Improving Operational Decision-Making in Hospital Pharmacies in the Presence of Disruptions

by

Lauren L. Czerniak

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Industrial and Operations Engineering)
in the University of Michigan
2023

Doctoral Committee:

Professor Mariel S. Lavieri, Co-Chair
Professor Mark S. Daskin, Co-Chair
Professor Eunshin Byon
Professor Jun Li
Professor Burgunda V. Sweet

Lauren L. Czerniak
czernL@umich.edu
ORCID iD: 0000-0003-4379-2004

ACKNOWLEDGMENTS

I am truly grateful for the love, support, and guidance throughout my PhD journey at the University of Michigan (UM). I would first like to thank my wonderful family. To my parents - David and Mary Beth, I am incredibly blessed to have you as my parents. From my Dad driving 9 hours round trip anytime I came home (solely because he enjoyed doing it) to my Mom and I's almost daily 45 minute phone calls. To my brother - Luke, thank you for bringing my stress level down a few notches and always making me laugh.

To my advisors - Mark Daskin and Mariel Lavieri, thank you for your guidance and support, both academically and personally. I have significantly grown as a researcher over the course of my PhD, and this is thanks to the two of you. It was an absolute honor to work with both of you, and it is clearly evident why the two of you are as successful as you are.

To my research collaborators - Gundy Sweet (Pharmacy), Jenn Erley-Leja (Pharmacy), Karl Renius (Pharmacy), Matthew Tupps (Pharmacy), Eunshin Byon (Industrial and Operations Engineering; IOE), Jun Li (Ross), Steve Broglio (Concussion), Gian Garcia (Concussion), Spencer Liebel (Concussion), Thomas McAllister (Concussion), Michael McCrea (Concussion), Paul Pasquina (Concussion), Chris Andrews (Glaucoma), Luke DeRoos (Glaucoma), Joshua Stein (Glaucoma), and Mark Van Oyen (Glaucoma), I am very appreciative of all of the help, support, and many answered questions. I would also like to express my appreciation for the funding that supported my research over the course of my PhD: National Science Foundation Graduate Research Fellowship (Grant DGE 1841052), Rackham Merit Fellowship, and the IOE Department at the UM (Graduate Student Instructor and Graduate Research Assistant).

To the IOE staff at the UM - thank you for your time, help, and kindness during my PhD. I would like to give a special thanks to Chris Konrad and Edwin Tang for always helping me troubleshoot my technology issues. I would also like to give a special thanks to Rod Capps, Akosua Dow, Matt Irelan, Valerie Martin, Mint Minto, Allyson Rodriguez, Tyra Sampson, Tina Picano Sroka, Jessalyn Tamez, and Cornelius Wright.

To my lab mates at the UM - Emily Tucker (Daskin), Gian Garcia (Lavieri), Wesley Marrero (Lavieri), Luke DeRoos (Lavieri), Daniel Otero (Lavieri), and Yili Wang (Lavieri), thank you for the mentorship and friendship that you have provided. I am grateful to have you as my academic brothers and sisters.

To my undergraduate researchers at the UM - Nicole Mor (Pharmacy), Halle Mohr (Pharmacy), Valeria Chavez (Pharmacy), Hannah Zhou (Concussion), Max Genthe (Concussion), and Yueyun Xia (Concussion), thank you for allowing me to learn as much from you as you did from me. I have no doubts that you will all be nothing but successful in your future careers.

To my mentees at the UM - Kevin Smith (2020-2021), Jacqueline Hannan (2021-2022), Tracy Chen (2022-2023), and Yili Wang (2023-2024), thank you for our friendship and allowing me to be a part of your PhD experience. I wish you all the best and I know that you will thrive in your future careers.

To my students at the UM - IOE 316 (Winter 2021) and IOE 202 (Fall 2022), thank you for always making my day brighter. Teaching is a passion of mine and it was a pleasure to have you as my students.

To my friends at the UM - I am blessed to have too many to name everyone. I would like to thank you all for the memories that I will treasure in my heart and thank you all for bringing light to my PhD experience. I would like to give a special thanks to my partner, Alex Valenze, who I met in Michigan. Thank you for all of your support, great times together, numerous ice cream trips, and always making me smile. I would also like to give a special thanks to Kelley Dugan, Kati Moug, Arlen Dean, Andrew Holen, and Emily Dumas. Thank you for the fun times filled with cards, board games, movies, walks, and many laughs. I would like to give a special thanks to my PhD cohort. I definitely chose the perfect year to start my PhD.

To my mentors at Pitt - Lisa Maillart, Natasa Vidic, Louis Luangkesorn, and Eugene Trofimov, thank you for the encouragement and support in my decision to pursue a PhD. I would like to give a special thanks to Lisa Maillart for being a great role model and allowing me to see myself in her.

To my friends at Pitt - Thank you for the wonderful undergraduate experience. I would like to thank Neil Patel and Joannie Schuman for our friendship that started during my undergraduate career, but continues to this day.

To my family - I would like to thank you all for your love and support. Thank you to my two “sisters": Christina Wallace and Alexa Szelong. I cannot thank you enough for the support and inspiration that you both provided me. I would like to sincerely thank my Pap Kutch who is watching my dissertation defense from heaven, and I would like to thank the Kaczmarek Family: Uncle Gene, Aunt Marcy, Uncle Andrew, Aunt Lori, Kinley, Cambree, and Kendall. Thank you to my “Michigan Family": Ryan, Brenna, Jack, and Ella Bartok. It was a blessing to have you so close for my first few years in Michigan and I truly appreciate you always being there if I needed anything. I would also like to thank the vacation crew: Axelsson Family, Harbaugh Family, Gorman Family, Reed Family, and Wallace Family. I am glad that attending the UM ruffled some feathers in the group since everyone seems to be Penn State fans. With this, all I have to say is “Go Blue!".

To close, I would like to thank God for the abundant blessings. In any challenge that I faced, he gave me the strength to overcome it; always coming out stronger on the other side. I am also a firm believer that:

If everything was easy, we would never see God work in our lives.

TABLE OF CONTENTS

LIST OF FIGURES

FIGURE

LIST OF TABLES

TABLE

LIST OF APPENDICES

ABSTRACT

The National Academies of Sciences, Engineering, and Medicine released a consensus study report in 2022 regarding concerns of the resiliency of medical supply chains. Narrowing the focus to pharmaceutical drugs, disruptions have plagued pharmaceutical supply chains long before 2022, but gaps in research prevail regarding how to improve inventory decision-making to help mitigate the negative effects of these disruptions. Supply chain disruptions impact the availability of the drug whereas demand disruptions impact the quantity of a drug needed. Pharmaceutical drugs are also a perishable product implying that hospital pharmacies cannot simply stock-up on a drug to be "safe". This dissertation makes contributions at the intersection of operations research and healthcare by creating a series of mathematical models to improve decision-making, provide insights, and challenge administrative policies in current practice for hospital pharmacy inventory systems.

This dissertation consists of four technical chapters where each chapter has its own model and research contributions. The first technical chapter consists of a simulation-optimization model that is used to solve for reorder point s and order-up-to level S periodic review inventory policies for a hospital pharmacy inventory system with supply chain disruptions and perishability. The model is quickly solved using a Binary Grid-Search algorithm.

From the first technical chapter, we recognize that not all hospital pharmacies have the resources to implement simulation-optimization models in practice. Therefore, we modify some of the assumptions made in the first technical chapter and provide a closed-form model in the second technical chapter. Specifically, the second technical chapter provides closed-form solutions for length of the review period R and order-up-to level S periodic review inventory policies in a hospital pharmacy inventory system with supply chain disruptions and perishability. The model provides the inventory policy quickly and is easy to implement.

Beyond supply chain disruptions and perishability, we recognize that hospital pharmacy inventory systems may face demand disruptions. The third technical chapter builds upon the second technical chapter by creating an adaptive inventory system that incorporates these demand disruptions. The research is designed to answer the following question: How does a drug's shortage-waste weighting along with the duration of and time between supply chain disruptions influence the benefits (or detriments) of adapting to demand disruptions? We also create a ranking procedure that provides a way of discerning which drugs are of most concern and illustrates which policies

to update given that a limited number of inventory policies can be updated.

We recognize that strict regulations in current practice generally prohibit hospital network pharmacies from sharing drugs outside of their network or make it very difficult for hospital network pharmacies to stay compliant when sharing drugs outside of their network. To this end, the fourth technical chapter builds upon the second technical chapter and includes a modeling framework to solve for continuous review order-up-to level inventory policies in a *two*-hospital network pharmacy inventory system with supply chain disruptions, perishability, and lateral transshipments (i.e., sharing of inventory; integrated inventory system).

The technical chapters of this dissertation capture the challenging characteristics of a hospital pharmacy inventory system like disruptions and perishability. These technical chapters also build upon one another which allows for more in-depth models and insights. Furthermore, patients depend on pharmaceutical drugs for treatment, and this dissertation is designed to help ensure that these patients have access to the treatment that they need.

CHAPTER 1

Introduction

Hospital pharmacy managers face a fundamental trade-off: drug waste versus drug shortages. Drug waste is costly; U.S. hospital pharmacies discard approximately $800 million worth of expired drugs a year (Allen, 2017). On the other hand, drug shortages are a public health crisis that increase costs for patients, increase medication errors, lead to inferior treatments, and lead to severe consequences such as death (Phuong et al., 2019). Drug shortages are often the result of fragile pharmaceutical supply chains (Tucker, 2020). Furthermore, there have been recent concerns regarding the resiliency of medical supply chains and with the current geo-political environment in which energy sources are being weaponized, it is very likely that drugs will follow soon if China or India elect to do so (i.e., drug shortages are likely to get worse than better) (NASEM, 2022).

To address both drug waste and drug shortages, three aspects need to be considered: (1) Perishability and supply chain disruptions make it difficult to select the appropriate levels of drug inventories to maintain. This leads to hospitals ending up with more drugs than needed [waste] or an insufficient number of drugs [shortage]. (2) Demand disruptions (i.e., declines or surges in demand) are prevalent in the inventory system, but few studies have investigated how inventory policies should adapt to decreases [often leads to waste] and increases [often leads to shortages] in demand. (3) Hospital networks do not work in isolation, but it is unclear if sharing drug inventory between hospital networks would help to mitigate drug waste and drug shortages.

This dissertation considers these aspects by (see Figure 1.1): (1) solving for periodic review inventory policies for a hospital pharmacy with perishability and supply chain disruptions, (2) creating an adaptive periodic review inventory system that endogenously detects and responds to significant shifts in demand to account for demand disruptions, while still accounting for perishability and supply chain disruptions, and (3) formulating and solving a multi-location model that quantifies the impact of an integrated inventory system across hospital networks with respect to drug waste and drug shortages. From model conception, design, and publication of this dissertation research, I, Lauren L. Czerniak, worked closely with a interdisciplinary team consisting of Mariel S. Lavieri and Mark S. Daskin (academic co-advisors at the University of Michigan [UM];

Industrial and Operations Engineering), Eunshin Byon (Industrial and Operations Engineering Associate Professor at the UM), Burgunda V. Sweet (Assistant Dean for Curriculum Assessment and Clinical Professor of Pharmacy at the UM), and Jennifer Erley-Leja, Matthew A. Tupps, and Karl Renius (licensed pharmacists and Central Pharmacy managers at the UM).

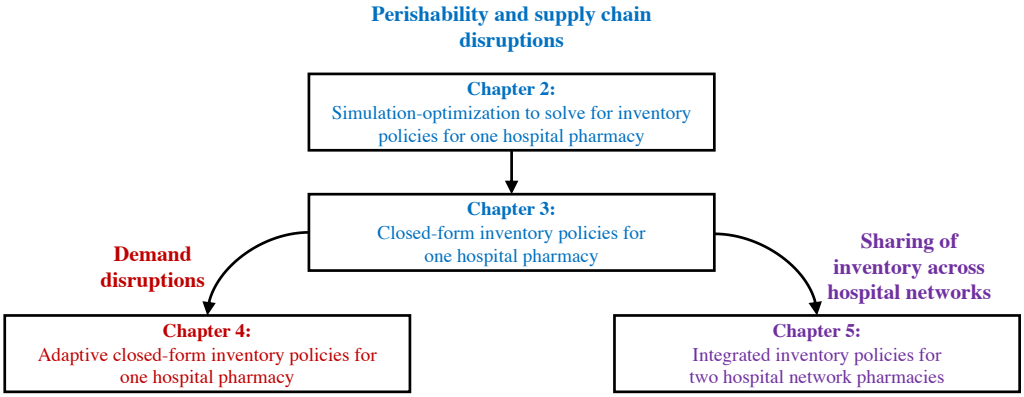

Figure 1.1: Depiction illustrating how this dissertation incorporates the three aspects that need consideration. Chapters 2-5 correspond to the four technical chapters of this dissertation.

We proceed to provide an overview of the four technical chapters of this dissertation (i.e., Chapters 2-5) where each chapter is a standalone paper. For each technical chapter, the key contributions are highlighted. At the end of this chapter, we provide a note on the application of the models from the four technical chapters and describe how this dissertation is organized.

1.1 Chapters 2-3: Inventory Policies with Perishability and Supply Chain Disruptions

To reduce drug waste and drug shortages, it is essential to consider perishability and supply chain disruptions when establishing how much inventory to have on-hand and when to place orders in a hospital pharmacy inventory system. However, hospital pharmacy managers are often responsible for thousands of different drugs (e.g., 2,500+ drugs at the UM). Hence, it is crucial to have models that provide quick computation times for inventory policies that account for perishability and supply chain disruptions. In Chapters 2-3 of this dissertation, we tackle this challenge by developing two models which quickly solve for inventory policies in a hospital pharmacy inventory system.

1.1.1 Chapter 2: Simulation-optimization Model

In Chapter 2, the reorder point s and order-up-to level S periodic review inventory policy are identified using simulation-optimization. The contributions of this research are as follows:

1. We develop a simulation-optimization model to determine near-optimal (s,S) periodic review inventory policies for a perishable system that faces supply chain disruptions, stochastic demand, and positive lead time. Further, there are no probability distribution restrictions for the uncertain parameters.

2. We implement a Binary Grid-Search algorithm which uses the structure of the objective function to solve for the (s,S) policies rather than taking an Exhaustive Grid-Search approach.

This research is published in the *Proceedings of the 2021 Winter Simulation Conference* and has the following co-authors: Lauren L. Czerniak, Mark S. Daskin, Mariel S. Lavieri, Burgunda V. Sweet, Jennifer Erley-Leja, and Matthew A. Tupps (Czerniak et al., 2021).

1.1.2 Chapter 3: Closed-form Model

The model presented in Chapter 2 provides quick computation times and offers great flexibility in terms of the system characteristics. Still, not all hospital pharmacies have the resources to implement simulation-optimization models in practice. Hence, in Chapter 3, we derive optimal length of the review period R and order-up-to level S periodic review inventory policies in closed-form by relaxing the stochastic demand and positive lead time assumption. The contributions of this research are as follows:

1. We derive exact closed-form solutions for a non-perishable and perishable lost-sales (R,S) periodic review inventory system with supply chain disruptions.

2. We provide insights regarding how model inputs and parameters influence the (R,S) inventory policies using the closed-form expressions and numerical sensitivity analysis.

3. Our models are applied to a hospital pharmacy inventory system to (i) illustrate the consequences of not implementing the proposed closed-form model in a perishable inventory system with supply chain disruptions, (ii) analyze the influence of supply chain disruption patterns, (iii) study the impact of stochastic demand, and (iv) investigate the trade-off between model simplicity and model accuracy.

This research is published in *INFOR: Information Systems and Operational Research* and has the following co-authors: Lauren L. Czerniak, Mark S. Daskin, Mariel S. Lavieri, Burgunda V. Sweet, Jennifer Erley-Leja, Matthew A. Tupps, and Karl Renius (Czerniak et al., 2023a).

1.2 Chapter 4: Adaptive Inventory Policies

In addition to perishability and supply chain disruptions, demand disruptions pose an additional challenge for hospital pharmacy managers when deciding how much inventory to have on-hand and when to place orders. Demand disruptions were prevalent during the Covid-19 pandemic (e.g., surges in demand to treat Covid-19 patients and declines in demand for other drugs due to the cancellation of elective surgeries). Failing to adapt to disruptions that increase or decrease demand often leads to drug shortages and drug waste, respectively. To this end, Chapter 4 answers the following question: how does a drug's shortage-waste weighting and supply chain disruption profile (i.e., duration of and time between supply chain disruptions) influence the benefits (or detriments) of adapting to demand disruptions? The contributions of this research are as follows:

1. We create an adaptive (R,S) periodic review inventory system that accounts for perishability, supply chain disruptions, and demand disruptions. All expressions are presented in closed-form providing quick solution times and easy implementation which is critical in a hospital pharmacy where managers are often responsible for thousands of different drugs.

2. The adaptive inventory system relies on the expected proportion of drugs wasted per day in a (R,S) periodic review perishable inventory system with supply chain disruptions and stochastic demand. To the best of our knowledge, we are the first to present this value in closed-form.

3. We use the adaptive inventory system to create a ranking procedure. The ranking procedure provides a way of discerning which drugs are of most concern and illustrates which policies to update given that a limited number of inventory policies can be updated.

4. We leverage simulation modeling and perform an extensive numerical analysis using real-world demand data from the University of Michigan's Central Pharmacy to distinguish how a drug's shortage-waste weighting and supply chain disruption profile influence the benefits (or detriments) of adapting to demand disruptions.

This paper has been submitted for publication and has the following co-authors: Lauren L. Czerniak, Mariel S. Lavieri, Mark S. Daskin, Eunshin Byon, Karl Renius, Burgunda V. Sweet, Jennifer Erley-Leja, and Matthew A. Tupps (Czerniak et al., 2023b).

1.3 Chapter 5: Integrated Inventory Policies

Strict regulations in current practice generally prohibit hospital network pharmacies (HNP; i.e., a central pharmacy that provides drugs for all hospitals in the healthcare network) from sharing drugs

outside of their network or make it very difficult for HNPs to stay compliant when sharing drugs outside of their network. However, it is unclear if sharing drug inventory between hospital networks (i.e., lateral transshipments [LT]; shipments between hospital networks) would help mitigate drug waste and drug shortages. In Chapter 5, we create a new modeling framework that solves for continuous review order-up-to level inventory policies in a two-HNP integrated inventory system with supply chain disruptions and perishability. The contributions of this research are as follows:

1. We create a modeling framework to solve for the continuous review inventory policies that minimize the expected cost per day in a two-HNP non-perishable integrated inventory system with independent suppliers experiencing supply chain disruptions. To incorporate perishability, we introduce a constraint on the probability drugs are wasted at each HNP. This modeling framework can also be applied to many inventory systems of interest (e.g., blood banks).

2. Working closely with our hospital pharmacy collaborators at the University of Michigan, we numerically study a two-HNP perishable integrated inventory system to (i) demonstrate when it is beneficial to allow two HNPs with independent suppliers experiencing supply chain disruptions to share inventory and (ii) provide insights on how the supply chain disruption characteristics of the two HNPs influence whether it is beneficial to share inventory.

This paper is published in *IISE Transactions on Healthcare Systems Engineering* and has the following co-authors: Lauren L. Czerniak, Mariel S. Lavieri, Mark S. Daskin, Burgunda V. Sweet, Jennifer Erley-Leja, Matthew A. Tupps, and Karl Renius (Czerniak et al., 2023c).

1.4 Application of the Models

The four technical chapters of this dissertation focus on hospital pharmacy inventory systems. We note that the application of this research may not be directly applicable to other pharmacy inventory systems (e.g., retail pharmacy inventory systems). For example, hospital pharmacies often stock thousands of different drugs (e.g., 2,500+ drugs at the UM) for a wide-range of treatments (e.g., surgical procedures, cancer treatment, daily care). However, retail pharmacies often only stock a small subset of the drugs that hospital pharmacies stock making quick computation times less of a concern. Furthermore, we discuss how strict regulations in current practice generally prohibit HNPs from sharing drugs outside of their network or make it very difficult for HNPs to stay compliant when sharing drugs outside of their network. In a retail pharmacy setting, sharing regulations may be less of a concern. As an example, consider a patient whose primary retail pharmacy (e.g., Walgreens in their neighborhood) does not have the medication that the patient needs

on-hand. The patient may either drive to a Walgreens outside of their neighborhood to get the medication that they need or have the prescription transferred to a different community pharmacy that has the medication in stock.

1.5 Organization of Dissertation

We proceed to describe how the remainder of this dissertation is organized. Chapters 2-5 consist of the four technical chapters. Each of these chapters have its own introduction and motivation, literature review, model, numerical analysis, and concluding remarks. Chapter 6 provides concluding thoughts and future research directions.

CHAPTER 2

Improving Simulation Optimization Run Time When Solving for Periodic Review Inventory Policies in a Pharmacy

2.1 Introduction

It is estimated the United States spends about $329 billion annually on drugs (Hartman et al., 2018). However, demand and supply uncertainties in a pharmacy inventory system make decision-making a challenging task. Having insufficient inventory can lead to shortages which increase cost for care, increase medication errors, and lead to inferior treatments (Phuong et al., 2019). In contrast, drugs are a perishable product, so holding too much inventory can lead to waste. Accounting for these different sources of uncertainty, perishability, and system characteristics (like positive lead time) make it computationally difficult to determine how much inventory to keep on hand and when orders should be placed.

We develop a simulation-optimization model to determine near-optimal (s,S) periodic review inventory policies where s denotes the reorder point and S denotes the order-up-to level. The objective is to minimize the expected cost per day and the model represents a drug at one hospital pharmacy. Simulation-optimization is often used for optimization problems where the objective function cannot be computed exactly, usually due to stochastic parameters in the problem of interest. Instead, replications of a simulation model with realized random variables are used to estimate the objective function (Fu, 2015). In the pharmacy system of interest, an inventory manager can be responsible for over 2,500 drugs. Hence, quickly solving the simulation-optimization model is crucial. To do this, we implement a Binary Grid-Search (GS) algorithm to solve for the near-optimal periodic review policies and compare its run time to an Exhaustive GS and a Convergent Optimization via Most-Promising-Area Stochastic Search (COMPASS). We refer to these solutions as near-optimal because we discretize the solution space. Also, throughout this chapter, we refer to this discretized solution space as a square grid where the rows correspond to the values of

s and the columns correspond to the values of S.

To illustrate our solution framework, we apply our methodology to solve for the (s,S) parameters in a periodic review inventory system for a drug test case. While this chapter looks at the problem from a pharmacy perspective, the models presented can be applied to other first-in-first-out perishable inventory systems such as chemicals and food warehouses. Also, the model places no restrictions on the probability distributions for the uncertain parameters, which provides greater flexibility for products that may be of interest. Below, we present related research from inventory management to motivate how the simulation-optimization model expands upon current literature. We also explicitly highlight the contributions of this chapter.

Past literature includes a variety of methodologies to solve for periodic review inventory policies: stochastic programming, reinforcement learning, simulation, closed-form solutions. With respect to (s,S) periodic review policies, it is important to recognize the valuable research that studies (s,S) policies from a gradient estimation perspective (e.g., Bashyam & Fu, 1991; Fu & Hu, 1997) and early research that does not consider a perishable and supply disrupted system like the one presented in this research (e.g., Sahin, 1982; Zheng & Federgruen, 1991). Recent perishable inventory research uses models such as stochastic programming and reinforcement learning to determine optimal periodic review policies in the presence of stochastic demand (Dillon et al., 2017; Rajendran & Ravindran, 2019, Rajendran & Srinivas, 2020; Kara & Dogan, 2018). Further, Nguyen & Chen (2019) use stochastic programming to model stochastic demand and stochastic yield supply. When looking at the supply disruption side, Atan & Rousseau (2016) consider supply disruptions in a perishable inventory system, but they do not account for uncertain demand. For supply disruptions in non-perishable inventory systems, some researchers include stochastic demand with supply disruptions (Schmitt et al., 2010) and some only account for supply disruptions (Skouri et al., 2014; Konstantaras et al., 2019), but again, these models do not account for perishability.

For research that solves for optimal inventory policies in a pharmacy system, Zhang et al. (2014) use a deterministic demand and supply model whereas Saedi et al. (2016) account for uncertain demand and supply disruptions, but they assume there is no lead time. Other researchers look into the lead time question like Li et al. (2018) who consider stochastic lead time, but no uncertainty in terms of supply and demand, while Franco & Alfonso-Lizarazo (2020) include stochastic demand and lead time, but they do not consider supply disruptions. In summary, there is a gap in periodic review policies for perishable products that face stochastic demand, supply disruptions, and positive lead time.

The two main contributions of this chapter are:

1. We develop a simulation-optimization model to determine near-optimal (s,S) periodic review inventory policies for a perishable system that faces stochastic demand, supply disrup-

tions, and positive lead time. Further, there are no probability distribution restrictions for the uncertain parameters.

2. We implement a Binary GS algorithm which uses the structure of the objective function to solve for the (s,S) policies rather than taking an Exhaustive GS approach.

The remainder of the chapter is organized as follows: Section 2.2 describes a (s,S) inventory system and how to model this system using simulation. Section 2.3 presents three methods for solving the simulation-optimization model: Exhaustive GS, COMPASS, and Binary GS. Section 2.4 provides numerical results using all three methods. Finally, Section 2.5 closes the chapter with concluding thoughts and future research.

2.2 (s,S) Periodic Review Inventory System

2.2.1 (s,S) Inventory Process

There exists a variety of inventory policies that any given system can follow. Consistent with the system of interest, this research focuses on (s,S) periodic review inventory policies. We define inventory position as the amount of inventory on hand plus the amount of inventory en route. When the inventory position falls below s, an order is placed for S minus the inventory position. Here, the review period is 1 day. Figure 2.1 illustrates (s,S) inventory policies in a perishable lost-sales inventory system where H is the total inventory on hand, H' is the total inventory on hand that expires at the end of day t, P is the inventory position, d is the demand, s is the reorder point, and S is the order-up-to level.

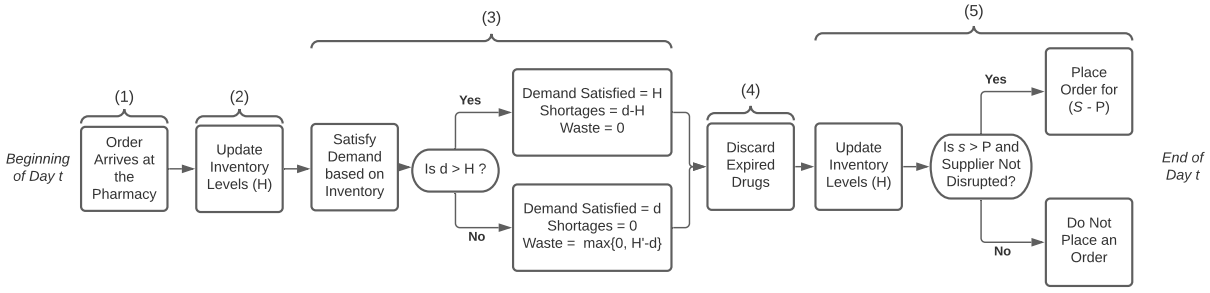

Figure 2.1: Process flow of (s,S) perishable inventory system ((#) correspond to model in Section 2.2.2.3).

2.2.2 Simulation Model

To develop a simulation-optimization model, the first step is to create a simulation model that represents a (s,S) periodic review inventory system and outputs the objective of the model given (s,S) as inputs. The objective of the model is to minimize the expected cost per day where cost is measured relative to purchasing cost. The costs of interest are shortage, waste, holding, and ordering.

2.2.2.1 Assumptions

The simulation model represents a drug at one hospital pharmacy and drug substitutes are not considered. Our modeling framework makes this no substitute assumption because in practice, there is often a preferred choice and even when substitutes are available, these substitutes often come with additional logistics like ensuring the patient does not have a certain comorbidity. Drugs typically have an expiration date that is with respect to the end of the month (e.g., July 31). Therefore, our model assumes all drugs that arrive in the same month come from the same production batch and have the same end of the month expiration date. Hence, drugs are only discarded at the end of the month. We assume the inventory on-hand and estimated days until inventory expires is known each day t. Also, the lead time is deterministic and non-negative ($l \in \mathbb{N}_0$) where $l = 0$ implies an order placed at the end of day $(t-1)$ arrives at the beginning of day t. Demand is stochastic and supply uncertainty is due to disruptions (i.e., no supply yield uncertainty). We assume these two sources of uncertainty are independent from one another. Last, we assume all demand not met is lost, first-in-first-out protocols are always followed, and there is no seasonal demand.

2.2.2.2 Input Parameters

For input parameters, we consider a planning horizon of length $|\mathbf{T}|$ days $t = 1,2,...,|\mathbf{T}|$. We denote l as the deterministic lead time ($l \in \mathbb{N}_0$), e as the shelf-life of the drug (months), $|\mathbf{R}|$ as the number of simulation replications (r denotes a particular replication), and b, z, h, and o as the shortage, waste, holding, and ordering cost relative to purchasing cost, respectively. For the discretized square grid, we denote $s_l = S_l$ as the minimum value for s and S, $s_u = S_u$ as the maximum value for s and S, and Δ as the grid increment for s and S. For example, $s_l = 25$, $s_u = 100$, $\Delta = 25 \implies s \in [25,50,75,100]$. For the uncertain parameters in the model, we define d_t as the demand on day t and y_t as a binary variable for the supply disruption status on day t. For the binary variable, $y_t = 0$ denotes supply is disrupted and $y_t = 1$ denotes supply is not disrupted. There are no restrictions on the probability distributions of these sources of uncertainty.

2.2.2.3 Simulation Model Description

Given these assumptions and input parameters, we describe the simulation model at a high level. Pseudocode for the simulation model can be found in Appendix A.1 and the steps in the simulation model are identified on Figure 2.1. Also, it is important to note the decision variables (s,S) are inputs to the simulation model and we discuss how to optimize these decision variables in Section 2.3.

The simulation model begins with the function StaticSim where for each replication r, the function generates a daily demand pattern and supply disruption pattern using the associated input parameters and probability distributions. Then, the InventoryProcess function is called to calculate the number of shortages, drugs wasted, orders placed, and drugs held during the planning horizon for this replication r. The function InventoryProcess calculates these values by taking the input parameters, demand pattern for replication r, supply pattern for replication r, s, and S and simulating how the system would operate with these parameters. For each day in the planning horizon, (1) orders arrive to the pharmacy at the beginning of the day, (2) inventory levels are updated, (3) inventory levels are updated based on the demand observed on day t, (4) if day t corresponds to the end of the month, expired drugs are discarded, and (5) an order is placed according to the (s,S) inputs if supply is not disrupted. This process continues until the end of the planning horizon is achieved. At this point, the objective value is calculated and InventoryProcess returns this output to the StaticSim function. All replications r follow the same process (computed in parallel to reduce computation time) and at the end, StaticSim calculates the average objective value over all replications r.

2.3 Simulation-optimization

When using the simulation model described in Section 2.2.2 to solve for the near-optimal (s,S) periodic review policies, we consider three approaches: Exhaustive GS, COMPASS, and Binary GS. Also, we use common random numbers regardless of the search procedure.

2.3.1 Near-Optimal Policies Using Exhaustive Grid-Search

Exhaustive GS takes a discretized grid for the parameters s and S and enumerates the objective value for all feasible solutions. In this inventory policy setting, a solution is only feasible if $s \leq S$. Since one is enumerating all solutions, the computation time can become excessive, especially for fine grid spaces.

2.3.2 Near-Optimal Policies Using COMPASS

COMPASS is a discrete simulation-optimization algorithm where at each iteration, a most promising area is defined which depends on the distance between a feasible solution and the best solution found thus far (Hong & Nelson, 2006). This most promising area is adaptive and to keep consistency with the other search procedures, we allocate the same number of simulation replications to all enumerated solutions. Also, the algorithm requires a uniform search procedure over the discretized grid, so we implement the RMD algorithm (Hong & Nelson, 2006).

2.3.3 Near-Optimal Policies Using Binary Grid-Search

Binary GS takes a discretized grid for the parameters s and S, but instead of enumerating all solutions, binary searches are completed successively on the rows and columns until no improvement to the objective value is made. For initialization, the algorithm enumerates all diagonal solutions on the discretized grid.

2.3.3.1 Structural Properties of the Objective Function

Our empirical observations support that we cannot assume the objective function is always monotonically decreasing and then monotonically increasing with respect to s for fixed S and S for fixed s (i.e., the negation of the objective function is unimodal). Consider the drug test case (described in Section 2.4) with 5,000 replications and solved using an Exhaustive GS (Section 2.3.1). Figure 2.2 displays (a) Expected Cost Per Day vs Reorder Point s and (b) Expected Cost Per Day vs Order-up-to Level S for five different values of S and s, respectively. For (a), one may consider why the objective function is flat up to some reorder point s. This occurs in cases where the reorder point, s, is too small and the order-up-to level, S, is too large resulting in there (almost surely) being a significant number of drugs that expire before there is sufficient demand to use these drugs (i.e., inventory position drops to 0). Therefore, an order is placed to bring the inventory level up to S regardless of the value of s and it is important to consider that all of these drugs have the same expiration date. Figure 2.2 illustrates the objective function is not always monotonically decreasing and then monotonically increasing, but that it does exhibit characteristics of this structure.

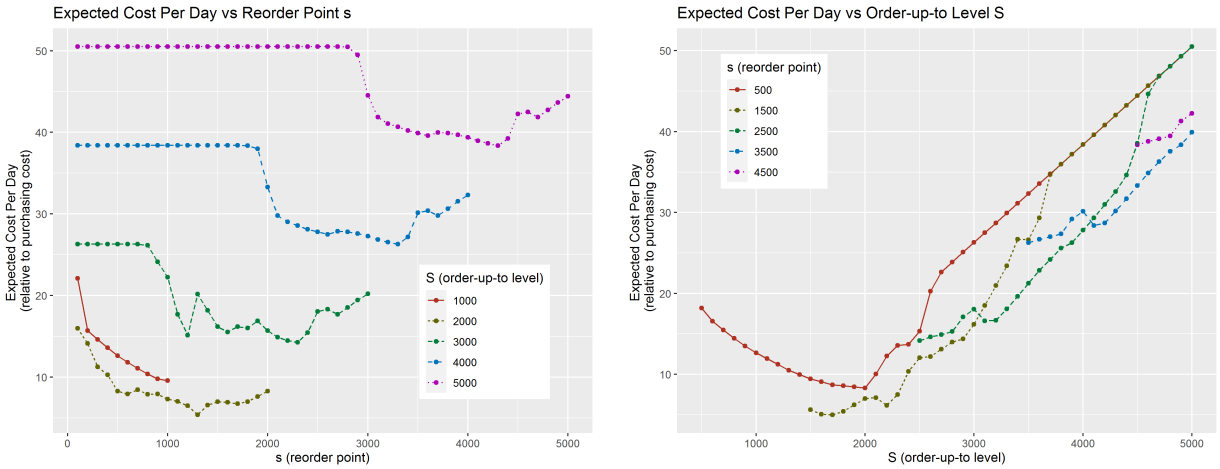

(a) Expected cost per day vs reorder point s.

(b) Expected cost per day vs order-up-to level S.

Figure 2.2: Objective function (expected cost per day) with respect to s for fixed S and S for fixed s.

We also consider if the monotonically decreasing and then increasing property does not hold for this test case due to simulation error. Regardless of the number of simulation replications completed, there will be some error with respect to the objective value since this output is a result of a stochastic process. When considering solutions that are significantly different from one another at a confidence level of 95%, 5.7% of the pairwise comparisons with respect to the columns of the grid and 3.9% of the pairwise comparisons with respect to the rows of the grid fail to meet the monotonically decreasing then increasing structure. These results demonstrate the proportion of time the monotone property does not hold is relatively small.

These findings motivate a regular binary search coupled with a forced binary search to ensure one is not caught in a local optimum. Forced binary search means we complete a binary search above, to the right, below, and to the left of the current solution regardless of the objective value of the neighboring solutions. The algorithm can be seen as a metaheuristic where we exploit when neighboring solutions improve the objective value. However, before terminating, we also accept worse neighbors with the aim of avoiding local optimality. We simply refer to the algorithm as Binary GS and details follow in the subsequent section.

2.3.3.2 Description of Binary Grid-Search

Throughout the explanation of the Binary GS algorithm, we will refer to the discretized grid described earlier (square grid matrix where the rows correspond to s values in increasing magnitude and the columns correspond to S values in increasing magnitude). A binary search, also known as a half-interval search, is when one successively divides a unimodal search interval in half to find

the global maximum (i.e., negation finds global min). For example, consider the objective value, V, for a fixed s across nine increasing S values: $V = (100, 99, 75, 72, 84, 88, 94, 100, 106)$. One first starts at the midpoint of the full interval (i.e., 84). Then, one considers the neighboring solutions (i.e., 72 and 88). Since $72 < 84$, one considers the midpoint (i.e., 75) of the halved interval $(100, 99, 75, 72, 84)$ and the same procedure continues.

To initialize the algorithm, (a) we enumerate all feasible solutions on the diagonal of the discretized grid with the number of replications desired by the user. From these initial enumerations, (b) we select the column that minimizes the objective function and perform a binary search on this column. (c) The feasible solution minimizing the objective function is chosen as the starting point.

Then, we begin an iterative process where (1) we enumerate the objective value to the left and right of our current solution. If both of these objective values are greater than the current solution, no change is made to the current solution. Otherwise, we carry out a binary search on this row of the grid and update our current solution with the new minimum found. (2) We enumerate the objective value above and below our current solution. If both of these objective values are greater than the current solution, no change is made to the current solution. Otherwise, we carry out a binary search on this column of the grid and update our current solution with the new minimum found. (3) We check if the solution found in the previous iteration is the same solution as this iteration. If not, the iterative process starts back at (1) with the current solution. If yes, then (4) a binary search is completed on the interval above, to the right, below, and to the left of the current solution to avoid local optimality (Section 2.3.3.1). As an example, the above interval means we consider the $S_{current}$ column and search the interval $s \in [s_l, s_{current}]$. These four interval searches are completed regardless of the value of the neighboring solutions of the current solution (i.e., the neighboring solutions may be worse). The key idea is that if the current solution is truly optimal, then performing a binary search on these four intervals should find a solution no better than the current solution. If the current solution remains optimal, the algorithm terminates. If not, then the iterative procedure starts back at (1) with the best solution found thus far as the current solution. Figure 2.3 below demonstrates the process where the light shaded circle with (*) denotes the solution and the fully shaded black circles denote the search area for that step. Pseudocode for the algorithm can be found in the Appendix A.2.

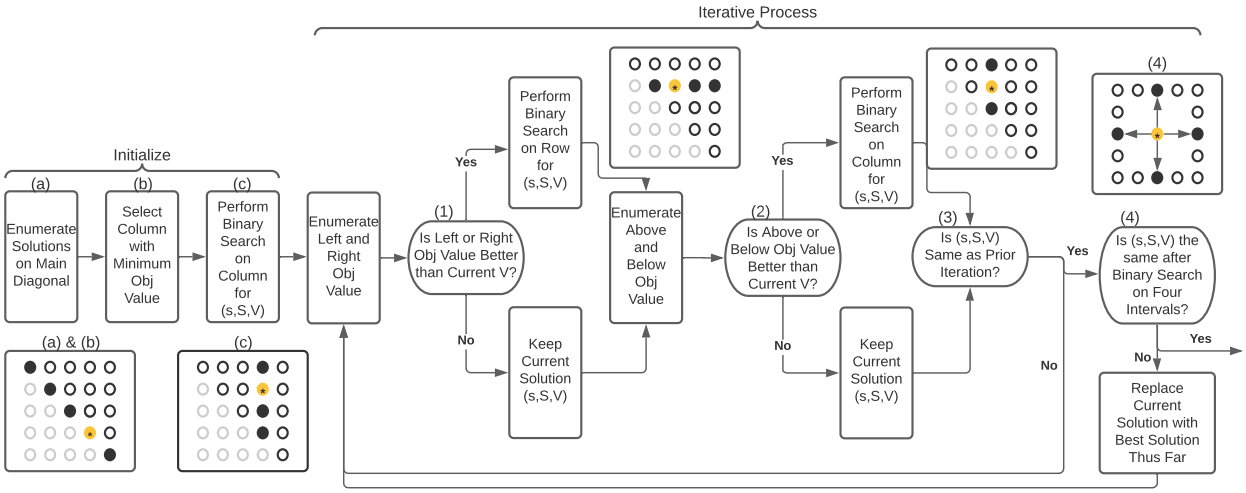

Figure 2.3: Process flow of binary grid-search (s: reorder point, S: order-up-to level, V: objective value).

2.4 Experimental Analysis

To evaluate the performance of the simulation-optimization model solved with an Exhaustive GS, COMPASS, and Binary GS, we present eight different cases where we vary the number of replications ("computational cost") for each method. We vary the replications from 100-10,000. We provide this range because $|\mathbf{R}| = 5,000$ replications ensures a confidence interval width, w, of at most 1. Specifically, one can determine $|\mathbf{R}|$ by ensuring $w > (2cd_{max})/(\sqrt{|\mathbf{R}|})$ where c is a user-specified confidence level (e.g., $c = 1.96$ for 95%), d_{max} is the maximum standard deviation from the simulation model when analyzed on a discretized grid, and w is the desired confidence interval width. For each of these eight cases, we present the near-optimal objective value, near-optimal s, near-optimal S, and run time when using (a) Exhaustive GS, (b) COMPASS, and (c) Binary GS. Objective values are rounded to the nearest tenths place. For COMPASS, we consider $m = \frac{1}{10}(no.\ of\ feasible\ solutions)$ solutions to randomly and uniformly search with replacement each iteration and consider a warm-up of $W = 10$ iterations for the uniform search. We terminate this algorithm when all neighboring solutions have been enumerated and the same solution has been found for 10 iterations in a row. All numerical experiments are completed in R (Version 4.0.3) on a Windows desktop (4.28GHz).

2.4.1 Pharmacy Model Input Parameters and Numerical Results

We use the input parameters across all eight cases where the cost (relative to purchasing cost) associated with waste, shortages, holding, and ordering are relevant to the pharmacy system of interest. We define $|\mathbf{T}| = 360$ Days (30 Days/Month), $l = 6$ Days, $e = 3$ Months, $b = 5$ (Haijema, 2013), $z = 1$ (Haijema, 2013), $h = 0.001$ (Jia & Zhao, 2017), $o = 0.5$ (Haijema, 2013), $s_l = S_l = 100$, $s_u = S_u = 5000$, $\Delta = 100$, Demand $\sim$ Poisson($\lambda = 25$/day), Days to Supply Disruption [Recovery] $\sim$ Geometric($p_1 = 0.01$) [Geometric($p_2 = \frac{1}{30}$)].

Figure 2.4 below displays (a) Expected Cost Per Day vs Number of Simulation Replications and (b) Run Time (hours) vs Number of Simulation Replications. For Figure 2.4 (a), we present the objective value with replications independent of the replications used to optimize the (s,S) policies to avoid optimization bias. Notice that in all cases, the Binary GS algorithm finds the exact same solution as the Exhaustive GS algorithm or a solution that is better in objective value when accounting for optimization bias. This also holds for the COMPASS algorithm. Furthermore, the Binary GS stays within an hour run time (maximum run time is approximately 30 minutes) whereas the Exhaustive GS increases to a computation time of approximately 10.5 hours at 10,000 replications. Another interesting aspect is the Binary GS algorithm always outperforms the COMPASS algorithm with respect to run time.

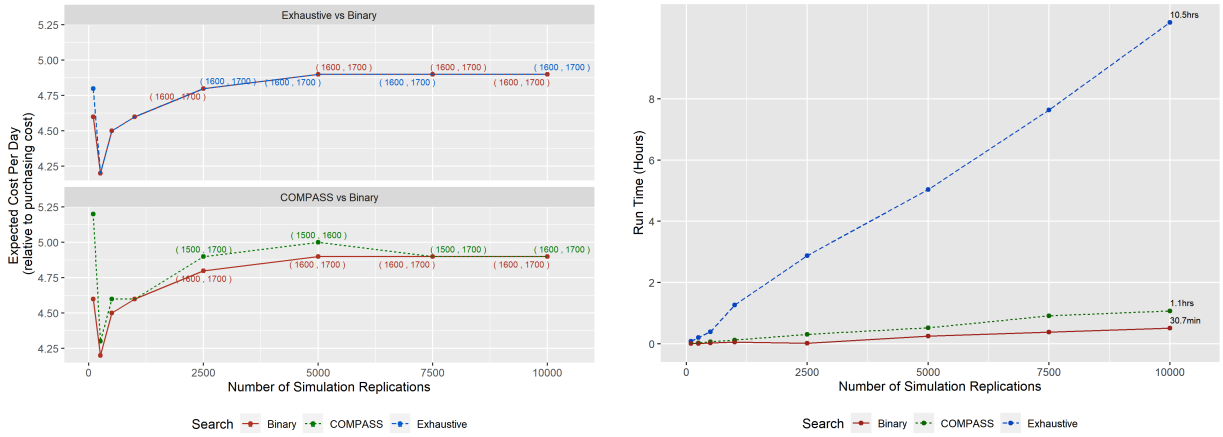

(a) Expected cost per day vs number of simulation replications (no. in () represent the (s,S) policies).

(b) Simulation run time vs number of simulation replications (run time indicated for 10,000 reps).

Figure 2.4: Numerical results with pharmacy model input parameters.

The Exhaustive GS guarantees the best solution is found (before considering optimization bias). However, it is not only computationally inefficient, but also unrealistic for the pharmacy system of interest. In particular, when considering the demand for all drugs at the University of Michigan's Central Pharmacy, the 20% highest demand drugs (545 drugs) make up about 85% of the total

demand. Hence, if one is interested in updating policies every 3 months (2,160 hours) using 10,000 replications, one would need $(545) * (10.5hrs) = 5,722.5hrs$, or about 7.95 months, when using the Exhaustive GS. Hence, even if one desktop computer was solely utilized to run the exhaustive algorithm for the entire 3 months, one could not solve for all policies within this time frame. However, if using the Binary GS, one could solve for the policies in $(545) * (0.5hrs) = 272.5hrs \approx$ 11.35 days.

2.5 Conclusion and Future Research

We create a simulation-optimization (s,S) periodic review inventory model that accounts for stochastic demand, supply disruptions, perishability, and positive lead time. We solve the simulation-optimization model using a Binary GS algorithm which uses the structural properties of the objective function. Based on our experimental results, we notice this algorithm improves computation time without sacrificing accuracy of the solution.

As for limitations, we assume a first-in-first-out system. If this is violated, then one can expect more waste and more shortages. However, through observations at the Central Pharmacy at the University of Michigan, first-in-first-out is a valid assumption with the protocols in place. We also assume demand and supply are independent of one another. When constructing the demand and supply patterns for a replication, r, one can consider the dependency of these uncertainties and we leave that to future work. Other assumptions such as lost-sales, deterministic lead time, and expired drugs being discarded at the end of the month can easily be adjusted in the model to tailor it to the characteristics of the system of interest.

In this research, we do not theoretically exploit the structure of the objective function and we do not provide theoretical guarantees on the Binary GS algorithm. Also, we find the Binary GS algorithm does not sacrifice the accuracy of the solution when defining the discretized grid with respect to s and S, but an alternative approach is defining the discretized grid with respect to s and $\delta = S - s$. Finally, we compare the Binary GS algorithm to an Exhaustive GS and the COMPASS algorithm, but other methods one can consider for comparison are stochastic gradient descent with random restart, random search, and Bayesian optimization. We leave these ideas for future research.

This research provides pharmacy inventory managers with a valuable tool that can aid decision-making in the complex pharmacy system.

Acknowledgments for Chapter 2

The Central Pharmacy at the University of Michigan provided information regarding protocols in place (e.g., first-in-first-out) and the number of drugs the pharmacy is responsible for. We would like to thank Bruce Chaffee, Seema Jetli, and Karl Renius for providing access to the data. We would also like to thank Nicole Mor for the helpful literature searches.

CHAPTER 3

Closed-form (R,S) Inventory Policies for Perishable Inventory Systems with Supply Chain Disruptions

3.1 Introduction

Hospital pharmacy managers make inventory decisions like how much to order and the frequency of orders for thousands of different drugs (e.g., 2,500+ drugs at the University of Michigan's Central Pharmacy). Commonly, these managers do not have the resources to implement advanced mathematical models (e.g., stochastic programs, simulation/simulation-optimization). Closed-form solutions for inventory policies in this setting are attractive because they provide the inventory policy quickly and are easy to implement (Borwein & Crandall, 2013), but existing closed-form solutions do not account for perishability and supply chain disruptions. Supply chain disruptions are defined as 'random events that cause a supplier or other element of the supply chain to stop functioning, either completely or partially, for a (typically random) amount of time' (Snyder et al., 2016). There have been recent concerns regarding the resiliency of medical supply chains (NASEM, 2022). It is critical to account for supply chain disruptions because having insufficient inventory can lead to drug shortages which increase the cost for care and lead to inferior/cancelled treatments (Phuong et al., 2019). Not accounting for perishability can lead to drug waste (HDMA, 2009; Allen, 2017).

We consider a (R,S) periodic review inventory system with perishability, supply chain disruptions, and deterministic demand. Here, R denotes the length of the review period and S denotes the order-up-to level (i.e., attempt to place an order up to S every R days). In a hospital pharmacy setting, R is typically small (e.g., $R \leq 7$ days) and it takes very little time to receive the drug when the supply chain is not disrupted (i.e., negligible/small lead times). However, when the supply chain is disrupted, the disruption can last weeks, months, or years (median duration of 14 months; Tucker et al., 2020a). From the perspective of a hospital pharmacy inventory system, we derive closed-form solutions that minimize the expected holding, ordering, and waste costs while ensuring a certain proportion of demand is satisfied for a (a) lost-sales (R,S) periodic review inventory

system with (b) supply chain disruptions and (c) perishability. For a lost-sales (R,S) periodic review inventory system (i.e., (a)), it is well known that higher fixed ordering costs encourage the system to place orders less frequently (i.e., have a larger R) and that higher holding costs encourage the system to stock less inventory on-hand (i.e., have a smaller S).

When incorporating supply chain disruptions into a lost-sales (R,S) periodic review inventory system (i.e., (a)-(b)), we expect to stock more inventory on-hand (i.e., have a higher S) and attempt to place orders more frequently (i.e., have a smaller R) due to the unreliability of the supply chain. We will attempt to place orders more frequently because (1) a greater number of order attempts implies a greater chance of placing an order when supply is not disrupted and (2) for any fixed S, a smaller R implies a larger expected inventory on-hand. We show that R and S depend on both the supply disruption parameters (e.g., disruption probability and recovery probability) and the type of the supply disruption process.

When adding perishability into a lost-sales (R,S) periodic review inventory system with supply chain disruptions (i.e., (a)-(c)), we expect that if the expiration lifetime of the drug is long enough, we will not waste any drug product if we simply follow a non-perishable (R,S) model with supply chain disruptions (i.e., (a)-(b)). However, how long is long enough? This ultimately depends on the order-up-to level S and deterministic daily demand q. We show that if the expiration lifetime is not long enough, following a non-perishable (R,S) inventory policy leads to unnecessary waste. Even worse, the expected proportion of demand satisfied is less than the demand satisfied when using a perishable (R,S) model.

The contributions of this chapter are as follows:

1. We derive exact closed-form solutions for a non-perishable and perishable lost-sales (R,S) periodic review inventory system with supply chain disruptions (two-state supply process; defined later).

2. We provide insights regarding how model inputs and parameters influence the (R,S) inventory policies using the closed-form expressions and numerical sensitivity analysis.

3. Our models are applied to a hospital pharmacy inventory system to (i) illustrate the consequences of not implementing the proposed closed-form model in a perishable inventory system with supply chain disruptions, (ii) analyze the influence of supply chain disruption patterns, (iii) study the impact of stochastic demand, and (iv) investigate the trade-off between model simplicity and model accuracy.

The remainder of this chapter is organized as follows: Section 3.2 presents literature relevant to this research. Section 3.3 provides the model notation and assumptions. Section 3.4 provides closed-form solutions for a non-perishable lost-sales (R,S) periodic review inventory system with

supply chain disruptions. In Section 3.5, we extend the model to incorporate perishability. Section 3.6 provides a numerical analysis of a hospital pharmacy inventory system to draw managerial insights and address (i)-(iv). Section 3.7 provides concluding thoughts and future research directions.

3.2 Literature Review

Literature reviews on inventory models with supply chain disruptions are presented in Paul et al. (2016), Snyder et al. (2016), Shen & Li (2017), and Ivanov et al. (2017). While this setting might seem similar to stochastic lead time models, we start by describing some of the literature and discuss key differences between a stochastic lead time model and supply chain disruption model. We proceed to define the two common ways supply chain disruptions are modeled in the literature and discuss the early research that considers supply chain disruptions in an inventory system. Then, we present the research most relevant to this chapter: non-perishable closed-form solutions and perishable closed-form solutions. We consistently use R as the length of the review period, S as the order-up-to level, s as the reorder point, and Q as the order quantity.

Supply Chain Disruptions versus Stochastic Lead Time

Our supply chain disruption setting is distinct from a stochastic lead time setting (e.g., Janakiraman & Roundy, 2004; Bischak et al., 2014; Tai et al., 2021; Duc et al., 2022 [(R,S) examples]). In our setting, during the supply chain disruption, no orders are successfully placed and no orders are en route. When the supply chain is properly functioning again on a review period day, an order is placed to raise the inventory position (i.e., inventory on-hand plus inventory en route) to the order-up-to level S. All of the drugs in this order arrive at the same time. In contrast, stochastic lead time implies that every R days, an order is always successfully placed to raise the inventory position to the order-up-to level S. The order is immediately en route, but there is some variability in when this order will arrive to the pharmacy (e.g., 5 days to receive the drug instead of 3 days). For every successful order with the supply chain disruption model, there could be multiple orders en route with the stochastic lead time model. We proceed to define the two common ways supply chain disruptions are modeled.

Two-State versus Bernoulli Supply Process

In the literature, researchers often model supply chain disruptions in two ways: two-state supply process and Bernoulli supply process. A two-state supply process is when the system alternates between a not disrupted state and disrupted state. Most researchers assume the length of time in each state is exponentially (or geometrically if discrete) distributed and thus form a Markov

chain (Snyder et al., 2016). Specifically, when time is discrete, researchers define α as the disruption probability (i.e., $\alpha = \Pr(\text{disrupted}|\text{not disrupted})$) and β as the recovery probability (i.e., $\beta = \Pr(\text{not disrupted}|\text{disrupted})$). Defining these parameters with respect to the length of the review period R implies that the probability the system has been disrupted for exactly j review periods is $\pi_j = \frac{\beta}{\alpha+\beta}$ when $j = 0$ and $\pi_j = \frac{\alpha\beta}{\alpha+\beta}(1-\beta)^{j-1}$ when $j \geq 1$ (Snyder & Shen, 2019). In the two-state supply process, the future state of the system depends on the current state of the system.

In contrast, a Bernoulli supply process assumes the future state of the system is independent of the current state of the system. In other words, the probability the system is disrupted in the next review period is the same, regardless of the current state of the system (Snyder et al., 2016). When disruptions follow a Bernoulli supply process, the number of review periods the system has been disrupted is a geometric random variable where $\pi_j = p(1-p)^j$, $j \geq 0$ (Skouri et al., 2014). In the definition, p ($0 \leq p \leq 1$) is the supply quality where $p = 1$ means perfect supply quality.

The long-run probability that the supply chain is not disrupted for a two-state and Bernoulli supply process are $\frac{\beta}{\alpha+\beta}$ and p, respectively, implying this probability is the same when $\frac{\beta}{\alpha+\beta} = p$ (Snyder & Shen, 2019; Skouri et al., 2014). An important question is whether supply should be modeled as a two-state or Bernoulli supply process. Even though supply chain disruptions in a pharmaceutical supply chain can be exogenous to the supply chain (e.g., natural disasters) or endogenous to the supply chain (e.g., drug recalls) (Konstantaras et al., 2019), we model supply as a two-state supply process as the duration of and time between supply chain disruptions in a pharmaceutical supply chain are well represented by the geometric probability distribution (Tucker et al., 2020b). To this end, all closed-form solutions presented in this chapter correspond to a two-state supply process. In the Appendix, we provide the general model for a two-state and Bernoulli supply process (in Appendix B.1) and the closed-form solutions for the Bernoulli supply process (in Appendix B.2). The Bernoulli supply process solutions are easier to interpret and implement as the closed-form expressions are much more condensed in comparison to the two-state supply process. We want to highlight that the Bernoulli supply process comes at a cost of not capturing the duration of and time between supply chain disruptions. Consider $\alpha_1 = 0.5$ and $\beta_1 = 0.5$ which corresponds to frequent supply chain disruptions, but short recovery times and $\alpha_2 = 0.01$ and $\beta_2 = 0.01$ which corresponds to infrequent supply chain disruptions, but long recovery times. Both have a long-run probability that supply is not disrupted of $\frac{\beta_1}{\alpha_1+\beta_1} = \frac{\beta_2}{\alpha_2+\beta_2} = \frac{1}{2}$ implying $p = \frac{1}{2}$ for both. However, the disruption profile is not captured in p. Later (in Section 3.6.3), we illustrate the importance of accounting for the duration of and time between supply chain disruptions when making inventory decisions. We also investigate (in Section 3.6.5) the consequences of using the Bernoulli solutions (i.e., simpler/more-condensed solutions) instead of the two-state solutions. We now proceed to present past literature relevant to our research.

Early Supply Chain Disruption Research

Research on inventory systems with complete supply chain disruptions was started by Parlar & Berkin (1991) with corrections made by Berk & Arreola-Risa (1994). In this research, the authors formulate an economic order quantity (EOQ) model with supply chain disruptions. The model is solved using line search techniques because of the exponential term in the formulation (i.e., $\frac{\lambda}{\lambda+\mu}(1-e^{-\frac{(\lambda+\mu)Q}{q}})$ where q is the deterministic daily demand, λ is the disruption rate, and μ is the recovery rate; $(1-e^{-\frac{(\lambda+\mu)Q}{q}})$ indicates that when the inventory on-hand reaches zero, the supply chain was not disrupted $\frac{Q}{q}$ days ago; Snyder, 2014; Snyder & Shen, 2019). Also, the researchers assume that orders are placed when the inventory position is zero (i.e., $s=0$). Parlar & Perry (1995) relax the zero reorder point assumption and treat the reorder point, s, as a decision variable in the model. All of this early research considers supply chain disruptions as a two-state supply process, but none of this research derives closed-form solutions nor considers perishability.

Non-perishable Closed-Form Solutions

We first present research that models supply chain disruptions as a two-state supply process and then present research that models supply chain disruptions as a Bernoulli supply process. For both supply disruption processes (i.e., two-state and Bernoulli), none of the past closed-form (R,S) models consider a complete lost-sales system, an objective that minimizes the expected cost with a constraint on the proportion of demand satisfied, and perishability.

For disruptions that follow a two-state supply process, Snyder (2014) replaces the exponential term (i.e., $\frac{\lambda}{\lambda+\mu}(1-e^{-\frac{(\lambda+\mu)Q}{q}})$) with a constant (i.e., $\frac{\lambda}{\lambda+\mu}r$ where $0<r\leq 1$) in the EOQ formulation by Berk & Arreola-Risa (1994) to derive approximate closed-form solutions for a non-perishable inventory system. Recall, this model assumes the reorder point, s, is zero, so only the order quantity is optimized. Qi et al. (2009) derive approximate closed-form solutions for the extended model that considers supply chain disruptions at both the supplier and retailer. These researchers assume that the reorder point is zero and that when a supply disruption occurs at the retailer, all inventory is lost. In contrast to the zero reorder point models, Heimann & Waage (2007) consider the non-zero reorder point EOQ model developed by Parlar & Perry (1995) and formulate approximate closed-form solutions for the reorder point and order quantity. Schmitt et al. (2010) present closed-form solutions for a periodic review (R,S) inventory system with an arbitrarily fixed $R=1$ that has deterministic demand, and the authors derive closed-form approximations for systems with stochastic demand or stochastic supply yield. Saithong & Lekhavat (2020) extend the deterministic demand model by deriving periodic review (R,S) policies with an arbitrarily fixed $R=1$ closed-form solutions with partial backordering and lost sales.

For disruptions that follow a Bernoulli supply process, Gullu et al. (1997) derive periodic review

(R,S) policies with an arbitrarily fixed $R = 1$ closed-form solutions for an inventory system with deterministic dynamic demand and stochastic supply that depends on the period in the planning horizon. Skouri et al. (2014) derive closed-form reorder point, s, and order quantity, Q, policies for an inventory system that minimizes the expected cost per day which consists of holding, back-ordering, and ordering costs. These authors assume deterministic demand so the problem can be directly redefined as a (R,S) policy (i.e., q is the deterministic daily demand, $R = \frac{Q}{q}$, $S = s + Q$). Further, the authors assume $S \leq qR$ implying orders cannot be placed unless the inventory position is at most zero (Konstantaras et al., 2019). This research has been extended by Salehi et al. (2016) and Taleizadeh (2017) who consider partial backorders with lost sales and partial backorders with advance payment, respectively.

Recall, Skouri et al. (2014) assume $S \leq qR$. More recent research by Konstantaras et al. (2019) relax this assumption and the authors derive closed-form solutions for the same non-perishable (R,S) inventory system. Taleizadeh et al. (2021) extend the model to consider partial backordering with lost-sales. These models minimize the expected holding, ordering, and backordering costs per day. For the lost-sales component in the partial backordering with lost-sales model, Taleizadeh et al. (2021) use a pessimistic approximation by including the total lost-sales cost in the objective function. In comparison to this chapter, this more recent research also assumes deterministic demand, but only models supply chain disruptions as a Bernoulli supply process. Also, this more recent research considers backordering or partial backordering whereas we focus on a complete lost-sales system with a constraint on the proportion of demand not satisfied. Zero backordering is a natural assumption in our hospital pharmacy setting of interest because patients cannot wait days, weeks, or months to receive a drug that they need at this very moment (e.g., post-surgery medication).

Perishable Closed-Form Solutions

We note that researchers often use methodologies such as stochastic programming (e.g., Dillon et al., 2017; Nguyen & Chen, 2019; Rajendran & Srinivas, 2020; Franco & Alfonso-Lizarazo, 2020) or simulation (e.g., Syawal & Alfares, 2020; Chapter 2) instead of considering closed-form models. For emphasis, in this chapter, we focus on closed-form solutions which provide the inventory policy quickly and are easy to implement. Atan & Rousseau (2016) derive closed-form solutions for the order-up-to level, S, in periodic review (R,S) policies with an arbitrarily fixed $R = 1$ perishable inventory system with supply chain disruptions (two-state supply process). The authors assume backordering is allowed, and they assume a periodic review $(R = 1,S)$ ordering policy is optimal. However, inventory control literature demonstrates that when the fixed ordering cost, $k > 0$, a (R,S) inventory policy where R is not fixed to 1 may be more appropriate since each order placed has a cost associated with it (Snyder & Shen, 2019). Kouki et al. (2020) consider

a $(s = S-1, S)$ perishable inventory system (either lost-sales or backordering) with Poisson demand, stochastic lead time, stochastic lifetime, and no supply uncertainty. The authors provide the explicit expression for the expected cost of the system.

3.3 Model Notation and Assumptions

We study a perishable lost-sales (R,S) inventory system with supply chain disruptions that minimizes the expected holding, ordering, and waste costs while ensuring the proportion of demand not satisfied is at most γ. We describe the modeling notation (see Table 3.1) and assumptions.

Table 3.1 Summary of the modeling notation.

Notation	Description
Decision Variables	
R	Length of the review period (i.e., attempt to place an order every R days; $R \geq 1$)
S	Order-up-to level (i.e., attempt to place an order up to S every R days)
Input Parameters	
h	Holding cost per day per drug ($h > 0$)
k	Fixed ordering cost (i.e., for each order attempted; $k > 0$)
q	Deterministic demand per day ($q > 0$)
e	Expiration lifetime in days ($e \geq 1$)
γ	Maximum proportion of demand not satisfied over infinite horizon ($0 < \gamma < 1$; lost-sales)
Other Parameters	
j	State indicating the number of consecutive review periods that supply has been disrupted
π_j	Probability supply has been disrupted for exactly j consecutive review periods
I_j^+	Average inventory held per day over a review period that starts in state j (in Section 3.4.1)
I_j^-	Average inventory short per day over a review period that starts in state j (in Section 3.4.1)
m	Number of review periods the order-up-to level, S, fully satisfies demand; $m = \lfloor \frac{S}{qR} \rfloor$ (in Section 3.4.2)
Two-state Supply	
α	Disruption probability with respect to the length of the review period ($0 < \alpha < 1$)
β	Recovery probability with respect to the length of the review period ($0 < \beta < 1$)
π_0	Probability supply is disrupted for 0 consecutive review periods (i.e., not disrupted); $\pi_0 = \frac{\beta}{\alpha+\beta}$
π_j	Probability supply is disrupted for exactly j consecutive review periods; $\pi_j = \frac{\alpha\beta}{\alpha+\beta}(1-\beta)^{j-1}, j \geq 1$

3.3.1 Assumptions

We consider an infinite planning horizon and treat the length of the review period R and order-up-to level S as continuous decision variables to allow the derivation of the solutions in closed-form. For supply, we assume supply is completely disrupted; once a supply chain disruption occurs, the current inventory on-hand is not lost, but no additional inventory can be acquired until the supply chain is properly functioning again. Consistent with pharmaceutical supply chains, we assume supply chain disruptions follow a two-state supply process (Tucker et al., 2020b). For the review period of length R, we assume orders can only be attempted in increments of R (e.g., $R = 3$ implies orders can only be attempted on days $1, 4, 7, ...$), regardless of the supply status between review periods, and that orders are only processed successfully when supply is not disrupted (i.e., the order is lost during a supply chain disruption). This implies that if an order is attempted at the start of a review period, say day t, but supply is disrupted, then the inventory on-hand is not raised to the

order-up-to level S, until the first day $t \in \{(t+R),(t+2R),...\}$ that has non-disrupted supply. In a hospital pharmacy setting, we want to highlight that the length of the review period R is typically much less than the length of the supply chain disruption, and hospital pharmacies often do not know the exact instance that the supply chain is not disrupted. These characteristics support the increments of R assumption. We assume the lead time is zero which is a common assumption in the application area of interest. Managers at the University of Michigan's Central Pharmacy report very small lead times (e.g., 36-72 hours) while supply chain disruptions can last weeks, months, or years (Tucker et al., 2020a). Due to these two prior assumptions, inventory on-hand (i.e., amount in stock) and inventory position (i.e., amount in stock plus orders en route) are the same, so we consistently refer to inventory as inventory on-hand for clarity. We assume daily demand is deterministic and constant, but demonstrate how much of an impact this assumption has on a system's performance in Section 3.6.4. Also, we assume the lifetime of the perishable product is deterministic, has no quality decay (Khan et al., 2014), and begins when the product arrives, as the expiration lifetime of drugs is well known. Finally, we assume first-in-first-out protocols are in place, which is in accordance with practice at the University of Michigan's Central Pharmacy.

3.4 Non-perishable (R,S) Inventory System

We derive closed-form solutions for a *non-perishable* lost-sales (R,S) periodic review inventory system with supply chain disruptions. We start by illustrating the procedure in Figure 3.1.

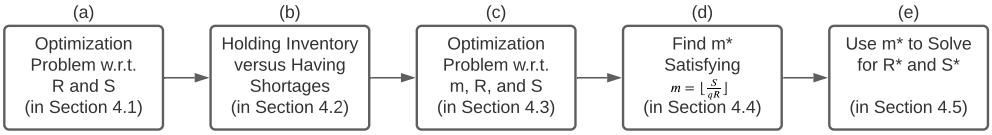

Figure 3.1: Summary of the non-perishable derivation.

Step (e) in Figure 3.1 presents the final closed-form solutions and implementation (in Section 3.4.5). We conclude this section by developing insights from the closed-form solutions (in Section 3.4.6). We provide the derivation details for this section in Appendix B.1.2.

3.4.1 Optimization Problem w.r.t. R and S

For the (R,S) inventory system, the model minimizes the expected (i) ordering cost and (ii) holding cost while (iii) ensuring the proportion of demand not satisfied is at most γ.

In regards to ordering cost, (i) a fixed cost of k is incurred for each order attempted, regardless of the supply status (i.e., cyclic delivery contract; Konstantaras et al., 2019). Hence, with a

review period of length R, the expected ordering cost per day is $\frac{k}{R}$. For systems that experience a cost for each order attempted (e.g., k_1) and each order received (e.g., k_2), the fixed cost k can be replaced with $k = k_1 + \pi_0 k_2$ to account for both types of cost where π_0 corresponds to the long-run probability that supply is not disrupted.

(ii) We only incur holding costs when we have positive inventory on-hand, so we first consider the inventory on-hand at the start of the review period (i.e., we just received an order recalling the zero lead time assumption). If we had no supply chain disruptions, then our current inventory on-hand would always be S at the start of the review period. However, we have supply chain disruptions. From Table 3.1, state j indicates a supply chain disruption that has lasted j consecutive review periods. Therefore, we let t indicate the time since the inventory on-hand has been reviewed to find that our current inventory on-hand is $(S - jqR - qt)^+ = \max\{0, (S - jqR - qt)\}$ at the start of the review period (i.e., $t = 0$). We define I_j^+ as the average inventory held per day over a review period that starts in state j. We have:

$$I_j^+ = \frac{1}{R} \int_0^R (S - jqR - qt)^+ dt \tag{3.1}$$

where $(S - jqR - qt)^+$ corresponds to holding cost only being incurred when inventory on-hand is positive (Konstantaras et al., 2019).

(iii) When ensuring the proportion of demand not satisfied is at most γ, we must consider the average number of shortages per day. Similar to holding, we define I_j^- as the average number of shortages per day over a review period that starts in state j. Given a daily demand of q, we have:

$$I_j^- = \frac{1}{R} \int_0^R [(S - jqR - qt)^- - (S - jqR - q(t-1))^-] dt \tag{3.2}$$

where $(S - jqR - qt)^- = -\min\{0, (S - jqR - qt)\}$ indicates the cumulative shortages with respect to t days. To get the shortages per day, we subtract the cumulative shortages with respect to $t - 1$ days (i.e., $(S - jqR - q(t-1))^-$). Then, dividing I_j^- by the daily demand ($q > 0$) indicates the expected proportion of shortages per day.

Using (i), (ii), and (iii) and recalling that π_j indicates the probability the system has been disrupted for exactly j review periods, we can write the optimization problem as Formulation (3.3).

$$\begin{aligned} \underset{R,S}{\text{minimize}} \quad & C(R,S) = \frac{k}{R} + \sum_{j=0}^{\infty} (h\pi_j I_j^+) \\ \text{subject to} \quad & \frac{\sum_{j=0}^{\infty} (\pi_j I_j^-)}{q} \leq \gamma \end{aligned} \tag{3.3}$$

In Formulation (3.3), the objective function minimizes the expected ordering and holding cost per day and the constraint ensures the proportion of demand not satisfied is at most γ. We consider this constrained optimization problem because hospital pharmacies often aim to achieve certain service levels (Gebicki et al., 2014).

3.4.2 Holding Inventory versus Having Shortages

For any given review period of length R, we may only hold inventory (Case 1), hold inventory and have shortages (Case 2), or only have shortages (Case 3). We define $m = \lfloor \frac{S}{qR} \rfloor$ which denotes how many review periods S fully satisfies demand (Konstantaras et al., 2019). Consider a daily demand of $q = 1$, review period of length $R = 4$ days, and order-up-to level of $S = 10$. Here, $m = \lfloor \frac{10}{4} \rfloor = 2$ implying S fully satisfies demand for 2 review periods (i.e., 8 days). For 3 review periods (i.e., 12 days), S satisfies demand for day 9 and day 10, but not for day 11 and day 12 explaining why S only fully satisfies demand for 2 review periods. We use m and consider the three cases that define the system's status when we start in state j (see Figure 3.2).

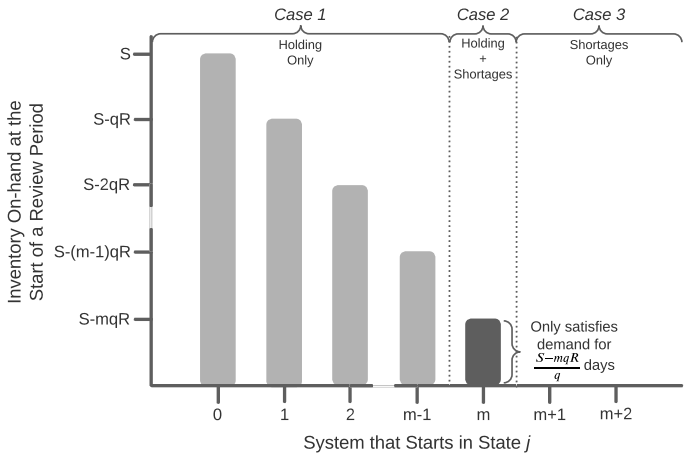

Figure 3.2: Holding inventory versus shortages.

Case 1: $j \leq m-1$

In this case, the system has been disrupted for less than m review periods. Therefore, S can fully satisfy the demand in this review period, so we never have shortages and only hold inventory (i.e.,

$I_j^+ \geq 0$ and $I_j^- = 0$). If we consider the difference between I_j^+ and I_{m-1}^+, we find:
$I_j^+ - I_{m-1}^+ = (S - jqR - qt) - (S - (m-1)qR - qt) = (m-1-j)qR$.

Case 2: $j = m$

In this case, the system has been disrupted for m review periods (i.e., starting inventory on-hand is $S - mqR$), but recall that S only fully satisfies demand for m review periods. Hence, we will have days that we hold inventory and days that we are short (i.e., $I_j^+ \geq 0$ and $I_j^- \geq 0$). Specifically, we hold inventory from the start of the review period until time $\frac{S-mqR}{q}$ (i.e., $S - mqR - qt = 0$). From time $\frac{S-mqR}{q}$ until the end of the review period, we are short.

Case 3: $j \geq m+1$

In this case, we have exhausted all of our inventory on-hand since S only fully satisfies demand for m review periods. Therefore, we always have shortages and never hold inventory (i.e., $I_j^+ = 0$ and $I_j^- \geq 0$). For $j \geq m+1$, we always have a shortage (i.e., lost-sales) of q per day. If we consider the difference between I_j^- and I_{m+1}^-, we find:
$I_j^- - I_{m+1}^- = q - q = 0$.

3.4.3 Optimization Problem w.r.t. m, R, and S

Defining $m = \lfloor \frac{S}{qR} \rfloor$ and using the holding versus shortage cases (in Section 3.4.2), we combine the findings to obtain the optimization problem with respect to m, R, and S in Formulation (3.4).

$$\begin{aligned} \underset{R,S,m}{\text{minimize}} \quad & C(R,S,m) = \frac{k}{R} + \overbrace{\sum_{j=0}^{m-1} h\pi_j I_j^+}^{\textit{Case 1}} + \overbrace{h\pi_m I_m^+}^{\textit{Case 2}} \\ \text{subject to} \quad & \frac{\overbrace{\pi_m I_m^-}^{\textit{Case 2}} + \overbrace{\sum_{j=m+1}^{\infty} \pi_j I_j^-}^{\textit{Case 3}}}{q} \leq \gamma \end{aligned} \tag{3.4}$$

Applying basic rules of addition and subtraction to Formulation (3.4), we notice:

$$\begin{aligned} \underset{R,S,m}{\text{minimize}} \quad & C(R,S,m) = \frac{k}{R} + \sum_{j=0}^{m-1} h\pi_j (I_j^+ - I_{m-1}^+) + \sum_{j=0}^{m-1} h\pi_j I_{m-1}^+ + h\pi_m I_m^+ \\ \text{subject to} \quad & \frac{\pi_m I_m^- + \sum_{j=m+1}^{\infty} \pi_j (I_j^- - I_{m+1}^-) + \sum_{j=m+1}^{\infty} \pi_j I_{m+1}^-}{q} \leq \gamma \end{aligned} \tag{3.5}$$

Using the findings from Section 3.4.2 in Formulation (3.5) and simplifying leads to:

$$\begin{aligned} \underset{R,S,m}{\text{minimize}} \quad & C(R,S,m) = \frac{k}{R} + h\sum_{j=0}^{m-2}\pi_j(m-1-j)qR + h\pi_m I_m^+ + hI_{m-1}^+\sum_{j=0}^{m-1}\pi_j \\ \text{subject to} \quad & \frac{\pi_m I_m^- + I_{m+1}^-\sum_{j=m+1}^{\infty}\pi_j}{q} \leq \gamma \end{aligned} \tag{3.6}$$

where:

$$I_m^+ = \frac{(S-mqR)^2}{2qR} \tag{3.7}$$

$$I_m^- = \frac{q}{R}\left(R - \frac{S-mqR}{q}\right) \tag{3.8}$$

$$I_{m-1}^+ = S - mqR + \frac{1}{2}qR \tag{3.9}$$

$$I_{m+1}^- = q \tag{3.10}$$

Using Equations (3.7)-(3.10) in Formulation (3.6) and simplifying using the definition of π_j (see Table 3.1), we have the optimization problem in Formulation (3.11).

$$\begin{aligned} \underset{R,S,m}{\text{minimize}} \quad & C(R,S,m) = \frac{k}{R} + h\left(S - \frac{1}{2}qR\right)\left(\frac{\beta}{\alpha+\beta} - \frac{\alpha\beta}{(1-\beta)(\alpha+\beta)}\right) + \\ & hqR\left(\frac{1}{\beta}\right)(m\beta - 1 + (1-\beta)^m)\left(\frac{\alpha}{(1-\beta)(\alpha+\beta)}\right) + h\frac{(S-mqR)^2}{2qR}(1-\beta)^{m-1}\left(\frac{\alpha\beta}{(\alpha+\beta)}\right) + \\ & h\left(S - mqR + \frac{1}{2}qR\right)(1-(1-\beta)^m)\left(\frac{\alpha}{(1-\beta)(\alpha+\beta)}\right) \\ \text{subject to} \quad & \left(\frac{1}{\beta} + m - \frac{S}{qR}\right)(1-\beta)^{m-1}\left(\frac{\alpha\beta}{(\alpha+\beta)}\right) \leq \gamma \end{aligned} \tag{3.11}$$

3.4.4 Find m^* Satisfying $m = \lfloor \frac{S}{qR} \rfloor$

For Any Fixed m, Optimize R and S

When considering the objective function in Formulation (3.11), we find that for any fixed m, the Hessian is positive semi-definite and we have that S and R are jointly convex. We can solve the model presented in Formulation (3.11) using a Lagrange multiplier (λ_1). From the Lagrange multiplier formulation, we find that the constraint in Formulation (3.11) must be satisfied with equality (i.e., the constraint must be tight) to satisfy the requirements of a two-state supply process. We use these findings to solve for m^*.

Solving For m^*

Since the constraint in Formulation (3.11) must be satisfied with equality, we can use the definition

of m to solve for the unique and optimal m^*.

Using the definition of m, we have:

$$m = \lfloor \frac{S}{qR} \rfloor \implies m \leq \frac{S}{qR} < m+1 \tag{3.12}$$

Considering the constraint in Formulation (3.11) satisfied with equality, we have:

$$\frac{S}{qR} = \frac{1}{\beta} + m - \frac{\gamma(\alpha+\beta)}{\alpha\beta(1-\beta)^{m-1}} \tag{3.13}$$

Combining Equation (3.12) and Equation (3.13) simplifies to:

$$\frac{\ln(\frac{\gamma(\alpha+\beta)(1-\beta)}{\alpha})}{\ln(1-\beta)} - 1 < m \leq \frac{\ln(\frac{\gamma(\alpha+\beta)(1-\beta)}{\alpha})}{\ln(1-\beta)} \tag{3.14}$$

Using Equation (3.14), we define m^* in Equation (3.15).

$$\text{(Solving for } m^*) \implies m^* = \left\lfloor \frac{\ln(\frac{\gamma(\alpha+\beta)(1-\beta)}{\alpha})}{\ln(1-\beta)} \right\rfloor \tag{3.15}$$

Considering Equation (3.15), we have four input parameter requirements:
(R1) $\gamma \leq \frac{\alpha}{\alpha+\beta}$; (R2) $\gamma > 0$; (R3) $\alpha > 0$; (R4) $0 < \beta < 1$

(R1) ensures $m \geq 1$ (i.e., S covers at least one review period) and this requirement implies that the maximum proportion of demand not satisfied must be less than or equal to the long-run probability that supply is disrupted (i.e., $\frac{\alpha}{\alpha+\beta}$). This requirement is not restrictive as it aligns with the intuition of the problem. Consider the extreme case of $\alpha = 0$. This case implies the supply chain is never disrupted, so we should be able to fully satisfy demand. Thus, it is appropriate to set $\gamma = 0$ (i.e., we satisfy all of the demand). In general, the requirement implies that the more confident we are about our supply (i.e., the smaller $\frac{\alpha}{\alpha+\beta}$ is), the smaller we need to set γ.

(R2)-(R4) ensure m is defined with the $\ln(\frac{\gamma(\alpha+\beta)(1-\beta)}{\alpha})$ term in the numerator and the $\ln(1-\beta)$ term in the denominator of Equation (3.15). All three requirements are also not restrictive. $\gamma > 0$ implies that the maximum proportion of demand not satisfied must be greater than 0. This is appropriate because we can never guarantee that all of our demand will be satisfied since we have supply uncertainty (i.e., supply chain disruptions). $\alpha > 0$ corresponds to a positive probability that the supply chain will transition to the disrupted state given it is currently not disrupted. Since α is a probability, the only other possibility is $\alpha = 0$ which implies supply is perfect. If this were the case, a supply chain disruption model would not be necessary. For $0 < \beta < 1$, $\beta = 0$ corresponds to the supply chain never recovering after a disruption (i.e., the disrupted state is an absorbing state in the two-state Markov chain). If this were the case, the purchaser would find another supplier

who is more reliable. On the other hand, $\beta = 1$ corresponds to the supply chain recovering exactly one review period after the disruption occurs. This implies supply recovery is deterministic and thus, a supply chain disruption model would not be necessary.

3.4.5 Use m^* to Solve for R^* and S^*

Using the optimal m^* (see Equation (3.15)), we can solve for the non-perishable closed-form (R,S) inventory policies as presented in Equations (3.16)-(3.17). We numerically verify that $A_1 > 0$ in Equation (3.16) for all $\gamma \in \{0.001, 0.002, ..., 0.999\}$, $\alpha \in \{0.001, 0.002, ..., 0.999\}$, and $\beta \in \{0.001, 0.002, ..., 0.999\}$ combinations that meet requirement (R1). For each combination, we use the corresponding m^* (see Equation (3.15)) for the verification.

$$\text{(Solving for } R^*) \implies R^* = \max\left\{1, \sqrt{\frac{2k\alpha\beta(\alpha+\beta)(1-\beta)^{m^*+1}}{qh(A_1)}}\right\} \tag{3.16}$$

where:

$$\begin{aligned} A_1 = \big(&-2\alpha^2\gamma - 2\beta^2\gamma + 4\beta^3\gamma - 2\beta^4\gamma + \alpha^2\gamma^2 + \beta^2\gamma^2 - 2\beta^3\gamma^2 + \beta^4\gamma^2 - 4\alpha\beta^2\gamma^2 - 2\alpha^2\beta\gamma^2 - 2\alpha^2\beta^2\gamma + 2\alpha\beta^3\gamma^2 \\ &- 4\alpha\beta\gamma + \alpha^2\beta^2\gamma^2 + 2\alpha\beta\gamma^2 + 8\alpha\beta^2\gamma + 4\alpha^2\beta\gamma - 4\alpha\beta^3\gamma + 2\alpha^2\beta^2\gamma(1-\beta)^{m^*} - 2\alpha\beta^2\gamma(1-\beta)^{m^*} - 2\alpha^2\beta\gamma(1-\beta)^{m^*} \\ &+ 2\alpha\beta^3\gamma(1-\beta)^{m^*} + \alpha^2(1-\beta)^{2m^*} + 2\alpha\beta(1-\beta)^{m^*} - 3\alpha\beta^2(1-\beta)^{m^*} - \alpha^2\beta(1-\beta)^{m^*} + \alpha\beta^3(1-\beta)^{m^*} \\ &+ \alpha^2\beta(1-\beta)^{2m^*} + \alpha^2\beta^2(1-\beta)^{m^*}\big) + \big(m^*(2\alpha\beta(1-\beta)^{m^*})(-\alpha\beta + \alpha\beta(1-\beta)^{m^*} + \beta + \alpha - \beta^2)\big) \end{aligned}$$

$$\text{(Solving for } S^*) \implies S^* = qR^*\left(\frac{1}{\beta} + m^* - \frac{\gamma(\alpha+\beta)}{\alpha\beta(1-\beta)^{m^*-1}}\right) \tag{3.17}$$

The two-state supply process assumes the future state of the supply chain depends on the current state of the supply chain (Snyder & Shen, 2019). This implies the future state of the supply chain is dependent on the length of the review period R, so the disruption probability and recovery probability parameters, α and β, must be with respect to the length of the review period. However, we do not know the length of the review period R beforehand. Therefore, to properly solve for the (R,S) policies, we consider an iterative algorithm. For the algorithm, we denote η as the iteration number, R^*_η as the length of the review period obtained on iteration η, S^*_η as the order-up-to level obtained on iteration η, $\alpha^{(i)}$ as the disruption probability with respect to a review period of length i, $\beta^{(i)}$ as the recovery probability with respect to a review period of length i, and $\mathbf{P}^{(i)}$ as the i-step transition probability matrix (Kulkarni, 2011) where:

$$\mathbf{P}^{(i)} = \begin{pmatrix} & \begin{matrix} not\ disrupted & disrupted \end{matrix} \\ \begin{matrix} not\ disrupted \\ disrupted \end{matrix} & \begin{pmatrix} \cdot & \alpha^{(i)} \\ \beta^{(i)} & \cdot \end{pmatrix} \end{pmatrix} = \begin{pmatrix} & \begin{matrix} not\ disrupted & disrupted \end{matrix} \\ \begin{matrix} not\ disrupted \\ disrupted \end{matrix} & \begin{pmatrix} (1-\alpha^{(1)}) & \alpha^{(1)} \\ \beta^{(1)} & (1-\beta^{(1)}) \end{pmatrix}^i \end{pmatrix}$$

The algorithm has 3-steps: initialization, iterative procedure, and ensure model requirements are satisfied. At a high level, the algorithm iteratively solves for the (R^*, S^*) inventory policy until

$\alpha^{(R^*)}$ and $\beta^{(R^*)}$ output R^* as the optimal length of the review period. We proceed to formally present the algorithm.

Algorithm 1 Iterative algorithm for two-state supply process.

Step 1: Initialization
Set $\eta = 0$, $R^*_{(-1)} = 0$, $R^*_0 = 1$, $\alpha = \alpha^{(1)}$, $\beta = \beta^{(1)}$; ensure $\gamma \leq \frac{\alpha}{\alpha+\beta}$, $\gamma > 0$, $\alpha > 0$, and $0 < \beta < 1$

Step 2: Iterative Procedure
while $R^*_{\eta-1} \neq R^*_\eta$
    $\eta = \eta + 1$ *% update iteration number*
    Solve for (R^*_η, S^*_η) using α and β *% solve for inventory policy (see Equations (3.16)-(3.17))*
    Set $\alpha = \alpha^{(R^*_\eta)}$, $\beta = \beta^{(R^*_\eta)}$ *% update α and β to correspond to the length of the review period*

    Termination to Avoid Infinite While Loop
    if $R^*_{\eta-2} = R^*_\eta$ and $R^*_{\eta-1} \neq R^*_\eta$ for 10 consecutive η values *% check if an alternating pattern arises*
        $R^* = \min\{R^*_{\eta-1}, R^*_\eta\}$; terminate while loop *% terminate while loop and define R^* as the minimum*
    end if
    if $round(R^*_\eta) = round(R^*_{\eta-1})$ for 10 consecutive η values *% check if integer solutions remain the same*
        $R^* = \min\{R^*_{\eta-1}, R^*_\eta\}$; terminate while loop *% terminate while loop and define R^* as the minimum*
    end if
end while

Step 3: Ensure Model Requirements are Satisfied
if $\gamma \leq \frac{\alpha}{\alpha+\beta}$
    output: (R^*, S^*) for two-state supply process.
else
    output: Model not feasible with given input parameters. Decrease γ and repeat procedure.
end if

3.4.6 Insights from the Non-perishable Solutions

From the closed-form solutions, we gain insights about the (R,S) inventory policies. From R^* (see Equation (3.16)), we notice an increase in ordering cost (i.e., k) corresponds to an increase in R^* (i.e., place orders less frequently). An increase in daily demand (i.e., q) leads to a decrease in R^*. Ordering cost is incurred per order attempted whereas holding cost is incurred for each item on-hand each day. Therefore, an increase in daily demand implies we want to place orders more frequently, incurring greater ordering costs and less holding costs. We notice that R^* is dependent on the supply disruption process parameters (i.e., α and β). R^* cannot be expressed as a function of $\frac{\beta}{\alpha+\beta}$ implying that only considering the long-run probability that supply is not disrupted is insufficient when solving for R^*.

From S^* (see Equation (3.17)), we notice a higher daily demand (i.e., q) and longer length of the review period (i.e., R^*) implies a larger S^* (i.e., stock more inventory). When considering the maximum proportion of demand not satisfied (i.e., γ), a larger γ implies we need to fulfill less demand, leading to a smaller S^*. A smaller γ implies we need to satisfy more demand, leading to a larger S^*. This conclusion is obtained using the negative term in the S^* expression (i.e., $-\frac{\gamma(\alpha+\beta)}{\alpha\beta(1-\beta)^{m^*-1}}$). Like R^*, we notice that S^* is dependent on the supply disruption process parameters (i.e., α and β). S^* cannot be expressed as a function of $\frac{\beta}{\alpha+\beta}$ implying that only considering the long-run probability that supply is not disrupted is insufficient when solving for S^*.

Using Equations (3.15)-(3.17), we numerically verify that the model reduces to the well-known standard EOQ model in the limiting case where the supply chain becomes perfect (i.e., no supply chain disruptions). We consider $\alpha \longrightarrow 0$ and $\beta \longrightarrow 1$ when studying the limiting case that the supply chain becomes perfect. With (R1), we also have that $\gamma \longrightarrow 0$. We find that $m^* \longrightarrow 1$, $R^* \longrightarrow \sqrt{\frac{2k}{qh}}$, and $S^* \longrightarrow \sqrt{\frac{2kq}{h}}$ which is equivalent to the standard EOQ model.

3.5 Perishable (R,S) Inventory System

In this section, we present closed-form solutions for a *perishable* lost-sales (R,S) periodic review inventory system with supply chain disruptions. We provide the derivation details for this section in Appendix B.1.4. Atan & Rousseau (2016) derive closed-form solutions for a perishable periodic review $(R=1,S)$ inventory system with supply chain disruptions where all demand not met is backordered. The objective of the model is to minimize the expected holding (h), backordering (b), and waste (z) costs. Extending this work to the lost-sales case (i.e., all demand not met is lost), we find the optimal policy has the form $S^* = \min\{f(h,b,\alpha,\beta), eq\}$ where $f(h,b,\alpha,\beta)$ indicates a closed-form expression dependent on the input parameters $(\cdot)$ (in Appendix B.1.3). This result shows that it is never optimal to set $S > eq$ (i.e., we never waste items). It is important to recall the deterministic demand assumption. When introducing a review period of length $R \geq 1$ into the problem, it is also never optimal to set $S > eq$. Ordering more than what is demanded in a lifetime cycle (i.e., eq) would have no impact on the proportion of demand satisfied. The excess inventory would simply be discarded (i.e., $S-eq$) and this excess inventory would contribute to an increase in holding costs. Therefore, reconsidering the question 'how long is long enough?' stated in Section 3.1, we claim e is long enough if $e \geq \frac{S}{q}$, in which case, it is sufficient to use a non-perishable (R,S) model with supply chain disruptions. We can analytically find the smallest e that is long enough by setting $e = \frac{S}{q}$ and solving for e. Using these findings, we provide the following procedure (see Figure 3.3) to extend the non-perishable closed-form (R,S) solutions presented in Section 3.4 to a perishable system.

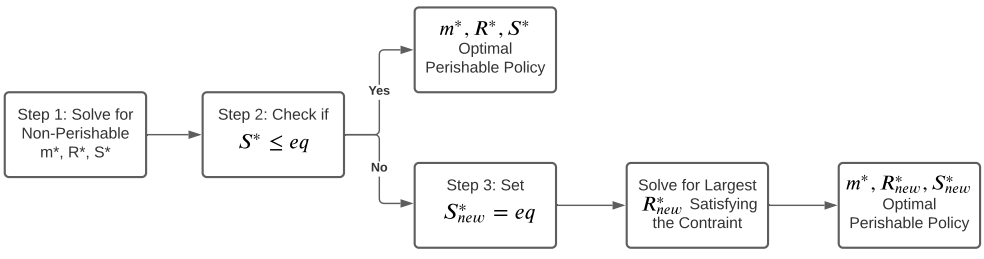

Figure 3.3: Summary of the perishable derivation.

Step 1: Solve for the Non-Perishable m^*, R^*, and S^* Inventory Policy
Solve for m^*, R^*, and S^*. See Equations (3.15), (3.16), and (3.17).

Step 2: Check if $S^* \leq eq$
If $S^* \leq eq$, then (m^*, R^*, S^*) denotes the optimal policy for the perishable inventory system. Otherwise, continue to Step 3.

Step 3: Enforce Perishability Condition by Setting $S^*_{new} = eq$
Since it is never optimal to have $S > eq$, we first modify S^* to be $S^*_{new} = eq$. Then, we are interested in finding the largest R^*_{new} such that our shortage constraint (see Formulation (3.11)) is satisfied when $S^*_{new} = eq$. Given we set $S^*_{new} = eq$, both holding and ordering costs (i.e., costs in the objective function) are the smallest when R^*_{new} is the largest. Hence, we are interested in finding the largest R^*_{new} because this solution will have the smallest objective value. Stating this as an optimization problem, we have Formulation (3.18).

$$\begin{aligned} \underset{R}{\text{maximize}} \quad & R \\ \text{subject to} \quad & \left(\frac{1}{\beta} + m - \frac{S^*_{new}}{qR}\right)(1-\beta)^{m-1}\left(\frac{\alpha\beta}{(\alpha+\beta)}\right) \leq \gamma \end{aligned} \tag{3.18}$$

Solution to the Optimization Problem
Rearranging the constraint in Formulation (3.18) and recalling $S^*_{new} = eq$, we have:

$$R \leq \frac{e}{(\frac{1}{\beta} + m) - \frac{\gamma(\alpha+\beta)}{\alpha\beta(1-\beta)^{m-1}}} \tag{3.19}$$

To maximize R, Constraint (3.19) must be satisfied with equality (i.e., the constraint is tight) and similar to Section 3.4.4, we use the relation that $m = \lfloor \frac{S^*_{new}}{qR} \rfloor = \lfloor \frac{e}{R} \rfloor \implies m \leq \frac{e}{R} < m+1$ to find:

$$(\text{Solving for } m^*_{new}) \implies m^*_{new} = m^* = \left\lfloor \frac{\ln(\frac{\gamma(\alpha+\beta)(1-\beta)}{\alpha})}{\ln(1-\beta)} \right\rfloor \tag{3.20}$$

with the requirements of (R1) $\gamma \leq \frac{\alpha}{\alpha+\beta}$, (R2) $\gamma > 0$, (R3) $\alpha > 0$, and (R4) $0 < \beta < 1$ to ensure

$m^*_{new} = m^*$ is defined. Having found $m^*_{new} = m^*$, we have:

$$\text{(Solving for } R^*_{new}) \implies R^*_{new} = \max\left\{1, \frac{e}{(\frac{1}{\beta} + m^*) - \frac{\gamma(\alpha+\beta)}{\alpha\beta(1-\beta)^{m^*-1}}}\right\} \tag{3.21}$$

$(m^*_{new} = m^*, R^*_{new}, S^*_{new} = eq)$ provides the final solution when the perishability condition is enforced. These (R^*_{new}, S^*_{new}) policies are solved iteratively like the procedure defined in Section 3.4.5. However, instead of starting with $R^*_0 = 1$, $\alpha = \alpha^{(1)}$, and $\beta = \beta^{(1)}$, we use the non-perishable length of the review period found in Step 1 to initialize the algorithm: $R^*_0 = R^*$, $\alpha = \alpha^{(R^*)}$, and $\beta = \beta^{(R^*)}$.

We require $R^*_{new} \geq 1$ in Equation (3.21) which is consistent with periodic review inventory policies. In the non-perishable case (see Equation (3.16)), we enforce $R^* \geq 1$ and we use this R^* to solve for S^*. In the perishable case, S^*_{new} is fixed to $S^*_{new} = eq$. Therefore, if R^*_{new} is forced to 1 due to the maximum function in Equation (3.21), then the maximum proportion of demand not fulfilled constraint (γ) is not satisfied. With the given inputs, one must increase γ to satisfy the constraint. We demonstrate in Section 3.6.2 and Section 3.6.4 examples of when this occurs.

3.5.1 Insights from the Solutions with Perishability Enforced (Step 3)

When enforcing the perishability condition by setting $S^*_{new} = eq$ (see Step 3), we notice that an increase in the expiration lifetime corresponds to an increase in R^*_{new} (i.e., place orders less frequently; see Equation (3.21)). Like the non-perishable closed-form solutions, we notice that R^*_{new} is dependent on the supply disruption process parameters (i.e., α and β). R^*_{new} cannot be expressed as a function of $\frac{\beta}{\alpha+\beta}$ implying that only considering the long-run probability that supply is not disrupted is insufficient when solving for R^*_{new}. Moving forward, we simply refer to the optimal inventory policy as (R^*, S^*).

3.6 Numerical Analysis

We proceed to analyze the perishable closed-form (R^*, S^*) policies (in Section 3.5) in a hospital pharmacy setting. The results generalize to other inventory systems of interest (e.g., blood banks, chemical plants). It is important to consider that the closed-form (R^*, S^*) policies treat the decision variables as continuous variables. To not obscure the results, we simply present the inventory policies as continuous variables. We also consistently refer to the disruption probability (α) and the recovery probability (β) with respect to a review period of length 1 day (i.e., $\alpha = \alpha^{(1)}$ and $\beta = \beta^{(1)}$).

We first present the input parameters for a hospital pharmacy drug base case where we consider Fentanyl 50mcg/mL 30mL (Fentanyl) which is a synthetic opioid used to manage pain (Centers for Disease Control and Prevention, 2021). In particular, the Central Pharmacy at the University of Michigan uses this drug for continuous pain relief treatment during surgeries making it a critical asset to hospital operations. This form of Fentanyl is classified as a 503B drug which pharmacists define as pre-compounded drugs available for commercial sale without a specific patient order.

We (i) illustrate the importance of implementing the closed-form solutions presented in this chapter for a perishable inventory system with supply chain disruptions (in Section 3.6.2). Next, we (ii) analyze how supply chain disruption patterns influence the inventory policies (in Section 3.6.3) and (iii) study the impact of stochastic demand (in Section 3.6.4). We also (iv) investigate the trade-off between model simplicity and model accuracy (in Section 3.6.5). For (i) and (iii), we utilize a simulation model of the inventory system where we round R^* down to the nearest whole number, consider a 12-month warm-up period (i.e., 360 days), consider a 5-year planning horizon (i.e., 1800 days), and we perform 500 simulation replications. The timing and duration of the supply chain disruptions are the random events in the model. 500 simulation replications ensures (with respect to the base case) asymptotic convergence for the expected number of drugs wasted per day, expected number of drug shortages per day, and expected cost per day (holding plus ordering; see Appendix B.3). We conclude this section by presenting a sensitivity analysis (in Section 3.6.6) for model validation and to gain additional managerial insights.

3.6.1 Base Case

Table 3.2 presents the model parameters and the optimal perishable (R^*, S^*) periodic review inventory policy for the hospital pharmacy base case. We consider a holding cost of $h = \$0.025$ (Fentanyl average wholesale price approximately \$25, IBM Watson Health, 2022; holding cost relative to drug price 0.001, Jia & Zhao, 2017) and an ordering cost of $k = \$250$ (ordering cost relative to drug price 10). Using 6 months of the most recent data available from the University of Michigan's Central Pharmacy (i.e., October 22, 2020-March 22, 2021; missing data imputed using linear interpolation with respect to the day of the week), we consider a daily demand of $q = 45$. Consistent with 503B drugs at the University of Michigan, we assume the expiration lifetime of the drug is $e = 90$ days. Drugs may have smaller (e.g., $e \in \{14, 30, 60\}$ days) or larger (e.g., $e = 120$ days) expiration lifetimes in comparison to the base case of $e = 90$ days. We consider these smaller and larger expiration lifetimes and perform a sensitivity analysis when drawing conclusions from the numerical analysis. Also, the base case ensures 95% of demand is satisfied (i.e., $\gamma = 0.05$) when supply is disrupted on average every three months (i.e., $\alpha = \frac{1}{90\,days}$) and remains disrupted for about 1 month (i.e., $\beta = \frac{1}{30\,days}$). These values satisfy the requirements presented in Section

3.4.4 (i.e., $\gamma \leq \frac{\alpha}{\alpha+\beta}$, $\gamma > 0$, $\alpha > 0$, and $0 < \beta < 1$).

Table 3.2 Base case input parameters.

Notation	Description
Input Parameters	
$h = \$0.025$	Holding cost per day per drug ($h > 0$)
$k = \$250$	Fixed ordering cost (i.e., for each order attempted; $k > 0$)
$q = 45$	Deterministic demand per day ($q > 0$)
$e = 90$ days	Expiration lifetime in days ($e \geq 1$)
$\gamma = 0.05$	Maximum proportion of demand not satisfied over infinite horizon ($0 < \gamma < 1$)
Two-state Supply	
$\alpha^{(1)} = \frac{1}{90}$	Disruption probability with respect to a review period of length 1 day (i.e., expected time not disrupted is 90 days)
$\beta^{(1)} = \frac{1}{30}$	Recovery probability with respect to a review period of length 1 day (i.e., expected time disrupted is 30 days)
Decision Variables[a]	
$R^* = 4.95$	Length of the review period (i.e., attempt to place an order every R days; $R \geq 1$)
$S^* = 2412.92$	Order-up-to level (i.e., attempt to place an order up to S every R days)

[a]Rounding to $(R^* = 5, S^* = 2400)$ changes the objective value and proportion of demand not satisfied by less than 1%

3.6.2 Importance of Considering Perishability and Supply Chain Disruptions

The closed-form (R^*, S^*) periodic review inventory policies presented in Section 3.5 are derived for a perishable inventory system with supply chain disruptions. We proceed to quantify the importance of accounting for perishability and supply chain disruptions (SCD) using four models: (a) No SCD/Perish, (b) Perish Only, (c) SCD Only, and (d) SCD/Perish (this research).

For models (a)-(d), Table 3.3 provides a model description and (R^*, S^*) when using the base case input parameters (see Table 3.2). We observe that (R^*, S^*) for models (a) and (b) are equivalent and (R^*, S^*) for models (c) and (d) are equivalent. The reason for the equivalence is that $S^* \leq eq$ implying the perishability condition is never enforced. We notice that the safety stock for models (a) and (b) is $S^* - R^* q = 0$ and that the safety stock for models (c) and (d) is $S^* - R^* q = 2190$. Recalling the maximum demand not satisfied constraint (i.e., γ), models (c) and (d) keep extra inventory on-hand to ensure this constraint remains satisfied even in the presence of supply chain disruptions. Models (a) and (b) have no safety stock because they do not consider supply chain disruptions.

Table 3.3 Description of the four models analyzed (Section 3.6.2).

Model	Solving for $(\mathbf{R}, \mathbf{S})$	$(\mathbf{R}^*, \mathbf{S}^*)$
(a) **No SCD/Perish**	Use the standard EOQ model (Snyder & Shen, 2019) which does not consider supply chain disruptions nor perishability (see Appendix B.4).	$(21.08, 948.68)$
(b) **Perish Only**	Use the standard EOQ model, but enforce the perishability condition by setting $S^* = eq$ when $S^*_{EOQ} > eq$.	$(21.08, 948.68)$
(c) **SCD Only**	Use the closed-form solutions for a non-perishable inventory system with supply chain disruptions presented in Section 3.4.	$(4.95, 2412.92)$
(d) **SCD/Perish (this research)**	Use the closed-form solutions for a perishable inventory system with supply chain disruptions presented in Section 3.5 (i.e., enforce the perishability condition by setting $S^* = eq$ when the non-perishable solution provides a $S^* > eq$).	$(4.95, 2412.92)$

Using the base case input parameters with varying expiration lifetimes, $e \in \{60, 90, 120\}$ days, and varying disruption probabilities, $\alpha \in \{0.01, ..., 0.25\}$, we solve for the (R^*, S^*) inventory policies using each of the four models (i.e., (a)-(d)). Figure 3.4 illustrates the expected proportion of drug shortages per day and expected proportion of drugs wasted per day. The proportion is measured relative to the daily demand of the base case (i.e., $q = 45$). These results are obtained by simulating the perishable inventory system with the (R^*, S^*) inventory policy using the defined expiration lifetime (e), disruption probability (α), and other input parameters defined in Table 3.2. The timing and duration of the supply chain disruptions are the random events in the simulation model. When $S > eq$ (i.e., waste will occur), we must record the age of the inventory to appropriately calculate the proportion of drug shortages and drugs wasted. Thus, we use a simulation model of the (R, S) inventory system to keep the calculations consistent across all models (a)-(d).

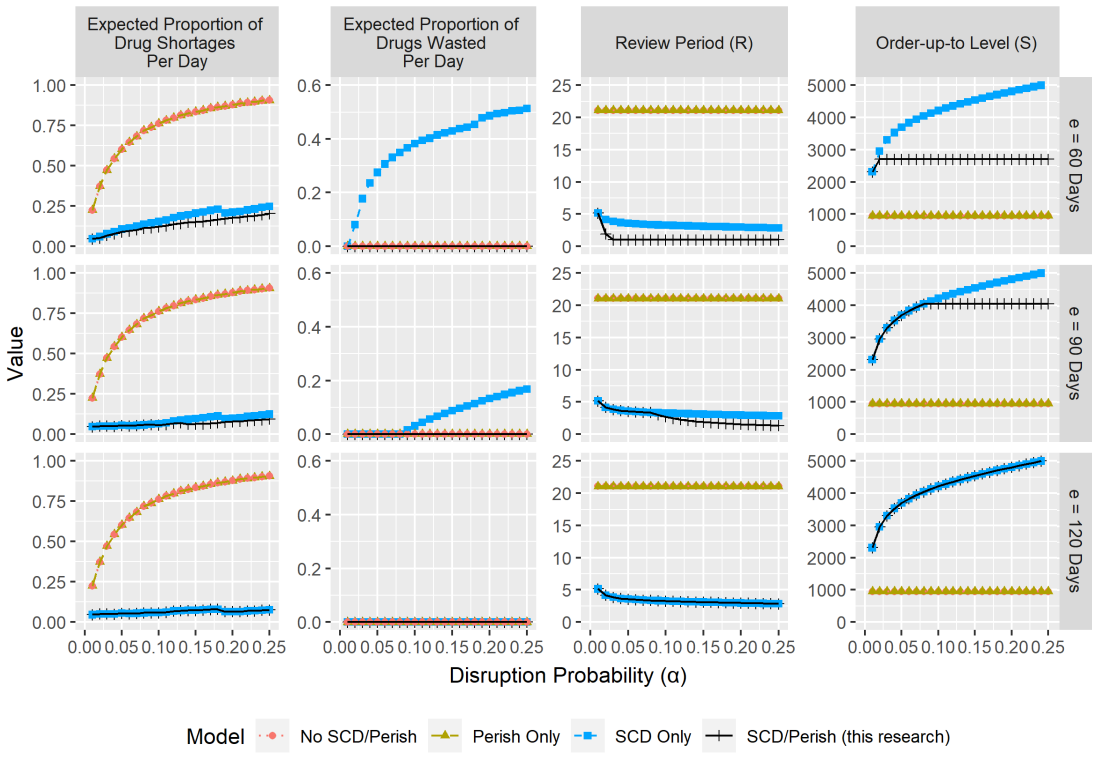

Figure 3.4: Importance of accounting for perishability and supply chain disruptions.

We first notice (see Figure 3.4) that the No SCD/Perish and Perish Only models perform the same for all disruption probabilities (α) and expiration lifetimes (e). This occurs because the order-up-to level S^* is at most eq for all input parameters considered when using the standard EOQ model. S^* does not exceed eq unless $e \leq 21$ ($S^*_{EOQ} = \sqrt{2(250)(45)/0.025} = 949$ and $q = 45 \implies e < \frac{949}{45} = 21.1$). When $e \leq 21$, the Perish Only model would set $S^* = eq$ (i.e., enforce perishability condition). We also notice that when the (R^*, S^*) inventory policy for the SCD Only model and SCD/Perish (this research) model are not the same (i.e., $\alpha \geq 0.02$ for $e = 60$ and $\alpha \geq 0.09$ for $e = 90$), this implies that the SCD/Perish (this research) model sets $S^* = eq$ (i.e., enforces perishability condition). We proceed to highlight the importance of accounting for both perishability and supply chain disruptions simultaneously.

If we only consider perishability (i.e., Perish Only model), we have an increase in the expected proportion of drug shortages per day in comparison to the SCD/Perish (this research) model. Using the base case disruption probability $\alpha = \frac{1}{90} \approx 0.01$ for all expiration lifetimes (e), the Perish Only model leads to demand not being satisfied 22.5% of the time in comparison to the SCD/Perish (this research) model of 4.6% (i.e., about 5 times the amount of shortages). The difference in the expected proportion of drug shortages per day between the two models increases as the disruption probability (α) increases. This aligns with our intuition as a larger disruption probability

corresponds to a larger long-run probability that supply is disrupted (i.e., $\frac{\alpha}{\alpha+\beta}$). For the expected proportion of drugs wasted, there is zero waste when using the Perish Only model and SCD/Perish (this research) model. These results show that only considering perishability instead of simultaneously considering perishability and supply chain disruptions increases the proportion of drug shortages, especially for drugs with large disruption probabilities (α).

The consequences of only considering supply chain disruptions (i.e., SCD Only model) are highlighted when $S^* > eq$ with the non-perishable model. The SCD/Perish (this research) model enforces the perishability condition by setting $S^* = eq$, but the SCD Only model keeps $S^* > eq$. When $S^* > eq$ with the non-perishable model, the SCD/Perish (this research) model leads to no waste, but the SCD Only model leads to a positive proportion of drugs wasted that increases as the disruption probability (α) increases. A larger disruption probability corresponds to a larger long-run probability that supply is disrupted (i.e., $\frac{\alpha}{\alpha+\beta}$) and thus, a larger order-up-to level (see the fourth column panel in Figure 3.4). For disruption probabilities $\alpha \geq 0.23$, the expected number of drugs wasted per day is at least 50% of the daily demand (see Figure 3.4). Even worse, the SCD Only model increases the expected proportion of drug shortages per day in comparison to the SCD/Perish (this research) model. For disruption probability $\alpha = 0.18$, there is a 65% increase in the proportion of drug shortages (see Figure 3.4). When $S^* \leq eq$ with the non-perishable model, the SCD Only and SCD/Perish (this research) model perform the same. These results show that only considering supply chain disruptions instead of simultaneously considering perishability and supply chain disruptions can increase the proportion of drugs wasted, but can also increase the proportion of drug shortages.

From Figure 3.4, we notice that the expected proportion of drug shortages with the SCD/Perish (this research) model is not always less than or equal to $\gamma = 0.05$ (i.e., the demand satisfaction constraint) for $e = 60$ and $e = 90$ days. These are cases where the perishability condition is enforced (i.e., $S^* = eq$; see the fourth column panel in Figure 3.4), but there is not a feasible $R^* \geq 1$ such that the proportion of demand not satisfied is at most $\gamma = 0.05$. We set $R^* = 1$ (in Section 3.5; see the third column panel in Figure 3.4) and while the γ constraint is no longer satisfied, the SCD/Perish (this research) model continues to outperform or perform the same as the other models considered from a drug shortage and drug waste perspective. Also, in Figure 3.4, we observe that the expected proportion of drug shortages per day decreases as α goes from 0.18 to 0.19 for $e \in \{60, 90, 120\}$ days. Recall, the closed-form solutions treat the length of the review period (R) and order-up-to level (S) as continuous decision variables. However, the simulation model rounds the optimal length of the review period R^* down to the nearest whole number to model a periodic review inventory system. This explains the decrease between the two values of α.

3.6.3 Influence of Supply Chain Disruption Patterns

Using the perishable closed-form model, we analyze how supply chain disruption patterns influence the inventory policies. We consider the disruption probability (α) and define the recovery probability (β) such that $\alpha = y\beta$ ($0 < \alpha < 1$; $y \in \mathbb{R}^+$; $0 < \beta < 1$). For any fixed value of y, the corresponding α and β values that satisfy the relation $\alpha = y\beta$ have the same long-run probability of being in the disrupted state (i.e., $\frac{\alpha}{\alpha+\beta}$). However, the supply chain disruption pattern may differ. Figure 3.5 illustrates the supply pattern interpretation between $y : \alpha = y\beta$ and α.

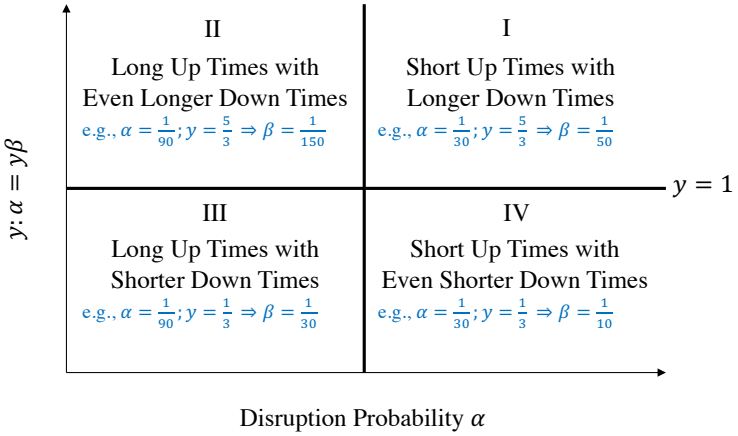

Figure 3.5: Illustration of different supply chain disruption patterns.

We study the expected cost per day (holding plus ordering), the length of the review period (R), and the order-up-to level (S) when considering a disruption probability (α) and recovery probability (β) such that $\alpha = y\beta$ ($0 < \alpha < 1$; $y \in \mathbb{R}^+$; $0 < \beta < 1$). Figure 3.6 presents the results for $\alpha \in \{0.01, 0.02, ..., 0.1\}$ and $y \in \{\frac{1}{3}, \frac{2}{3}, 1, \frac{4}{3}, \frac{5}{3}\}$ with all other inputs fixed (see Table 3.2). When the disruption probability (α) is small and/or y is large, the supply chain is less reliable (see quadrants I-III in Figure 3.5). In these cases, it is optimal to enforce the perishability condition by setting $S^* = eq$. Further, when $y = \frac{5}{3}$, it is optimal to enforce the perishability condition by setting $S^* = eq$ with the same R^* for both $\alpha = 0.01$ and $\alpha = 0.02$. When $\alpha = 0.01$, we have $\beta = 0.006$ and when $\alpha = 0.02$, we have $\beta = 0.012$. For both cases, the long-run probability that supply is disrupted is $\frac{\alpha}{\alpha+\beta} = 0.625$. However, when $\alpha = 0.02$ and $\beta = 0.012$, the supply chain recovers more quickly than the $\alpha = 0.01$ and $\beta = 0.006$ case. This implies we have a greater chance of holding inventory (or equivalently, not stocking out) when $\alpha = 0.02$ over $\alpha = 0.01$ explaining why the expected cost per day increases in Figure 3.6. Parallel arguments hold for $y = \frac{5}{3}$ and $\alpha \in \{0.03, 0.04\}$.

More generally, the results illustrate that for any value of y, there exists a disruption probability α' such that for $\alpha \geq \alpha'$, it becomes optimal to follow the (R^*, S^*) using the non-perishable model (e.g., $\alpha' = 0.03$ when $y = 1$). Another key takeaway is that the inventory policies differ based on

the value of α and y. This illustrates that it is not sufficient to solely use the long-run probability that supply is not disrupted (i.e., $\frac{\beta}{\alpha+\beta}$) when making inventory decisions and instead, both the duration of and time between supply chain disruptions should be considered.

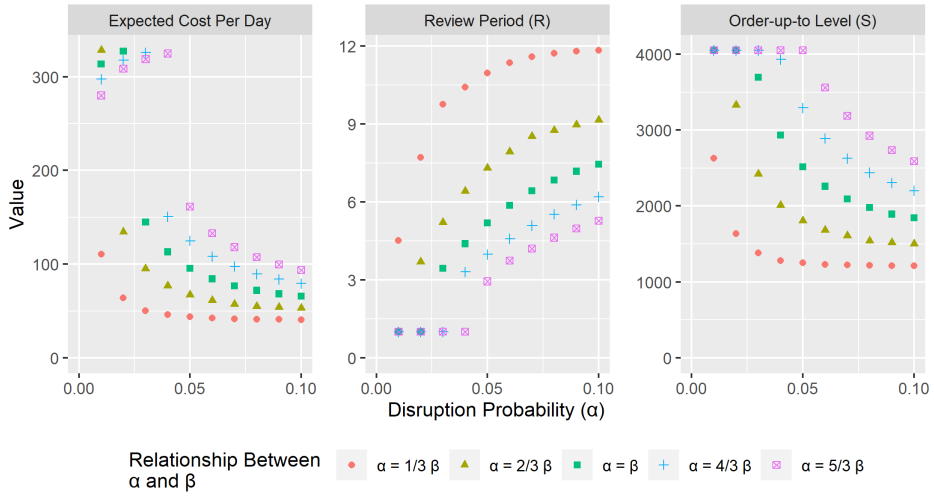

Figure 3.6: Results for different supply chain disruption patterns.

3.6.4 Impact of Stochastic Demand

Using the perishable closed-form model, we study the impact of stochastic demand that is normally distributed. With the base case input parameters (see Table 3.2), we solve for the optimal (R^*, S^*) inventory policies using expiration lifetimes of $e \in \{14, 30, 60, 90, 120\}$ days. Using these (R^*, S^*) inventory policies, we simulate how the inventory system would perform using the (R^*, S^*) policies with stochastic demand. Specifically, we consider daily demand that is normally distributed with mean, $q = 45$, and standard deviation, $\sigma \in \{0, 2, 4, ..., 20\}$ (with the demand in any realization forced to 0 if a negative demand is generated). We consider this range of standard deviations because the most recent 6 months of Fentanyl data has a standard deviation of 15. We calculate the proportion of drug shortages and drugs wasted per day as a fraction of the average daily demand for the simulation run. Figure 3.7 presents the expected cost per day (holding plus ordering), expected proportion of drug shortages per day, and expected proportion of drugs wasted per day versus the standard deviation of daily demand.

The change in standard deviation (i.e., variability in daily demand) has a negligible impact on the expected cost per day which consists of holding and ordering costs. With respect to shortages, the increase in variability in daily demand leads to a negligible increase in the expected proportion of drug shortages for $e \in \{60, 90, 120\}$ days, and the (R^*, S^*) policies continue to satisfy the $\gamma = 0.05$ constraint for all standard deviation values (σ). For $e \in \{14, 30\}$ days, the expected

proportion of drug shortages has a negligible increase as demand variability increases. These expiration lifetimes are examples where given the perishability condition is enforced with $S^* = eq$, there is not a feasible $R^* \geq 1$ such that the proportion of demand not satisfied is at most $\gamma = 0.05$. Therefore, all shortage results for $e \in \{14, 30\}$ days are greater than 0.05. Also, no feasible R^* implies the length of the review period is set to $R^* = 1$ (in Section 3.5). Consequently, orders are placed every day, which explains the large difference in the expected cost between $e = \{14, 30\}$ days and $e \in \{60, 90, 120\}$ days. With respect to waste, the expected proportion of drugs wasted linearly increases with an increase in daily demand variability, but only impacts drugs with an expiration lifetime of $e \in \{14, 30, 60\}$ days.

We believe the negligible increase in the proportion of drug shortages in contrast to a significant increase in the proportion of drug shortages for all expiration lifetimes is a result of the inventory system holding extra inventory on-hand to guard against supply chain disruptions. We also believe the extra inventory on-hand leads to a negligible change in the expected cost per day for all expiration lifetimes. For small expiration lifetimes (i.e., $e \in \{14, 30\}$ days), we suspect the increase in the expected proportion of drugs wasted is a result of the inventory system having an order-up-to level S^* equal to the maximum inventory level with the defined perishability condition (i.e., $S^* = eq$). Other small lifetimes like $e = 60$ days require a large standard deviation (i.e., $\sigma = 14$) to see an increase in waste. The (R^*, S^*) policies are insensitive to changes in the standard deviation when demand is normally distributed, but when a drug's expiration lifetime is small, waste increases linearly as the standard deviation increases.

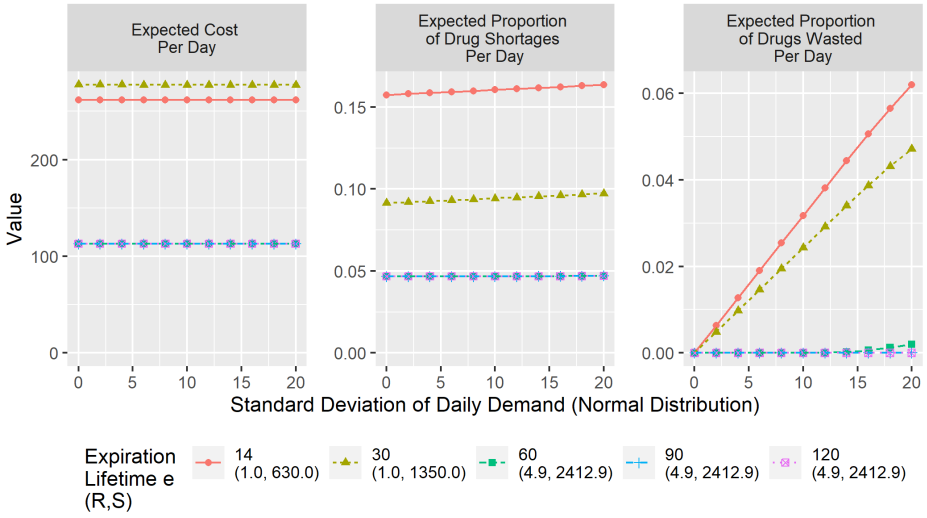

Figure 3.7: Impact of not considering stochastic demand.

3.6.5 Model Simplicity versus Model Accuracy

We investigate the trade-off between model simplicity and model accuracy. In Section 3.6.3, we showed it is critical to account for the duration of and time between supply chain disruptions which are captured when modeling the inventory system as a two-state supply process. These are not captured when modeling the inventory system as a Bernoulli supply process. The Bernoulli closed-form solutions are much more condensed in comparison to a two-state supply process (see Appendix B.2). We consider the base case input parameters with the relationship $\alpha = \frac{1}{3}\beta$ (see Figure 3.6 in Section 3.6.3). We solve for the (R^*, S^*) inventory policy using the two-state closed-form solutions (Chapter 3) and Bernoulli closed-form solutions (see Appendix B.2). When solving for the (R^*, S^*) inventory policy with the Bernoulli closed-form solutions, we use the long-run probability that the supply chain is not disrupted as the input parameter (i.e., $p = \frac{\beta}{\alpha+\beta} = \frac{3}{4}$). We calculate the expected proportion of drug shortages per day using the left-hand side of the constraint in Formulation (3.11) with the (R^*, S^*) inventory policy and the corresponding $\alpha^{(R)}$ and $\beta^{(R)}$. Figure 3.8 presents the results.

With the disruption probability (α) fixed, we first observe that the (R^*, S^*) inventory policies are not always equivalent for the two-state supply process and Bernoulli supply process. We conclude that R^* and S^* are dependent on the supply chain disruption process. We find that as the disruption probability (α) increases, the two systems have similar inventory policies and thus, similar performance. However, for small disruption probabilities, the (R^*, S^*) inventory policies can be significantly different leading to an increase in the expected proportion of drug shortages per day when using the Bernoulli closed-form solutions instead of the two-state closed-form solutions. Specifically, the two-state supply process model and Bernoulli supply process model with the base case disruption probability $\alpha = \frac{1}{90} \approx 0.01$ leads to demand not being satisfied 5% and 16% of the time, respectively (i.e., 3 times the number of shortages with the Bernoulli model). For waste, the two-state and Bernoulli closed-form solutions both ensure $S^* \leq eq$ which guarantees no waste in the system. Hence, there is no impact on waste when using the Bernoulli closed-form solutions instead of the two-state closed-form solutions.

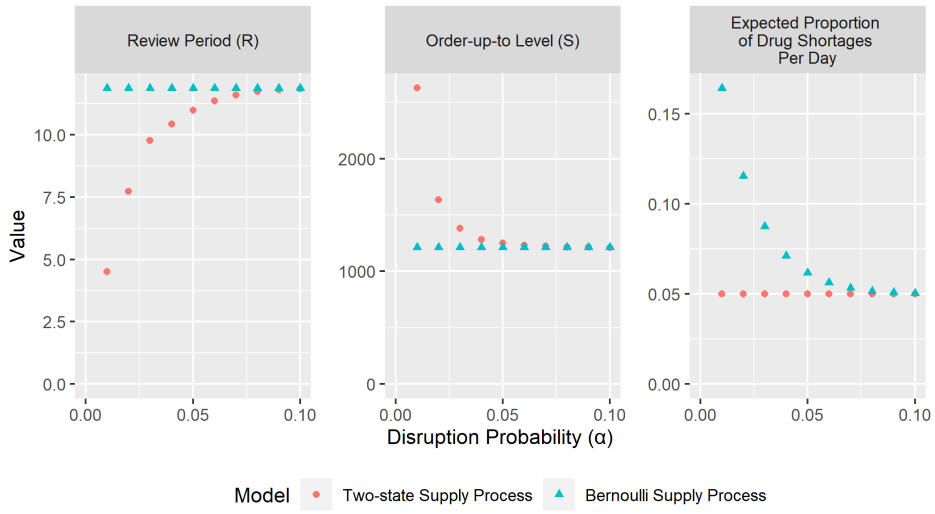

Figure 3.8: Results for two-state versus Bernoulli supply process where $\alpha = \frac{1}{3}\beta$.

3.6.6 Managerial Insights

We study the sensitivity of the model outputs (i.e., the expected cost per day [holding plus ordering], R, and S) with respect to changes in the input parameters: expiration lifetime (e), disruption probability (α), recovery probability (β), maximum proportion of demand not satisfied (γ), holding cost (h), and ordering cost (k). We present the key findings as (a)-(f) (see Appendix B.5 for the detailed analysis along with the corresponding figures).

(a) As the expiration lifetime (e) increases, the length of the review period (R) and the order-up-to level (S) increase. Also, there exists a lifetime e' beyond which it becomes optimal to always follow the non-perishable policies for $e \geq e'$. For $e < e'$, the expected cost per day increases.

(b) As the disruption probability (α) increases, the long-run probability that supply is disrupted increases leading to an increasing expected cost per day, decreasing length of the review period (R), and increasing order-up-to level (S).

(c) As the recovery probability (β) increases, the long-run probability that supply is disrupted decreases leading to a decreasing expected cost per day, increasing length of the review period (R), and decreasing order-up-to level (S).

(d) As the maximum proportion of demand not satisfied (γ) increases, the constraint in the model becomes less restrictive resulting in a decreasing expected cost per day, increasing length of the review period (R), and decreasing order-up-to level (S).

(e) As the holding cost (h) increases, the expected cost per day increases, the length of the review period (R) decreases, and the order-up-to level (S) decreases.

(f) As the ordering cost (k) increases, the expected cost per day increases, the length of the review period (R) increases, and the order-up-to level (S) increases.

3.7 Conclusion

A key finding from this chapter is that it is critical to account for perishability and supply chain disruptions simultaneously, especially in hospital pharmacy inventory systems where shortages can be life-threatening. In particular, we show that following a non-perishable (R,S) inventory policy with supply chain disruptions can lead to unnecessary waste and even worse, a decrease in the expected proportion of demand satisfied in comparison to a perishable (R,S) inventory policy with supply chain disruptions (in Section 3.6.2). This implies that an inventory manager cannot simply choose a supply chain disruption model (without perishability) to be 'safe'. This chapter also demonstrates that it is crucial to account for the duration of and time between supply chain disruptions, as the long-run probability that the supply chain is not disrupted is insufficient for inventory decision-making (in Section 3.6.3). We illustrate the consequences of ignoring the duration of and time between supply chain disruptions by simply modeling the inventory system as a Bernoulli supply process (in Section 3.6.5). We find that making this assumption has no impact on waste, but increases the expected proportion of drug shortages per day, especially for small disruption probabilities holding the long-run probability that the supply chain is not disrupted fixed. This shows that a simple model (i.e., more condensed closed-form solutions) that does not accurately depict the supply chain disruption process (i.e., Bernoulli supply process instead of a two-state supply process) comes at the cost of model performance.

It is worth noting how the closed-form model in this chapter may be implemented in our hospital pharmacy setting of interest. As a first step, we would work closely with our hospital pharmacy collaborators at the University of Michigan to conduct a pilot study of our closed-form model with a subset of drugs (e.g., 503B drugs). For the pilot study, the closed-form model could be coded in a programming language (e.g., R, Python), and ideally, this code would be embedded in an application that has a "user-friendly" front end (e.g., Excel). The front end would require a hospital pharmacy manager to enter the proportion of demand not satisfied for a particular drug (or category of drugs) and the other input parameters (i.e., demand, costs, expiration lifetime, supply chain disruption parameters) would be directly computed/retrieved on the back end from databases. A hospital pharmacy manager would interact with this front end, the application would output the optimal (R,S) inventory policy, and the (R,S) inventory policy would be manually entered into the

ordering system (e.g., BD Logistics). We would proceed this way with our pilot study because in current practice, it is challenging to integrate applications directly into the ordering system. However, given the pilot study is successful, we could foresee working with inventory automation vendors to embed the application into the ordering system to eliminate the manual (R,S) entry. Furthermore, if the pilot study is successful, we could foresee extending the implementation of the closed-form model to all drugs at the University of Michigan, other hospitals in the United States, and worldwide.

For perishability, we assume that the expiration lifetime of the product is deterministic and begins when it arrives at the inventory system. Future research can study stochastic expiration lifetimes (e.g., exponential lifetimes) and/or stochastic arrival lifetimes (i.e., the remaining lifetime of the product when it arrives at the inventory system is stochastic). Another avenue for future research is considering the capacity of the inventory system. For example, a hospital pharmacy may have limited storage to hold a particular drug (e.g., a medicine refrigerator). Constraining the amount of inventory held may be appropriate in this setting, and we leave this for future research.

The easy to implement closed-form solutions in this chapter provide the inventory policy quickly for non-perishable and perishable lost-sales (R,S) periodic review inventory systems with supply chain disruptions. This can help policy and decision makers avoid consequences of perishability, like waste, and supply chain disruptions, like shortages.

Acknowledgments for Chapter 3

We would like to give a special thanks to Christine Betts (University of Michigan Health Pharmacy Services) for providing access to the Fentanyl daily demand data. The research team would also like to thank the two anonymous reviewers, associate editor, and area editor at *INFOR: Information Systems and Operational Research* for their feedback on the chapter.

CHAPTER 4

The Benefits (or Detriments) of Adapting to Demand Disruptions in a Hospital Pharmacy with Supply Chain Disruptions

4.1 Introduction

Hospital pharmacy managers are responsible for determining how much inventory to have on-hand and when to place orders. Key aspects that need consideration when making these inventory decisions are supply chain disruptions and demand disruptions. A supply chain disruption is a random amount of time such that the hospital pharmacy is unable to receive a particular drug (Snyder et al., 2016). The resiliency of pharmaceutical supply chains is an active area of concern (NASEM, 2022). Insufficient inventory during these supply chain disruptions leads to drug shortages which can increase costs, increase drug errors, and decrease the quality of care (Phuong et al., 2019). Shortages in a food inventory system may cause you to change your recipe and shortages in a concert inventory system may cause you to miss your favorite band. However, drug shortages in a hospital pharmacy inventory system can cause significant drawbacks such as patients experiencing denied/delayed care (e.g., cancelled surgeries) or being provided with sub-optimal treatment (e.g., substitute drugs with less efficacy or more side effects). A simple solution is stocking more inventory than is needed, but holding too much inventory leads to waste as drugs have a finite shelf life.

Balancing the drug shortage versus drug waste trade-off is further complicated by demand disruptions. Demand disruptions occur when the demand mean is different from "normal/baseline" (i.e., it increases or decreases) for a random or extended amount of time. Demand disruptions were prevalent during the Covid-19 pandemic (e.g., surges in demand to treat Covid-19 patients and declines in demand for other drugs due to the cancellation of elective surgeries). Failing to adapt to disruptions that increase or decrease demand often leads to drug shortages and drug waste, respectively. However, drugs differ in their shortage-waste weighting (i.e., concern for shortages versus

concern for waste) and supply chain disruption profile (i.e., duration of and time between supply chain disruptions). We are interested in studying how these differences influence the benefits (or detriments) of adapting to demand disruptions.

A drug's shortage-waste weighting defines the shortage concern versus waste concern for a particular drug of interest. For the supply chain disruption profile, we analyze both the duration of supply chain disruptions and time between supply chain disruptions. With the varying drug shortage-waste weightings and supply chain disruption profiles, simulation modeling is a viable method to assess the performance of the inventory system given different inventory policies, such as adaptive inventory policies (i.e., inventory policies change over time), are in place.

To analyze the performance of adaptive inventory policies, we recognize that hospital pharmacy managers decide how frequently to place orders and how much inventory to have on-hand. To this end, we develop an adaptive lost-sales (R,S) periodic review inventory system where R denotes the length of the review period (i.e., attempt to place an order every R days) and S denotes the order-up-to level (i.e., attempt to place an order up to S every R days). By adaptive, we imply that the (R,S) inventory policy changes over time to reflect the shifts in the demand mean. We design the adaptive inventory system such that it (1) solves for the (R,S) inventory policy in a hospital pharmacy with supply chain disruptions, (2) endogenously detects when the inventory policy needs to be updated due to a demand disruption, and (3) appropriately updates the (R,S) inventory policy. Contracts, logistics, and resources can make it challenging to update the (R,S) inventory policy very frequently. We use the drug's shortage-waste weighting, the change in the expected proportion of drug shortages per day, and the change in the expected proportion of drugs wasted per day to support when the (R,S) inventory policy needs to be updated. Furthermore, we recognize that hospital pharmacy managers are often responsible for thousands of different drugs. To provide quick solution times and easy implementation, we create an adaptive inventory system that consists solely of closed-form expressions. Although, hospital pharmacy managers may only be able to adapt a limited number of (R,S) inventory policies at any given time. To give these managers a sense of which drugs are of most concern, we use the adaptive inventory system and provide a procedure to rank drugs based on multiple characteristics of a drug (e.g., expiration lifetime, shortage concern, waste concern, demand for the drug).

This chapter makes the following contributions:

1. We create an adaptive (R,S) periodic review inventory system that accounts for perishability, supply chain disruptions, and demand disruptions. All expressions are presented in closed-form providing quick solution times and easy implementation which is critical in a hospital pharmacy where managers are often responsible for thousands of different drugs.

2. The adaptive inventory system relies on the expected proportion of drugs wasted per day

in a (R,S) periodic review perishable inventory system with supply chain disruptions and stochastic demand. To the best of our knowledge, we are the first to present this value in closed-form.

3. We use the adaptive inventory system to create a ranking procedure. The ranking procedure provides a way of discerning which drugs are of most concern and illustrates which policies to update given that a limited number of inventory policies can be updated.

4. We leverage simulation modeling and perform an extensive numerical analysis using real-world demand data from the University of Michigan's Central Pharmacy to distinguish how a drug's shortage-waste weighting and supply chain disruption profile influence the benefits (or detriments) of adapting to demand disruptions.

The remainder of the chapter is organized as follows: Section 4.2 provides literature relevant to this research. Section 4.3 presents the adaptive inventory system and the ranking procedure. Section 4.4 presents the simulation models of the multiple inventory systems we consider in this chapter. In Section 4.5, we conduct a numerical analysis using real-world demand data from the University of Michigan's Central Pharmacy. Section 4.6 closes the chapter and provides future research directions.

4.2 Literature Review

Disruptions in a supply chain, whether from demand, supply, or transportation, are not a new problem. Paul et al. (2016), Snyder et al. (2016), Shen & Li (2017), and Ivanov et al. (2017) provide insightful review articles on the topic. To demonstrate how this chapter expands upon past research, we present literature in three areas: perishable inventory models without disruptions, perishable inventory models with disruptions, and adaptive inventory models. We close this section by using past literature to motivate adaptive inventory policies in a hospital pharmacy setting.

In Table 4.1, we present (a) perishable inventory models without disruptions, (b) perishable inventory models with disruptions, (c) adaptive inventory models, and (d) this research. For each paper, we characterize the research by perishability (NP: non-perishable, P: perishable), demand (D: deterministic, S: stochastic, DD: demand disruptions), supply (SSC: stochastic supply capacity, SLT: stochastic lead time, SCD: supply chain disruptions), adaptive (i.e., inventory policy/control parameters change over time or research uses thresholds to incorporate dynamic decision-making), and methodology. We note that some perishable inventory models have NP when characterizing their perishability (e.g., Zhang et al., 2014). This implies that the authors consider a perishable inventory system, but the authors do not capture perishability in their models (i.e., the authors do not capture the finite lifetime of the product).

Table 4.1 Summary of relevant literature.

Paper	NP/P	Demand	Supply	Adaptive	Methodology
(a) Perishable Inventory Models without Disruptions					
Zhang et al. (2014)	NP	D			Simulation-optimization
Little & Coughlan (2008)	NP	S			Constraint program
Neve & Schmidt (2022)	NP	S			Cost/service level optimization
Dillon et al. (2017)	P	S			Stochastic program
Kara & Dogan (2018)	P	S			Reinforcement learning
Rajendran & Srinivas (2020)	P	S			Stochastic program
Syawal & Alfares (2020)	P	S			Simulation-optimization
Xu & Szmerekovsky (2022)	P	S			Stochastic program
Li et al. (2018)	P	D	SLT		Nonlinear program and closed-form
Franco & Alfonso-Lizarazo (2020)	P	S	SLT		Stochastic program
Nguyen & Chen (2019)	P	S	SSC		Stochastic program
Nguyen & Chen (2022)	P	S	SSC		Stochastic program
(b) Perishable Inventory Models with Disruptions					
Atan & Rousseau (2016)	P	D	SCD		Closed-form
Chapter 3	P	D	SCD		Closed-form
Azghandi (2019)	NP	S	SCD		Simulation model
Saedi et al. (2016)	P	S	SCD		Continuous time Markov chain
Chapter 2	P	S	SCD		Simulation-optimization
He & Wang (2012)	P	DD			Analytical model
Rana et al. (2021)	P	DD			Analytical model
Uthayakumar & Karuppasamy (2019)	P	DD			Economic order quantity model
(c) Adaptive Inventory Models					
Eilon & Elmaleh (1970)	NP	S	SLT	✓	Forecasting procedure
Peterson et al. (1972)	NP	S		✓	Heuristics for dynamic program
Kim et al. (2005)	NP	S		✓	Reinforcement learning
Schmitt et al. (2017)	NP	S	SCD	✓	Simulation model
Li et al. (2016)	P	S		✓	Dynamic program
(d) This Research					
Chapter 4	P	DD	SCD	✓	Closed-form framework

NP: non-perishable, P: perishable, other abbreviations defined in Section 4.2

In comparison to the literature, we would like to motivate adaptive inventory policies and the importance of assessing the performance of the inventory policies for a hospital pharmacy inventory system. From a hospital pharmacy perspective, Azghandi (2019) creates a simulation model to study a multi-echelon pharmaceutical supply chain faced with supply chain disruptions and finds that adaptive inventory policies (i.e., inventory policies change over time) may help decrease drug shortages and drug costs for an inventory system with supply chain disruptions. However, demand disruptions are not taken into consideration. On the contrary, our research incorporates supply chain disruptions when solving for the optimal inventory policy and adapts the inventory policy over time to account for demand disruptions to overcome drug shortages and drug waste.

From a disruption perspective, Snyder et al. (2016) discuss the need for integrating proactive (i.e., guard against future uncertainties) and reactive (i.e., implemented when unexpected events occur) strategies when overcoming disruptions (Ivanov et al., 2017). In this chapter, we account for supply chain disruptions when solving for the optimal (R^*, S^*) inventory policy [proactive] as well as endogenously detect demand disruptions and update the (R^*, S^*) inventory policy [reactive]

to tackle the drug shortage and drug waste challenges in a hospital pharmacy inventory system. Furthermore, by detecting demand disruptions and appropriately updating the (R^*, S^*) inventory policy, we are able to adapt to changing conditions and consequently, improve the viability of the inventory system (Ivanov, 2021).

Also, we leverage simulation modeling to assess the performance of multiple inventory systems, including adaptive inventory systems, for varying drug shortage-waste weightings (i.e., concern for shortages versus concern for waste) and supply chain disruption profiles (i.e., duration of and time between supply chain disruptions). Gebicki et al. (2014) encourages the use of the criticality of the drug (i.e., concern for shortages defined by a drug's shortage-waste weighting in this chapter) and availability of the drug (i.e., supply chain disruption profile in this chapter) in inventory decision-making to improve the outcomes of a hospital pharmacy. But, they do not study adaptive inventory policies.

4.3 Adaptive (R, S) Inventory System

We proceed to present the adaptive (R,S) inventory system. We start by presenting how to solve for a (R,S) inventory policy for a perishable inventory system with supply chain disruptions (in Section 4.3.1) and how to adapt this (R,S) inventory policy over time to respond to demand disruptions (in Section 4.3.2). We close this section by presenting a ranking procedure (in Section 4.3.3) which provides a way of discerning which drugs are of most concern and illustrates which inventory policies to update given that a limited number of inventory policies can be updated.

4.3.1 (R,S) Inventory Policies with Supply Chain Disruptions

We start by modeling a lost-sales (R,S) perishable inventory system with supply chain disruptions as done in Chapter 3. The closed-form expressions derived in Chapter 3 provide quick-to-solve and easy-to-implement (R,S) periodic review inventory policies. We simply refer to this model as the (R,S) model. We refer the reader to Chapter 3 for the full model and implementation details. In Table 4.2, we present the notation for the (R,S) model and then we proceed to present the model assumptions.

Table 4.2 Summary of the modeling notation for the (R,S) model (Chapter 3).

Notation	Description
Input Parameters	
k	Fixed ordering cost (i.e., for each order attempted; $k > 0$)
h	Holding cost per day per drug ($h > 0$)
q	Deterministic demand per day ($q > 0$)
e	Expiration lifetime in days ($e \geq 1$)
γ	Maximum proportion of drug shortages per day over the infinite horizon ($0 < \gamma < 1$)
$\alpha^{(R)}$	Supply chain disruption probability with respect to the length of the review period R ($0 < \alpha^{(R)} < 1$)
$\beta^{(R)}$	Supply chain recovery probability with respect to the length of the review period R ($0 < \beta^{(R)} < 1$)
Variables	
R	Length of the review period ($R \geq 1$)
S	Order-up-to level ($S \geq 0$)
m	Number of review periods S fully covers ($m = \lfloor \frac{S}{qR} \rfloor$; $m \geq 1$)

The (R,S) model minimizes the expected ordering and holding cost, while constraining the proportion of drug shortages per day to be at most γ over an infinite horizon. The constraint on the proportion of drug shortages per day is tight with the optimal (R^*,S^*) inventory policy (i.e., the proportion of drug shortages per day is γ over an infinite horizon). The model also enforces zero waste in the inventory system by always ensuring $S^* \leq eq$ where e is the expiration lifetime in days and q is the deterministic daily demand. With this upper bound on S^*, there may be instances where the maximum proportion of drug shortages constraint cannot be satisfied. When the maximum proportion of drug shortages constraint cannot be satisfied, the optimal inventory policy is always $(R^* = 1, S^* = eq)$ (i.e., the smallest R^* and the largest S^*) as this maximizes the expected inventory on-hand.

The (R,S) model assumes that the lead time is zero which is consistent with hospital pharmacies that tend to see very small lead times. The Central Pharmacy at the University of Michigan experiences lead times of 36-72 hours. Also, the model assumes orders can only be attempted at times that are integer multiples of R (regardless of the not disrupted versus disrupted status of the supply chain between review periods) and that orders are only successfully placed when the supply chain is not disrupted. Given these assumptions, we can simply refer to inventory position as inventory on-hand. When applying the (R,S) model in this chapter, we assume that we have full and accurate knowledge of the inventory on-hand. Furthermore, we note that supply chain disruptions often last weeks, months, or years (Tucker et al., 2020a) which is much longer than the typical time between orders (i.e., R).

The (R,S) model assumes that supply chain disruptions follow a two-state supply process which is consistent with pharmaceutical supply chains (Tucker et al., 2020b). A two-state supply process is modeled as a two-state discrete time Markov chain where the time the supply chain is in the not disrupted state is a geometric random variable with parameter $\alpha^{(1)}$ (i.e., disruption probability; mean up time of $\frac{1}{\alpha^{(1)}}$) and the time the supply chain is in the disrupted state is a geometric random variable with parameter $\beta^{(1)}$ (i.e., recovery probability; mean down time of $\frac{1}{\beta^{(1)}}$) (Snyder et al.,

2016). A two-state supply process accounts for the duration of and time between supply chain disruptions. For perishability, the (R,S) model assumes that the drugs have no quality decay (Khan et al., 2014) and have a deterministic lifetime that starts when the drug arrives at the pharmacy. Also, the model assumes first-in-first-out protocols are in place which is consistent with practice at the University of Michigan's Central Pharmacy. For demand, the (R,S) model assumes that the daily demand, $q > 0$, is deterministic and static. For emphasis, our hospital pharmacy setting of interest does not have deterministic or static demand. In Section 4.3.2, we discuss how to use the (R,S) model in the adaptive framework to account for the real-world hospital pharmacy system that has stochastic and variable demand.

4.3.1.1 Supply Chain Disruption Parameters in the $(\mathbf{R},\mathbf{S})$ Model

In the (R,S) model, $\alpha^{(R)}$ and $\beta^{(R)}$ represent the supply chain disruption and recovery probability, respectively, for a review period of length R days. An important observation is that these values depend on the length of the review period R (i.e., a decision variable). Chapter 3 illustrates how to use the closed-form expressions to solve for the optimal (R^*,S^*) with only $\alpha^{(1)}$ and $\beta^{(1)}$ as input by leveraging $\mathbf{P}^{(i)}$ which is the i-step transition probability matrix (see Equation (4.1); Kulkarni, 2011).

$$\mathbf{P}^{(i)} = \begin{pmatrix} \cdot & \alpha^{(i)} \\ \beta^{(i)} & \cdot \end{pmatrix} = \begin{pmatrix} (1-\alpha^{(1)}) & \alpha^{(1)} \\ \beta^{(1)} & (1-\beta^{(1)}) \end{pmatrix}^{i} \tag{4.1}$$

In practice, we can find the values of $\alpha^{(1)}$ and $\beta^{(1)}$ by answering the following questions:

1. What proportion of the time is the drug short? Define this value as Q_1 ($Q_1 \in (0,1)$; 0 implies never short and 1 implies always short).

2. If the drug goes short, how long do you think the shortage will last (in days)? Define this value as Q_2 ($Q_2 > 1$).

With $Q_2 > \frac{Q_1}{1-Q_1}$ (due to the (R,S) model requirement of $\alpha^{(R)} \leq \alpha^{(1)} < 1$; see Table 4.2) and $Q_2 > 1$ (due to the (R,S) model requirement of $\beta^{(R)} \leq \beta^{(1)} < 1$; see Table 4.2), we can solve for $\alpha^{(1)}$ and $\beta^{(1)}$ using the relation that $Q_1 = \frac{\alpha^{(1)}}{\alpha^{(1)}+\beta^{(1)}}$ and $Q_2 = \frac{1}{\beta^{(1)}}$. Here, Q_1 corresponds to the long-run probability that the supply chain is disrupted (Snyder & Shen, 2019). Using these equations, we have the results presented in Equations (4.2)-(4.3).

$$\alpha^{(1)} = \frac{Q_1}{Q_2(1-Q_1)} \tag{4.2}$$

$$\beta^{(1)} = \frac{1}{Q_2} \tag{4.3}$$

Through discussions with our hospital pharmacy collaborators, the value of Q_2 is very difficult to quantify in practice, but it is an important parameter as it defines $\alpha^{(1)}$ and $\beta^{(1)}$ in Equation (4.2) and Equation (4.3), respectively. With this, we provide the expression for the ratio of the expected proportion of drug shortages per day given a value Q_2 and the true value Q_2^* in Appendix C.1.1.

4.3.2 Adapting the (R,S) Inventory Policy Over Time

We want to emphasize that in practice, contracts, logistics, and resources can make it challenging to update the (R,S) inventory policy very frequently (e.g., every day/week/month). A more common approach is updating the inventory policy every 3 or 6 months (i.e., update every B days where $B \in \{90, 180\}$) where the average demand from the past B days is used to update the (R,S) inventory policy. We refer to this approach as a benchmark inventory system. Unlike the benchmark inventory system that updates the inventory policy on long fixed intervals, we create an adaptive inventory system that endogenously detects when the (R,S) inventory policy needs to be updated at any point in time. We endogenously detect a change in the inventory policy is needed using a shortage threshold δ_s and waste threshold δ_w. These thresholds represent the change in the proportion of drug shortages per day and drugs wasted per day, respectively, that the inventory system is willing to tolerate. Specifically, these thresholds guard against shortages when an increasing demand disruption and waste when a decreasing demand disruption occur, respectively. When considering demand, as mentioned in Section 4.3.1, the (R,S) model assumes that demand is deterministic and static. However, our hospital pharmacy system of interest has stochastic and variable demand. For the deterministic aspect of the (R,S) model, we replace the deterministic daily demand (q) with the expected daily demand ($\bar{q}_{current}$) when solving for the (R,S) inventory policy; the term *current* implies that this expected daily demand defines the current (R,S) inventory policy. For the static aspect of the (R,S) model, we update this (R,S) inventory policy using a new expected daily demand when we endogenously detect a demand disruption.

We define the additional notation necessary for the adaptive inventory system in Table 4.3. We also present a depiction of how the adaptive demand parameters and estimates relate to one another in Figure 4.1. We then proceed to define the expected proportion of drug shortages and drugs wasted per day for the (R,S) inventory system. We also explain how to endogenously detect when the inventory policy needs to be updated.

Table 4.3 Summary of the modeling notation for the adaptive inventory system.

Notation	Description
Thresholds	
δ_s	Change in the proportion of drug shortages per day that signals a change in the (R,S) inventory policy when exceeded ($0 < \delta_s < 1$; input parameter)
δ_w	Change in the proportion of drugs wasted per day that signals a change in the (R,S) inventory policy when exceeded ($0 < \delta_w < 1$; input parameter)
Adaptive	
N	Number of past daily demand observations to consider for the adaptive inventory system daily demand estimates (input parameter)
q_t	Daily demand observed on day t (real-world hospital pharmacy data)
$\bar{q}_{current}$	Expected daily demand used for the current (R,S) policy
$\bar{q}_{new}$	Expected daily demand calculated using the average of the most recent N daily demand observations
σ_{new}	Standard deviation of daily demand calculated using the most recent N daily demand observations
$P_{short\|(\bar{q},R,S)}$	Expected proportion of drug shortages per day when following a (R,S) inventory policy for an inventory system that has an expected daily demand of $\bar{q}$
$P_{waste\|(\bar{q},\sigma,R,S)}$	Expected proportion of drugs wasted per day when following a (R,S) inventory policy for an inventory system that has an expected daily demand of $\bar{q}$ and standard deviation of daily demand σ
Benchmark	
B	Number of days the inventory system follows the same (R,S) inventory policy where the average of the last B days is used to update the inventory policy (input parameter)

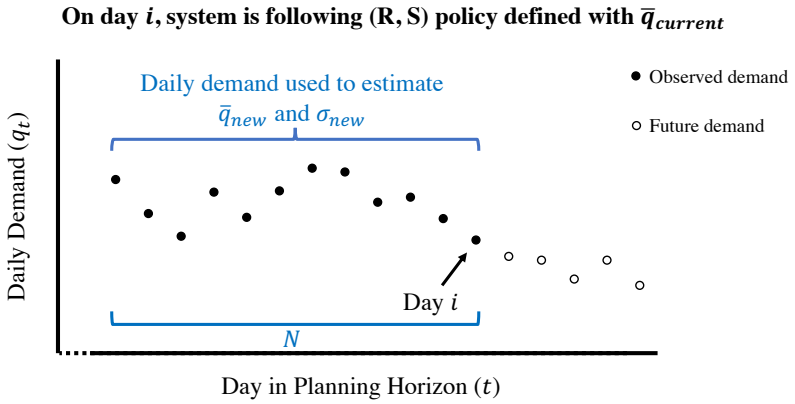

Figure 4.1: Depiction of the adaptive demand parameters and estimates. $\bar{q}_{new}$ and σ_{new} are calculated at the end of day i after demand has been observed.

4.3.2.1 Expected Proportion of Drug Shortages and Drugs Wasted Per Day

We proceed to present the closed-form expressions for the expected proportion of drug shortages per day and drugs wasted per day. When calculating these proportions, we assume that the daily demand is normally distributed. We make this assumption because of the quantile-quantile plots (i.e., QQ-plots) with the real-world unique training daily demand observations.

Expected Proportion of Drug Shortages Per Day

Using the (R,S) model (see Section 4.3.1), Equation (4.4) represents the expected proportion of drug shortages per day when following a (R,S) inventory policy for an inventory system that has an

expected daily demand of $\bar{q}$ (in Appendix C.2.1). Proportion is measured relative to the expected daily demand $\bar{q}$. From Equation (4.4), the expected proportion of drug shortages per day depends on the length of the review period (R), the order-up-to level (S), the expected daily demand ($\bar{q}$), the disruption probability ($\alpha^{(R)}$), and the recovery probability ($\beta^{(R)}$). It is worth noting that the (R,S) policy is calculated using $\bar{q}_{current}$ which ensures that $m = \lfloor \frac{S}{\bar{q}_{current}R} \rfloor \geq 1$. However, with an expected daily demand of $\bar{q}$, the $m \geq 1$ requirement may no longer be satisfied which explains the two cases presented in Equation (4.4). Furthermore, Chapter 3 illustrates that stochastic demand that is normally distributed has a negligible impact on the expected proportion of drug shortages per day with the (R,S) model. The reason for this finding is that the (R,S) model incorporates supply chain disruptions which encourage the inventory system to hold extra inventory on-hand. As a result, there is a negligible impact on the expected proportion of drug shortages per day when demand is stochastic. Hence, we simply consider a deterministic expected daily demand $\bar{q}$.

(4.4)

$$P_{short|(\bar{q},R,S)} = \begin{cases} \frac{\alpha^{(R)}\beta^{(R)}(1-\beta^{(R)})^{m-1}}{(\alpha^{(R)}+\beta^{(R)})}\left(m+1-\frac{S}{\bar{q}R}\right) + \left(\frac{\alpha^{(R)}(1-\beta^{(R)})^{m}}{(\alpha^{(R)}+\beta^{(R)})}\right); \ m \geq 1 \text{ where } m = \lfloor \frac{S}{\bar{q}R} \rfloor \\ \\ \frac{\beta^{(R)}}{\alpha^{(R)}+\beta^{(R)}}\left(\frac{\bar{q}R-S}{\bar{q}R}\right) + \left(1-\frac{\beta^{(R)}}{\alpha^{(R)}+\beta^{(R)}}\right); \ m < 1 \text{ where } m = \lfloor \frac{S}{\bar{q}R} \rfloor \end{cases}$$

Expected Proportion of Drugs Wasted Per Day

Equation (4.5) represents the expected proportion of drugs wasted per day when following a (R,S) inventory policy for an inventory system that has an expected daily demand of $\bar{q}$ and standard deviation of daily demand σ (in Appendix C.2.2). Proportion is measured relative to the expected number of drugs ordered. To our knowledge, we are the first to present this value in closed-form. The expected proportion of drugs wasted per day in Equation (4.5) depends on the length of the review period (R), the order-up-to level (S), the expected daily demand ($\bar{q}$), the standard deviation of daily demand (σ), the disruption probability ($\alpha^{(R)}$; in π_j), the recovery probability ($\beta^{(R)}$; in π_j), and the expiration lifetime (e). Z denotes a standard normal random variable and Appendix C.2.2 provides the closed-form expressions for the summations including π_j. π_j is the probability that the supply chain is disrupted for exactly j consecutive review periods ($\pi_0 = \frac{\beta^{(R)}}{\alpha^{(R)}+\beta^{(R)}}$; $\pi_j = \frac{\alpha^{(R)}\beta^{(R)}}{(\alpha^{(R)}+\beta^{(R)})(1-\beta^{(R)})}(1-\beta^{(R)})^j, \ j \geq 1$). Equation (4.5) accounts for stochastic demand where we assume that the daily demand is independent and normally distributed with mean $\bar{q}$ and standard deviation σ. Chapter 3 illustrates that the expected proportion of drugs wasted with the (R,S) model is sensitive to stochastic demand that is normally distributed, especially for drugs with short

expiration lifetimes.

$$P_{waste|(\bar{q},\sigma,R,S)} = \tag{4.5}$$

$$\begin{cases} \frac{E_w}{S};\ \lceil \frac{e}{R} \rceil = 1 \\ \\ \frac{E_w}{(\lceil \frac{e}{R} \rceil R\bar{q}+E_w)\pi_0+\sum\limits_{j=\lceil \frac{e}{R} \rceil -1}^{\infty} \pi_j(S+R\bar{q}(\frac{(\lceil \frac{e}{R} \rceil -1)}{2}))};\ \lceil \frac{e}{R} \rceil = 2 \\ \\ \frac{E_w}{(\lceil \frac{e}{R} \rceil R\bar{q}+E_w)(\pi_0+\sum\limits_{j=1}^{\lceil \frac{e}{R} \rceil -2} \pi_j-\frac{1}{\lceil \frac{e}{R} \rceil}\sum\limits_{j=1}^{\lceil \frac{e}{R} \rceil -2} j\pi_j)+\frac{S}{\lceil \frac{e}{R} \rceil}\sum\limits_{j=1}^{\lceil \frac{e}{R} \rceil -2} j\pi_j+\frac{R\bar{q}}{\lceil \frac{e}{R} \rceil}\sum\limits_{j=1}^{\lceil \frac{e}{R} \rceil -2} j^2\pi_j+\sum\limits_{j=\lceil \frac{e}{R} \rceil -1}^{\infty} \pi_j(S+R\bar{q}(\frac{(\lceil \frac{e}{R} \rceil -1)}{2}))};\ \lceil \frac{e}{R} \rceil \geq 3 \end{cases}$$

$$\text{where } E_w = \left(S \cdot \Pr(Z < \frac{S-e\bar{q}}{\sqrt{e\sigma^2}}) - e\bar{q} - \frac{1}{2\pi}\Big(-e^{-\frac{(S-e\bar{q})^2}{2e\sigma^2}} + e^{-\frac{(-e\bar{q})^2}{2e\sigma^2}} \Big)\sqrt{e\sigma^2} \right)$$

4.3.2.2 Detecting When the Inventory Policy Needs to Be Updated

On any day t, we follow a particular (R,S) inventory policy which is defined with respect to an expected daily demand of $\bar{q}_{current}$. In the presence of increasing and/or decreasing demand disruptions, we may need to update the inventory policy to avoid excessive drug shortages and/or drug waste. The key idea is that we consider how a shift in the expected daily demand impacts the expected proportion of drug shortages per day and expected proportion of drugs wasted per day. For the shift in the expected daily demand, we estimate the new expected daily demand $\bar{q}_{new}$ by averaging the most recent N daily demand observations and we compare this to $\bar{q}_{current}$; the expected daily demand used to define the current (R,S) inventory policy. We detect a change in the inventory policy is necessary if (a) $\bar{q}_{new} \geq \bar{q}_{current}$ and the shift in the expected daily demand causes the change in the expected proportion of drug shortages per day to exceed δ_s or (b) $\bar{q}_{new} < \bar{q}_{current}$ and the shift in the expected daily demand causes the change in the expected proportion of drugs wasted per day to exceed δ_w. If (a) or (b) holds, we update the (R,S) inventory policy using the new expected daily demand $\bar{q}_{new}$. It is worth noting that if $\bar{q}_{new} < \bar{q}_{current}$, the estimated standard deviation of demand (i.e., σ_{new}) will influence the expected proportion of drugs wasted per day as illustrated in Equation (4.5). We illustrate an overview of the procedure in Figure 4.2.

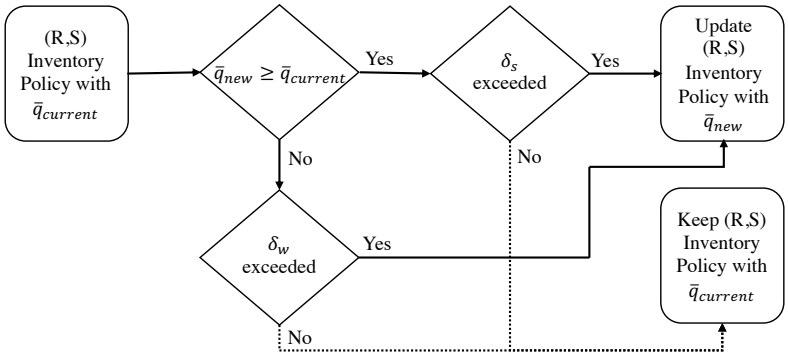

Figure 4.2: Overview of the adaptive inventory system.

We proceed to describe the daily demand estimates and expected proportions needed for the adaptive inventory system. We then formally define the conditions that detect when the inventory policy needs to be updated.

Daily Demand Estimates and Expected Proportions

We estimate the expected daily demand $\bar{q}_{new}$ by averaging the most recent N daily demand observations. We initially considered other prediction approaches (e.g., ARIMA), but we found that a simple average approach performed just as well or better. We hypothesize that an averaging approach performs well due to the variable and noisy real-world daily demand data. We also found that knowing everything about the future (i.e., knowing the future N daily demand observations exactly) had a small or negligible impact on the performance. For drugs with seasonal data over a period of daily demand observations (e.g., demand is dependent on the day of the week), we suggest selecting N such that it is a multiple of the seasonal period (e.g., $N \in \{7, 14, 21, 28, 35, 42, 49, 56\}$ days for weekly seasonality). For drugs that experience longer periods of seasonality (e.g., yearly), we note that the adaptive inventory system will endogenously detect changes in the demand over this longer period of time and appropriately update the (R,S) inventory policy. For the standard deviation of daily demand, we estimate the standard deviation of daily demand (i.e., σ_{new}) using the most recent N daily demand observations.

We first calculate the expected proportion of drug shortages per day when following a particular (R,S) inventory policy. We calculate this proportion for an inventory system that has an expected daily demand of $\bar{q}_{new}$ (i.e., $P_{short|(\bar{q}_{new},R,S)}$, see Equation (4.4)) which corresponds to the proportion estimate with the new demand mean. We also calculate this proportion for an inventory system that has an expected daily demand of $\bar{q}_{current}$ (i.e., $P_{short|(\bar{q}_{current},R,S)}$, see Equation (4.4)) which corresponds to the proportion estimate with the demand mean that defines the current (R,S) inventory policy. Next, we calculate the expected proportion of drugs wasted per day when

following a particular (R,S) inventory policy. We calculate this proportion for an inventory system that has an expected daily demand of $\bar{q}_{new}$ and standard deviation of daily demand σ_{new} (i.e., $P_{waste|(\bar{q}_{new},\sigma_{new},R,S)}$, see Equation (4.5)) which corresponds to the proportion estimate with the new demand mean and new standard deviation of daily demand. We also calculate this proportion for an inventory system that has an expected daily demand of $\bar{q}_{current}$ and standard deviation of daily demand σ_{new} (i.e., $P_{waste|(\bar{q}_{current},\sigma_{new},R,S)}$, see Equation (4.5)) which corresponds to the proportion estimate with the demand mean that defines the current (R,S) inventory policy and new standard deviation of daily demand.

Conditions that Detect When the Inventory Policy Needs to Be Updated

Given $\bar{q}_{new} \geq \bar{q}_{current}$, we have an increase or no change in the expected daily demand. In this case, we compare $P_{short|(\bar{q}_{new},R,S)} - P_{short|(\bar{q}_{current},R,S)}$ to the shortage threshold δ_s; the change in the proportion of drug shortages per day that signals a change in the (R,S) inventory policy when exceeded. δ_s provides a numerical value to quantify the shortage concern. A smaller δ_s is recommended for drugs that have a high shortage concern. We note that when $\bar{q}_{new} = \bar{q}_{current}$, $P_{short|(\bar{q}_{new},R,S)} - P_{short|(\bar{q}_{current},R,S)}$ will equal zero, but we include this scenario to simply break the analysis into two cases.

Given $\bar{q}_{new} < \bar{q}_{current}$, we have a decrease in the expected daily demand. In this case, we compare $P_{waste|(\bar{q}_{new},\sigma_{new},R,S)} - P_{waste|(\bar{q}_{current},\sigma_{new},R,S)}$ to the waste threshold δ_w; the change in the proportion of drugs wasted per day that signals a change in the (R,S) inventory policy when exceeded. δ_w provides a numerical value to quantify the waste concern. A smaller δ_w is recommended for drugs that have a high waste concern.

Formally, the adaptive inventory system detects that the inventory policy needs to be updated if either condition in Equation (4.6) is satisfied. We provide a detailed discussion on the selection of the input parameters δ_s and δ_w in Section 4.5.2.

$$\begin{cases} P_{short|(\bar{q}_{new},R,S)} - P_{short|(\bar{q}_{current},R,S)} > \delta_s; & \bar{q}_{new} \geq \bar{q}_{current} \\ P_{waste|(\bar{q}_{new},\sigma_{new},R,S)} - P_{waste|(\bar{q}_{current},\sigma_{new},R,S)} > \delta_w; & \bar{q}_{new} < \bar{q}_{current} \end{cases} \tag{4.6}$$

Given we detect that the (R,S) inventory policy needs to be updated, we set $\bar{q}_{current} = \bar{q}_{new}$ and solve for the new (R,S) inventory policy using the (R,S) model.

4.3.3 Ranking Drugs

Managers at the University of Michigan's Central Pharmacy are responsible for making inventory decisions for 2,500+ drugs. We create a ranking procedure to address the following questions:

1. Out of the 2,500+ drugs, which (if any) drugs should the hospital pharmacy be most concerned about?

2. If only a limited number of inventory policies can be updated, which drugs (if any) should the hospital pharmacy focus on?

We present a ranking procedure that depends on the:

(a) storage/holding cost for the drug (h)

(b) expiration lifetime of the drug (e)

(c) supply chain disruption profile of the drug ($\alpha^{(R)}$ and $\beta^{(R)}$)

(d) shortage concern (δ_s; change in the proportion of drug shortages per day that the decision-maker is willing to tolerate)

(e) waste concern (δ_w; change in the proportion of drugs wasted per day that the decision-maker is willing to tolerate)

(f) demand for the drug ($\bar{q}_{current}$ and $\bar{q}_{new}$)

(g) demand variability of the drug (σ_{new})

For the ranking procedure, we observe that the adaptive inventory system keeps a record of the change in the expected proportion of drug shortages per day and the change in the expected proportion of drugs wasted per day (see Equation (4.6)). We are interested in ranking the drugs in order of decreasing concern. We introduce a proportion exceedance metric denoted P_{metric} which measures how much the shortage threshold δ_s is exceeded given $\bar{q}_{new} \geq \bar{q}_{current}$ and waste threshold δ_w is exceeded given $\bar{q}_{new} < \bar{q}_{current}$. We define P_{metric} in Equation (4.7) where it is important to note that the adaptive inventory system only indicates that the (R,S) inventory policy needs to be updated when $P_{metric} > 0$.

$$P_{metric} = \begin{cases} \max\{0, P_{short|(\bar{q}_{new},R,S)} - P_{short|(\bar{q}_{current},R,S)} - \delta_s\}; & \bar{q}_{new} \geq \bar{q}_{current} \\ \max\{0, P_{waste|(\bar{q}_{new},\sigma_{new},R,S)} - P_{waste|(\bar{q}_{current},\sigma_{new},R,S)} - \delta_w\}; & \bar{q}_{new} < \bar{q}_{current} \end{cases} \tag{4.7}$$

P_{metric} encompasses characteristics (a)-(g). We rank the drugs in order of decreasing concern by sorting the drugs from largest to smallest using P_{metric}. There are multiple ways that decision-makers can implement the ranking procedure in practice (e.g., update drugs based on P_{metric} every day, update drugs based on the average value of P_{metric} over a fixed interval of days). In Section 4.5.5, we illustrate one way of implementing the ranking procedure in practice and we analyze the results.

4.4 Simulation Models

With the varying drug shortage-waste weightings and supply chain disruption profiles, we use simulation modeling to assess the performance of multiple inventory systems. We create simulation models of four inventory systems: (A) Adaptive Inventory System (in Section 4.4.1), (B) Adaptive with Buyback Inventory System (in Section 4.4.2), (C) Benchmark Inventory System (in Section 4.4.3), and (D) Static Inventory System (in Section 4.4.3). We let t denote the day in the planning horizon where the decision-maker can select the initialization of t. We let r denote the number of days remaining in the review period until the next order is attempted, b denote the number of days since the (R,S) inventory policy has been updated, I_i denote the inventory on-hand with a lifetime remaining of i days ($i = 1,...,e$), and I_{tot} denote the total inventory on-hand. To assess the performance of the system, we record the number short each day t (i.e., $short_t$), the number wasted each day t (i.e., $waste_t$), the number held each day t (i.e., h_t), and the number successfully ordered each day t (i.e., o_t).

4.4.1 (A) Adaptive Inventory System

Ignoring the bold text, Figure 4.3 provides a step-by-step description of the simulation model representing a periodic review inventory system with adaptive inventory policies. To initialize the model, we (a) define $\bar{q}_{current}$ using the first B daily demand observations and solve for the optimal (R^*,S^*) inventory policy. We then (b) initialize $r = R^*$ (i.e., there are R^* days remaining in the review period until an order should be attempted), $b = 0$ (i.e., it has been 0 days since the inventory policy has been updated), $I_i = 0 \ \forall i = 1,...,e$ (i.e., the inventory system has zero inventory on-hand across all lifetimes), $I_{tot} = 0$ (i.e., the total inventory on-hand is zero), and $o_{t-1} = S^*$ (i.e., S^* drugs will arrive on day t).

After initializing the model, an iterative procedure begins. At the beginning of day t, (c) an order placed on day $t-1$ arrives at the hospital pharmacy (when supply is not disrupted) recalling the (R,S) model zero lead time assumption. The inventory levels are appropriately updated. Then, we (d) observe the real-world daily demand (i.e., q_t). At the end of day t, we (e) record the number of drug shortages on day t (i.e., $short_t$), record the number of drugs wasted on day t (i.e., $waste_t$), and discard these wasted drugs. Then, we (f) subtract one day from the number of days remaining in the review period (i.e., $r = r-1$), add one day to the number of days since the inventory policy has been updated (i.e., $b = b+1$), and appropriately update the inventory levels. Then, (g) if $r = 0$ (i.e., there are zero days remaining in the review period), an order should be attempted. Thus, we proceed to step (h). Otherwise, no order attempt is necessary so we (n) record the number of drugs held (i.e., $h_t = I_{tot}$) and ordered (i.e., $o_t = 0$) on day t, and start the process back at step (d).

Given an order should be attempted, we (h) use the most recent N daily demand observations

to estimate the expected daily demand $\bar{q}_{new}$ and standard deviation of daily demand σ_{new}. Then, we (i) determine if the (R,S) inventory policy needs to be updated by seeing if either condition in Equation (4.6) is satisfied (i.e., $\bar{q}_{new} \geq \bar{q}_{current}$ & $P_{short|(\bar{q}_{new},R^*,S^*)} - P_{short|(\bar{q}_{current},R^*,S^*)} > \delta_s$ or $\bar{q}_{new} < \bar{q}_{current}$ & $P_{waste|(\bar{q}_{new},\sigma_{new},R^*,S^*)} - P_{waste|(\bar{q}_{current},\sigma_{new},R^*,S^*)} > \delta_w$). If the (R,S) inventory policy needs to be updated, we (j) update the inventory policy by setting $\bar{q}_{current} = \bar{q}_{new}$ and solve for the corresponding optimal (R^*,S^*) inventory policy. After step (j), we (k) set $b = 0$ to indicate that it has been zero days since the inventory policy has been updated. Finally, we (l) attempt to place an order based on the (R^*,S^*) inventory policy and set $r = R^*$. When supply is not disrupted, the order is successful. We (m) record the number of drugs held (i.e., $h_t = I_{tot}$) and ordered (i.e., $o_t = S^* - I_{tot}$) on day t, and start the process back at step (c). When supply is disrupted, the order is unsuccessful. We (n) record the number of drugs held (i.e., $h_t = I_{tot}$) and ordered (i.e., $o_t = 0$) on day t, and start the process back at step (d).

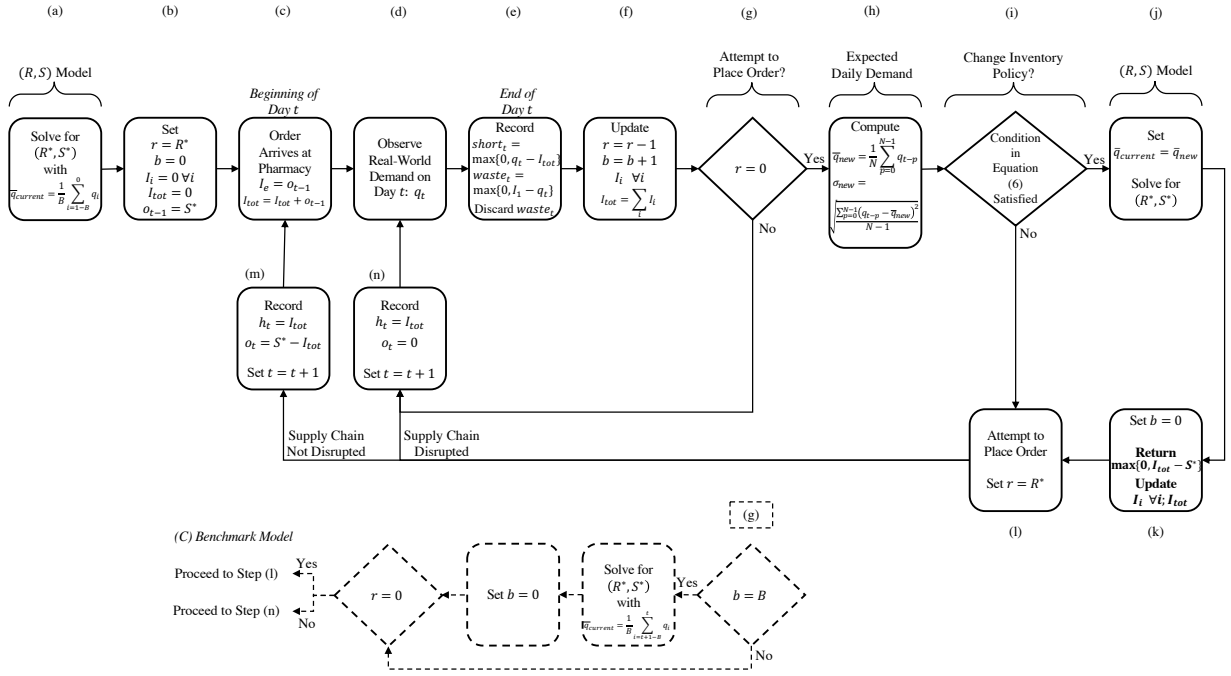

Figure 4.3: Simulation model road map for the (A) Adaptive, (B) Adaptive with Buyback, and (C) Benchmark inventory systems. For (A), follow steps (a)-(n) and omit the bold text. For (B), follow steps (a)-(n) and implement the bold text to incorporate buyback. For (C), starting at step (g), replace the process with the dashed shapes/lines.

4.4.2 (B) Adaptive Inventory System with Buyback

Through discussions with our hospital pharmacy collaborators at the University of Michigan's Central Pharmacy, some contracts allow the pharmacy to return drugs if the hospital pharmacy has

too much inventory on-hand. With these buyback programs in mind, we create a simulation model of a periodic review inventory system with adaptive inventory policies where the hospital pharmacy can return drugs to the supplier given the inventory on-hand exceeds the optimal order-up-to level S^*. Taking notice to the bold text, Figure 4.3 presents the step-by-step description of the simulation model. We add an extra operation (see bold text) in step (k) where we allow the hospital pharmacy to return all drugs that cause the inventory on-hand to exceed the optimal order-up-to level S^*. We assume that the newest inventory (i.e., longest remaining lifetime) is returned and that the hospital pharmacy receives full compensation for the returned drugs.

4.4.3 (C) Benchmark Inventory System and (D) Static Inventory System

The simulation model for the benchmark inventory system and static inventory system is the same as the adaptive inventory system (i.e., (A) Adaptive) except the (R,S) inventory policy is only updated when $b = B$ (see dashed shapes/lines in Figure 4.3) and the (R,S) inventory policy is never updated, respectively.

4.5 Numerical Analysis

We use daily demand data from the University of Michigan's Central Pharmacy (October 2019-November 2021). This two year period captures demand before the Covid-19 pandemic and fluctuations in demand during the Covid-19 pandemic. We analyze how a drug's shortage-waste weighting and supply chain disruption profile influence the benefits (or detriments) of adapting to demand disruptions. We present the data (in Section 4.5.1), shortage-waste weightings (in Section 4.5.2), and input parameters (in Section 4.5.3). We proceed to study the benefits (or detriments) of adapting to demand disruptions in Section 4.5.4. Then, we analyze the ranking procedure (see Section 4.3.3) in Section 4.5.5. Throughout the numerical analysis, we denote days $t \leq 0$ as the training horizon and days $t > 0$ as the testing horizon. Also, the (R,S) model treats R and S as continuous decision variables. We take a conservative approach by rounding R down to the nearest whole number and S up to the nearest whole number. For all other computations that require an integer value, we round to the nearest whole number (e.g., $\bar{q}_{current}$). We also ensure that all daily demand values are positive as the (R,S) model requires a positive daily demand ($\bar{q}_{current} > 0$) and demand is always non-negative in practice.

4.5.1 Real-world Data

The Central Pharmacy at the University of Michigan manages and keeps records for 2,500+ drugs (e.g., surgical, cancer, daily care for inpatients). Except when analyzing the ranking procedure

(in Section 4.5.5), we focus on 503B drugs which are pre-compounded drugs that arrive to the pharmacy in ready-to-use presentations (Jones, 2020). If a 503B drug experiences a shortage, hospital pharmacies will often substitute the drug with the form that requires compounding before administration. This form of the drug requires additional pharmacy resources and has a very small expiration lifetime once compounded (e.g., 24 hours). If a 503B drug is wasted due to expiration, hospital pharmacies experience a higher waste cost because a 503B drug is often more expensive than the form that requires compounding before administration. These are additional reasons why it is critical to avoid drug shortages and drug waste for this class of drugs.

We start by analyzing two 503B drugs. The first drug is (a) Rocuronium 10 mg/1mL (Rocuronium) which is a paralyzing agent that is critical to have on-hand and is most often used for rapid sequence intubation. The second drug is (b) Labetalol 5mg/1mL (Labetalol) which is a critical drug used for blood pressure reduction for several indications. Figure 4.4 presents the weekly demand where the red horizontal lines denote the corresponding mean weekly demand that minimizes the sum of squared errors for the daily demand data. It is worth noting that the red horizontal lines for Rocuronium resemble an increasing demand disruption and the red horizontal lines for Labetalol resemble a decreasing demand disruption. Furthermore, for both Rocuronium and Labetalol, the two groups of daily demand observations corresponding to the red horizontal lines are statistically different at a 0.05 significance level when applying a non-parametric Mann Whitney U test to the two groups. The training horizon (i.e., $t \leq 0$) consists of the first 56 unique daily demand observations (i.e., 8 weeks) stacked four times to have 224 observations. Then, we only select observations $45 - 224$ to obtain 180 daily demand training observations (i.e., about 6 months). We replicate and stack these daily demand observations when constructing the training data set to have a training data set that is not impacted by the Covid-19 pandemic. Beyond Rocuronium and Labetalol, we consider four additional 503B drugs at the end of Section 4.5.4.

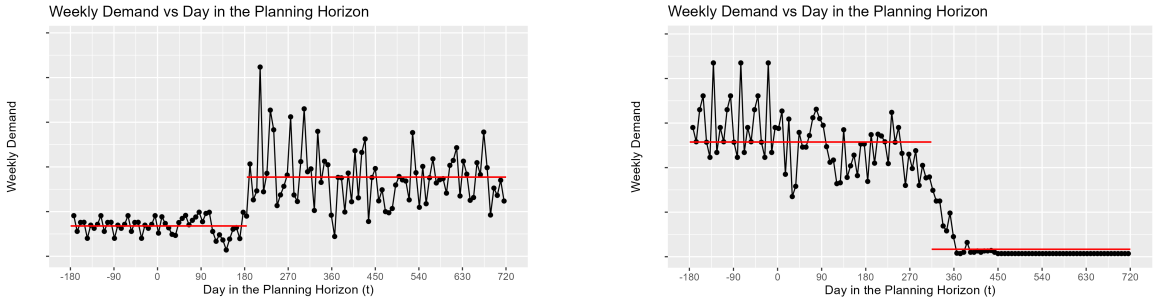

(a) Rocuronium.

(b) Labetalol.

Figure 4.4: Weekly demand versus day in the planning horizon. Numerical values are removed on the y-axis for data confidentiality.

4.5.2 Shortage-Waste Weighting

A drug's shortage-waste weighting defines the shortage concern (defined by δ_s) and waste concern (defined by δ_w) for a particular drug of interest. Working closely with our hospital pharmacy collaborators, we study a range of (δ_s, δ_w) values such that $\delta_s + \delta_w = 0.1$ where $\delta_s \in [0.01, 0.09]$ and $\delta_w \in [0.01, 0.09]$ to capture the trade-off between these two measures. In the following section, we select $\gamma = 0.05$. This implies that the (R,S) model ensures that shortages occur 5% of the time given the shortage constraint can be satisfied in the (R,S) model (see Section 4.3.1). Therefore, our range of $\delta_s \in [0.01, 0.09]$ values detects when shortages occur more than $6\% - 14\%$ of the time, respectively. It is important to note that a shift from 5% of the time to 6% of the time is a 20% increase in shortages. Also, the adaptive inventory system considers the percentage of drugs wasted with the mean demand used to define the current (R^*, S^*) inventory policy (i.e., $100 * P_{waste|(\bar{q}_{current}, \sigma_{new}, R^*, S^*)}\%$) and the new mean demand (i.e., $100 * P_{waste|(\bar{q}_{new}, \sigma_{new}, R^*, S^*)}\%$) given a standard deviation of daily demand σ_{new}. Therefore, our range of $\delta_w \in [0.01, 0.09]$ values detects when the difference in the percentage of drugs wasted is more than $1\% - 9\%$, respectively. Figure 4.5 illustrates a drug's shortage-waste weighting where we also provide drug examples in black *italics*. We note that when $(\delta_s, \delta_w) = (0.05, 0.05)$, the model is equally sensitive to drug shortages and drug waste.

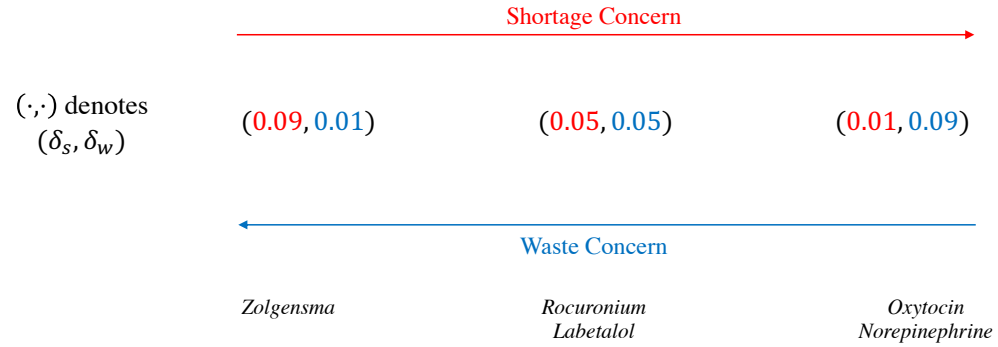

Figure 4.5: A drug's shortage-waste weighting. From left to right, the concern for shortages increases (i.e., δ_s decreases) and the concern for waste decreases (i.e., δ_w increases).

4.5.3 Model Input Parameters

We provide a summary of the input parameters for the numerical analysis in Table 4.4. For the (R,S) model, we consider an ordering cost (k) relative to drug price of 10 (i.e., $k = 10 \cdot$(drug price)). We consider a daily holding cost (h) relative to drug price of 0.001 (i.e., $h = 0.001 \cdot$(drug price); Jia & Zhao, 2017). Through discussions with our hospital pharmacy collaborators, Rocuronium and Labetalol have a wholesale price of about \$12 and \$7 per dose, respectively. We

consider an expiration lifetime of 90 days (i.e., $e = 90$) which is consistent with the class of 503B drugs. When solving for the optimal (R^*, S^*) policy, we constrain shortages to occur only 5% of the time (i.e., $\gamma = 0.05$). Also, we consider four supply chain disruption profiles, $(\alpha^{(1)}, \beta^{(1)}) = \{(\frac{1}{30}, \frac{1}{10}), (\frac{1}{90}, \frac{1}{30}), (\frac{1}{270}, \frac{1}{90}), (\frac{1}{810}, \frac{1}{270})\}$, where $\frac{1}{x}$ corresponds to an expected duration of x days. Our hospital pharmacy collaborators at the University of Michigan have observed a variety of supply chain disruption lengths in practice (e.g., 1-3 months and 8-9 months). These $(\alpha^{(1)}, \beta^{(1)})$ supply chain disruption profiles all have the same long-run probability that the supply chain is disrupted: $\frac{\alpha^{(1)}}{\alpha^{(1)}+\beta^{(1)}} = 0.25$.

Drug demand often depends on the day of the week, so we consider $N = 56$ past daily demand observations when estimating the expected daily demand and standard deviation of daily demand for the adaptive inventory system (i.e., 56 is divisible by 7). We provide additional details on the selection of N in Section 4.5.6. We consider a benchmark inventory system that follows the same inventory policy for $B = 90$ days (i.e., about 3 months).

For all simulation models, we initialize $\bar{q}_{current}$ using the first $B = 90$ training daily demand observations as shown in Section 4.4. However, to warm-up the simulation model, we replicate and stack the entire set of 180 training daily demand observations 4 times to have a total warm-up period of $4(180)$ days (i.e., ≈ 2 years). Following the warm-up period, we always consider a testing horizon of 720 days (i.e., ≈ 2 years). We first note that we consider a long warm-up period to ensure that the inventory on-hand is appropriately balanced across the potential lifetimes of the drug (i.e., not all of the inventory on-hand is "new"). Also, a sufficient warm-up period coupled with 1,000 simulation replications makes it likely that all possible supply chain disruption patterns are well represented. We specifically select 1,000 simulation replications as this ensures a 95% confidence interval half-width of at most 0.01 for the proportion of drug shortages per day for Rocuronium and Labetalol with the (A) Adaptive model when $(\delta_s, \delta_w) = (0.05, 0.05)$ and $(\alpha^{(1)}, \beta^{(1)}) = (\frac{1}{270}, \frac{1}{90})$.

For the adaptive models (i.e., (A) Adaptive and (B) Adaptive with Buyback), we use the most recent $N = 56$ daily demand observations to endogenously detect demand disruptions and update the (R, S) inventory policy. It is worth noting that we initialize all the simulation models using the first $B = 90$ training daily demand observations which keeps consistency across all four inventory systems of interest: (A) Adaptive, (B) Adaptive with Buyback, (C) Benchmark, and (D) Static. Furthermore, we implement a long warm-up period that uses the first 180 training daily demand observations stacked 4 times for the simulation models (i.e., ≈ 2 years). But, the sole purpose of consistent initialization and long warm-up periods is to compare the performance of the systems without bias. When implementing these inventory systems in practice, the (A) Adaptive and (B) Adaptive with Buyback inventory systems only require N daily demand observations for initialization. The (C) Benchmark inventory system only requires B daily demand observations

for initialization. The (D) Static inventory system only requires N or B daily demand observations (based on the decision-maker) for initialization.

Table 4.4 Numerical analysis input parameters.

Notation	Description
Input Parameters	
$k = 10$·(drug price)	Fixed ordering cost (i.e., for each order attempted)
$h = 0.001$·(drug price)	Holding cost per day per drug
$e = 90$	Expiration lifetime in days
$\gamma = 0.05$	Maximum proportion of drug shortages per day over the infinite horizon
$\alpha^{(1)} \in \{\frac{1}{30}, \frac{1}{90}, \frac{1}{270}, \frac{1}{810}\}$	Supply chain disruption probability with respect to 1 day
$\beta^{(1)} \in \{\frac{1}{10}, \frac{1}{30}, \frac{1}{90}, \frac{1}{270}\}$	Supply chain recovery probability with respect to 1 day
$\delta_s \in [0.01, 0.09]$	Change in the proportion of drug shortages per day that signals a change in the (R,S) inventory policy when exceeded
$\delta_w \in [0.01, 0.09]$	Change in the proportion of drugs wasted per day that signals a change in the (R,S) inventory policy when exceeded
$N = 56$	Number of past daily demand observations to consider for the adaptive inventory system daily demand estimates
$B = 90$	Number of days the inventory system follows the same (R,S) inventory policy where the average of the last B days is used to update the inventory policy
Warm-up= $4(180)$	Length of the warm-up period in days for the simulation models
Reps= $1,000$	Number of simulation replications

4.5.4 Drug Case Studies

We proceed to present the results for the drug case studies. Specifically, we provide the results for (a) Rocuronium and (b) Labetalol. At the end of this section, we consider four additional 503B drugs and summarize the results in a table. For all analyses, we use the real-world demand data provided by the Central Pharmacy at the University of Michigan and we simulate the performance of the (A) Adaptive, (B) Adaptive with Buyback, (C) Benchmark, and (D) Static models. For clarity, from a shortage perspective, we claim that adaptive inventory policies are beneficial if the adaptive model (e.g., (A) Adaptive and (B) Adaptive with Buyback) leads to a smaller proportion of drug shortages per day in comparison to the (D) Static model. We claim that adaptive inventory policies are detrimental if the adaptive model has a larger proportion of drug shortages per day in comparison to the (D) Static model. From a waste perspective, we claim that adaptive inventory policies are beneficial if the adaptive model (e.g., (A) Adaptive and (B) Adaptive with Buyback) leads to a smaller proportion of drugs wasted per day in comparison to the (D) Static model. We claim that adaptive inventory policies are detrimental if the adaptive model has a larger proportion of drugs wasted per day in comparison to the (D) Static model. We also provide additional insights using the (C) Benchmark model.

4.5.4.1 Rocuronium and Labetalol Case Studies

We present the results for Rocuronium (see Figure 4.6) and Labetalol (see Figure 4.7). We present the overall expected proportion of drug shortages per day and proportion of drugs wasted per day over the testing horizon. The testing horizon is 720 days (i.e., ≈ 2 years). We also present the expected number of inventory policy changes made over the course of the testing horizon. For the shortage-waste weightings (see x-axis in Figures 4.6 and 4.7), we consider $\delta_w \in [0.01, 0.09]$ such that $\delta_s + \delta_w = 0.1$. For the supply chain disruption profile, we focus on supply chain disruptions with an expected duration of 90 days: $(\alpha^{(1)}, \beta^{(1)}) = (\frac{1}{270}, \frac{1}{90})$. Also, for the supply chain disruption profile, we provide the disruption probability with respect to 1 day (i.e., $\alpha^{(1)}$), recovery probability with respect to 1 day (i.e., $\beta^{(1)}$), and optimal (R^*, S^*) inventory policy when the mean of the first $B = 90$ training daily demand observations (i.e., initial $\bar{q}_{current}$) are used for the expected daily demand $\bar{q}_{current}$ (see Figure 4.3 step (a)).

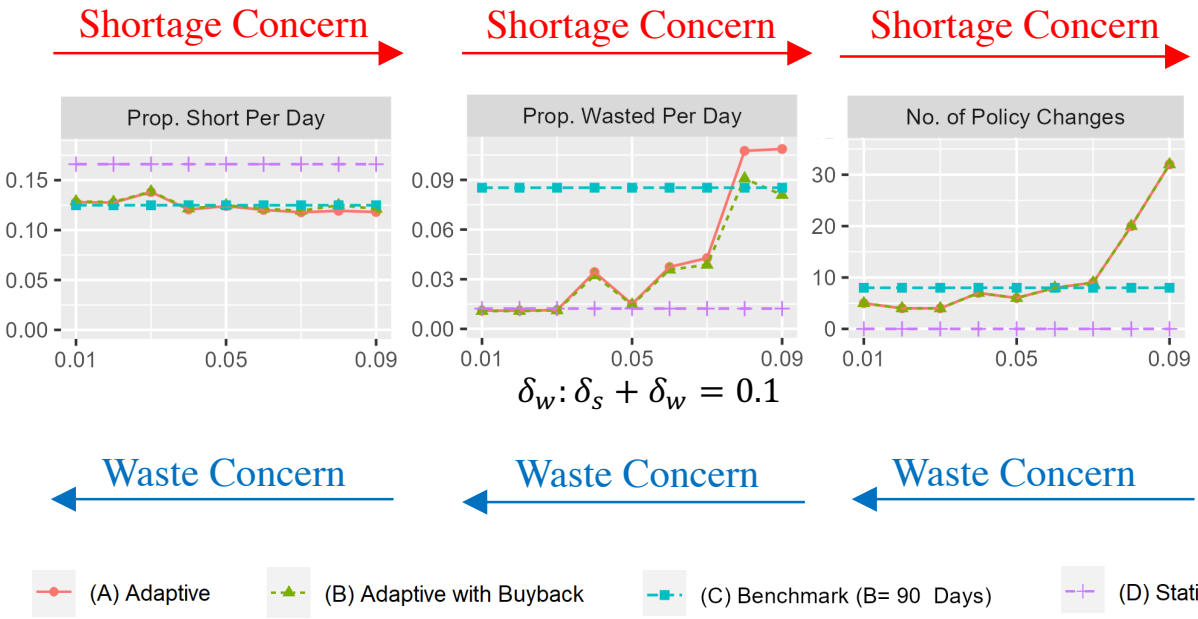

Figure 4.6: Rocuronium results with models (A)-(D). We consider long supply chain disruption durations where $\alpha^{(1)} = \frac{1}{270}$ and $\beta^{(1)} = \frac{1}{90}$. We have that $(R^*, S^*) = (1, 8910)$ with an initial $\bar{q}_{current} = 98$. The x-axis illustrates δ_w such that $\delta_w + \delta_s = 0.1$ implying δ_w increases from left to right and δ_s decreases from left to right. Recall that δ_w and δ_s represent the change in the proportion of drugs wasted per day and drug shortages per day that the inventory system is willing to tolerate. See Figure 4.4(a) for the observed demand over the planning horizon.

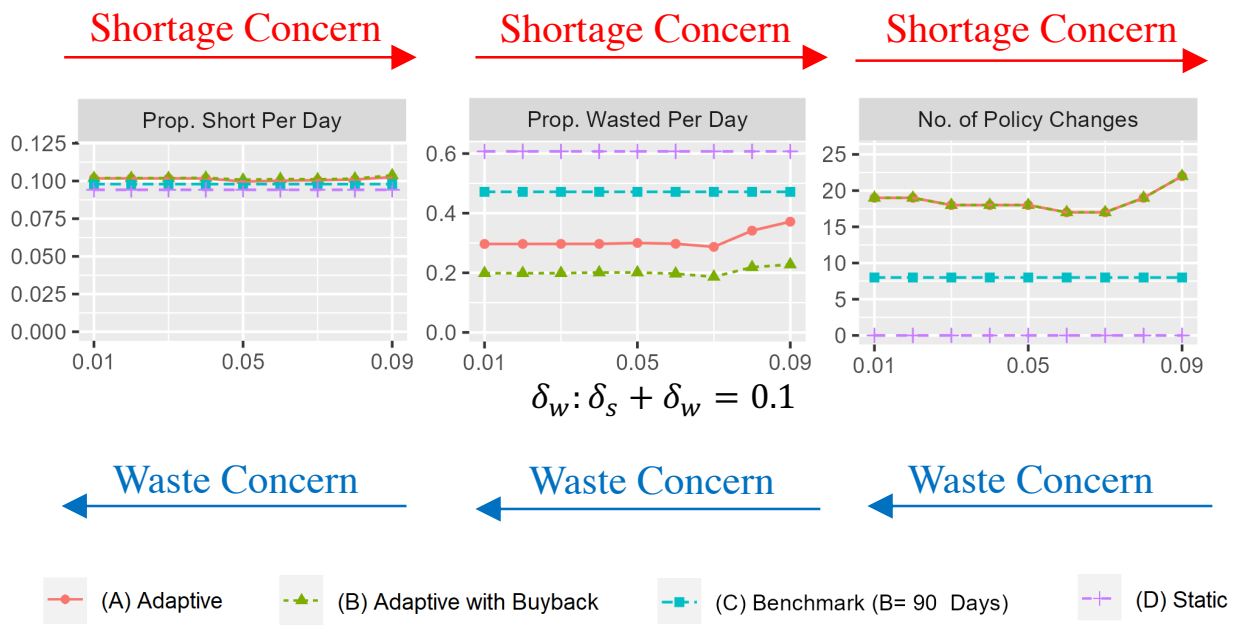

Figure 4.7: Labetalol results with models (A)-(D). We consider long supply chain disruption durations where $\alpha^{(1)} = \frac{1}{270}$ and $\beta^{(1)} = \frac{1}{90}$. We have that $(R^*, S^*) = (1, 3780)$ with an initial $\bar{q}_{current} = 43$. The x-axis illustrates δ_w such that $\delta_w + \delta_s = 0.1$ implying δ_w increases from left to right and δ_s decreases from left to right. Recall that δ_w and δ_s represent the change in the proportion of drugs wasted per day and drug shortages per day that the inventory system is willing to tolerate. See Figure 4.4(b) for the observed demand over the planning horizon.

From Figures 4.6 and 4.7, we find that (a) for a fixed supply chain disruption profile, a drug's shortage-waste weighting dictates the magnitude of the benefits (or detriments) of adaptive inventory policies. For example, consider the proportion of drugs wasted per day for Rocuronium (see Figure 4.6 column 2). When the model is more concerned with shortages implying δ_w is large since $\delta_w + \delta_s = 0.1$ (see x-axis), the detriments of an adaptive model (i.e., (A) Adaptive and (B) Adaptive with Buyback) in comparison to the static model (i.e., (D) Static) are greater than a model that is more concerned with waste implying δ_w is small. We note that the (D) Static model always performs well in terms of the proportion of drugs wasted. We most likely observe this because Rocuronium resembles an increasing demand disruption (see Figure 4.4). The (D) Static model never changes the inventory policy when the increase in demand occurs implying the (D) Static model is always under-ordering. As a result, the (D) Static model performs well from a waste perspective (see Figure 4.6 column 2), but not a shortage perspective (see Figure 4.6 column 1). Consider the proportion of drugs wasted per day for Labetalol (see Figure 4.7 column 2). When the model is more concerned with shortages implying δ_w is large (see x-axis), the benefits of an adaptive model (i.e., (A) Adaptive and (B) Adaptive with Buyback) in comparison to the static model (i.e., (D) Static) are slightly less than a model that is more concerned with waste implying

δ_w is small.

Selecting a shortage-waste weighting with a high shortage [waste] concern often performs better than a shortage-waste weighting with a low shortage [waste] concern with respect to the proportion of drug shortages [drugs wasted] per day. However, it is worth noting that we do not always observe a monotone property for the proportion of drug shortages per day and the proportion of drugs wasted per day as we vary (δ_s, δ_w). For example, consider the proportion of drugs wasted per day for Rocuronium (see Figure 4.6 column 2). There are instances where a shortage-waste weighting that is less concerned with waste (e.g., $\delta_w = 0.05$) has a smaller proportion of drugs wasted per day in comparison to a shortage-waste weighting that is more concerned with waste (e.g., $\delta_w = 0.04$). Taking a deeper look, we also find that the 95% confidence intervals around the point estimates for the proportion of drugs wasted per day for $\delta_w = 0.04$ and $\delta_w = 0.05$ do not intersect. However, it is important to note that the difference in the expected proportion of drugs wasted per day with $\delta_w = 0.04$ and $\delta_w = 0.05$ is very small (i.e., difference of 0.019). Furthermore, if we consider all non-monotone instances across the six 503B drugs of interest in this chapter, we continue to find a very small difference in the expected proportion of drug shortages per day or drugs wasted per day (i.e., difference ≤ 0.03). From a practical perspective, these very small differences may not be significant. We suspect that these non-monotone instances are a combination of (1) simulation sampling error and (2) the drug shortage concern (i.e., δ_s) and drug waste concern (i.e., δ_w) causing small differences on when an inventory policy is updated and the frequency of an inventory policy being updated.

When looking at the number of policy changes, we find that (b) the number of policy changes with the (A) Adaptive and (B) Adaptive with Buyback models is largely influenced by the drug shortage concern (i.e., δ_s) and drug waste concern (i.e., δ_w). Consider the number of policy changes for Rocuronium (see Figure 4.6 column 3). As the model becomes more concerned with shortages (i.e., δ_w increases), the number of policy changes generally increase. Consider the number of policy changes for Labetalol (see Figure 4.7 column 3). The number of policy changes is fairly stable for the multiple (δ_s, δ_w) shortage-waste weightings. But, the number of policy changes starts to increase as the model becomes more concerned with shortages.

We next take a deeper look at the (C) Benchmark model which updates the inventory policy every $B = 90$ days using the average daily demand from the most recent $B = 90$ days. In the upcoming Section 4.5.6, we emphasize that (c) the (C) Benchmark model can lead to very poor performance. If we consider Figures 4.6 and 4.7, we can start to see this finding. The proportion of drug shortages per day with the (A) Adaptive, (B) Adaptive with Buyback, and (C) Benchmark models are very similar. But, the proportion of drugs wasted per day with the (C) Benchmark model is much greater than the proportion of drugs wasted per day with the (A) Adaptive and (B) Adaptive with Buyback models for almost all shortage-waste weightings.

We also find that (d) when considering the proportion of drugs wasted per day particularly for a drug that resembles a decreasing demand disruption (e.g., Labetalol; see Figure 4.7), the (B) Adaptive with Buyback model outperforms the (A) Adaptive model (see Figure 4.7 column 2). Recall that the (B) Adaptive with Buyback model allows the hospital pharmacy to return drugs when the order-up-to level S^* is less than the inventory on-hand. The improvement in the proportion of drugs wasted per day with zero to negligible impact on the proportion of drug shortages per day encourages the use of such buyback programs.

Other Drugs of Interest and Statistical Significance

In addition to Rocuronium and Labetalol, we summarize the results (see Table 4.5) for four other 503B drugs of interest: Avastin 1.25mg/0.05mL (chemotherapy with several indications; wholesale price of about $55 per dose), Oxytocin 30 units/500mL (induction of labor; wholesale price of about $10 per dose), Cefazolin 2gm/100mL (antibiotic; wholesale price of about $14 per dose), Norepinephrine 16mg/250mL (vasopressor used to increase blood pressure; wholesale price of about $25 per dose). The wholesale prices were obtained through discussions with our hospital pharmacy collaborators. The weekly demand data are provided in Appendix C.3.1. In Table 4.5, we present the ratio of the proportion of drug shortages per day with the (D) Static model to the proportion of drug shortages per day with the (A) Adaptive model for varying shortage-waste weightings and supply chain disruption profiles. We present the ratio to the hundredths place. A ratio that is greater than 1 implies that the (A) Adaptive model is beneficial (i.e., decreases the proportion of drug shortages per day) and a ratio that is less than 1 implies that the (A) Adaptive model is detrimental (i.e., increases the proportion of drug shortages per day). We do the same for the proportion of drugs wasted per day. Numbers with asterisks indicate that the proportion with the (A) Adaptive model and (D) Static model are statistically different at a 0.05 (*) and 0.01 (**) significance level when applying a non-parametric Wilcoxon signed-rank paired test to the 1,000 simulation replications. Furthermore, when the proportion with the (A) Adaptive model or (D) Static model is zero, we present the difference between the two proportions (i.e., (A) Adaptive - (D) Static) and indicate this with a “D" before the numerical value.

In Table 4.5, we present three shortage-waste weightings: $(\delta_s, \delta_w) \in \{(0.075, 0.025), (0.05, 0.05), (0.025, 0.075)\}$. When viewing Table 4.5, we note that the drug shortage concern increases and the drug waste concern decreases when viewing the shortage-waste weightings from left to right. In Table 4.5, we present four supply chain disruption profiles: $(\alpha^{(1)}, \beta^{(1)}) = \{ (\frac{1}{30}, \frac{1}{10}), (\frac{1}{90}, \frac{1}{30}), (\frac{1}{270}, \frac{1}{90}), (\frac{1}{810}, \frac{1}{270})\}$. All of the supply chain disruption profiles have the same long-run probability that the supply chain is disrupted (i.e., $\frac{\alpha^{(1)}}{\alpha^{(1)}+\beta^{(1)}} = 0.25$). Fixing the long-run probability that the supply chain is disrupted, from top to bottom, the four supply chain disruption profiles are listed in order of increasing supply chain disruption duration.

In Table 4.5, we find that (e), for a fixed shortage-waste weighting and long-run probability that the supply chain is disrupted, if adaptive inventory policies are beneficial from a shortage [waste] perspective, the benefits generally decrease as the supply chain disruption duration increases. Similarly, if adaptive inventory policies are detrimental from a shortage [waste] perspective, the detriments generally decrease as the supply chain disruption duration increases. As an example, consider Labetalol with $(\delta_s, \delta_w) = (0.05, 0.05)$. The shortage ratio is less than 1 implying there are detriments in implementing the (A) Adaptive model in comparison to the (D) Static model. As the supply chain disruption duration increases (move down the rows of the table), we notice that the shortage ratio increases implying the detriments decrease. It is likely that the (D) Static model performs better than the (A) Adaptive model for Labetalol from a shortage perspective because this drug resembles a decreasing demand disruption (see Figure 4.4(b)). Therefore, if we never change the inventory policy as done in the (D) Static model, we will always be over-ordering. Consequently, the (D) Static model leads to less shortages in comparison to the (A) Adaptive model that updates the inventory policy when changes in the mean demand occur. Keeping in mind that we are over-ordering with the (D) Static model, as we would expect, the (A) Adaptive model leads to a smaller proportion of drugs wasted per day in comparison to the (D) Static model which is seen with the waste ratio being greater than 1. We notice that as the supply chain disruption duration increases (move down the rows of the table), the waste ratio decreases implying the benefits decrease.

It is also worth noting that when a drug's inventory policy shifts from beneficial to detrimental as the supply chain disruption duration increases (e.g., see Avastin $(\delta_s, \delta_w) = (0.05, 0.05)$), the same conclusion (e) holds. Specifically, for these instances, as the supply chain disruption duration increases, the benefits of adaptive inventory policies decrease until the adaptive inventory policies become detrimental. When the adaptive inventory policies become detrimental, the detriments decrease as the supply chain disruption duration increases. Furthermore, when stating finding (e), we use the term generally because there are instances where the relationship does not hold exactly, but the difference in performance is very small (e.g., see Labetalol with $(\delta_s, \delta_w) = (0.025, 0.075)$ where $1.89 > 1.83$). Although, this is most likely a result of simulation sampling error.

We find that (f) as the expected supply chain disruption duration increases, the trade-off between drug shortages and drug waste becomes more apparent as improving one often negatively impacts the other. To see this finding, consider $(\delta_s, \delta_w) = (0.05, 0.05)$ for Rocuronium. We notice that the (A) Adaptive model is always beneficial from a shortage perspective as the shortage ratio is greater than 1. However, as the supply chain disruption duration increases, we notice that the (A) Adaptive model becomes detrimental from a waste perspective as the ratio is less than 1 illustrating the trade-off between drug shortages and drug waste.

Table 4.5 Case study results for all 503B drugs of interest.

$(\alpha^{(1)}, \beta^{(1)})$	Disruption Duration Description	(δ_s, δ_w) (0.075, 0.025) more concerned with waste		(0.05, 0.05)		(0.025, 0.075) more concerned with shortages	
		Shortage Ratio	**Waste Ratio**	**Shortage Ratio**	**Waste Ratio**	**Shortage Ratio**	**Waste Ratio**
				Rocuronium			
(1/30, 1/10)	Short	2.15**	D0	2.68**	D0	2.94**	D0
(1/90, 1/30)	Moderate	1.89**	D0	2.53**	D0	2.14**	D0
(1/270, 1/90)	Long	1.24**	1.1**	1.34**	0.82**	1.44**	0.31**
(1/810, 1/270)	Very Long	0.96**	1**	1.11**	0.9**	1.13**	0.87**
				Labetalol			
(1/30, 1/10)	Short	0.88**	2.99**	0.88**	2.99**	0.88**	2.99**
(1/90, 1/30)	Moderate	0.93**	2.46**	0.93**	2.46**	0.93**	2.46**
(1/270, 1/90)	Long	0.93**	2.05**	0.94**	2.02**	0.94**	1.83**
(1/810, 1/270)	Very Long	0.98**	1.84**	0.98**	1.82**	0.98**	1.89**
				Avastin			
(1/30, 1/10)	Short	1.12**	D0	1.12**	D0	1.12**	D0
(1/90, 1/30)	Moderate	1.08**	D0	1.08**	D0	1.08**	D0
(1/270, 1/90)	Long	0.86**	2.35**	0.88**	2.09**	0.95**	1.03**
(1/810, 1/270)	Very Long	0.95**	1.36**	0.95**	1.3**	0.97**	0.84**
				Oxytocin			
(1/30, 1/10)	Short	3.04**	D0	3.04**	D0	4.05**	D0
(1/90, 1/30)	Moderate	1.9**	D0	1.9**	D0	2.4**	D0
(1/270, 1/90)	Long	1.3**	0.36**	1.31**	1.18**	1.43**	0.23**
(1/810, 1/270)	Very Long	0.98**	0.99**	0.98**	0.99**	1.11**	0.45**
				Cefazolin			
(1/30, 1/10)	Short	8.42**	D0	8.36**	D0	9.56**	D0
(1/90, 1/30)	Moderate	4.53**	D0	5.05**	D0	6.21**	D0
(1/270, 1/90)	Long	1.44**	D0.01	1.59**	D0	1.66**	D0.07
(1/810, 1/270)	Very Long	1.02**	D0	1.12**	D0	1.2**	D0.02
				Norepinephrine			
(1/30, 1/10)	Short	2.06**	D0	2.5**	D0	2.7**	D0
(1/90, 1/30)	Moderate	1.25**	D0.02	1.7**	D0	1.39**	D0.01
(1/270, 1/90)	Long	0.97**	0.04**	1.02**	0.06**	0.99*	0.03**
(1/810, 1/270)	Very Long	0.99**	1**	0.99**	1**	1	0.12**

Values represent the ratio of the proportion with the (D) Static model to the proportion with the (A) Adaptive model. Significance levels of 0.05 (*) and 0.01 (**) are indicated. When the proportion with the (A) Adaptive model or (D) Static model is zero, the value represents the difference between the two proportions (i.e., (A) Adaptive - (D) Static) and this is indicated with a "D" before the numerical value.

4.5.5 Ranking Analysis

In Section 4.3.3, we present a P_{metric} ranking procedure that prioritizes drugs and indicates which drugs are of most concern. We proceed to describe a way of implementing the ranking procedure in practice and we analyze the performance when applying this method to the 500 highest (unit price)·(demand) drugs at the University of Michigan's Central Pharmacy. We consider the total yearly demand when selecting the 500 highest (unit price)·(demand) drugs. When considering the highest (unit price)·(demand) drugs, we only consider drugs that have demand data available for the time period of interest (i.e., October 2019-November 2021) and we ensure that the data only includes drugs (e.g., removed medical supplies). 370 drugs meet this criteria.

For implementation, we consider that the Central Pharmacy at the University of Michigan keeps record of P_{metric} every day for all drugs. We consider that every 30 days, the Central Pharmacy decides which drugs to update using the average value of P_{metric} over the most recent 30 days.

However, instead of updating all drugs that have an average value of P_{metric} greater than 0 (i.e., a value that signals a change in the (R,S) inventory policy is needed), we only allow the Central Pharmacy to update at most $M^{(\%)}$ percent of drugs: $M^{(\%)} \in \{0, 2.5, 5, ..., 15\}$. If more than $M^{(\%)}$ percent of drugs need to updated, we sort the average value of P_{metric} from largest to smallest for all drugs, and we only update the top $M^{(\%)}$ percent of drugs. For the analysis, we consider the input parameters defined in Section 4.5.3. But, we assume that all 503B drugs have an expiration lifetime of 90 days (i.e., about 3 months) and all non-503B drugs have an expiration lifetime of 360 days (i.e., about 1 year). Through discussions with our hospital pharmacy collaborators, non-503B drugs have a variety of expiration lifetimes and much longer lifetimes than 503B drugs. We select 360 days for these drugs because it is a conservative expiration lifetime estimate. Also, we consider $(\delta_s, \delta_w) = (0.05, 0.05)$ for the shortage-waste weighting, $(\alpha^{(1)}, \beta^{(1)}) = (\frac{1}{90}, \frac{1}{30})$ for the supply chain disruption profile, and we focus on the (A) Adaptive model where updates can only occur every 30 days (i.e., no buyback is in place). We present the average proportion of drug shortages per day across all drugs and the average proportion of drugs wasted per day across all drugs. We also present the cost of wasting a drug relative to the cost of a drug shortage that is required for the reduction in drug shortages to equally outweigh the detriments of drug waste. To calculate this value, we take the absolute difference between the average number of drug shortages with $M^{(\%)}$ percent and 0 percent updates across all drugs divided by the absolute difference between the average number of drugs wasted with $M^{(\%)}$ percent and 0 percent updates across all drugs. When the cost of wasting a drug relative to the cost of a drug shortage exceeds this value, it implies that the reduction in drug shortages does not outweigh the detriments of drug waste. Figure 4.8 illustrates the results.

We find that (g) there is a decreasing marginal benefit from a drug shortage perspective as the maximum percent of drugs that can be updated increases. This implies that a decision-maker needs to update a very small proportion of drugs at any point in time to get the greatest benefits of adaptive inventory policies. From a drug waste perspective, there is a decreasing marginal detriment as the maximum percent of drugs that can be updated increases. However, the detriments from a drug waste perspective are very small practically speaking. Furthermore, with the results illustrating the trade-off between drug shortages and drug waste (i.e., decrease in drug shortages but increase in drug waste), we take a deeper look at the cost of wasting a drug relative to the cost of a drug shortage that is required for the reduction in drug shortages to equally outweigh the detriments of drug waste (see Figure 4.8 row 3). We find that (h) the cost of wasting a drug would need to be far greater than the cost of a drug shortage for the reduction in drug shortages to not outweigh the detriments of drug waste. For example, consider $M^{\%} = 2.5$. The cost of wasting a drug would need to be about 3.5 times the cost of a drug shortage for the reduction in drug shortages to not outweigh the detriments of drug waste. Drug shortages can increase the cost of care, increase drug

errors, and delay/cancel treatment (Phuong et al., 2019). Therefore, except for the few very high priced drugs purchased on an order-by-order basis, it is highly unlikely that the cost of wasting a drug would be greater than a drug shortage let alone 3.5 times the cost of a drug shortage. It is worth noting that when we consider the 500 highest priced drugs and the 500 highest demanded drugs (total yearly demand), findings (g) and (h) continue to hold.

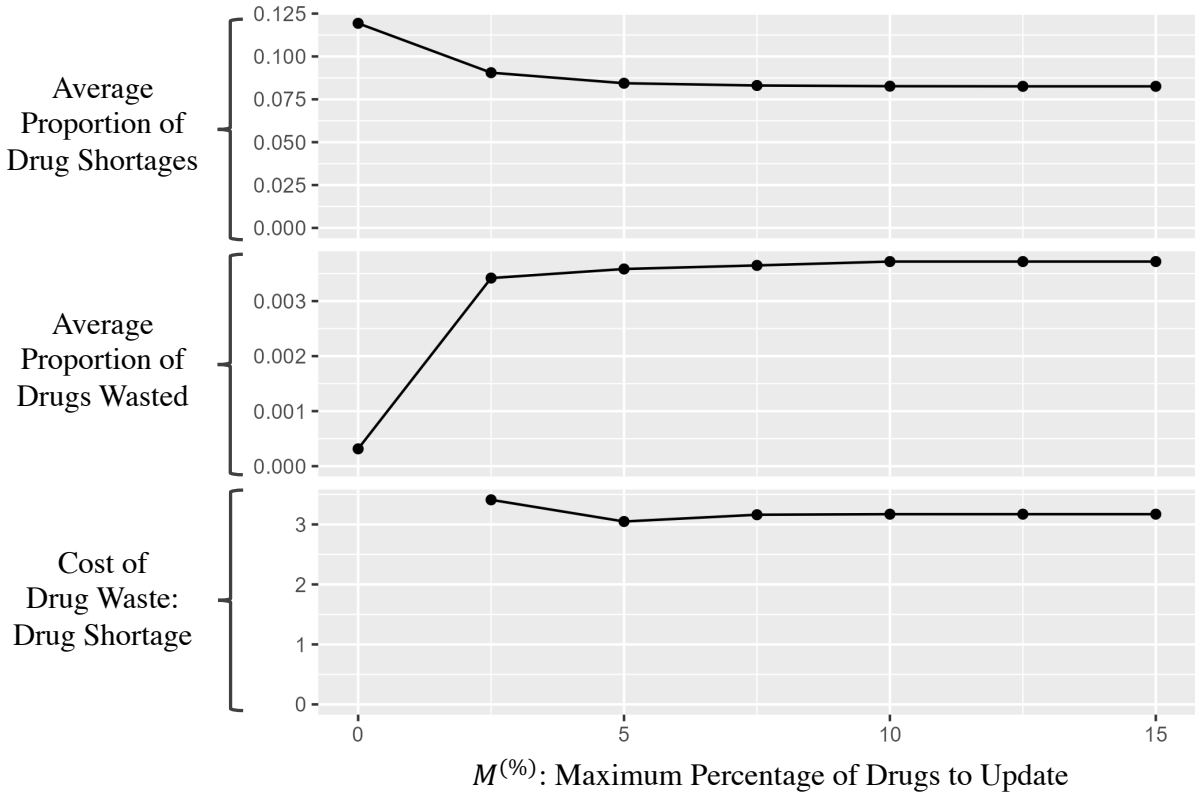

Figure 4.8: Proportion of drug shortages, proportion of drugs wasted, and the cost of wasting a drug relative to the cost of a drug shortage that is required for the reduction in drug shortages to equally outweigh the detriments of drug waste. The analysis focuses on the highest (unit price)·(demand) drugs at the University of Michigan's Central Pharmacy and for the analysis, a decision-maker can update at most $M^{(\%)}$ percent of drugs every 30 days.

4.5.6 Additional Takeaways

In the Appendix, we gain additional takeaways by analyzing (T1) the sensitivity of N (i.e., number of past daily demand observations used to estimate the expected daily demand; in Section C.3.2) and (T2) the length of demand disruptions (in Section C.3.3). We find that:

(T1) It is important to have a large enough N to avoid high shortage and/or waste detriments, but not too large such that the benefits of an adaptive model start to decrease. Furthermore, a

drug with a higher variability in demand (i.e., a larger standard deviation) often requires a larger N to avoid high shortage and/or waste detriments. As a rule of thumb, we suggest using about 50 daily demand observations when selecting N (e.g., $N = 49$, 7 weeks; $N = 56$, 8 weeks).

(T2) Given we vary the length of the demand disruption (for both increasing and decreasing demand disruptions), the benchmark model that only updates the inventory policy every $B = 90$ days using the average demand from the recent $B = 90$ days can lead to very poor performance. Furthermore, the adaptive model with buyback often performs better than the adaptive model without buyback for varying demand disruption lengths encouraging the use of such programs.

4.6 Conclusion

This research considers a perishable inventory system with supply chain disruptions and demand disruptions. This chapter leverages simulation modeling to distinguish how a drug's shortage-waste weighting (i.e., concern for shortages versus concern for waste) along with the duration of and time between supply chain disruptions influences the benefits (or detriments) of adapting to demand disruptions. We find that when fixing the mean duration of and mean time between supply chain disruptions, a drug's shortage-waste weighting dictates the magnitude of the benefits (or detriments) of adaptive inventory policies. We also find that for a fixed shortage-waste weighting and long-run probability that the supply chain is disrupted, if adaptive inventory policies are beneficial from a shortage [waste] perspective, the benefits generally decrease as the supply chain disruption duration increases. Similarly, if adaptive inventory policies are detrimental from a shortage [waste] perspective, the magnitude of the negative impact of the adaptive inventory policies decreases as the supply chain disruption duration increases. Furthermore, when fixing the long-run probability that the supply chain is disrupted, we find that long supply chain disruption durations are the most sensitive to the trade-off between drug shortages and drug waste as improving one often negatively impacts the other. The results also suggest that hospital pharmacies should avoid updating the inventory policy on long fixed intervals and negotiate for buyback programs when establishing contracts with their suppliers/wholesalers. For the latter, this is under the assumption that a buyback program adds no additional costs to the hospital pharmacy inventory system. When looking at a large collection of drugs, we find that a decision-maker needs to update a very small proportion of drugs (e.g., $< 5\%$) at any point in time to get the greatest benefits of adaptive inventory policies.

Our research illustrates that it is often beneficial for hospital pharmacies to participate in buy-

back programs. We do not consider buyback programs from the supplier perspective and we assume that the hospital pharmacy receives full compensation for the returned drug. If we assume that the hospital pharmacy does not receive full compensation for the returned drug, then the supplier may greatly benefit from buyback programs because they can possibly sell this collected drug to other hospital pharmacies at full price. However, there is no guarantee that another hospital pharmacy will purchase this returned drug before it expires. Hence, it would be interesting to study if the supplier should only accept a certain number of drugs or drugs with a certain remaining shelf life from the hospital pharmacy to reduce costs and waste at the supplier. We leave these studies for future research. In addition, we consider (R,S) inventory policies which are commonly implemented in practice. Future research can study other inventory policies (e.g., (s,S) inventory policies) and how adapting these policies impact the performance of the inventory system. Furthermore, when considering (R,S) inventory policies, we assume full and accurate knowledge of the inventory on-hand. Future research can consider the case where inventory records are inaccurate (Neve & Schmidt, 2022). A final direction for future research is using the adaptive inventory system presented in this chapter and analyzing how dynamic supply chain disruption parameters (i.e., $\alpha^{(1)}$ and $\beta^{(1)}$) influence the results.

Decision-making is difficult in hospital pharmacy inventory systems due to perishability, supply chain disruptions, and demand disruptions. A drug's shortage-waste weighting and supply chain disruption profile influence the benefits (or detriments) of adapting to demand disruptions, and hospital pharmacy managers should consider these characteristics when implementing such policies in practice.

Acknowledgments for Chapter 4

The research team would like to thank the Central Pharmacy at the University of Michigan for providing pharmaceutical drug data for the numerical analysis. We would like to thank Christine Betts (University of Michigan Health Pharmacy Services) for providing access to the data. We would also like to thank Jun Li (University of Michigan Ross School of Business) for the helpful discussions on the model and Nicole Mor (University of Michigan) for the helpful literature searches.

CHAPTER 5

When Is It Worth It for Two Hospital Network Pharmacies to Operate as an Integrated Inventory System In the Presence of Supply Chain Disruptions?

5.1 Introduction

Supply chain disruptions are prevalent in many perishable inventory systems (e.g., hospital pharmacies, blood banks, chemical warehouses). However, perishability limits how much the inventory system can "stock up" to avoid the negative effects of shortages caused by these supply chain disruptions. In a hospital pharmacy inventory system, it is crucial to avoid shortages as they increase the potential for medication errors, result in the use of inferior treatments, and increase the overall cost of care (Phuong et al., 2019). Further, there has been much interest in the resiliency of pharmaceutical supply chains (NASEM, 2022). We define supply chain disruptions as random periods of time where the supplier completely stops functioning (Snyder et al., 2016). A potential solution to alleviate the negative effects of supply chain disruptions is lateral transshipments (i.e., sharing of inventory); the shipment of products between retailers in the same echelon of the supply chain (Paterson et al., 2011). Literature from non-perishable inventory systems illustrates that lateral transshipments reduce shortages, increase service levels, and decrease costs (Herer et al., 2006; Özdemir et al., 2013; Noham & Tzur, 2014). These thoughts give rise to two research questions. (1) How can decision-making be optimized in perishable inventory systems with supply chain disruptions and lateral transshipments and (2) when do lateral transshipments benefit perishable inventory systems with independent suppliers experiencing supply chain disruptions?

A hospital network pharmacy (HNP) is a central pharmacy that provides drugs for all hospitals in the healthcare network. Strict regulations in current practice generally prohibit HNPs from sharing drugs outside of their network or make it very difficult for HNPs to stay compliant when sharing drugs outside of their network. However, it has been hypothesized that relaxing these strict regulations may improve the performance of the inventory systems. Such improvements are needed

to ensure that patients have access to the treatment that they need. Further, in a HNP setting, it is critical to have inventory policies that are accurate and quick-to-solve since hospital pharmacy managers are often responsible for thousands of different drugs.

We model a two-HNP integrated inventory system using a continuous time Markov chain. By integrated, we mean that the two HNPs participate in lateral transshipments when one HNP has positive inventory on-hand, and the other HNP has zero inventory on-hand (due to supply chain disruptions). Each HNP follows a continuous review order-up-to level inventory policy (i.e., HNP 1 [HNP 2] always has an inventory on-hand of S_1 [S_2] when the supply chain for HNP 1 [HNP 2] is not disrupted) and each HNP has their own independent supplier that experiences supply chain disruptions. We first create an exact model which relies on solving the system of balance equations to find the long-run probability of being in each state of the Markov chain. The long-run probabilities are used to calculate the expected cost per day: holding cost, shortage cost, and lateral transshipment cost. To solve for the optimal order-up-to levels in the integrated inventory system (i.e., S_1^* and S_2^*), we perform an exhaustive search over all feasible inventory policies. Second, we create an approximate model which formulates the Markov chain balance equations by conditioning on the inventory on-hand at one of the HNPs. The approximation achieves high accuracy and our numerical studies illustrate that the expected cost per day calculated with the approximate model is within 2.6% of the expected cost per day calculated with the exact model. With the approximate long-run probabilities, we quickly solve for the optimal order-up-to levels for the integrated inventory system (i.e., S_1^* and S_2^*) solely using closed-form expressions and a subset of feasible inventory policies. Then, we extend the exact and approximate models to incorporate perishability by defining a perishability condition and we perform a two-HNP numerical analysis. We note that for our base case numerical analysis, the approximate model solves for the optimal order-up-to levels in only 10 seconds when perishability does not influence the order-up-to levels [23 seconds when perishability influences the order-up-to levels]. Further, the approximate model with these optimal order-up-to levels calculates an expected cost per day that is only 1% different [2% different] than the expected cost per day when simulating the integrated inventory system with these optimal order-up-to levels (in Appendix D.2.1).

In the two-HNP integrated inventory system, we assume that the cost of a lateral transshipment is less than the cost of a shortage. Otherwise, it is always optimal to operate as a non-integrated inventory system. However, is this assumption sufficient to conclude that it is always optimal to operate as an integrated inventory system? Unfortunately, it is not as shown in Section 5.5.1. The cost of a lateral transshipment must be sufficiently less than the cost of a shortage where sufficient is largely influenced by the supply chain disruption profile (i.e., disruption rate and recovery rate) for each HNP. It is worth noting that sufficiently less is typically satisfied in a pharmaceutical setting where shortage costs tend to be high. Furthermore, we assume that each HNP has their own

supply chain disruption profile (i.e., disruption rate and recovery rate). When selecting a HNP to partner with, is it important for a HNP to consider the supply chain disruption profile of the other HNP? Yes, as shown in Section 5.5.2, it is critical for a HNP to consider the duration of and time between supply chain disruptions when selecting a HNP with which to share inventory.

The contributions of this chapter are as follows:

1. We create a modeling framework to solve for the continuous review inventory policies that minimize the expected cost per day in a two-HNP non-perishable integrated inventory system with independent suppliers experiencing supply chain disruptions. To incorporate perishability, we introduce a constraint on the probability drugs are wasted at each HNP. This modeling framework can also be applied to many inventory systems of interest (e.g., blood banks).

2. Working closely with our hospital pharmacy collaborators at the University of Michigan, we numerically study a two-HNP perishable integrated inventory system to (i) demonstrate when it is beneficial to allow two HNPs with independent suppliers experiencing supply chain disruptions to share inventory and (ii) provide insights on how the supply chain disruption characteristics of the two HNPs influence whether it is beneficial to share inventory.

The remainder of this chapter is organized in the following manner: Section 5.2 provides literature related to this research. Section 5.3 describes the modeling notation and assumptions. Section 5.4 presents the exact model and approximate model for the two-HNP non-perishable integrated inventory system, and how to extend the models to incorporate perishability. In this section, we also present other inventory systems for comparison. In Section 5.5, we provide a two-HNP numerical analysis. In Section 5.6, we provide concluding remarks and outline opportunities for future research.

5.2 Literature Review

Paterson et al. (2011) and Kumari et al. (2022) provide good reviews of lateral transshipments in supply chains. We highlight that lateral transshipments in perishable inventory systems with supply chain disruptions (e.g., hospital pharmacies) are not well understood. We proceed to define how supply chain disruptions are modeled and define common lateral transshipment terminology. Then, we present past research relevant to our perishable integrated inventory models: (a) multi-retailer non-perishable inventory systems with lateral transshipments and imperfect supply, and (b) multi-retailer perishable inventory systems. We use the term retailer for generality, but in our hospital network pharmacy setting of interest, a retailer refers to a HNP. We finish this section

by presenting how our research is distinct from past research.

Modeling Supply Chain Disruptions

Supply chain disruption research usually defines disruptions as a Bernoulli or two-state supply process. In a Bernoulli supply process, the status of the supplier on day t (i.e., disrupted or not disrupted) is independent of the status of the supplier on day $t-1$ (Snyder et al., 2016). A drawback of modeling supply chain disruptions in this manner is that it fails to capture the duration of and time between supply chain disruptions, which is critical in inventory decision-making (Chapter 3). Instead, we model supply chain disruptions as a two-state supply process where the time in each state (i.e., disrupted or not disrupted) is exponentially distributed (Snyder et al., 2016). We define λ_1 [λ_2] as the disruption rate (i.e., the rate of transitioning from the not disrupted state to the disrupted state) and μ_1 [μ_2] as the recovery rate (i.e., the rate of transitioning from the disrupted state to the not disrupted state) for the supplier of retailer 1 [retailer 2]. A two-state supply process is consistent with pharmaceutical supply chains (Tucker et al., 2020b).

Supply chain disruptions are a type of supply uncertainty, but supply uncertainty can also be modeled as supply capacity (i.e., a finite amount can be ordered from the suppliers each period; deterministic or stochastic). We want to highlight that supply chain disruptions and supply capacity are distinctly different as supply chain disruptions cause no inventory to be received for a random period of time whereas supply capacity limits the amount of inventory that can be received (Yao & Minner, 2017). Also, we want to highlight that supply chain disruptions are distinct from stochastic lead time. Supply chain disruptions imply that orders are only successfully placed when the supply chain is not disrupted whereas stochastic lead time implies that orders are always successfully placed, but the time at which the order arrives is variable. It is also worth noting that in a pharmaceutical supply chain, disruptions often last weeks, months, or years (Tucker et al., 2020a), far longer than the typical time between orders.

Lateral Transshipment Terminology

There are two types of lateral transshipment models: proactive and reactive. Proactive lateral transshipments are made to redistribute stock to maximize the performance of the system whereas reactive lateral transshipments are made in response to a stock out (Paterson et al., 2011). The two-retailer integrated inventory system modeled in this chapter implements reactive lateral transshipments by allowing one retailer that has positive inventory on-hand to share with the other if it has zero inventory on-hand (due to supply chain disruptions). We present past research that considers either type of lateral transshipments as they both quantify the importance of sharing in a multi-retailer inventory system. In addition, lateral transshipment models fall within one of two categories: centralized or decentralized (Paterson et al., 2011). Centralized models optimize the

overall performance of the integrated inventory system where as decentralized models optimize the performance at each individual retailer in the integrated inventory system. The two-retailer integrated inventory system modeled in this chapter minimizes the total cost of the system (i.e., centralized) because in our setting, we are interested in quantifying the overall benefit of allowing two HNPs to share inventory. We present past literature focusing on this category as decentralized systems introduce competition into the model which is beyond the scope of this chapter (decentralized; see Hu et al., 2007; Lee & Park, 2016; Avci & Selim, 2016; Arikan & Silbermayr, 2018).

Relevant Past Research

Table 5.1 presents past research most relevant to our perishable integrated inventory models. We separate the research summary into (a) multi-retailer non-perishable inventory systems with lateral transshipments and imperfect supply, (b) multi-retailer perishable inventory systems, and (c) this research. We characterize the research by perishability (NP: non-perishable, P: perishable), the number of retailers/suppliers (No.: number), the number of echelons in the supply chain, demand (D: deterministic, S: stochastic), supply (DSC: deterministic supply capacity, SSC: stochastic supply capacity, SCD: supply chain disruptions), lateral transshipments (LT; R: reactive, P: proactive), and model/methodology.

Table 5.1 Summary of relevant literature.

Paper	NP/P	No. Retailers	No. Suppliers	3+ Echelons	Demand	Supply	LT	Model
(a) Multi-retailer Non-perishable Inventory Systems with Lateral Transshipments and Imperfect Supply								
Özdemir et al. (2013)	NP	2+	1		S	DSC	R	Stochastic program
Hu et al. (2008)	NP	2	2		S	SSC	R	Dynamic program
Firouz et al. (2017)	NP	2+	2+		S	SSC, SCD	P	Stochastic program
Nasr et al. (2012)	NP	2	1		D	SCD	R	Non-linear program
Avci (2019)	NP	2+	2+		S	SCD	P	Simulation-optimization
Ghomi-Avili et al. (2017)	NP	2+	2+	✓	D	SCD	R	Stochastic program
Jabbarzadeh et al. (2018)	NP	2+	2+	✓	S	SCD	R	Stochastic program
(b) Multi-retailer Perishable Inventory Systems								
Puranam et al. (2017)	P	1	2+		S	SSC	-	Dynamic program
Rajendran & Ravindran (2019)	P	2+	1		S	-	-	Stochastic program
Shih & Rajendran (2020)	P	2+	1		S	SSC	-	Stochastic program
Dolgui et al. (2018)	P	2+	2+	✓	D	DSC	-	Mixed-integer program
Dehghani & Abbasi (2018)	P	2	1		S	-	R	Heuristic that uses PDE
Jin & Agirbas (2013)	P	2+	1		S	-	P	Transshipment decision rule
Nakandala et al. (2017)	P	2+	2+	✓	S	-	R	Transshipment decision rule
Dehghani et al. (2021)	P	2+	1		S	-	P	Stochastic program
Zhang et al. (2022)	P	2	1		S	-	R	Markov decision process
Wei et al. (2022)	P	2+	1		S	-	P	Dynamic program
Shokouhifar et al. (2021)	P	2+	2+		S	SSC	P	Whale optimization algorithm
Arani et al. (2021)	P	2+	2+	✓	S	SSC	R	Stochastic program
Larimi et al. (2019)	P	2+	2+	✓	S	DSC	R	Stochastic program
Zhou et al. (2021)	P	2+	2+		S	SSC	R	Chance-constrained program
Wang & Ma (2015)	P	2+	-		S	SCD	R	Mixed-integer program
Bozkir et al. (2022)	P	2+	1		S	SCD	P	Closed-form service-level metrics
Aldrighetti et al. (2019)	P	2+	1		S	SCD	R	Simulation
(c) This Research								
Chapter 5	P	2	2		S	SCD	R	Markov chain/Closed-form solutions

Distinction of Our Research

The two-HNP perishable integrated inventory system in this chapter captures stochastic demand (Poisson distribution), supply chain disruptions (two independent suppliers), lateral transshipments (reactive), and perishability, which to our knowledge, has not yet been studied in the literature. Looking at the past research presented (non-perishable and perishable), most researchers use mathematical programming, dynamic programming, or simulation-optimization. There are not always resources available to implement these methodologies in practice since they are typically solved using an optimization solver (e.g., Gurobi, CPLEX), backward induction, or simulation and they can also require long computation times. In this chapter, we use a continuous time Markov chain to analytically study a two-HNP integrated inventory system following continuous review order-up-to level inventory policies. Calculating the exact expected cost per day can be computationally intensive as the problem size increases. We create an approximate model, which achieves high accuracy, where all expressions are presented in closed-form. The approximate model provides easy-to-implement and quick-to-solve inventory policies.

Using a continuous time Markov chain is a common approach in the queueing inventory system stream of literature (recent review; see Karthikeyan & Sudhesh, 2016). This stream of literature is different from traditional inventory management because it assumes that each customer is serviced before leaving the system, and consequently, reduces the inventory on-hand by one upon departure. This implies that when there are a positive number of customers in the queue (in service plus waiting), inventory is used according to the service rate rather than the arrival rate (Karthikeyan & Sudhesh, 2016). Traditional inventory management literature also employs Markov chains (e.g., Ross et al., 2008; Vicil & Jackson, 2016; Saedi et al., 2016; Saha & Ray, 2019; Fathi et al., 2021; Poormoaied & Demirci, 2021). None of the past research considers our non-perishable nor perishable setting: two-HNP integrated inventory system following continuous review order-up-to level inventory policies with stochastic demand, lateral transshipments, and supply chain disruptions. For example, Saedi et al. (2016) consider a hospital pharmacy inventory system, but focus on substitute products and use a continuous time Markov chain to optimize a single reorder point fixed order quantity (r,Q) inventory policy. We create a continuous time Markov chain to determine the order-up-to levels for two distinct HNPs (i.e., S_1^* for HNP 1 and S_2^* for HNP 2). We also highlight that in this chapter, we use the long-run probability of being in each state of the continuous time Markov chain to solve for the optimal (S_1^*, S_2^*) integrated inventory policy.

We note the difference between our research and a past research paper that analytically studies an inventory system with lateral transshipments. Li & Zhang (2012) study a two-retailer integrated inventory system following order-up-to level inventory policies with reactive lateral transshipments by expressing the cost function and requirements for optimality in closed-form. These researchers look at the lateral transshipment problem from a stochastic demand perspective and do

not consider supply uncertainty nor perishability. In this chapter, we consider stochastic demand (Poisson distribution), supply chain disruptions, lateral transshipments, and perishability. Also, our research focuses on lateral transshipments between HNPs to improve the performance of the inventory system. This is in contrast to past research that studies the partial or complete removal (i.e., direct-ship) of wholesalers in a pharmaceutical supply chain to improve the performance of the system (Niziolek et al., 2012).

A final important note is that our research contributes to the area of literature focused on viable supply chains (e.g., Ruel et al., 2021; Ivanov, 2022). In particular, by allowing the two HNPs to participate in lateral transshipments, we (i) intersect the supply chains of each system (i.e., supplier + HNP). This intersection introduces flexible redundancy into the integrated inventory system because it allows a HNP that has a disrupted supplier and zero inventory on-hand to receive inventory from the other HNP during this time of need (intertwined supply networks; Ivanov & Dolgui, 2020). Furthermore, the intersection (ii) captures each area of a reconfigurable supply chain at the micro-level: structural redundancy (structural variety), process risk pooling (process flexibility), inventory buffers (parametric redundancy), and coordination (execution visibility; see Dolgui et al., 2020). With our two-HNP integrated inventory system that incorporates characteristics (i) and (ii), we illustrate the reduction in drug shortages by operating as an integrated inventory system when two HNPs with independent suppliers experience supply chain disruptions. Equivalently, we illustrate the viability of the two-HNP integrated inventory system; "ability to maintain itself and survive in a changing environment over a long period of time through a redesign of the structures and replanning of economic performance with long-term impacts" (Ivanov, 2022).

5.3 Model Notation and Assumptions

We consider a two-HNP perishable integrated inventory system (i.e., lateral transshipments are permitted) where each HNP has their own independent supplier experiencing supply chain disruptions. The integrated inventory system is centralized and the objective of the integrated inventory system is to minimize the expected sum of holding, shortage, and lateral transshipment costs per day with a constraint on the probability drugs are wasted at each HNP. We present the notation for the integrated inventory system in Table 5.2, and we proceed to describe the modeling assumptions.

Table 5.2 Summary of the modeling notation.

Notation	Description
Decision Variables	
S_1	Order-up-to level for HNP 1 ($S_1 \geq 1$; i.e., continuously attempt to place an order up to S_1)
S_2	Order-up-to level for HNP 2 ($S_2 \geq 1$; i.e., continuously attempt to place an order up to S_2)
Input Parameters	
h_1	Holding cost per day per drug for HNP 1 ($h_1 > 0$)
h_2	Holding cost per day per drug for HNP 2 ($h_2 > 0$)
b	[a]Shortage cost per unsatisfied demand ($b > 0$; lost-sales)
t_{12}	Lateral transshipment cost from HNP 1 to HNP 2 per drug ($t_{12} > 0$)
t_{21}	Lateral transshipment cost from HNP 2 to HNP 1 per drug ($t_{21} > 0$)
x	[a]Expiration lifetime in days ($x \geq 1$)
q_1	Daily demand rate for HNP 1 ($q_1 > 0$; Poisson distribution)
q_2	Daily demand rate for HNP 2 ($q_2 > 0$; Poisson distribution)
λ_1	Supply chain disruption rate with respect to days for HNP 1 ($\lambda_1 > 0$; exponential distribution)
μ_1	Supply chain recovery rate with respect to days for HNP 1 ($\mu_1 > 0$; exponential distribution)
λ_2	Supply chain disruption rate with respect to days for HNP 2 ($\lambda_2 > 0$; exponential distribution)
μ_2	Supply chain recovery rate with respect to days for HNP 2 ($\mu_2 > 0$; exponential distribution)
Markov Chain	
I_1	Inventory and supply status at HNP 1 ($0 \leq I_1 \leq S_1 + 1; I_1 \in \mathbb{Z}_0^+$) *$0 \leq I_1 \leq S_1$ denotes inventory on-hand is I_1 and the supply chain is disrupted *$I_1 = S_1 + 1$ denotes inventory on-hand is S_1 and the supply chain is not disrupted
I_2	Inventory and supply status at HNP 2 ($0 \leq I_2 \leq S_2 + 1$; $I_2 \in \mathbb{Z}_0^+$) *$0 \leq I_2 \leq S_2$ denotes inventory on-hand is I_2 and the supply chain is disrupted *$I_2 = S_2 + 1$ denotes inventory on-hand is S_2 and the supply chain is not disrupted
X_t	State at time t ($t \geq 0$) defined as the tuple (I_1, I_2)
P_{ij}	Exact long-run probability that the integrated inventory system is in state $X = (I_1 = i, I_2 = j)$ (i.e., $P_{ij} = \Pr(I_1 = i,\ I_2 = j)$)
$P_i^{(p,s)}$	Approximate long-run probability that the primary HNP p has inventory and supply status i (i.e., $P_i^{(p,s)} = \Pr(I_p = i)$); s denotes the secondary HNP where $p \neq s$
$P_{j\|i}^{(p,s)}$	Approximate long-run probability that the secondary HNP s has inventory and supply status j given the primary HNP p has inventory and supply status i (i.e., $P_{j\|i}^{(p,s)} = \Pr(I_s = j \| I_p = i)$)
$P_{ij}^{(p,s)}$	Approximate long-run probability that the integrated inventory system is in state $X = (I_p = i, I_s = j)$ (i.e., $P_{ij}^{(p,s)} = \Pr(I_p = i,\ I_s = j) = P_{j\|i}^{(p,s)} P_i^{(p,s)}$)
Perishability Condition	
δ	[a]Upper limit (i.e., tolerance) on the probability that drugs are wasted at each HNP

[a]With very minor changes, the shortage cost b, expiration lifetime x, and tolerance δ can be made hospital dependent.

5.3.1 Model Assumptions

In the two-HNP perishable integrated inventory system, we assume that each HNP follows an order-up-to level inventory policy (i.e., S_1 and S_2) for the same perishable drug. We assume that each HNP operates as a continuous review inventory system and that when supply is not disrupted at HNP 1 [HNP 2], the inventory level at HNP 1 [HNP 2] is instantaneously replenished (i.e., zero lead time) to S_1 [S_2]. Therefore, we simply refer to inventory as inventory on-hand since inventory on-hand and inventory-position are equivalent in this case. We assume that each HNP has a static daily demand rate that follows a Poisson distribution (i.e., q_1 and q_2). Poisson distributed demand implies that the time between demand observations is exponentially distributed. For perishability, we assume that the drug has no quality decay (Khan et al., 2014) and a deterministic lifetime of x days since expiration dates for drugs are well regulated by pharmaceutical manufacturers. Also,

we assume the expiration lifetime begins when the drug arrives at the HNP and that first-in-first-out protocols are in place like they are at the University of Michigan's Central Pharmacy.

For supply, we assume that each HNP has a supplier (independent of the other HNP's supplier) that experiences supply chain disruptions that follow a two-state supply process (Tucker et al., 2020b). In a HNP setting, purchasing contracts are well established between a healthcare network and their supplier/wholesaler, and there often exists restrictions on purchasing from suppliers/wholesalers outside of the purchasing contract. We note that truly independent suppliers require knowledge of the entire supply chain which may not exist in pharmaceutical supply chains nor any other type of supply chain. Further, since some drugs are made by a single pharmaceutical company (e.g., vinblastine and vincristine sulfate [generic oncology drugs]; Tucker et al., 2020b), we also consider the case where the two HNPs use the same supplier/wholesaler in Section 5.4.3.

By integrated inventory system, we assume that the HNPs participate in lateral transshipments when one HNP has zero inventory on-hand (due to a supply chain disruption), and the other HNP has positive inventory on-hand (i.e., reactive lateral transshipments, Paterson et al., 2011). We assume that the transport time for lateral transshipments between the two HNPs is zero and that a transshipment cost (i.e., t_{12} or t_{21} depending on the direction of the transshipment) is incurred for each drug transshipped. In our HNP setting of interest, these two assumptions imply that if HNP 1 has zero inventory on-hand and has a patient that needs the drug, the patient immediately receives the drug from HNP 2 given HNP 2 has positive inventory on-hand. Consequently, a cost of t_{21} is incurred. Also, when both HNPs have zero inventory on-hand, a shortage cost (i.e., b) is incurred for each unsatisfied demand. We assume that unsatisfied demand on any given day is completely lost (i.e., lost-sales) and that this unsatisfied demand does not impact the demand rate on another day (i.e., independent). In our numerical analysis (in Section 5.5), we consider a pre-compounded form of Fentanyl that is given to patients during surgery. Given this form of Fentanyl is unavailable, the University of Michigan's Central Pharmacy prescribes one of three alternatives (based on their availability) and still completes the surgery as scheduled (i.e., lost-sales). Since surgeries only happen one time, the unsatisfied demand for this form of Fentanyl has no impact on the future demand for this form of Fentanyl (i.e., independent). For the alternatives, we want to note that the first alternative requires additional resources as it involves compounding the drug in-pharmacy whereas the second and third alternative are morphine and hydromorphone, respectively. Through conversations with the pharmacists at the University of Michigan's Central Pharmacy, the second and third alternative are difficult changes to make which further highlights the importance of having the first-choice drug on-hand when needed. Furthermore, in this case, the shortage cost (i.e., b) can be seen as the incremental cost of having to use one of the three alternatives instead of the first-choice drug. It is worth noting that when substitutes are not available for a first-choice drug, the shortage cost (i.e., b) needs to capture the impacts of delaying and/or cancelling treatment. In

this case, the shortage cost needs to be high enough to capture the detriments of sub-optimal care and ultimately encourage extra inventory on-hand (i.e., design for resiliency), but not too high such that excess inventory is held on-hand (i.e., design for efficiency; see Ivanov, 2021). We assume that the cost of a shortage exceeds the cost of a lateral transshipment in either direction (i.e., $b > t_{12}$ and $b > t_{21}$) as drug shortages can be life-threatening in our pharmaceutical setting. If $b \leq t_{12}$ and $b \leq t_{21}$, it would always be optimal to follow a non-integrated inventory policy (i.e., act as independent HNPs).

5.4 Two-HNP Integrated Inventory System

We model a two-HNP integrated inventory system with independent suppliers experiencing supply chain disruptions (see Figure 5.1). First, we consider a two-HNP non-perishable integrated inventory system where we model the inventory system as a continuous time Markov chain (in Section 5.4.1.1) and express the expected cost per day (in Section 5.4.1.2). Then, we present the exact model (in Section 5.4.1.3) and approximate model (in Section 5.4.1.4) to solve for the inventory policies that minimize the exact and approximate expected cost per day, respectively. In Section 5.4.1.4, we also evaluate the accuracy of the approximate model. Then, we define a perishability condition which provides a way of extending the two-HNP *non-perishable* integrated inventory policies to a two-HNP *perishable* integrated inventory system (in Section 5.4.2). We close this section by presenting two inventory systems that we are interested in comparing to the two-HNP integrated inventory system with two independent suppliers (in Section 5.4.3).

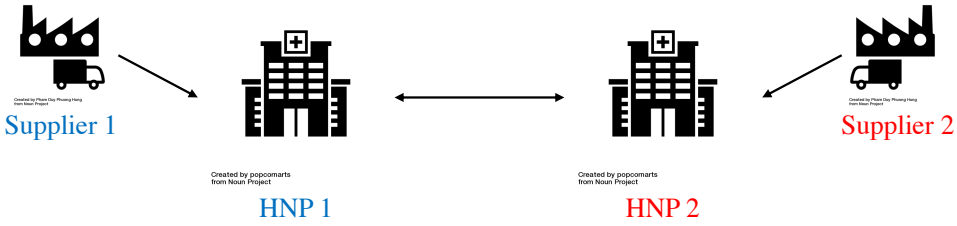

Figure 5.1: Two-HNP integrated inventory system with independent suppliers.

5.4.1 Two-HNP Non-perishable Integrated Inventory System

5.4.1.1 Non-perishable Integrated Inventory System Markov Chain

For the two-HNP non-perishable integrated inventory system, at HNP 1 and HNP 2, we assume that demand follows a Poisson distribution with static daily demand rates of q_1 and q_2, respectively. HNP 1 [HNP 2] experiences supply chain disruptions according to a two-state supply process

where the time in each state is exponentially distributed with a corresponding disruption rate, λ_1 [λ_2], and recovery rate, μ_1 [μ_2]. Therefore, we can model the integrated inventory system as a continuous time Markov chain as depicted in Figure 5.2. The structure of the Markov chain presented in Figure 5.2 assumes that each order-up-to level (i.e., S_1 and S_2) is at least 1. This is appropriate as our HNP setting of interest tends to stock at least 1 drug dose for any particular drug. A few drugs are ordered on a patient-by-patient basis, but these drugs are outside the scope of this chapter.

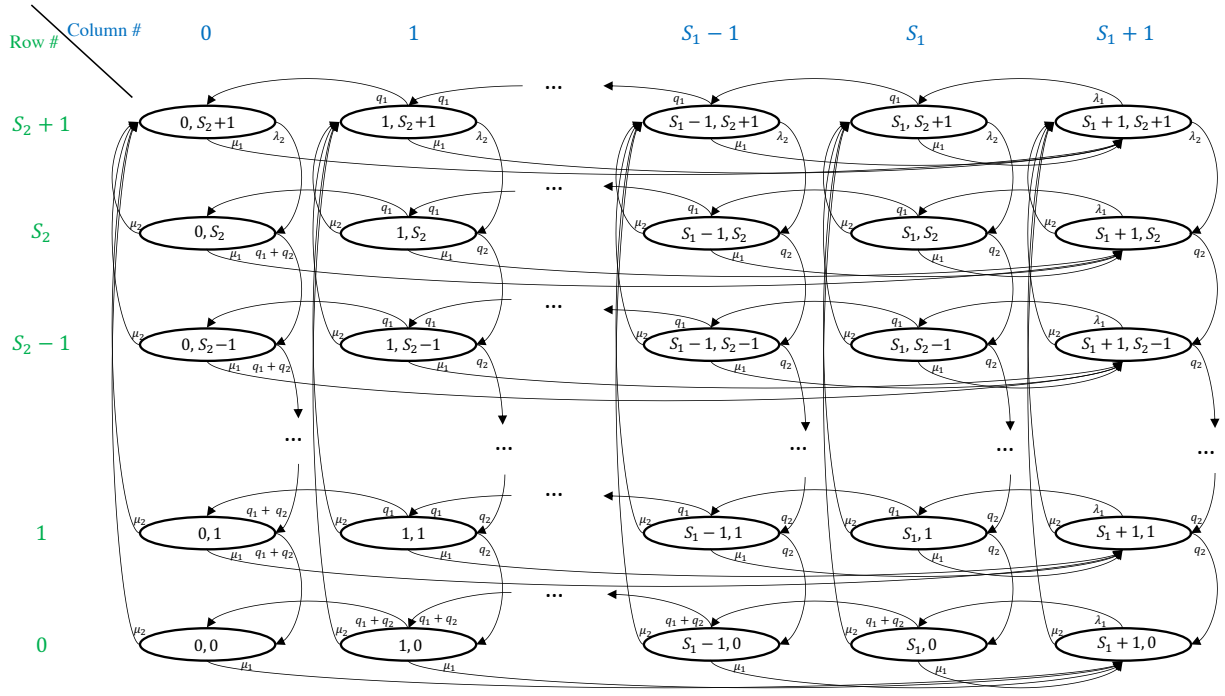

Figure 5.2: Integrated inventory system Markov chain with independent suppliers. Columns $k = 0, \ldots, S_1$ imply that HNP 1 has k inventory on-hand and the supply chain is disrupted. Column $S_1 + 1$ implies that HNP 1 has S_1 inventory on-hand and the supply chain is not disrupted. Row $S_2 + 1$ implies that HNP 2 has S_2 inventory on-hand and the supply chain is not disrupted. Rows $v = S_2, \ldots, 0$ imply that HNP 2 has v inventory on-hand and the supply chain is disrupted.

In Figure 5.2, the state at any time t (X_t; $t \geq 0$) is defined as the tuple (I_1, I_2) ($0 \leq I_1 \leq S_1 + 1$; $0 \leq I_2 \leq S_2 + 1$; $I_1, I_2 \in \mathbb{Z}_0^+$). For $0 \leq I_1 \leq S_1$; $0 \leq I_2 \leq S_2$, we let I_1 and I_2 represent the inventory on-hand at HNP 1 and HNP 2, respectively, when the supply chain is *disrupted* at the respective HNP. We introduce $I_1 = S_1 + 1$ to denote that the inventory on-hand at HNP 1 is S_1 plus the supply chain is *not disrupted*. Similarly, we introduce $I_2 = S_2 + 1$ to denote the inventory on-hand at HNP 2 is S_2 plus the supply chain is *not disrupted*. For emphasis, S_1 and S_2 are the order-up-to levels for the continuous review inventory system and S_1 and S_2 denote the maximum inventory on-hand at any point in time for HNP 1 and HNP 2, respectively.

5.4.1.2 Non-perishable Expected Cost Per Day

The non-perishable expected cost per day consists of holding, shortage, and lateral transshipment costs. When both HNPs have positive inventory on-hand, only holding costs are incurred. When one HNP has positive inventory on-hand and the other HNP has zero inventory on-hand, holding costs and lateral transshipment costs are incurred. When both HNPs have zero inventory on-hand, only shortage costs are incurred. Figure 5.3 illustrates when each cost is incurred based on the inventory on-hand at each HNP in the integrated inventory system.

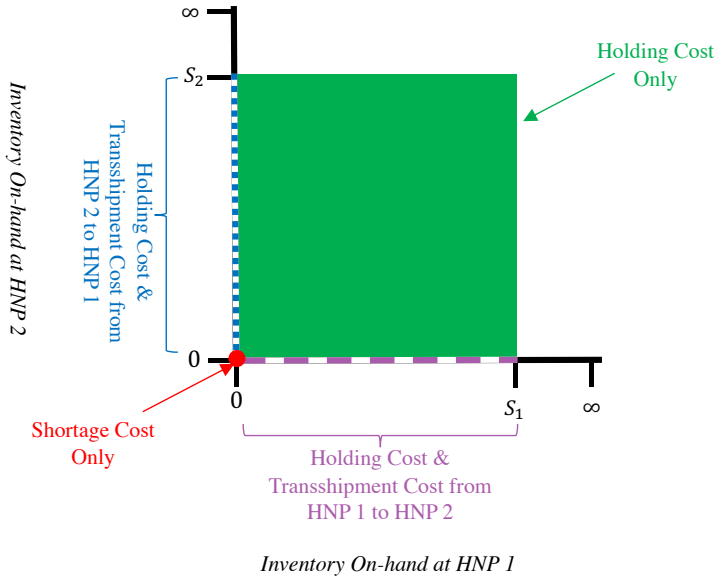

Figure 5.3: Costs incurred in the inventory system.

Using the cost parameters (see Table 5.2) and the long-run probability that the inventory system is in each state $X = (I_1 = i, I_2 = j)$ (i.e., P_{ij}), we have the expected cost per day as presented in Equation (5.1). The expected cost per day is comprised of four components: holding cost only, holding cost and transshipment cost from HNP 1 to HNP 2, holding cost and transshipment cost from HNP 2 to HNP 1, and shortage cost only.

$$\mathbb{E}[\text{cost per day}] = \tag{5.1}$$

$$\overbrace{(h_1S_1+h_2S_2)P_{(S_1+1)(S_2+1)}+\sum_{i=1}^{S_1}(h_1i+h_2S_2)P_{i(S_2+1)}+\sum_{j=1}^{S_2}(h_1S_1+h_2j)P_{(S_1+1)j}+\sum_{i=1}^{S_1}\sum_{j=1}^{S_2}(h_1i+h_2j)P_{ij}}^{\text{Holding Cost Only}}$$

$$+\overbrace{\left(h_1S_1+t_{12}q_2\right)P_{(S_1+1)0}+\sum_{i=1}^{S_1}\left(h_1i+t_{12}q_2\right)P_{i0}}^{\text{Holding Cost and Transshipment Cost from HNP 1 to HNP 2}}+\overbrace{\left(h_2S_2+t_{21}q_1\right)P_{0(S_2+1)}+\sum_{j=1}^{S_2}\left(h_2j+t_{21}q_1\right)P_{0j}}^{\text{Holding Cost and Transshipment Cost from HNP 2 to HNP 1}}$$

$$+\overbrace{b(q_1+q_2)P_{00}}^{\text{Shortage Cost Only}}$$

5.4.1.3 Exact Model: Non-perishable Integrated Inventory Policies

The integrated inventory system Markov chain (see Figure 5.2) is irreducible and it has a finite number of states where all states are positive recurrent. Thus, the limiting distribution exists and is unique (Sericola, 2013). Using Markov chain theory, we can write the balance equations for the integrated inventory system Markov chain to solve for the long-run probability (i.e., steady-state probability) that the integrated inventory system is in each state $X = (I_1, I_2)$ (Kulkarni, 2011). The system of equations to solve for the long-run probability of being in each state of the Markov Chain in Figure 5.2 (i.e., P_{ij}) is provided in Appendix D.1.1.1.

Exact Model: Solving for the Optimal Order-up-to Levels

To solve for the optimal integrated (S_1^*, S_2^*) inventory policy, we exhaustively calculate the expected cost per day using Equation (5.1) for all feasible integer order-up-to level combinations (S_1, S_2): $S_1 \in [S_{1(min)}, S_{1(max)}]$ and $S_2 \in [S_{2(min)}, S_{2(max)}]$. $S_{k(min)}$ and $S_{k(max)}$ are bounds that denote the minimum and maximum order-up-to level for HNP k, respectively. In Appendix D.1.2.5, we illustrate how to strategically select $S_{1(min)}$, $S_{1(max)}$, $S_{2(min)}$, and $S_{2(max)}$. To strategically select these bounds, we use the input parameters for both HNPs and leverage the optimal order-up-to level for a single HNP inventory system with one supplier. We note that it is important to strategically select these bounds instead of making naive selections like $S_{k(min)} = 1$ and $S_{k(max)} = 10^6$ for $k = 1, 2$ to reduce the computation time of the model. We select (S_1^*, S_2^*) corresponding to the solution with the smallest expected cost per day.

Exhaustively calculating the expected cost per day can be time consuming, but the bigger issue is that the system of equations to solve for the long-run probability of being in each state of the Markov chain consists of $(S_1+2)\cdot(S_2+2)$ variables. This large number of variables makes the calculation computationally intensive as the problem size increases. A potential solution is reducing the demand rates and adding a correction to the cost parameters (e.g., when $q_1 = q_2$,

set $q_1' = 1$, $q_2' = 1$, $h_1' = h_1 q_1$, $h_2' = h_2 q_2$, $t_{21}' = t_{21} q_1$, $t_{12}' = t_{12} q_2$, $b' = b q_1 = b q_2$) as this will reduce the order-up-to levels. But, "batching" the demand reduces the variance of the demand as it follows a Poisson distribution (i.e., variance equals the demand rate; Kulkarni, 2011). As an alternative, we present an approximate model in the following section that provides all calculations in closed-form. Further, the approximate model requires a search over a subset of feasible integer order-up-to level combinations instead of all feasible combinations. These two characteristics make the approximate model quick-to-solve which is critical in a hospital pharmacy setting that stocks thousands of different drugs. Fortunately, the approximate model achieves high accuracy.

5.4.1.4 Approximate Model: Non-perishable Integrated Inventory Policies

When creating the approximate model, we make two observations. First, we notice that the expected holding cost per day at HNP 1 and HNP 2 depends on the summation of each column and each row of long-run probabilities in the Markov chain (see Figure 5.2), respectively. This implies that the approximate model needs to accurately estimate the summation of each row and column. Second, the expected shortage cost per day depends on the long-run probability P_{00} and the lateral transshipment costs per day depend on the summation of column 0 and the summation of row 0 of long-run probabilities in the Markov chain (see Figure 5.2) excluding P_{00}. This implies that the approximate model needs to accurately estimate P_{00} and the summation of the other probabilities in column 0 as well as the summation of the other probabilities in row 0 of the Markov chain. It is worth noting that in our hospital pharmacy setting, the demand rates are far greater than the disruption and recovery rates (i.e., $q_1 > \lambda_1 + \lambda_2 + \mu_1 + \mu_2$ and $q_2 > \lambda_1 + \lambda_2 + \mu_1 + \mu_2$) as the vast majority of drugs have an expected daily demand of at least one, the time between supply chain disruptions is far greater than one day, and supply chain disruptions often last weeks, months, or years (Tucker et al., 2020a). When one or both suppliers are disrupted, this encourages movement to three of the corner states of the Markov chain: $(0,0)$, $(S_1+1,0)$, $(0,S_2+1)$.

With these two observations, we use conditional probability where we let $P_i^{(p,s)} = \Pr(I_p = i)$ and $P_{j|i}^{(p,s)} = \Pr(I_s = j | I_p = i)$ to approximate the long-run probabilities for the Markov chain depicted in Figure 5.2. We use the notation p and s to denote the primary HNP (i.e., the HNP conditioned on) and secondary HNP, respectively. Using conditional probability in this way segments the Markov chain by rows and columns and as a result, we lose some of the interaction between the states of the full Markov chain. Hence, the long-run probabilities are *approximate*. However, we find that the conditional probability approach accurately estimates the summation of each row and each column of the exact long-run probabilities (i.e., probabilities that influence the holding cost). Also, since the demand rates push the system to the three corner states when one or both suppliers are disrupted, we find that the conditional probability approach provides good estimates for P_{00} and the summation of the other probabilities in column 0 as well as the summation of the other

probabilities in row 0 of the Markov chain (i.e., probabilities that influence the shortage and lateral transshipment costs). Further, we find that the accuracy of the approximate model improves when conditioning on the HNP p that leads to the larger estimate for P_{00} (see Equations (5.4) and (5.9)). We analyze the accuracy of the approximate model at the end of this Section 5.4.1.4.

We use Equations (5.2)-(5.9) to calculate the approximate joint long-run probability that the integrated inventory system is in state $X = (I_p = i, I_s = j)$: $P_{ij}^{(p,s)} = P_{j|i}^{(p,s)} P_i^{(p,s)}$. We analytically verify that $\sum_{i=0}^{S_p+1} \sum_{j=0}^{S_s+1} P_{j|i}^{(p,s)} P_i^{(p,s)} = 1$.

$$P_{S_p+1}^{(p,s)} = \frac{\mu_p}{\mu_p + \lambda_p} \tag{5.2}$$

$$P_i^{(p,s)} = \frac{\mu_p}{\mu_p + \lambda_p} \left(\frac{\lambda_p}{\mu_p + q_p + q_s P_{0|S_p}^{(p,s)}} \right) \left(\frac{q_p + q_s P_{0|S_p}^{(p,s)}}{\mu_p + q_p + q_s P_{0|S_p}^{(p,s)}} \right)^{S_p - i}; \ i = 1, ..., S_p \tag{5.3}$$

$$P_0^{(p,s)} = \left(\frac{q_p + q_s P_{0|S_p}^{(p,s)}}{\mu_p} \right) P_1^{(p,s)} \tag{5.4}$$

$$P_{S_s+1|i}^{(p,s)} = \frac{\mu_s}{\mu_s + \lambda_s}; \ i = 0, 1, ..., S_p + 1 \tag{5.5}$$

$$P_{j|i}^{(p,s)} = \frac{\mu_s}{\mu_s + \lambda_s} \left(\frac{\lambda_s}{q_s + \mu_s} \right) \left(\frac{q_s}{q_s + \mu_s} \right)^{S_s - j}; \ i = 1, ..., S_p + 1, \ j = 1, ..., S_s \tag{5.6}$$

$$P_{0|i}^{(p,s)} = \left(\frac{q_s}{\mu_s} \right) P_{1|i}^{(p,s)}; \ i = 1, ..., S_p + 1 \tag{5.7}$$

$$P_{j|0}^{(p,s)} = \frac{\mu_s}{\mu_s + \lambda_s} \left(\frac{\lambda_s}{q_p + q_s + \mu_s} \right) \left(\frac{q_p + q_s}{q_p + q_s + \mu_s} \right)^{S_s - j}; \ j = 1, ..., S_s \tag{5.8}$$

$$P_{0|0}^{(p,s)} = \left(\frac{q_p + q_s}{\mu_s} \right) P_{1|0}^{(p,s)} \tag{5.9}$$

Approximate Model: Solving for the Optimal Order-up-to Levels

We replace the exact long-run probability P_{ij} in Equation (5.1) with the approximate long-run probability $P_{ij}^{(p,s)} = P_{j|i}^{(p,s)} P_i^{(p,s)}$ using Equations (5.2)-(5.9). Simplifying, we have the expected cost per day as shown in Equation (5.10). In Equation (5.10), c_p (corresponds to the constant term), e_p (corresponds to the exponential term), and $0 < r_p < 1$ (corresponds to the ratio term) are fixed values that only depend on the input parameters h_p, h_s, b, t_{ps}, t_{sp}, q_p, q_s, λ_p, μ_p, λ_s, μ_s (see Table 5.2) and the order-up-to level S_s. We use the subscript p to clarify that these are the fixed values when approximating the expected cost per day given we condition on HNP p. We also introduce the fixed values θ_p (long-run probability that HNP p is disrupted), θ_s (long-run

probability that HNP s is disrupted), α_p, and β_p to condense the expressions in Equation (5.10).

$$\mathbb{E}[\text{cost per day}|P_{ij}^{(p,s)}] = c_p + h_p S_p + e_p r_p^{S_p} \tag{5.10}$$

$$c_p = h_s(S_s + (\frac{q_s\theta_s}{\mu_s})(\alpha_p - 1)) - h_p(\frac{\theta_p}{\mu_p})(q_p + q_s\theta_s\alpha_p) + t_{ps}q_s\theta_s\alpha_p$$

$$e_p = h_s\theta_s\theta_p(\frac{-q_s\alpha_p}{\mu_s} - \frac{q_p}{\mu_s} + (\frac{q_p+q_s}{\mu_s})\beta_p) + h_p(\frac{\theta_p}{\mu_p})(q_p + q_s\theta_s\alpha_p) - t_{ps}q_s\theta_s\theta_p\alpha_p + t_{sp}q_p\theta_p(1-\theta_s\beta_p) + b(q_p+q_s)\theta_s\theta_p\beta_p$$

$$r_p = \frac{q_p + q_s\theta_s\alpha_p}{\mu_p + q_p + q_s\theta_s\alpha_p}$$

$$\theta_p = \frac{\lambda_p}{\mu_p+\lambda_p};\quad \theta_s = \frac{\lambda_s}{\mu_s+\lambda_s};\quad \alpha_p = (\frac{q_s}{q_s+\mu_s})^{S_s};\quad \beta_p = (\frac{q_p+q_s}{\mu_s+q_p+q_s})^{S_s}$$

Given Equation (5.10), we can solve for the optimal value of S_p (see Equation (5.11)) with all input parameters and the order-up-to level S_s fixed. In Equation (5.10), when $e_p > 0$, the constant function, linear function, and function $e_p r_p^{S_p}$ are convex. Thus, Equation (5.10) is convex as the sum of convex functions is convex. In the convex case, we set the first derivative equal to 0 to find the optimal value of $S_p = \frac{\ln(\frac{-h_p}{e_p \ln(r_p)})}{\ln(r_p)}$; $h_p > 0$, $e_p > 0$, and $0 < r_p < 1$ implies the numerator and denominator are always defined. The maximum function in Equation (5.11) comes from the assumption that $S_p \geq 1$. When $e_p \leq 0$, the constant function, linear function, and function $e_p r_p^{S_p}$ are concave implying Equation (5.10) is concave as the sum of concave functions is concave. In the concave case, we use the boundary condition that $S_p \geq 1$ to find the optimal value of S_p. Equation (5.10) treats S_p as a continuous decision variable, so we round up to the nearest integer (i.e., conservative).

$$S_p = \begin{cases} \max\{1, \frac{\ln(\frac{-h_p}{e_p \ln(r_p)})}{\ln(r_p)}\};\ e_p > 0 \\ 1;\ e_p \leq 0 \end{cases} \tag{5.11}$$

To solve for the optimal order-up-to levels in the integrated inventory system with the approximate long-run probabilities, we consider two properties of the problem. First, the accuracy of the approximate model improves when conditioning on the HNP p that leads to the larger estimate for P_{00}. Second, given we condition on HNP p, we can solve for the optimal value of S_p in closed-form. We use these properties to construct a search algorithm to solve for the optimal order-up-to levels (S_1^*, S_2^*) in the integrated inventory system. The search algorithm consists solely of closed-form expressions and requires a search over a subset of feasible integer order-up-to level combinations making the approximate model quick-to-solve; complexity $\mathcal{O}(S_{2(max)} - S_{2(min)} + S_{1(max)} - S_{1(min)})$. We proceed to formally present the search algorithm.

Algorithm Search algorithm for approximate model.

Input: h_1, h_2, b, t_{12}, t_{21}, q_1, q_2, λ_1, μ_1, λ_2, μ_2 (see Table 5.2)

Step A.1: Fix S_2 and solve for S_1^*
for S_2 in $[S_{2(min)}, S_{2(max)}]$ *% Iterate through feasible S_2 values*
  Solve for S_1^* given S_2 *% See equation* (5.11)
  Store combination $(\lceil S_1^* \rceil, S_2)$ in $\mathbb{S}_{subset}$ *% Store combination with S_1^* rounded up to the nearest integer*
end for

Step A.2: Fix S_1 and solve for S_2^*
for S_1 in $[S_{1(min)}, S_{1(max)}]$ *% Iterate through feasible S_1 values*
  Solve for S_2^* given S_1 *% See equation* (5.11)
  Store combination $(S_1, \lceil S_2^* \rceil)$ in $\mathbb{S}_{subset}$ *% Store combination with S_2^* rounded up to the nearest integer*
end for

Step A.3: Calculate P_{00} estimate for the subset of feasible (S_1, S_2) policies stored in $\mathbb{S}_{subset}$
for (S_1, S_2) in $\mathbb{S}_{subset}$ *% Iterate through the subset of feasible (S_1, S_2) policies*
  Calculate $P_{00}^{(1,2)} = P_{0|0}^{(1,2)} P_0^{(1,2)}$ *% See equations* (5.4) *and* (5.9)
  Calculate $P_{00}^{(2,1)} = P_{0|0}^{(2,1)} P_0^{(2,1)}$ *% See equations* (5.4) *and* (5.9)
  if $P_{00}^{(1,2)} \geq P_{00}^{(2,1)}$ *% Check if it is better to condition on HNP 1*
    Calculate $\mathbb{E}[\text{cost per day} | P_{ij}^{(1,2)}]$ *% See equation* (5.10)
  else
    Calculate $\mathbb{E}[\text{cost per day} | P_{ij}^{(2,1)}]$ *% See equation* (5.10)
  end if
end for

Output: $(S_1^*, S_2^*) = \min_{(S_1,S_2) \in \mathbb{S}_{subset}} \mathbb{E}[\text{cost per day}]$ denotes the optimal non-perishable integrated inventory policy

Approximate Model: Evaluating the Accuracy

We first evaluate the accuracy of the two-HNP integrated approximate model with two independent suppliers using the expected cost per day for the non-perishable integrated inventory system. We evaluate the accuracy by (1) calculating the expected cost per day given a fixed (S_1^*, S_2^*) when using the system of equations for the exact model and (2) calculating the expected cost per day given a fixed (S_1^*, S_2^*) when using the closed-form expressions for the approximate model (in Section 5.4.1.4). Our numerical studies suggest comparable results (difference in expected total cost per day $\leq 2.6\%$; in Appendix D.2.1) supporting the finding that the approximate model achieves high accuracy.

Second, we evaluate the accuracy of the two-HNP integrated approximate model with two independent suppliers using the optimal order-up-to levels for the non-perishable integrated inventory system. We evaluate the accuracy by (1) solving for the optimal order-up-to levels (S_1^*, S_2^*) for the non-perishable integrated inventory system using the exact model in Section 5.4.1.3 where we consider increments of 5 for the exhaustive search and (2) solving for the optimal order-up-to levels (S_1^*, S_2^*) for the non-perishable integrated inventory system using the approximate model in Section 5.4.1.4. We then calculate the expected cost per day with these two sets of (S_1^*, S_2^*) policies using the system of equations for the exact model. Our numerical studies suggest comparable results (difference in expected total cost per day $\leq 0.2\%$; in Appendix D.2.2) further supporting the conclusion that the approximate model achieves high accuracy.

5.4.2 Two-HNP Perishable Integrated Inventory System

5.4.2.1 Defining the Perishability Condition

Next, we examine the case in which the drug has a finite and deterministic lifetime of x days. To capture this, we define a perishability condition which can be seen as a chance constraint at each HNP. The perishability condition addresses the question: Are the order-up-to levels (S_1^*, S_2^*) too large for the expiration lifetime (i.e., x)? Formally, when following the (S_1^*, S_2^*) integrated inventory policy, we want the probability that drugs are wasted at each HNP to be at most δ.

Supply is Never Disrupted at HNP p

We consider when supply is never disrupted at HNP p for $p = 1, 2$ (i.e., no supply chain disruptions; $I_p = S_p + 1$). Using Equations (5.5)-(5.7), we calculate $P^{(p,s)}_{j|S_p+1}$ for $j = 0, 1, ..., S_s + 1$. We note that since we are conditioning on supply never being disrupted at HNP p, the value $P^{(p,s)}_{j|S_p+1}$ for $j = 0, 1, ..., S_s + 1$ from the approximate model is exact in this case. When $j = 1, ..., S_s + 1$, this implies that HNP s has positive inventory-hand. Hence, HNP s is not depending on HNP p for inventory (i.e., HNP p has a demand rate of q_p). When $j = 0$, this implies that HNP s is depending on HNP p through lateral transshipments (i.e., HNP p has a demand rate of $q_p + q_s$). Next, we consider the demand that occurs in a time period of x days (i.e., expiration lifetime). Formally, we define the random variables Q_{x,q_p} and Q_{x,q_p+q_s} as the number of demand arrivals in a time period of x days when the Poisson demand rate is q_p and $q_p + q_s$, respectively. We present the probability that drugs are wasted at HNP p given supply at HNP p is never disrupted (ND) and an expiration lifetime x in Equation (5.12).

$$P_{\text{waste HNP } p|(\text{ND},\, x)} = \left(\sum_{j=1}^{S_s+1} P^{(p,s)}_{j|S_p+1} \right) \Pr(Q_{x,q_p} \leq S_p - 1) + P^{(p,s)}_{0|S_p+1} \Pr(Q_{x,q_p+q_s} \leq S_p - 1) \tag{5.12}$$

Using the Poisson cumulative distribution function, we have the probability that drugs are wasted at HNP p given supply at HNP p is never disrupted and an expiration lifetime x in Equation (5.13).

$$P_{\text{waste HNP } p|(\text{ND},\, x)} = \left(\sum_{j=1}^{S_s+1} P^{(p,s)}_{j|S_p+1} \right) \left(e^{-q_p x} \sum_{i=0}^{S_p-1} \frac{(q_p x)^i}{i!} \right) + P^{(p,s)}_{0|S_p+1} \left(e^{-(q_p+q_s)x} \sum_{i=0}^{S_p-1} \frac{((q_p+q_s)x)^i}{i!} \right) \tag{5.13}$$

Perishability Condition as a Chance Constraint at Each HNP

When following the (S_1^*, S_2^*) integrated inventory policy, we want the probability that drugs are wasted at HNP p ($p = 1, 2$) given supply at HNP p is never disrupted to be at most δ. For $p = 1, 2$, we use $P^{(p,s)}_{j|S_p+1}$ ($j = 0, ..., S_s + 1$) and Equation (5.13) to formally define the perishability condition

in Equation (5.14). When the perishability condition is satisfied, (S_1^*, S_2^*) provides the optimal integrated inventory policy for the perishable system. When the perishability condition is violated (i.e., $P_{\text{waste HNP 1}|(\text{ND}, x)} > \delta$ or $P_{\text{waste HNP 2}|(\text{ND}, x)} > \delta$), we enforce the perishability condition as shown in Section 5.4.2.2.

$$P_{\text{waste HNP 1}|(\text{ND}, x)} = \overbrace{\left(\frac{\mu_2}{\mu_2+\lambda_2}\left(1+\frac{\lambda_2}{\mu_2}\left(1-(\frac{q_2}{q_2+\mu_2})^{S_2^*}\right)\right)\right)\left(e^{-q_1 x}\sum_{i=0}^{S_1^*-1}\frac{(q_1 x)^i}{i!}\right)}^{\text{Only HNP 1 is using the inventory on-hand at HNP 1}}$$
$$+\overbrace{\left(\frac{q_2}{\mu_2}\right)\left(\frac{\lambda_2}{q_2+\mu_2}\right)\left(\frac{\mu_2}{\mu_2+\lambda_2}\right)\left(\frac{q_2}{\mu_2+q_2}\right)^{S_2^*-1}\left(e^{-(q_1+q_2)x}\sum_{i=0}^{S_1^*-1}\frac{((q_1+q_2)x)^i}{i!}\right)}^{\text{HNP 1 and HNP 2 are using the inventory on-hand at HNP 1}} \leq \delta$$

and (5.14)

$$P_{\text{waste HNP 2}|(\text{ND}, x)} = \overbrace{\left(\frac{\mu_1}{\mu_1+\lambda_1}\left(1+\frac{\lambda_1}{\mu_1}\left(1-(\frac{q_1}{q_1+\mu_1})^{S_1^*}\right)\right)\right)\left(e^{-q_2 x}\sum_{j=0}^{S_2^*-1}\frac{(q_2 x)^j}{j!}\right)}^{\text{Only HNP 2 is using the inventory on-hand at HNP 2}}$$
$$+\overbrace{\left(\frac{q_1}{\mu_1}\right)\left(\frac{\lambda_1}{q_1+\mu_1}\right)\left(\frac{\mu_1}{\mu_1+\lambda_1}\right)\left(\frac{q_1}{\mu_1+q_1}\right)^{S_1^*-1}\left(e^{-(q_1+q_2)x}\sum_{j=0}^{S_2^*-1}\frac{((q_1+q_2)x)^j}{j!}\right)}^{\text{HNP 1 and HNP 2 are using the inventory on-hand at HNP 2}} \leq \delta$$

In Equation (5.14), we condition on supply never being disrupted at HNP 1 and HNP 2. First, we condition on each HNP individually to capture that the HNPs participate in lateral transshipments when one HNP has positive inventory on-hand and the other has zero inventory on-hand. Second, we condition on supply never being disrupted because we want to know if the order-up-to levels (S_1^*, S_2^*) are too large for the expiration lifetime (i.e., x). Alternatively, we could have conditioned on supply never being disrupted at each HNP and then multiplied by the respective long-run probability that the supply chain is not disrupted for each HNP (i.e., $\frac{\mu_1}{\lambda_1+\mu_1}$ and $\frac{\mu_2}{\lambda_2+\mu_2}$ for HNP 1 and 2, respectively). However, performing this multiplication would imply that when the supply chain is very unreliable (i.e., $\frac{\mu_k}{\lambda_k+\mu_k}$ for $k = 1, 2$ is very small), the probability that drugs are wasted may be very small simply because we multiplied by $\frac{\mu_k}{\lambda_k+\mu_k}$. As an example, consider a simple single HNP inventory system where demand is exactly 1 per day (i.e., deterministic), the drug has an expiration lifetime of 90 days, and supply chain disruptions occur according to the following process: the supply chain is not disrupted for 1 day and then the supply chain is disrupted for 99 days (i.e., deterministic; the long-run probability that the supply chain is not disrupted is $\frac{1}{100} = 0.01$). Consider you select an order-up-to level of $S = 100$ for this single HNP inventory system. The probability that drugs are wasted when supply is never disrupted for this single HNP inventory system is 1. Multiplying this probability by the long-run probability that the supply chain is not disrupted (i.e., 0.01) results in $1 \cdot 0.01 = 0.01$, which is very small. Calculating the probability

that drugs are wasted in this way hides the fact that every 100 days, you order-up-to 100, satisfy demand for exactly 90 days, and then discard 10 drugs due to the 90-day lifetime of the drug. The order-up-to level (i.e., S) is too high for the demand that occurs in a time period of x days (i.e., expiration lifetime) and this information could have been lost given we multiplied by the long-run probability that the supply chain is not disrupted.

5.4.2.2 Enforcing the Perishability Condition

When the perishability condition is violated (i.e., $P_{\text{waste HNP 1}|(\text{ND}, x)} > \delta$ or $P_{\text{waste HNP 2}|(\text{ND}, x)} > \delta$; see Equation (5.14)) with the optimal order-up-to levels (S_1^*, S_2^*) for the non-perishable integrated inventory system, we enforce the perishability condition to ensure that the order-up-to levels (S_1^*, S_2^*) are appropriate for the expiration lifetime of x days. We notice that $P_{\text{waste HNP 1}|(\text{ND}, x)}$ and $P_{\text{waste HNP 2}|(\text{ND}, x)}$ depend on both S_1^* and S_2^*. Therefore, to enforce the perishability condition, we decrease S_1^* and S_2^* iteratively until the probability that drugs are wasted at HNP p ($p = 1, 2$) given supply is never disrupted at HNP p is at most δ. When iteratively decreasing S_1^* and S_2^*, we always start with S_1^*. We proceed to formally present the algorithm.

Algorithm Enforcing the perishability condition.

Input:
Parameters: h_1, h_2, b, t_{12}, t_{21}, q_1, q_2, λ_1, μ_1, λ_2, μ_2 (see Table 5.2)
Optimal non-perishable policy: (S_1^*, S_2^*)

Step A: Decrease S_1^* and/or S_2^* Until the Perishability Condition is Satisfied
while $P_{\text{waste HNP 1}|(\text{ND}, x)} > \delta$ or $P_{\text{waste HNP 2}|(\text{ND}, x)} > \delta$ *% Perishability condition not satisfied; See equation* (5.14)
    Step A.1: Consider Decreasing S_1^*
    if $P_{\text{waste HNP 1}|(\text{ND}, x)} > \delta$ *% Waste constraint not satisfied at HNP 1*
        $S_1^* = \max\{1, S_1^* - 1\}$ *% Decrease the order-up-to level at HNP 1 by 1 if the boundary condition is not met*
        Calculate $P_{\text{waste HNP 1}|(\text{ND}, x)}$ and $P_{\text{waste HNP 2}|(\text{ND}, x)}$ with (S_1^*, S_2^*) *% See equation* (5.14)
    end if

    Step A.2: Consider Decreasing S_2^*
    if $P_{\text{waste HNP 2}|(\text{ND}, x)} > \delta$ *% Waste constraint not satisfied at HNP 2*
        $S_2^* = \max\{1, S_2^* - 1\}$ *% Decrease the order-up-to level at HNP 2 by 1 if the boundary condition is not met*
        Calculate $P_{\text{waste HNP 1}|(\text{ND}, x)}$ and $P_{\text{waste HNP 2}|(\text{ND}, x)}$ with (S_1^*, S_2^*) *% See equation* (5.14)
    end if

    Step A.3: Check if the Order-up-to Level Boundary Condition is Met at Each HNP
    if $S_1^* = 1$ and $S_2^* = 1$ *% Boundary condition is met at each HNP*
        break while loop (note: perishability condition may still be violated)
    end if
end while

Output: (S_1^*, S_2^*) denotes the optimal perishable integrated inventory policy

5.4.3 Other Inventory Systems for Comparison

In this chapter, we focus on a two-HNP integrated inventory system with independent suppliers. For comparison, we are also interested in studying a (a) two-HNP non-integrated inventory system and (b) two-HNP integrated inventory system with one supplier (see Figure 5.4).

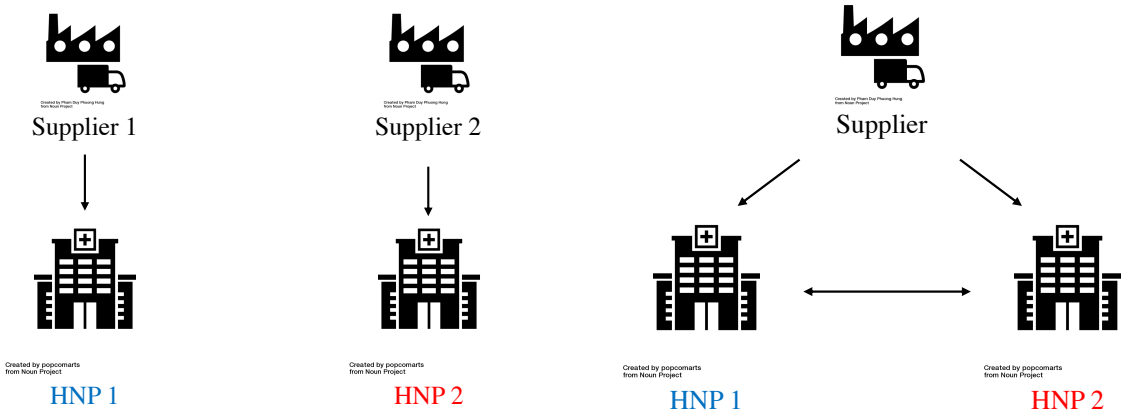

(a) Non-integrated inventory system.

(b) Integrated inventory system with one supplier.

Figure 5.4: Other inventory systems for comparison.

5.4.3.1 Two-HNP Non-Integrated Inventory System

The non-integrated inventory system largely resembles the strict regulations in current practice that generally prohibit HNPs from sharing drugs outside of their network or make it very difficult for HNPs to stay compliant when sharing drugs outside of their network. In this inventory system, the two HNPs act independently of one another, so the non-perishable order-up-to levels are obtained by considering a single HNP single supplier non-perishable inventory system. The optimal order-up-to level S_k for HNP k ($k = 1, 2$) is expressed in Equation (5.15) (see Appendix D.1.2). Equation (5.15) treats S_k as a continuous decision variable, so we round up to the nearest integer (i.e., conservative).

$$S_k = \max\{1, \frac{\ln\left(\frac{-h_k}{(\frac{\lambda_k q_k}{\mu_k+\lambda_k})(\frac{h_k}{\mu_k}+b)\ln(\frac{q_k}{q_k+\mu_k})}\right)}{\ln(\frac{q_k}{q_k+\mu_k})}\} \tag{5.15}$$

To incorporate perishability, we decrease S_k ($k = 1, 2$) until the probability that drugs are wasted at HNP k is at most δ given supply is never disrupted at HNP k and an expiration lifetime x as shown in Equation (5.16) (see Appendix D.1.3).

$$P_{\text{waste HNP } k|(\text{ND}, x)} = e^{-q_k x} \sum_{i=0}^{S_k-1} \frac{(q_k x)^i}{i!} \leq \delta \tag{5.16}$$

5.4.3.2 Two-HNP Integrated Inventory System with One Supplier

We also study a two-HNP integrated inventory system with one supplier. In this case, when the supplier is disrupted, both HNP 1 and HNP 2 are unable to replenish their drug inventory. For the

two-HNP non-perishable integrated inventory system with one supplier, we express the expected cost per day in closed-form (see Appendix D.1.4.1-D.1.4.2), but to solve for the optimal order-up-to levels, we must exhaustively calculate the expected cost per day for all feasible integer order-up-to level combinations (S_1, S_2): $S_1 \in [S_{1(min)}, S_{1(max)}]$ and $S_2 \in [S_{2(min)}, S_{2(max)}]$.

To incorporate perishability, we introduce a constraint on the probability drugs are wasted at each HNP. A one supplier system implies that when supply is never disrupted at one HNP (recall we condition on supply never being disrupted for the perishability condition), supply is also never disrupted at the second HNP since they have the same supplier. As a result, the probability that drugs are wasted only depends on the individual HNP's order-up-to level like the non-integrated inventory system. Therefore, we decrease S_k $(k = 1, 2)$ until the probability that drugs are wasted at HNP k is at most δ given supply is never disrupted at HNP k and an expiration lifetime x as shown in Equation (5.16) (see Appendix D.1.5). We provide a summary of all inventory systems in Table 5.3.

Table 5.3 Description of the inventory systems.

Notation	Description
Two Supplier Integrated Non-perishable	Two independent supplier integrated inventory system without enforcing the perishability condition (in Section 5.4.1)
Two Supplier Integrated Perishable	Two independent supplier integrated inventory system with the perishability condition enforced when appropriate (in Section 5.4.2)
Non-Integrated Non-perishable	Non-integrated inventory system where each HNP acts independently without enforcing the perishability condition (in Section 5.4.3.1 and Appendix D.1.2)
Non-Integrated Perishable	Non-integrated inventory system where each HNP acts independently with the perishability condition enforced when appropriate (in Section 5.4.3.1 and Appendix D.1.3)
One Supplier Integrated Non-perishable	One supplier integrated inventory system without enforcing the perishability condition (in Section 5.4.3.2 and Appendix D.1.4)
One Supplier Integrated Perishable	One supplier integrated inventory system with the perishability condition enforced when appropriate (in Section 5.4.3.2 and Appendix D.1.5)

5.5 Numerical Analysis

We consider our approximate model and numerically study a two-HNP perishable integrated inventory system. We consider Fentanyl 50mcg/mL 30mL (Fentanyl) which is a synthetic opioid used to provide pain relief (Centers for Disease Control and Prevention, 2021). This drug is essential in hospital operations as the University of Michigan's Central Pharmacy prescribes the drug for continuous pain relief during surgeries. This drug is defined as a 503B drug which is a pre-compounded drug that is ready-for-use when delivered to the HNP (Jones, 2020). The average wholesale price of the drug is approximately $25 per dose (IBM Watson Health, 2022).

Table 5.4 provides the input parameters for the Fentanyl base case. For both HNPs, we consider equal holding costs (i.e., $h_1 = h_2 = \$0.025$; holding cost relative to drug price 0.001, Jia & Zhao,

2017). We select a shortage cost twice the cost of the drug (i.e., $b = \$50$). We consider equivalent lateral transshipment costs (i.e., $t_{12} = t_{21}$) since the HNPs are sharing the same drug and from discussions with our hospital pharmacy collaborators, we assume that lateral transshipment costs are 25% the cost of a shortage (i.e., $t_{12} = t_{21} = \$12.50$). We also assume that the two HNPs are comparable in size supporting equivalent daily demand rates (i.e., $q_1 = q_2$). We use the average daily demand from the most recent six months of available data at the University of Michigan's Central Pharmacy (i.e., October 22, 2020-March 22, 2021; missing data imputed using linear interpolation with respect to the day of the week) to get a daily demand rate estimate of $q_1 = q_2 =$ 45/day. We assume the expiration lifetime of the drug is $x = 90$ days which is consistent with the class of 503B drugs. We set the upper limit on the probability that drugs are wasted at each HNP to 0.05 (i.e., $\delta = 0.05$). For both suppliers, we assume that supply chain disruptions occur about every three months (i.e., 90 days; $\lambda_1 = \lambda_2 = \frac{1}{90}$). For both suppliers, we also assume that supply chain disruptions last for about 1 month (i.e., 30 days; $\mu_1 = \mu_2 = \frac{1}{30}$; perishability condition does not need to be enforced) or 3 months (i.e., 90 days; $\mu_1 = \mu_2 = \frac{1}{90}$; perishability condition needs to be enforced). Our hospital pharmacy collaborators have observed disruption durations of 1-3 months. But, due to the difficulty in selecting the disruption rate and recovery rate input parameters, we consider multiple combinations and ranges of these input parameters throughout the numerical analysis (see Table 5.4).

Table 5.4 Numerical analysis input parameters.

Notation	Base Case	Sensitivity	Description
h_1	\$0.025		Holding cost per day per drug for HNP 1 ($h_1 > 0$; Jia & Zhao, 2017)
h_2	\$0.025		Holding cost per day per drug for HNP 2 ($h_2 > 0$; Jia & Zhao, 2017)
b	\$50		Shortage cost per unsatisfied demand ($b > 0$; lost-sales)
t_{12}	\$12.50	$[\$0, \$50]$	Lateral transshipment cost from HNP 1 to HNP 2 per drug ($t_{12} > 0$)
t_{21}	\$12.50	$[\$0, \$50]$	Lateral transshipment cost from HNP 2 to HNP 1 per drug ($t_{21} > 0$)
x	90		Expiration lifetime in days ($x \geq 1$)
q_1	45		Daily demand rate for HNP 1 ($q_1 > 0$; Poisson distribution)
q_2	45		Daily demand rate for HNP 2 ($q_2 > 0$; Poisson distribution)
δ	0.05		Upper limit (i.e., tolerance) on the probability that drugs are wasted at each HNP
λ_1	$\frac{1}{90}$	$[0.01, 0.3]$	Supply chain disruption rate with respect to days for HNP 1
μ_1	$\frac{1}{30}, \frac{1}{90}$	$[0.01, 0.1]$	Supply chain recovery rate with respect to days for HNP 1
λ_2	$\frac{1}{90}$	$[0.001, 0.3]$	Supply chain disruption rate with respect to days for HNP 2
μ_2	$\frac{1}{30}, \frac{1}{90}$	$[0.001, 0.1]$	Supply chain recovery rate with respect to days for HNP 2

In Sections 5.5.1-5.5.2, we provide the expected total, holding, shortage, and lateral transshipment cost per day from the closed-form approximate model. We note that the closed-form expressions are derived for a non-perishable inventory system. However, our HNP setting of interest requires a small probability of waste (i.e., $\delta = 0.05$) and consequently, there is a small chance of waste when following the optimal perishable order-up-to levels. As a result, there is a negligible difference between the expected cost per day (sum of holding, shortage, and lateral transshipment costs) for a non-perishable and perishable inventory system when using the optimal perishable order-up-to levels (e.g., difference with respect to each cost component $\leq 0.3\%$ when simulating

the two inventory systems with a warm-up period of 500 days, a planning horizon of $10,000$ days, and $1,000$ simulation replications for the base case with $\mu_1 = \mu_2 = \frac{1}{90}$ which enforces the perishability condition). We also present the probability that drugs are wasted given supply is never disrupted at each respective HNP and the optimal perishable order-up-to levels (S_1^*, S_2^*) at each HNP.

5.5.1 Benefit of Sharing Inventory

We show when it is beneficial to allow two HNPs with independent suppliers experiencing supply chain disruptions to share inventory. In Figure 5.5, we consider the Two Supplier Integrated Perishable and Non-Integrated Perishable inventory systems with the base case input parameters (see Table 5.4), but varying lateral transshipment costs (i.e., $t_{12} = t_{21} \in [\$0, \$50]$). The results correspond to infrequent supply chain disruptions (i.e., $\lambda_1 = \lambda_2 = \frac{1}{90}$) and either short supply chain disruptions (i.e., $\mu_1 = \mu_2 = \frac{1}{30}$) or long supply chain disruptions (i.e., $\mu_1 = \mu_2 = \frac{1}{90}$). Since we are using the same input parameters for both HNPs, the approximate model leads to identical order-up-to levels (see column 4) and consequently, the same probability of waste (see column 3).

We find that (a) the greatest benefit of an integrated inventory system (i.e., Two Supplier Integrated Perishable) is when the system has small lateral transshipment costs relative to the cost of a shortage (see row 1 column 1). From a shortage perspective, we find that (b) an integrated inventory system (i.e., Two Supplier Integrated Perishable) always reduces the expected shortage cost regardless of the cost of a lateral transshipment (see row 1 column 2). From a waste perspective, we find that (c) when the integrated inventory system (i.e., Two Supplier Integrated Perishable) and non-integrated inventory system (i.e., Non-Integrated Perishable) have the same optimal order-up-to levels, the probability that drugs are wasted with the integrated inventory system is slightly less since sharing between the HNPs occur (see triangular curves in column 3). We also find that when the order-up-to levels for the integrated inventory system (i.e., Two Supplier Integrated Perishable) are larger than the order-up-to levels for the non-integrated inventory system (i.e., Non-Integrated Perishable), the probability that drugs are wasted may be greater with the integrated inventory system compared to the non-integrated inventory system (see $t_{12} = t_{21} = \$50$ and circular curves in column 3).

The most interesting finding is that (d) it is not always beneficial to operate as an integrated inventory system given the cost of a lateral transshipment is less than the cost of a shortage. Instead, there exists a lateral transshipment cost $t^* \leq b$ such that it becomes optimal to operate as a non-integrated inventory system given $t_{12} = t_{21}$. This t^* may be strictly less than b. This finding holds in integrated inventory systems without supply chain disruptions (e.g., Needham & Evers, 1998) and integrated inventory systems with supply chain disruptions and one supplier (e.g., Nasr et al.,

2012). We proceed to elaborate on the intuition behind why this finding holds in our two-HNP integrated inventory system with independent suppliers experiencing supply chain disruptions. The intuition stems from the number of shortages plus the number of lateral transshipments from the time that both suppliers are simultaneously disrupted to the time that both suppliers are simultaneously not disrupted. During this time frame, in the integrated inventory system with two independent suppliers, the number of lateral transshipments plus the number of shortages is always greater than or equal to the number of shortages in the non-integrated inventory system (recall that there are no lateral transshipments in the non-integrated inventory system). As a concrete example, consider the case that both suppliers are currently disrupted, HNP 1 has 1 unit of inventory on-hand, and HNP 2 has 0 units of inventory on-hand. Next, assume that the next sequence of events is (A) demand occurs at HNP 2, (B) demand occurs at HNP 1, (C) the supplier for HNP 1 recovers, and (D) the supplier for HNP 2 recovers. From the time that both suppliers are simultaneously disrupted to the time that both suppliers are simultaneously not disrupted, the integrated inventory system experiences 1 lateral transshipment and 1 shortage. The non-integrated inventory system only experiences 1 shortage.

In Figure 5.5, the expected total cost illustrates this finding (see row 1 column 1) where $t^* = \$42.50$ for $\mu_1 = \mu_2 = \frac{1}{90}$ (see triangular curves). Given $t_{12} \neq t_{21}$, the result still holds, but there exists lateral transshipment costs $t_{12}^* \leq b$ and $t_{21}^* \leq b$ such that it becomes optimal to operate as a non-integrated inventory system. In the remainder of the chapter, we claim the lateral transshipment costs are sufficiently less than the cost of a shortage if $t_{12} < t_{12}^*$ and $t_{21} < t_{21}^*$ (i.e., it is optimal to operate as an integrated inventory system). In our HNP setting, drugs are critical to patient care and can be life-saving. Hence, the lateral transshipment cost is generally sufficiently less than the cost of a shortage.

Building upon our finding (d), we are interested in gaining a better understanding of the sensitivity of t^*. Specifically, with our interest in supply chain disruptions, we analyze how the disruption and recovery rate influence t^*. Recall that t^* ($t^* \leq b$) is the value of the lateral transshipment cost such that it becomes optimal to operate as a non-integrated inventory system given $t_{12} = t_{21}$. An alternative interpretation is that for all lateral transshipment costs $t_{12} = t_{21} < t^*$, it is beneficial to operate as an integrated inventory system as it results in cost savings. In Appendix D.3.1, we present the full analysis where we find that (e) the value of t^* is largely influenced by the disruption rate and recovery rate. More specifically, when $\lambda_1 = \lambda_2$ and $\mu_1 = \mu_2$, t^* decreases as the long-run probability that the supply chain is disrupted increases. Consider the case when both suppliers have a very large long-run probability that the supply chain is disrupted. This finding implies that the lateral transshipment cost often needs to be much less than the cost of a shortage for it to remain beneficial to operate as an integrated inventory system (e.g., the lateral transshipment cost must be less than 40% the cost of a shortage for it to be beneficial to operate as an integrated inventory

system; see Appendix D.3.1). We also find that (f) regardless of the disruption and recovery rate combination, we reach the same conclusion that the lateral transshipment cost must be sufficiently less than the cost of a shortage for it to be beneficial to operate as an integrated inventory system.

In contrast to strict regulations in current practice that generally prohibit HNPs from sharing drugs outside of their network or make it very difficult for HNPs to stay compliant when sharing drugs outside of their network, results (a)-(c) along with lateral transshipment costs that are sufficiently less than the cost of a shortage (result (d)) show that policy makers may consider relaxing these restrictions to improve the performance of the HNPs.

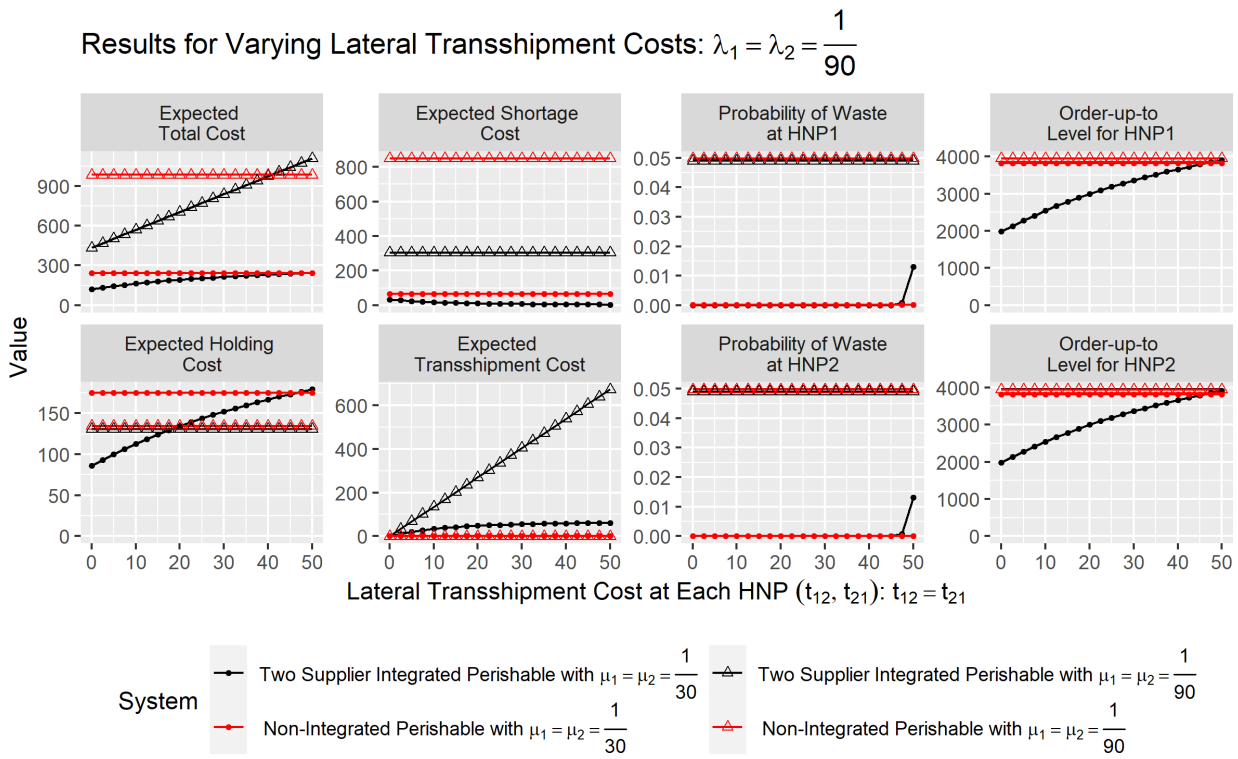

Figure 5.5: Results for the benefit of sharing inventory.

5.5.2 Supply Chain Disruption Characteristics

We now provide insights into how the supply chain disruption characteristics of the two HNPs influence whether it is beneficial to share inventory. In Figure 5.6, we present the Two Supplier Integrated Perishable and Non-Integrated Perishable inventory systems with the base case input parameters (see Table 5.4), but varying disruption and recovery rates at HNP 1 and HNP 2. With the base case input parameters, we want to emphasize that $t_{12} = t_{21} = \$12.50$ is sufficiently less than the cost of a shortage $b = \$50$ for all disruption and recovery rate combinations in this analysis to result in cost savings when using the integrated inventory system (see result (d) in Section

5.5.1). We consider varying recovery rates (i.e., $\mu_1 \in [0.01, 0.1]$) at HNP 1 such that $\lambda_1 = \frac{1}{3}\mu_1$. For all (λ_1, μ_1) combinations, the long-run probability that the supply chain is not disrupted is $\frac{\mu_1}{\mu_1+\lambda_1} = \frac{3}{4} = 0.75$. Also, all (λ_1, μ_1) combinations correspond to expected disruption durations (i.e., down time) that are shorter than the expected time until a disruption given the supply chain is currently not disrupted (i.e., up time). For the second HNP, we consider the same long-run probability that the supply chain is not disrupted (i.e., $\lambda_2 = \frac{1}{3}\mu_2$; $\frac{\mu_2}{\mu_2+\lambda_2} = 0.75$), but only distinguish two recovery rates: long disruptions (i.e., $\mu_2 = 0.01$) and short disruptions (i.e., $\mu_2 = 0.1$).

In Figure 5.6, we notice that the HNPs have identical order-up-to levels (see column 4) and consequently, the same probability of waste (see column 3) when the input parameters for both HNPs are identical. Identical in Figure 5.6 occurs when (i) $\mu_1 = 0.01$ (see x-axis) and $\mu_2 = 0.01$ (see circular curves) and (ii) $\mu_1 = 0.1$ (see x-axis) and $\mu_2 = 0.1$ (see triangular curves). We find that (a) an integrated inventory system is beneficial as it minimizes the expected total cost and shortage cost per day regardless of the supply chain disruption patterns. Also, similar to past research (Chapter 3), we find that (b) when solving for the optimal order-up-to levels, it is insufficient to only consider the long-run probability that the supply chain is not disrupted. Both the disruption rate and recovery rate impact the optimal order-up-to levels (see column 4) and consequently, the cost metrics.

We next consider the expected length of a disruption (i.e., $\frac{1}{\mu_1}$ and $\frac{1}{\mu_2}$) to verify that it is beneficial to operate as an integrated inventory system even when the expected length of a disruption is long in comparison to the expiration lifetime x. We find that (c) when the expected length of a disruption is long in comparison to the expiration lifetime x for HNP 1 (e.g., $\mu_1 = 0.01$; $\frac{1}{\mu_1} = 100$ days $> x = 90$ days), the decrease in the expected total cost with the Two Supplier Integrated Perishable inventory system (see black curves) in comparison to the Non-Integrated Perishable inventory system (see red curves) shows that it is beneficial to operate as an integrated inventory system. This statement holds when the expected length of a disruption for HNP 2 is long in comparison to the expiration lifetime x (e.g., $\mu_2 = 0.01$; $\frac{1}{\mu_2} = 100$ days $> x = 90$ days; see circular curves) and short in comparison to the expiration lifetime x (e.g., $\mu_2 = 0.1$; $\frac{1}{\mu_2} = 10$ days $< x = 90$ days; see triangular curves).

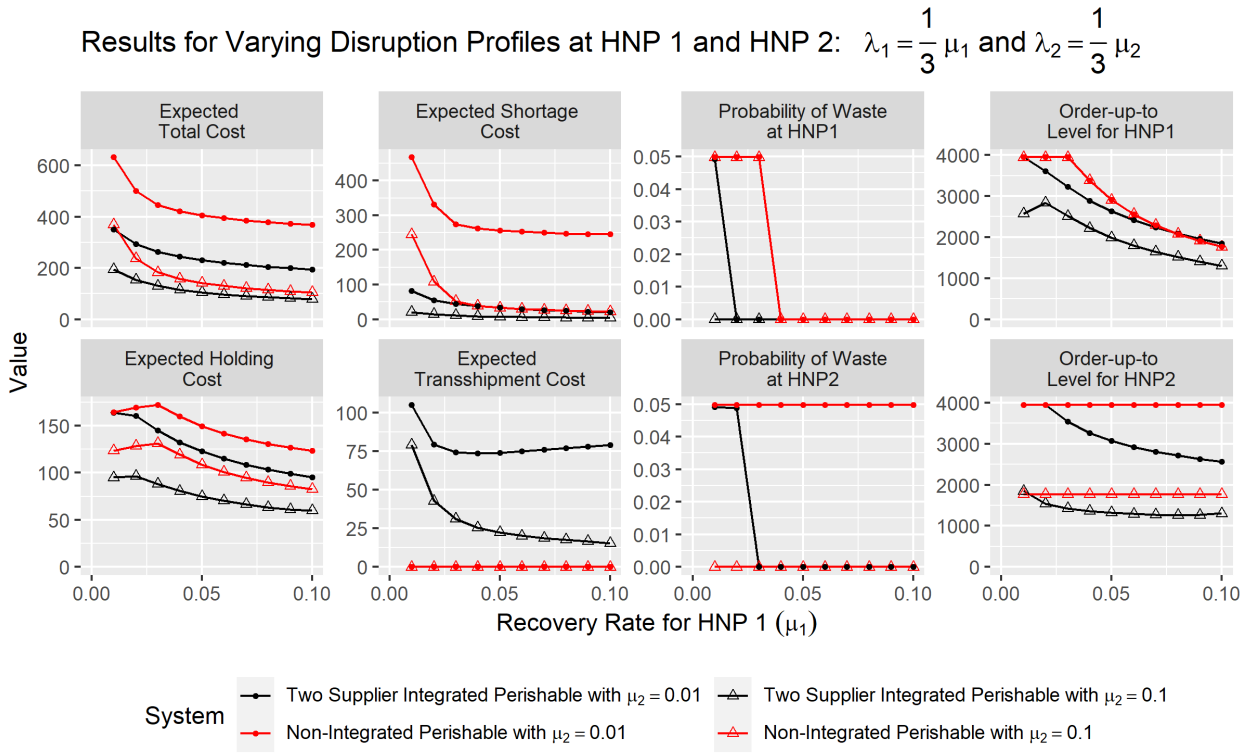

Figure 5.6: Results for supply chain disruption characteristics.

The final interesting finding is that (d) when selecting a partner HNP, it is critical to consider the supply chain disruption profile (i.e., disruption rate and recovery rate) of the partner HNP. Other research supports that a system's performance and decision-making depends on the supply chain disruption profile (e.g., Chapter 3; Rozhkov et al., 2022; Li et al., 2023). But, we hone in on how the disruption rates and recovery rates impact the performance of the HNPs and how these rates influence whether the HNPs should operate as an integrated or non-integrated inventory system. From the results in Figure 5.6, it is recommended to select a second HNP with similar disruption and recovery rates to the first HNP. For example, consider $\mu_1 = 0.1$ (see x-axis) and $\mu_2 = 0.1$ (see triangular curves). Recalling that both systems have equivalent holding costs, lateral transshipment costs, and demand rates, both HNPs experience the same benefit of an integrated inventory system from a decrease in order-up-to level perspective. Furthermore, while not shown in the figure, both HNPs have an equivalent reduction in shortage costs and holding costs by operating as an integrated inventory system instead of a non-integrated inventory system. Again, this implies that the two HNPs experience the same benefit of an integrated inventory system.

It is worth noting that as the long-run probability that the supply chain is disrupted increases, there are instances where failing to consider the disruption rate and recovery rate of the partner HNP can cause one HNP to be "worse-off". In Appendix D.3.2, we provide an example of such an

instance. The example considers two HNPs and compares the cost metrics when operating as an integrated versus a non-integrated inventory system. In the example, one HNP reaps benefits (i.e., expected shortage cost with the integrated inventory system is about 4% the shortage cost of the non-integrated inventory system). However, the second HNP is "worse-off" (i.e., expected shortage cost with the integrated inventory system is almost double the shortage cost of the non-integrated inventory system). To take a deeper look at this finding, we start by fixing the disruption rate and recovery rate for HNP 1 such that HNP 1 has infrequent supply chain disruptions (i.e., $\lambda_1 = \frac{1}{90}$) and short supply chain disruptions (i.e., $\mu_1 = \frac{1}{30}$). We vary the disruption rate (i.e., $\lambda_2 \in [0.001, 0.1]$) and recovery rate (i.e., $\mu_2 \in [0.001, 0.1]$) for HNP 2. For all other input parameters, we continue to consider the base case input parameters (see Table 5.4). We then distinguish which HNPs have a reduction in shortages with the integrated inventory system by calculating the percent change in shortages with the Two Supplier Integrated Perishable and Non-Integrated Perishable inventory systems (see Figure 5.7). Figure 5.7(a) presents the results for HNP 1. Figure 5.7(a) illustrates that there are instances where HNP 1 is "worse-off" by operating as an integrated inventory system because HNP 1 experiences an increase in shortages (see areas where percent change > 0). Figure 5.7(b) presents the results for HNP 2. Figure 5.7(b) illustrates that HNP 2 always has a reduction in shortages with an integrated inventory system since the percent change in shortages is always negative. In summary, the results illustrate that HNPs need to take into consideration the disruption rate and recovery rate of the partner HNP's supplier before deciding if it is beneficial to operate as an integrated inventory system.

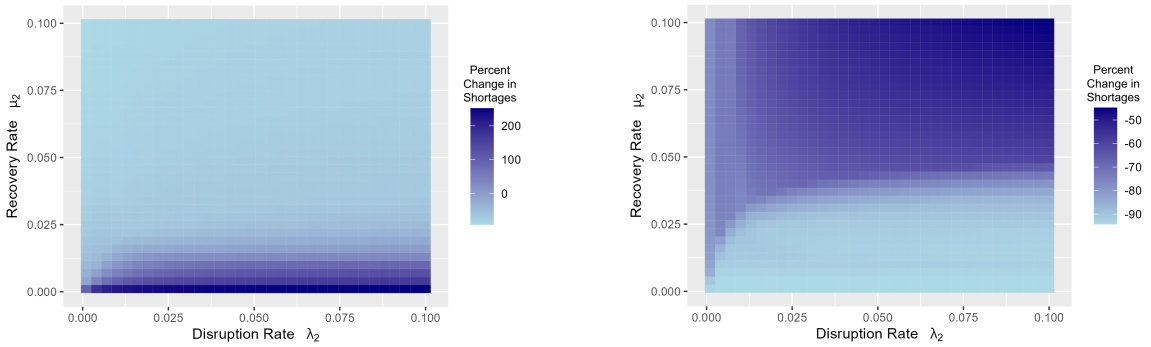

(a) % change in shortages at HNP 1.

(b) % change in shortages at HNP 2.

Figure 5.7: Percent change in shortages at HNP 1 and HNP 2 with the Two Supplier Integrated Perishable and Non-Integrated Perishable inventory systems. Percent change is calculated as follows:
Percent change = $\frac{\text{(Shortages with Integrated Perishable) - (Shortages with Non-Integrated Perishable)}}{\text{Shortages with Non-Integrated Perishable}} * 100$

5.5.3 Additional Takeaways

In the Appendix, we gain additional takeaways by studying (T1) two independent suppliers versus one supplier (in Appendix D.3.3.1) and (T2) the impact of HNPs not sharing inventory given both suppliers are disrupted (i.e., "hoard" inventory; in Appendix D.3.3.2). We find:

(T1) Operating as an integrated inventory system when the two HNPs share the same supplier provides negligible to no benefit. This implies that hospital pharmacy managers should focus on drugs that provide the opportunity to partner with a HNP that has a supplier independent of it's own supplier.

(T2) Given both suppliers are disrupted, it is beneficial to always allow the HNPs to share inventory (i.e., have lateral transshipments). Our numerical studies suggest that hoarding inventory (i.e., no lateral transshipments) when both suppliers are disrupted can cause significantly more shortages in comparison to the integrated inventory system that always allows lateral transshipments.

5.6 Summary and Conclusions

We provide a modeling framework that optimizes decision-making in a perishable inventory system with supply chain disruptions and lateral transshipments. The modeling framework first solves for the optimal (S_1^*, S_2^*) integrated inventory policy for a non-perishable inventory system. To incorporate perishability, we introduce a perishability condition which ensures the probability that drugs are wasted at each HNP does not exceed some tolerance level. When the perishability condition is violated, we illustrate how to enforce this perishability condition. It is worth noting that it is insufficient to ignore the optimal (S_1^*, S_2^*) integrated inventory policy for the non-perishable inventory system and simply consider the largest (S_1^*, S_2^*) integrated inventory policy that satisfies the perishability condition. As one example, there are cases where the optimal (S_1^*, S_2^*) integrated inventory policy for the non-perishable inventory system satisfies the perishability condition. Therefore, the perishability condition does not need to be enforced. In these cases, simply solving for the largest (S_1^*, S_2^*) integrated inventory policy that satisfies the perishability condition can result in a sub-optimal (S_1^*, S_2^*) integrated inventory policy. With our modeling framework, we illustrate when it is beneficial to operate as an integrated inventory system.

Strict regulations in current practice generally prohibit HNPs from sharing drugs outside of their network or make it very difficult for HNPs to stay compliant when sharing drugs outside of their network. However, through discussions with our hospital pharmacy collaborators, there have been instances where hospital pharmacies have shared with other hospital pharmacies outside of

their network. This sharing of drugs required extra time, energy, and resources to stay compliant. But, had this sharing not taken place, patients would have been denied care which emphasizes the benefits of and need for sharing. Our modeling framework takes this idea one step further by not only shedding light on when it is beneficial to share drugs and how beneficial it is to share drugs, but also how to incorporate sharing into the inventory decision-making process. Furthermore, while real-world systems may not satisfy every assumption in our modeling framework, our results show that operating as a non-integrated inventory system leads to overall higher shortage costs in comparison to an integrated inventory system. With U.S. hospitals spending more than $350 million annually due to the labor costs alone associated with drug shortages (Vizient, 2019), operating as an integrated inventory system can lead to significant cost savings.

Our research sheds light on some interesting findings. We first find that the cost of a lateral transshipment must be sufficiently less than the cost of a shortage for it to be beneficial to operate as an integrated inventory system. Sufficient is largely influenced by the supply chain disruption profile (i.e., disruption rate and recovery rate) for each HNP. The intuition behind this finding is that in the integrated inventory system with two independent suppliers, the number of lateral transshipments plus the number of shortages is always greater than or equal to the number of shortages in the non-integrated inventory system (recall that there are no lateral transshipments in the non-integrated inventory system). We also verify this finding for a broad range of disruption and recovery rates due to the difficulty in selecting these input parameters. We find that it is critical for HNPs to consider the duration of and time between supply chain disruptions when selecting a HNP to share inventory with. Given equivalent input parameters at the HNPs, we show that both HNPs experience similar benefits of an integrated inventory system when the HNPs have similar disruption and recovery rates. We also discuss how failing to consider the duration of and time between supply chain disruptions can cause one HNP to be "worse-off". Furthermore, when considering operating as an integrated inventory system, hospital pharmacy managers should partner with a HNP that has a supplier independent from it's own supplier and avoid hoarding inventory. Hoarding inventory can cause significantly more shortages.

We consider a two-HNP integrated inventory system where each HNP has their own independent supplier and for comparison, we also study a two-HNP integrated inventory system with one supplier. Future research can consider the case where two suppliers exist, but have supply chain disruptions that are correlated to one another. For example, consider two suppliers/wholesalers that receive their drug from the same manufacturer. Even though the two suppliers/wholesalers receive drug from the same manufacturer, if a supply chain disruption occurs at the manufacturer, this does not necessarily imply that the suppliers/wholesalers will stock out at the exact same time due to different inventory practices and storage capabilities at the suppliers/wholesalers. We expect that the results for a two-HNP integrated inventory system with correlated suppliers will fall

between the two independent supplier and one supplier models, but we leave the formal analysis for future research. Also, we consider a two-HNP integrated inventory system with independent suppliers, but future research can extend this to a larger number of HNPs. Although, it is unlikely that more than three independent suppliers exist for a particular drug. Furthermore, in the two-HNP integrated inventory system with independent suppliers, we assume that the lead time is zero for shipments from the supplier to the respective hospital network pharmacy. Future research can consider positive deterministic or stochastic lead times. In these instances, demand variability will increase and the order-up-to levels for each HNP will often increase given the perishability condition remains satisfied. When providing insights on selecting a HNP to partner with (in Section 5.5.2), we assume that we have full knowledge of the supply chain disruption profile (i.e., disruption rate and recovery rate) for both HNPs in the two-HNP integrated inventory system. Future research can take our modeling framework and study how sensitive the holding costs, shortage costs, lateral transshipments costs, and probability of waste are to changes in these input parameters. Future research can also study non-exponential times for these input parameters. In our HNP setting, exponential times are consistent with the current pharmaceutical supply chain literature (Tucker et al., 2020b). Furthermore, we emphasize that a HNP needs to consider the duration of and time between supply chain disruptions when selecting a partner HNP. There are instances where failing to do so can cause one HNP to be "worse-off" (e.g., higher shortage costs when operating as an integrated inventory system instead of a non-integrated inventory system). This is under the assumption that the integrated inventory system is centralized (i.e., minimize the expected total cost of the two-HNP integrated inventory system). Future research can consider the model from a decentralized perspective to overcome these "worse-off" consequences. In the decentralized case, our finding that a HNP can be "worse-off" from a shortage cost perspective encourages the use of a decision-rule for lateral transshipments. We recommend that future researchers consider that it may only be optimal for HNP 1 to share inventory if the inventory on-hand is at least I_1 and similarly, it may only be optimal for HNP 2 to share inventory if the inventory on-hand is at least I_2.

The integrated inventory policies justify that perishable inventory systems with supply chain disruptions may benefit from sharing inventory. In our HNP setting of interest, this contradicts the strict regulations in current practice that generally prohibit HNPs from sharing drugs outside of their network or make it very difficult for HNPs to stay compliant when sharing drugs outside of their network. Our findings encourage policy makers to consider the restrictions for drug sharing across HNPs, especially for drugs that can benefit from integration.

Acknowledgments for Chapter 5

The research team would like to thank the Central Pharmacy at the University of Michigan for sharing protocols for lateral transshipments between hospital network pharmacies and providing input parameters for the numerical analysis. We would like to give a special thanks to Christine Betts (University of Michigan Health Pharmacy Services) for providing access to the Fentanyl daily demand data and Nicole Mor (University of Michigan) for the helpful literature searches. The research team would also like to thank the anonymous reviewer and associate editor at *IISE Transactions on Healthcare Systems Engineering* for their feedback on the chapter.

CHAPTER 6

Conclusion

Pharmaceutical drugs are essential to patient care. This makes it crucial to have a large enough quantity of drugs on-hand to avoid shortages, but not too much as this can result in drug waste. In this dissertation, we make contributions at the intersection of operations research and healthcare by creating a series of mathematical models to improve decision-making, provide insights, and challenge administrative policies currently used in practice for hospital pharmacy inventory systems. Regarding insights, the results suggest that decision-making is not always intuitive.

We proceed to briefly summarize the four technical chapters (i.e., Chapters 2-5) of this dissertation. We close this chapter by providing thoughts for future research.

6.1 Summary of Technical Chapters

In Chapters 2-3, we address how perishability and supply chain disruptions make it difficult to select the appropriate levels of drug inventories to maintain. In Chapter 2, we create a simulation-optimization model and develop a Binary Grid-Search algorithm which exploits the structure of the objective function to solve for reorder point s and order-up-to level S periodic review inventory policies in a hospital pharmacy. The research provides an efficient method to solve for near-optimal (s,S) periodic review inventory policies which is essential in a hospital pharmacy inventory system that stocks thousands of different drugs.

In Chapter 3, we derive exact closed-form solutions for length of the review period R and order-up-to level S periodic review inventory policies in a hospital pharmacy. Closed-form solutions imply that the model provides the inventory policy quickly and is easy to implement. This is the first research to derive closed-form solutions for a lost-sales (R,S) periodic review inventory system with supply chain disruptions and perishability. The results using the closed-form model suggest that decision-making is not always intuitive. For example, hospital pharmacy managers must consider supply chain disruptions and perishability simultaneously. Only considering supply chain disruptions can cause the inventory system to be “worse-off" from both a drug waste and

drug shortage perspective.

In Chapter 4, we address how demand disruptions are prevalent in a hospital pharmacy inventory system, but few studies have investigated how inventory policies should adapt to increases and decreases in demand. In this technical chapter, we create a new adaptive periodic review inventory system solely comprised of closed-form expressions to illustrate how a drug's shortage-waste weighting along with the duration of and time between supply chain disruptions influence the benefits (or detriments) of adapting to demand disruptions. We also use the adaptive inventory system to create a ranking procedure. The ranking procedure provides a way of discerning which drugs are of most concern and illustrates which policies to update given that a limited number of inventory policies can be updated.

In Chapter 5, we address how hospital networks do not work in isolation, but it is unclear if sharing drug inventory between hospital networks would help to mitigate drug waste and drug shortages. In this technical chapter, we create a new modeling framework that solves for continuous review order-up-to level inventory policies in a two-HNP integrated inventory system with supply chain disruptions and perishability. This is the first research to analytically model an integrated two-HNP inventory system with supply chain disruptions and perishability. This is also the first research to provide a closed-form approximation that is quick-to-solve. Like Chapter 3, the results using the modeling framework suggest that decision-making is not always intuitive. For example, the lateral transshipment cost must be sufficiently less than the cost of a shortage for it to be beneficial to operate as an integrated inventory system; sufficient largely influenced by the duration of and time between supply chain disruptions. Also, hoarding inventory (i.e., no sharing of inventory when both suppliers are disrupted) can cause significantly more shortages. Furthermore, placing this model into perspective, this technical chapter illustrates that there are instances where perishable inventory systems with supply chain disruptions can benefit from sharing inventory. In a HNP setting, this contradicts the strict regulations in current practice that generally prohibit HNPs from sharing drugs outside of their network or make it very difficult for HNPs to stay compliant when sharing drugs outside of their network.

6.2 Future Research

This dissertation creates models that improve decision-making, provide insights, and challenge administrative policies in hospital pharmacy inventory systems. It is worth noting that this dissertation captures many complexities of a hospital pharmacy inventory system, but additional complexities remain. There exists storage capacity constraints, substitute drugs, interest in the resiliency of inventory policies, data challenges, and many echelons when expanding to the entire pharmaceutical supply chain network. These thoughts provide promising directions for future research and we

proceed to elaborate on these future research directions.

All four technical chapters in this dissertation (see Chapters 2-5) consider that hospital pharmacies are often responsible for thousands of different drugs, making it crucial to have quick-to-solve inventory policies. However, when hospital pharmacies stock these thousands of different drugs on-hand, there exists storage capacities. Therefore, it may not be possible to stock the optimal order-up-to level for every drug held on-hand in the hospital pharmacy. Future research can consider models that solve for all inventory policies simultaneously (e.g., constraint programming; see Little & Coughlan, 2008). This approach will allow storage capacity constraints to be considered in the model. It is also likely that this approach will shed light on which drug factors (e.g., expiration lifetime, duration of supply chain disruptions, time between supply chain disruptions) influence whether the true order-up-to level should be stocked or if a reduced order-up-to level should be stocked to satisfy the storage capacity constraints.

Throughout this dissertation, we aim to reduce drug shortages to ensure that patients have access to the care that they need. It is worth noting that when a patient requires a drug that is currently out of stock (i.e., short), hospital pharmacists may be able to prescribe a substitute drug. However, this substitute drug may cause additional side effects, may be less effective, and may require additional pharmacy resources (e.g., first-choice drug may be pre-compounded whereas the substitute drug may need to be compounded by a hospital pharmacist). These consequences emphasize the importance of having the first-choice drug on-hand when needed. We also note that prescribing the substitute drug while the first-choice drug is out of stock (i) may require an update to the current inventory policy for the substitute drug due to the possible increase in demand and (ii) may increase the likelihood of the substitute drug being out of stock because the supplier/wholesaler may not have the capacity to keep up with the possible increase in demand for the substitute drug. Given these ideas, an interesting future research direction is capturing substitute drugs in inventory decision-making and analyzing how the dependency between drugs influence the optimal inventory policies.

The research in this dissertation relies on expert opinion, data, and sensitivity analyses for the model input parameters. Two input parameters in use throughout all four technical chapters of this dissertation are the supply chain disruption probability (rate) and supply chain recovery probability (rate) in the discrete (continuous) case. These parameters drive the time between supply chain disruptions and duration of supply chain disruptions, respectively. An interesting future research direction is to reverse engineer the models developed in this dissertation to gain insights on the resiliency of an inventory system. For example, Chapter 3 derives closed-form solutions for a (R,S) periodic review inventory system with supply chain disruptions and perishability. Time is treated as discrete for this model and the model includes a constraint to guarantee that a certain proportion of demand is satisfied. Given a hospital pharmacy has a drug with a particular (R,S) policy, it would

be interesting to solve for the supply chain disruption probabilities and supply chain recovery probabilities that ensure the proportion of demand satisfied constraint is met. Taking a reverse engineer approach may shed insights on how resilient a (R,S) inventory policy is for a particular drug of interest.

Another thought is that this dissertation aims to improve decision-making in hospital pharmacy inventory systems particularly in the presence of supply chain disruptions. Taking a step back from the modeling and methodology perspective, another area for future research is improving the tracking of a drug's supply status and the data collected in a hospital pharmacy inventory system. In terms of tracking a drug's supply status at the hospital pharmacy level, it would be beneficial for hospital pharmacies to keep record of all instances where the supplier/wholesaler is unable to supply the drug. Although, it would be ideal for hospital pharmacies to form close connections with their supplier/wholesaler such that they know all instances (e.g., days) that the supplier/wholesaler is unable to supply the drug (i.e., the days that the supplier is disrupted). Furthermore, daily demand data in hospital pharmacies often correlates to how many drug units are administered on a given day. However, daily demand data does not often capture if this was the true demand for the drug or if this was the number of drug units administered simply because that is all of the inventory that the hospital pharmacy had on-hand. With this in mind, it would be beneficial for hospital pharmacies to keep record of their lost-sales for all drugs, even in the case where a substitute can be used. Improving the tracking of a drug's supply status and the data collected can provide additional guidance when selecting input parameters for models of hospital pharmacy inventory systems.

A pharmaceutical supply chain typically consists of an active pharmaceutical ingredient (API) manufacturer, final drug product manufacturer, supplier/wholesaler, and hospital pharmacy. All four technical chapters in this dissertation create models that focus on the supplier/wholesaler and hospital pharmacy levels of the supply chain. Future research can consider decision-making along every level of the supply chain to reduce both drug waste and drug shortages. From a modeling perspective, simulation-optimization provides a viable method to optimize decision-making for this complex system. But, it would be necessary to find current algorithms or develop new algorithms to quickly solve the simulation-optimization model. Also, to provide intuition and insights, it would be beneficial to model a simplified version of the multi-echelon supply chain and provide an analytical model (e.g., Markov chain model, closed-form model) that can be used to optimize decision-making. Furthermore, when considering an entire pharmaceutical supply chain network consisting of an API manufacturer, final drug product manufacturer, supplier/wholesaler, and hospital pharmacy, an interesting characteristic is that the location and identity of the API manufacturer is not generally known (NASEM, 2022). Future research can (i) quantify the benefit of making the location and identity of the API manufacturer known and (ii) gain a better under-

standing of how long it takes a supply chain disruption at the API manufacturer to propagate down to the hospital pharmacy level. Simulation modeling is a viable method to address (i) and (ii).

In conclusion, this dissertation consists of four technical chapters where each chapter has its own model and research contributions. We hope that the contributions at the intersection of operations research and healthcare help ensure that patients have access to the care that they need.

APPENDIX A

Chapter 2 Appendix

We proceed to present the pseudocode for the (s, S) inventory policy simulation model and binary grid-search algorithm in Appendix A.1 and Appendix A.2, respectively.

A.1 Pseudocode for (s,S) Inventory Policy Simulation Model

function StaticSim

for r=1:$|\mathbf{R}|$ *% iterate through all replications*
  Generate Demand d_t for $t = 1,2,...,|\mathbf{T}|$ *% generate demand pattern using appropriate probability distribution*
  Generate Supply y_t for $t = 1,2,...,|\mathbf{T}|$ *% generate supply pattern using appropriate probability distributions*
  ObjValue[r] = InventoryProcess *% calculate objective value using InventoryProcess function*
  ————————
  function InventoryProcess
  for t=1:$|\mathbf{T}|$ *% iterate through each day in planning horizon*
    Step 1: Update inventory based on amount that arrives at the beginning of day t
    Inventory[e]=Inventory[e]+InventoryArrival[1] *% store the amount that arrives as the "newest" inventory*
    Step 2: Update inventory en-route levels
    if $l = 0$ *% lead time is $l = 0$*
      InventoryArrival[1]=0 *% inventory to arrive tomorrow updated when order decision is made below*
    else
      InventoryArrival[i]=InventoryArrival[$i+1$] $\forall i = 1,...,l$ *% shift all inventory en route by one day*
      InventoryArrival[$l+1$]=0 *% inventory to arrive in l days updated when order decision made below*
    end if
    Step 3: Update inventory levels based on the amount of demand observed on day t
    NeedsFilled= d_t *% keep track of amount of demand that needs filled on day t*
    Short[t]=max$\{0, d_t - sum(\text{Inventory})\}$ *% keep track of the amount short on day t*
    $i = 1$ *% keep track of remaining lifetime index to enforce FIFO policy (lifetime is in months)*
    while NeedsFilled $\neq 0$ *% iterate until all demand filled or all inventory used*
      if Short[t]>0 *% check if one is short on day t*
        NeedsFilled=0 *% all inventory is used*
        Inventory[i]=0 $\forall i = 1,2,...,e$ *% update that all inventory is used*
      else if NeedsFilled $\geq$ Inventory[i] *% drug with lifetime remaining i has less inventory than required*
        NeedsFilled=NeedsFilled-Inventory[i] *% update how much demand needs filled after satisfying demand with "oldest" stock*
        Inventory[i]=0 *% indicate all inventory with remaining lifetime i is used*
      else
        Inventory[i]=Inventory[i]-NeedsFilled *% update inventory for this remaining lifetime i*
        NeedsFilled=0 *% no more demand for day t needs filled*
      end if
      $i = i+1$ *% increment lifetime remaining index*
    end while
    Step 4: Discard any expired drugs after observing demand on day t (this only occurs at the end of the month)
    if $t \in \{30,60,90,120,150,180,210,240,270,300,330,360\}$ *% last day of the month*
      Waste[t]=max$\{0,$Inventory$[1]\}$ *% calculate waste in day t*
      Inventory[1]=0 *% update inventory after discarding expired*
    end if
    Step 5: Place order based on (s,S) inventory policy
    if $[sum(\text{Inventory})+sum(\text{InventoryArrival})] < s$ *% place an order because inventory position falls below reorder point s*
      Arrival=S-$[sum(\text{Inventory})+sum(\text{InventoryArrival})]$ *% place an order for S minus the current inventory position*
      Arrival=(y_t)Arrival *% inventory only arrives if supply is not disrupted*
      Order[t]=1(y_t) *% indicate an order is placed on day t if supply is not disrupted*
    else
      Arrival=0 *% do not place an order*
    end if
    InventoryArrival[$l+1$]=Arrival *% update inventory en route based on most recent order decision*
    Step 6: Update inventory levels by age (this only occurs at the end of the month)
    if $t \in \{30,60,90,120,150,180,210,240,270,300,330,360\}$ *% last day of the month*
      Inventory[i]=Inventory[$i+1$] $\forall i = 1,...,(e-1)$ *% subtract 1 month from remaining lifetime*
      Inventory[e]=0 *% "newest" inventory set to 0 until order arrives next day*
    end if
    Hold[t]=$sum(\text{Inventory})$ *% store how many drugs are held on day t*
  end for
  Output: ObjRep[r]= $\sum_{t=31}^{|\mathbf{T}|} \frac{b(Short[t])+z(Waste[t])+o(Order[t])+h(Hold[t])}{(b+z+o+h)(|\mathbf{T}|-30)}$ *% begin at day 31 to remove intialization bias*
  ————————
end for
Output: Expected Cost Per Day = $\frac{1}{|\mathbf{R}|}\sum_{r=1}^{|\mathbf{R}|}$ ObjValue[r]

A.2 Pseudocode for Binary Grid-Search

function BinaryOptGrid

Step 1: Initialize the Grid for s and S
sGrid=$[s_l, s_u]$; by=Δ *% initialize grid for s*
SGrid=$[S_l, S_u]$; by=Δ *% initialize grid for S*
Step 2: Form Matrix and Denote Infeasible Solutions
Matrix is |sGrid| by |SGrid|, but referred to by (s,S) for clarity
ObjMatrix=matrix($|sGrid|, |SGrid|$) *% initialize matrix to store objective values*
ObjMatrix[s,S]=$Infeasible \quad \forall s,S$ such that $s > S$ *% denote infeasible solutions*
Step 3: Determine Objective Value for Main Diagonal
ObjMatrix[s,S]=StaticSim($s=c$, $S=c$) $\quad \forall c \in$ sGrid *% store objective value from StaticSim function given the s and S inputs*
CurrentOpt=[minimum sol_c, $s_{min\,c}$, $S_{min\,c}$] *% store the current minimum objective value from the main diagonal solution and the associated (s,S) in a vector*
CurrentOpt=[$s = \max\{s_l, \lfloor \frac{S_{current}}{2} \rfloor\}$, $S = S_{current}$] *% update the current solution using the midpoint value of s for this column*
Step 4: Perform a Binary Search on the Column With the Current Solution
CurrentOpt=BinarySearch(column, CurrentOpt) *% Perform binary search on the column; store new objective value and associated (s,S)*
OptFound=FALSE *% initialize an indicator variable for iterative process*
MaxIter *% define the maximum number of iterations*
while OptFound==FALSE
    Step 5: Check Solutions to the Left and Right
    % Left of Current Solution
    ObjMatrix[s,S]=StaticSim(s'=$s_{current}$, $S = \max\{s', \max\{S_l, S_{current} - \Delta\}\}$)
    % Right of Current Solution
    ObjMatrix[s,S]=StaticSim(s'=$s_{current}$, $S = \max\{s', \min\{S_u, S_{current} + \Delta\}\}$)
    if CurrentOpt $\leq$ LeftSolution & CurrentOpt $<$ RightSolution *% no improvement to objective value*
        continue algorithm
    else
        % moving improves the objective value
        CurrentOpt=BinarySearch(row, CurrentOpt) *% Perform binary search on the row; store new objective value and associated (s,S)*
    end if
    Step 6: Check Solutions Above and Below
    % Above Current Solution
    ObjMatrix[s,S]=StaticSim(s'=$\max\{s_l, s_{current} - \Delta\}$, $S = S_{current}$)
    % Below Current Solution
    ObjMatrix[s,S]=StaticSim(s'=$\min\{s_u, s_{current} + \Delta\}$, $S = S_{current}$)
    if CurrentOpt $\leq$ AboveSolution & CurrentOpt $<$ BelowSolution *% no improvement to objective value*
        continue algorithm
    else
        % moving improves the objective value
        CurrentOpt=BinarySearch(column, CurrentOpt) *%Perform binary search on column; store new objective value and associated (s,S)*
    end if
    Step 7: Check If Solution Remains Unchanged from Prior Iteration
    if $(sol_{current}, s_{current}, S_{current})$==$(sol_{prior\ iteration}, s_{prior\ iteration}, S_{prior\ iteration})$
        % Perform Forced Binary Search When True
        Perform binary search on the search interval above, to the right, below, and to the left of $(s_{current}, S_{current})$ regardless of the value of the neighboring solutions
        if All solutions are inferior to $(s_{current}, S_{current})$
            OptFound==TRUE *% near-optimal solution found and algorithm completed*
        else
            % a new near-optimal solution is found
            Replace $(s_{current}, S_{current})$ with the new optimal found and continue algorithm
        end if
    else
        continue algorithm
    end if
    if MaxIter is exceeded
        terminate algorithm early
        % if the algorithm is not terminating, increase the number of simulation replications
    end if
end while
Output: Near-Optimal Solution = $(sol_{current}, s_{current}, S_{current})$

APPENDIX B

Chapter 3 Appendix

B.1 General Model for Two-state and Bernoulli Supply Process

B.1.1 General Model Description

In the Chapter 3, we focus on a two-state supply process when modeling supply chain disruptions. However, past research often models supply chain disruptions as a two-state supply process or Bernoulli supply process. Therefore, we generalize the notation using the constants c_0, c_1, and c_2 so that the results presented in Appendix B.1 can be applied to either supply chain disruption process. When considering the probability the system has been disrupted for exactly j consecutive review periods (i.e., π_j), the structure of this value differs for the two-state and Bernoulli supply process. We let $\pi_j = c_1(1-c_0)^j;\ j \geq 1$ and let $\pi_0 = c_1 + c_2$ where c_0, c_1, and c_2 are constants. For the two-state supply process, we have $c_0 = \beta$, $c_1 = \frac{\alpha\beta}{(1-\beta)(\alpha+\beta)}$, and $c_2 = \frac{\beta}{\alpha+\beta} - \frac{\alpha\beta}{(1-\beta)(\alpha+\beta)}$. For the Bernoulli supply process, we have $c_0 = p$, $c_1 = p$, and $c_2 = p(1-p)^0 - p = 0$. We present all model parameters, the two-state supply process parameters, Bernoulli supply process parameters, and the general notation in Table B1.

Table B1 Summary of the modeling notation with both supply processes.

Notation	Description
Decision Variables	
R	Length of the review period (i.e., attempt to place an order every R days; $R \geq 1$)
S	Order-up-to level (i.e., attempt to place an order up to S every R days)
Input Parameters	
h	Holding cost per day per drug ($h > 0$)
k	Fixed ordering cost (i.e., for each order attempted; $k > 0$)
q	Deterministic demand per day ($q > 0$)
e	Expiration lifetime in days ($e \geq 1$)
γ	Maximum proportion of demand not satisfied over infinite horizon ($0 < \gamma < 1$)
Other Parameters	
j	State indicating the number of consecutive review periods that supply has been disrupted
π_j	Probability supply has been disrupted for exactly j consecutive review periods
I_j^+	Average inventory held per day over a review period that starts in state j (in Section 3.4.1)
I_j^-	Average inventory short per day over a review period that starts in state j (in Section 3.4.1)
m	Number of review periods the order-up-to level, S, fully satisfies demand; $m = \lfloor \frac{S}{qR} \rfloor$ (in Section 3.4.2)
Two-state Supply	
α	Disruption probability with respect to the length of the review period ($0 < \alpha < 1$)
β	Recovery probability with respect to the length of the review period ($0 < \beta < 1$)
π_0	Probability supply is disrupted for 0 consecutive review periods (i.e., not disrupted); $\pi_0 = \frac{\beta}{\alpha+\beta}$
π_j	Probability supply is disrupted for exactly j consecutive review periods; $\pi_j = \frac{\alpha\beta}{\alpha+\beta}(1-\beta)^{j-1}, j \geq 1$
Bernoulli Supply	
p	Supply quality ($0 < p < 1$; a larger p implies a better supply quality)
π_j	Probability supply is disrupted for exactly j consecutive review periods; $\pi_j = p(1-p)^j, j \geq 0$
General Notation for Supply Chain Disruptions	
c_0	Term in $\pi_j = c_1(1-c_0)^j; j \geq 1$
c_1	Term in $\pi_j = c_1(1-c_0)^j; j \geq 1$
c_2	Constant representing $c_2 = \pi_0 - c_1$

B.1.2 (R,S) Non-perishable Lost-Sales Inventory System

We present the optimization problem for the (R,S) non-perishable lost-sales inventory system. We proceed to provide the derivation details to solve for the optimal length of the review period R^* and order-up-to level S^* in closed-form.

Optimization Problem w.r.t. m, R, and S

For the non-perishable (R,S) periodic review inventory system with supply chain disruptions, we have the optimization problem presented in Formulation (B.1.1).

$$\begin{aligned} \underset{R,S,m}{\text{minimize}} \quad & C(R,S,m) = \frac{k}{R} + \sum_{j=0}^{m-2} h\pi_j(m-1-j)qR + \pi_m(hI_m^+) + hI_{m-1}^+ \sum_{j=0}^{m-1} \pi_j \\ \text{subject to} \quad & \frac{\pi_m(I_m^-) + I_{m+1}^- \sum_{j=m+1}^{\infty} \pi_j}{q} \leq \gamma \end{aligned} \tag{B.1.1}$$

where:

$$I_m^+ = \frac{1}{R}\int_0^R (S - mqR - qt)^+ dt = \frac{(S-mqR)^2}{2qR} \tag{B.1.2}$$

$$I_m^- = \frac{1}{R}\int_0^R [(S - mqR - qt)^- - (S - mqR - q(t-1))^-] dt = \frac{q}{R}\left(R - \frac{S-mqR}{q}\right) \tag{B.1.3}$$

$$I_{m-1}^+ = \frac{1}{R}\int_0^R (S - (m-1)qR - qt) dt = S - mqR + \frac{1}{2}qR \tag{B.1.4}$$

$$I_{m+1}^- = \frac{1}{R}\int_0^R [(S - (m+1)qR - qt)^- - (S - (m+1)qR - q(t-1))^-] dt = q \tag{B.1.5}$$

Using Equations (B.1.2)-(B.1.5), the definition of π_j, properties of the geometric series, and simplifying, we find:

$$\begin{aligned} \underset{R,S,m}{\text{minimize}} \quad & C(R,S,m) = \frac{k}{R} + h(m-1)qRc_2 + hqRc_1\left(\frac{1}{c_0^2}\right)(mc_0 - 1 + (1-c_0)^m) + \\ & (c_1)(1-c_0)^m\left(h\frac{(S-mqR)^2}{2qR}\right) + h\left(S - mqR + \frac{1}{2}qR\right)c_1\left(\frac{1-(1-c_0)^m}{c_0}\right) + \\ & h\left(S - mqR + \frac{1}{2}qR\right)c_2 \\ \text{subject to} \quad & (c_1)(1-c_0)^m\left(\frac{1}{R}\left(R - \frac{S-mqR}{q}\right)\right) + \left(1 - c_2 - c_1\left(\frac{1-(1-c_0)^{m+1}}{c_0}\right)\right) \leq \gamma \end{aligned} \tag{B.1.6}$$

Find m^* Satisfying $\lfloor \frac{S}{qR} \rfloor$ Derivations

For Any Fixed m, Optimize R and S Derivations

When considering the optimization problem presented in Formulation (B.1.6), we find that for any fixed m, the Hessian is positive semi-definite and thus we have that S and R are jointly convex. When calculating the Hessian, we find:

$$\left[\frac{\partial^2 C(R,S,m)}{\partial S^2}\right]\left[\frac{\partial^2 C(R,S,m)}{\partial R^2}\right] - \left[\frac{\partial^2 C(R,S,m)}{\partial R \partial S}\right]^2 = \left((c_1)(1-c_0)^m\left(h(\frac{1}{qR})\right)\frac{2k}{R^3}\right) \geq 0 \tag{B.1.7}$$

For any fixed m, we can solve the optimization problem presented in Formulation (B.1.6) by optimizing R and S using a Lagrange multiplier (λ_1). Equation (B.1.8) represents the objective function for the optimization problem. Equation (B.1.9) represents the constraint in the optimization problem with the slack variable t^2 since the constraint is an inequality.

$$\begin{aligned} C(R,S,m,t) = & \frac{k}{R} + h(m-1)qRc_2 + hqRc_1\left(\frac{1}{c_0^2}\right)(mc_0 - 1 + (1-c_0)^m) + (c_1)(1-c_0)^m\left(h\frac{(S-mqR)^2}{2qR}\right) \\ & + h\left(S - mqR + \frac{1}{2}qR\right)c_1\left(\frac{1-(1-c_0)^m}{c_0}\right) + h\left(S - mqR + \frac{1}{2}qR\right)c_2 \end{aligned} \tag{B.1.8}$$

$$g(R,S,m,t) = \gamma - (c_1)(1-c_0)^m\left(\frac{1}{R}\left(R - \frac{S-mqR}{q}\right)\right) - \left(1 - c_2 - c_1\left(\frac{1-(1-c_0)^{m+1}}{c_0}\right)\right) - t^2 \tag{B.1.9}$$

Using Equations (B.1.8)-(B.1.9), we present the Lagrange Multiplier Formulation which consists of Equations (B.1.10)-(B.1.13).

$$\begin{aligned}&\frac{\partial C(R,S,m,t)}{\partial S} = \lambda_1 \frac{\partial g(R,S,m,t)}{\partial S} \\ &\implies (c_1)(1-c_0)^m\left(h\left(\frac{S}{qR} - m\right)\right) + hc_1\left(\frac{1-(1-c_0)^m}{c_0}\right) + hc_2 = \lambda_1 c_1 (1-c_0)^m\left(\frac{1}{qR}\right)\end{aligned} \tag{B.1.10}$$

$$\begin{aligned}&\frac{\partial C(R,S,m,t)}{\partial R} = \lambda_1 \frac{\partial g(R,S,m,t)}{\partial R} \\ &\implies \frac{-k}{R^2} + h(m-1)qc_2 + hqc_1\left(\frac{1}{c_0^2}\right)(mc_0 - 1 + (1-c_0)^m) + (c_1)(1-c_0)^m h\left(\frac{-S^2}{2qR^2} + \frac{m^2 q}{2}\right) \\ &+ h(-mq + \frac{1}{2}q)c_1\left(\frac{1-(1-c_0)^m}{c_0}\right) + h(-mq + \frac{1}{2}q)c_2 = \lambda_1 c_1 (1-c_0)^m\left(\frac{-S}{qR^2}\right)\end{aligned} \tag{B.1.11}$$

$$\frac{\partial C(R,S,m,t)}{\partial t} = \lambda_1 \frac{\partial g(R,S,m,t)}{\partial t} \implies 0 = \lambda_1(-2t) \tag{B.1.12}$$

$$\gamma - (c_1)(1-c_0)^m\left(\frac{1}{R}\left(R - \frac{S-mqR}{q}\right)\right) - \left(1 - c_2 - c_1\left(\frac{1-(1-c_0)^{m+1}}{c_0}\right)\right) - t^2 = 0 \tag{B.1.13}$$

Solving For m^* Derivations

The Lagrange Multiplier Formulation consisting of Equations (B.1.10)-(B.1.13) provides the requirements for the optimality of R and S. Before optimizing R and S, we use Equations (B.1.10)-(B.1.13) to find the unique and optimal m^* ($m = \lfloor \frac{S}{qR} \rfloor$) that satisfies these equations. From Equation (B.1.12), we can see we have two cases: $t = 0$ or $\lambda_1 = 0$. To ensure that S covers at least one review period (i.e., $m \geq 1$) and to ensure that all numerators and denominators are defined, we find that only case $t = 0$ is feasible for both a two-state supply process and Bernoulli supply process. $t = 0$ implies that the constraint in Formulation (B.1.6) is satisfied with equality (i.e., the constraint is tight).

Feasible Case: $t = 0$

Using Equation (B.1.13) and setting $t = 0$ leads to:

$$\gamma - (c_1)(1-c_0)^m\left(\frac{1}{R}\left(R - \frac{S-mqR}{q}\right)\right) - \left(1 - c_2 - c_1\left(\frac{1-(1-c_0)^{m+1}}{c_0}\right)\right) - t^2 = 0 \tag{B.1.14}$$

$$\implies \frac{S}{qR} = \frac{-\gamma + c_1(1-c_0)^m(m+1) + (1-c_2-c_1(\frac{1-(1-c_0)^{m+1}}{c_0})}{c_1(1-c_0)^m} \tag{B.1.15}$$

Using the definition of m, we have:

$$m = \lfloor \frac{S}{qR} \rfloor \implies m \leq \frac{S}{qR} < m+1 \tag{B.1.16}$$

Using Equation (B.1.15) and Equation (B.1.16), we have:

$$\text{(Left bound)} \implies m \leq \frac{-\gamma + c_1(1-c_0)^m(m+1) + (1-c_2-c_1(\frac{1-(1-c_0)^{m+1}}{c_0}))}{c_1(1-c_0)^m} \tag{B.1.17}$$

$$\implies m \leq \frac{\ln\left(\frac{\gamma-(1-c_2-\frac{c_1}{c_0})}{\frac{c_1}{c_0}}\right)}{\ln(1-c_0)} \tag{B.1.18}$$

$$\text{(Right bound)} \implies m+1 > \frac{-\gamma + c_1(1-c_0)^m(m+1) + (1-c_2-c_1(\frac{1-(1-c_0)^{m+1}}{c_0}))}{c_1(1-c_0)^m} \tag{B.1.19}$$

$$\implies m > \frac{\ln\left(\frac{\gamma-(1-c_2-\frac{c_1}{c_0})}{\frac{c_1}{c_0}}\right)}{\ln(1-c_0)} - 1 \tag{B.1.20}$$

With the left bound (Equation (B.1.18)) and right bound (Equation (B.1.20)) results, we can simply write:

$$m^* = \left\lfloor \frac{\ln(\frac{\gamma-(1-c_2-\frac{c_1}{c_0})}{\frac{c_1}{c_0}})}{\ln(1-c_0)} \right\rfloor \tag{B.1.21}$$

Four Requirements for the Closed-form Solutions:

(R1) $\frac{\ln\left(\frac{\gamma-(1-c_2-\frac{c_1}{c_0})}{\frac{c_1}{c_0}}\right)}{\ln(1-c_0)} \geq 1 \implies (1-c_0) \geq \frac{\gamma-(1-c_2-\frac{c_1}{c_0})}{\frac{c_1}{c_0}} \implies \gamma \leq 1-c_1-c_2.$

(R2) $\gamma - (1-c_2-\frac{c_1}{c_0}) > 0 \implies \gamma > 1-c_2-\frac{c_1}{c_0}.$

(R3) $c_1 > 0.$

(R4) $0 < c_0 < 1.$

(R1) ensures $m \geq 1$ (i.e., S covers at least one review period). (R2)-(R4) ensure the numerator and denominator in Equation (B.1.21) are defined.

Use m^* to Solve for R^* and S^* Derivations

Having found m^* using Equation (B.1.21), we can use this m^* to solve for S^* and R^*. Recall, we require $\gamma \leq 1-c_1-c_2$, $\gamma > 1-c_2-\frac{c_1}{c_0}$, $c_1 > 0$, and $0 < c_0 < 1$. Also, we have that $t = 0$. Using Equations (B.1.10)-(B.1.13), we find:

$$\text{(Solving for } R^*) \implies R^* = \max\left\{1, \sqrt{\frac{2k(1-c_0)^{m^*}c_0^2c_1}{qh(A_1)}}\right\} \tag{B.1.22}$$

where:

$$
\begin{aligned}
A_1 = &\big(m^*(2c_0^2c_1^2(1-c_0)^{2m^*} + 2c_0c_1^2(1-c_0)^{m^*} + 2c_0^2c_1c_2(1-c_0)^{m^*})\big) \\
&+ \big((c_1^2(1-c_0)^{2m^*} - 2c_0^2\gamma + c_0^2\gamma^2 - 2c_0^2c_1\gamma(1-c_0)^{m^*} + c_0^2 - c_1^2 - c_0^2c_2^2 - c_0c_1^2(1-c_0)^{m^*} \\
&+ 2c_0^2c_1(1-c_0)^{m^*} - 2c_0c_1c_2 + c_0c_1^2(1-c_0)^{2m^*} + 2c_0c_1c_2(1-c_0)^{m^*} - 3c_0^2c_1c_2(1-c_0)^{m^*})\big)
\end{aligned}
$$

$$
\text{(Solving for } S^*) \implies S^* = qR^*\left(\frac{-\gamma + c_1(1-c_0)^{m^*}(m^*+1) + (1-c_2-c_1(\frac{1-(1-c_0)^{m^*+1}}{c_0})}{c_1(1-c_0)^{m^*}}\right) \tag{B.1.23}
$$

B.1.3 $(R=1,S)$ Perishable Lost-Sales Inventory System

Atan & Rousseau (2016) present closed-form solutions for a perishable $(R=1,S)$ inventory system with supply chain disruptions where all demand not met is backordered. The objective of the model is to minimize the expected holding (h), backordering (b), and waste (z) costs. Extending this research to the lost-sales case (i.e., all demand not met is lost), we can write the objective function as presented in Equations (B.1.24)-(B.1.25).

Case 1: $S < eq$

$$
C(R=1,S) = h\sum_{i=0}^{\frac{S}{q}-1}\pi_i(S-(i+1)q) + bq\sum_{i=\frac{S}{q}}^{\infty}\pi_i \tag{B.1.24}
$$

Case 2: $S \geq eq$

$$
C(R=1,S) = \frac{1}{e}\left[h\sum_{i=0}^{e-1}\pi_i(e-i-1)S + z(S-eq)\sum_{i=0}^{e-1}\pi_i + bq\sum_{i=e}^{\infty}\pi_i\right] \tag{B.1.25}
$$

Taking the same approach as Atan & Rousseau (2016), the optimal solution to the lost-sales $(R=1,S)$ model with supply chain disruptions is:

$$
S^* = \min\left\{q\frac{\ln(\frac{h(c_2+\frac{c_1}{c_0})}{bc_1+h\frac{c_1}{c_0}})}{\ln(1-c_0)}, eq\right\} \tag{B.1.26}
$$

B.1.4 (R,S) Perishable Lost-Sales Inventory System (Step 3)

From Appendix B.1.3, it is never optimal to set $S > eq$ given $R = 1$, and this claim also holds given $R > 1$ (in Section 3.5). Therefore, if the non-perishable model has an optimal order-up-to level $S^* > eq$, we first enforce the perishability condition by setting S^* to $S^*_{new} = eq$. Then, we are interested in finding the largest R^*_{new} such that our shortage constraint (see Formulation (B.1.6) in Appendix B.1.2) is satisfied when $S^*_{new} = eq$. Given we set $S^*_{new} = eq$, both holding and ordering costs (i.e., costs in the objective function) are the smallest when R^*_{new} is the largest. Hence, we are interested in finding the largest R^*_{new} because this solution will have the smallest objective value.

Formulating this as an optimization problem, we are interested in:

$$\begin{aligned} &\underset{R}{\text{maximize}} \quad R \\ &\text{subject to} \quad (c_1)(1-c_0)^m\left(\frac{1}{R}\left(R-\frac{S^*_{new}-mqR}{q}\right)\right)+\left(1-c_2-c_1\left(\frac{1-(1-c_0)^{m+1}}{c_0}\right)\right)\leq \gamma \end{aligned} \tag{B.1.27}$$

Solving the optimization problem in Formulation (B.1.27) with the approach described in Appendix 3.5, we find:

$$(\text{Solving for } m^*_{new}) \implies m^*_{new}=m^*=\left\lfloor \frac{\ln(\frac{\gamma-(1-c_2-\frac{c_1}{c_0})}{\frac{c_1}{c_0}})}{\ln(1-c_0)} \right\rfloor \tag{B.1.28}$$

$$(\text{Solving for } R^*_{new}) \implies R^*_{new}=\max\left\{1, \frac{e}{\frac{-\gamma+c_1(1-c_0)^{m^*}(m^*+1)+(1-c_2-c_1(\frac{1-(1-c_0)^{m^*+1}}{c_0}))}{c_1(1-c_0)^{m^*}}}\right\} \tag{B.1.29}$$

$(m^*_{new}, R^*_{new}, S^*_{new}=eq)$ provides the final solution when the perishability condition is enforced.

B.2 Bernoulli Supply Process Model

For a Bernoulli supply process, we provide the closed-form solutions for the non-perishable (R,S) inventory system. Given $S^* \leq eq$, these closed-form solutions provide the optimal inventory policy for the perishable inventory system. Given $S^* > eq$, we provide Step 3 which extends the non-perishable closed-form solutions to account for perishability.

B.2.1 Bernoulli Supply Process (R,S) Non-perishable Closed-form Solutions

For a Bernoulli supply process, we have $c_0 = p$, $c_1 = p$, and $c_2 = 0$. The four requirements are:
(R1) $\gamma \leq 1-c_1-c_2 \implies \gamma \leq (1-p)$.
(R2) $\gamma > 1-c_2-\frac{c_1}{c_0} \implies \gamma > 0$.
(R3) $c_1 > 0 \implies p > 0$.
(R4) $0 < c_0 < 1 \implies 0 < p < 1$.

With requirements (R1)-(R4), we can use the definition of c_0, c_1, and c_2 in Equations (B.1.21), (B.1.22), and (B.1.23) to solve for m^*, R^*, and S^* for the Bernoulli supply process model:

$$(\text{Solving for } m^*) \implies m^*=\left\lfloor \frac{\ln(\gamma)}{\ln(1-p)} \right\rfloor \tag{B.2.1}$$

$$(\text{Solving for } R^*) \implies R^*=\max\left\{1, \sqrt{\frac{2kp(1-p)^{m^*}}{qh(A_1)}}\right\} \tag{B.2.2}$$

where:

$$A_1 = m^*(2p(1-p)^{m^*} + 2p^2(1-p)^{2m^*}) + (-2\gamma + \gamma^2 - 2\gamma p(1-p)^{m^*} + p(1-p)^{2m^*} + (1-p)^{2m^*} + p(1-p)^{m^*})$$

$$\text{(Solving for } S^*) \implies S^* = qR^*\left(\frac{1}{p} + m^* - \frac{\gamma}{p(1-p)^{m^*}}\right) \tag{B.2.3}$$

We numerically verify that $A_1 > 0$ in Equation (B.2.2) for all $\gamma \in \{0.001, 0.002, ..., 0.999\}$ and $p \in \{0.001, 0.002, ..., 0.999\}$ combinations that meet requirement (R1). For each combination, we use the corresponding m^* (see Equation (B.2.1)) for the verification.

B.2.2 Bernoulli Supply Process (R,S) Perishable Closed-form Solutions (Step 3)

Given $S^* > eq$, we set $S^*_{new} = eq$. Using c_0, c_1, and c_2 for a Bernoulli supply process in Equations (B.1.28)-(B.1.29), we have m^*_{new} and R^*_{new} as presented in Equations (B.2.4)-(B.2.5), respectively, along with the definition of S^*_{new} presented in Equation (B.2.6).

$$\text{(Solving for } m^*_{new}) \implies m^*_{new} = m^* = \left\lfloor \frac{\ln(\gamma)}{\ln(1-p)} \right\rfloor \tag{B.2.4}$$

$$\text{(Solving for } R^*_{new}) \implies R^*_{new} = \max\left\{1, \frac{e}{(\frac{1}{p} + m^*) - \frac{\gamma}{p(1-p)^{m^*}}}\right\} \tag{B.2.5}$$

$$\text{(Solving for } S^*_{new}) \implies S^*_{new} = eq \tag{B.2.6}$$

B.3 Number of Simulation Replications

For the numerical analysis in Chapter 3 (in Section 3.6), we utilize a simulation model of the inventory system where we round R^* down to the nearest whole number, consider a 12-month warm-up period (i.e., 360 days), consider a 5-year planning horizon (i.e., 1800 days), and we perform 500 simulation replications. We perform 500 simulation replications as this number ensures (with respect to the base case) asymptotic convergence for the expected number of drugs wasted per day, expected number of drug shortages per day, and expected cost per day (holding plus ordering; see Figure B1).

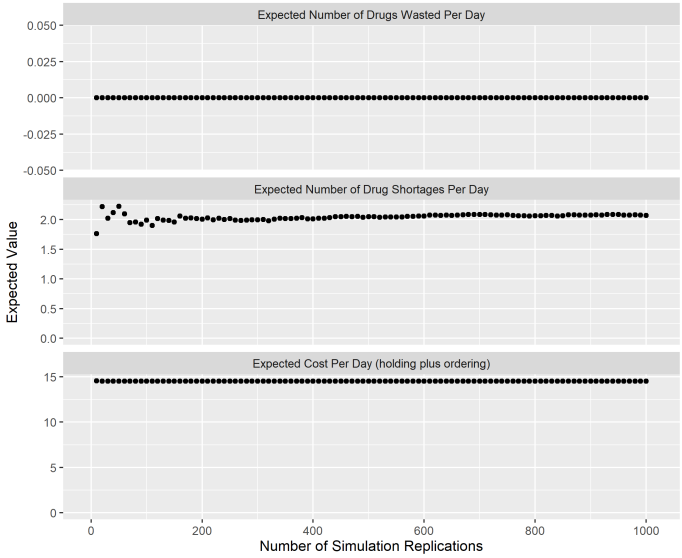

Figure B1: Number of simulation replications analysis.

B.4 Standard EOQ Model

In Section 3.6.2 of Chapter 3, we consider the standard EOQ model (Snyder & Shen, 2019) to illustrate the importance of implementing the perishable closed-form (R,S) inventory policies (in Section 3.5) for a perishable inventory system with supply chain disruptions. For the EOQ model, we solve for the (R,S) inventory policies using Equations (B.4.1)-(B.4.2).

$$\text{(Solving for } S^*_{EOQ}) \implies S^*_{EOQ} = \sqrt{\frac{2kq}{h}} \tag{B.4.1}$$

$$\text{(Solving for } R^*_{EOQ}) \implies R^*_{EOQ} = \frac{S^*_{EOQ}}{q} \tag{B.4.2}$$

B.5 Managerial Insights

We study the sensitivity of the model outputs (i.e., the expected cost per day [holding plus ordering], R, and S) with respect to changes in the input parameters: expiration lifetime (e), disruption probability (α), recovery probability (β), maximum proportion of demand not satisfied (γ), holding cost (h), and ordering cost (k).

B.5.1 Sensitivity Analysis for Expiration Lifetime (e)

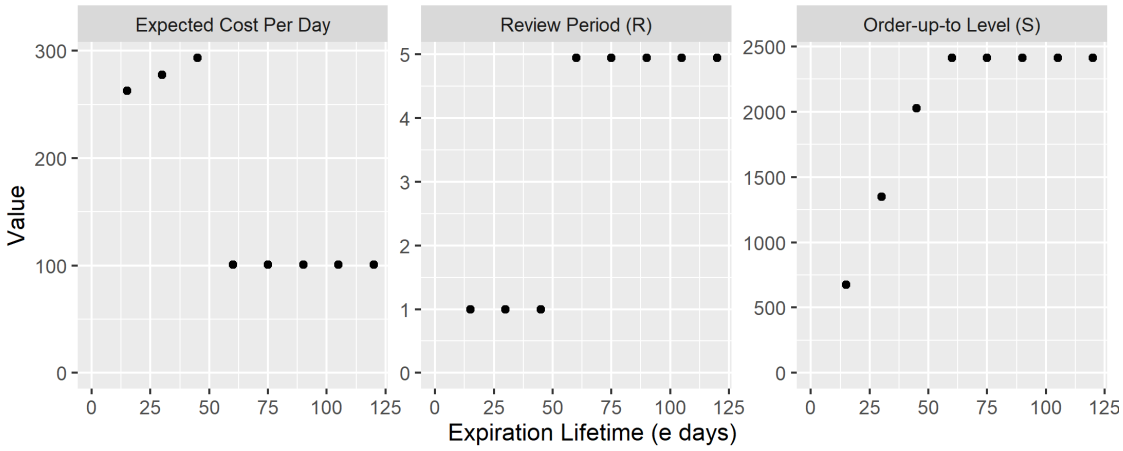

Figure B2: Sensitivity analysis for expiration lifetime (e).

Figure B2 illustrates that as the expiration lifetime (e) increases with the other inputs and parameters fixed, the length of the review period and order-up-to level increase. In particular, the order-up-to level and the length of the review period increase until $e = 60$ days. At this point, it is always optimal to follow the (R,S) policy found for with the non-perishable model making the expected cost per day for $e \geq 60$ the same. For $e < 60$, it is optimal to enforce the perishability condition (i.e., $S^* = eq$) leading to an increase in the expected cost per day due to the additional holding cost as e increases (ordering cost equivalent with $R = 1$ for $e < 60$).

B.5.2 Sensitivity Analysis for Disruption Probability (α)

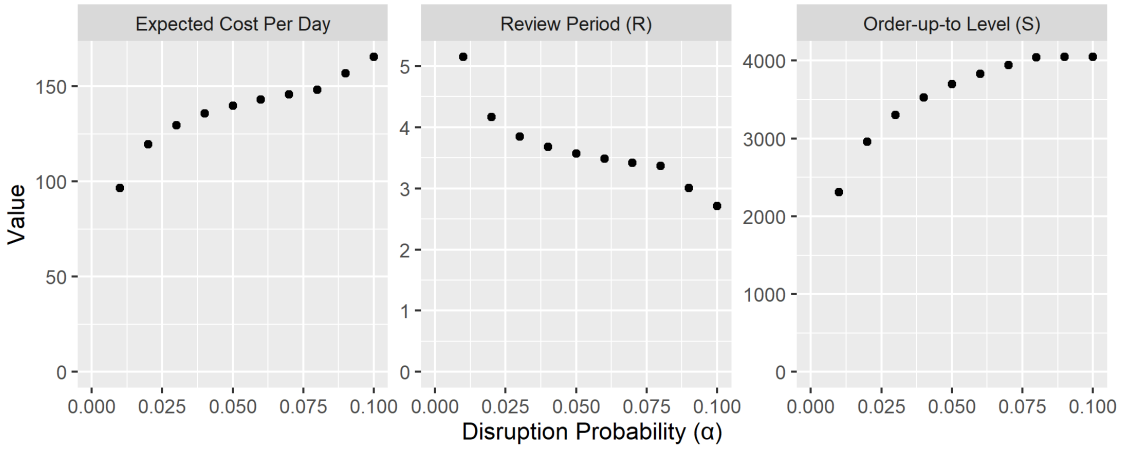

Figure B3: Sensitivity analysis for disruption probability (α).

Figure B3 illustrates how increasing the disruption probability (α) with all other inputs and parameters fixed, impacts the expected cost per day, order-up-to level, and the length of the review period. In the long-run, the probability that supply is disrupted is $\frac{\alpha}{\alpha+\beta}$ (Snyder & Shen, 2019). Therefore, as α increases, supply becomes less reliable leading to an increasing expected cost per day, decreasing length of the review period, and increasing order-up-to level.

B.5.3 Sensitivity Analysis for Recovery Probability (β)

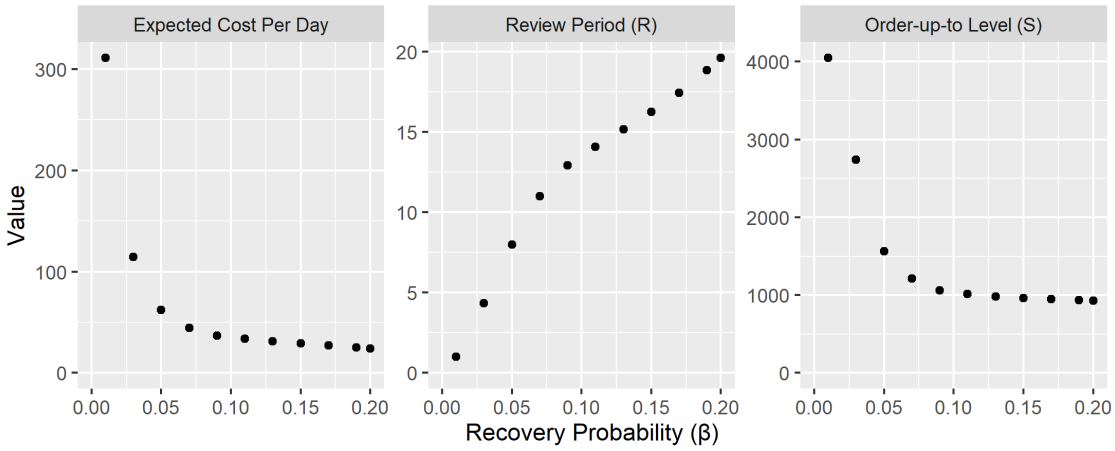

Figure B4: Sensitivity analysis for recovery probability (β).

Figure B4 illustrates how increasing the recovery probability (β) with all other inputs and parameters fixed, impacts the expected cost per day, the length of the review period, and order-up-to level. In the long-run, the probability that supply is not disrupted is $\frac{\beta}{\alpha+\beta}$ (Snyder & Shen, 2019). Therefore, as β increases, supply becomes more reliable leading to a decreasing expected cost per day, increasing length of the review period, and decreasing order-up-to level. Also, when considering the results for a change in β, it is important to consider that when β is very small (e.g., $\beta \leq 0.01$), supply is usually disrupted, so it is optimal to set $S = eq = 4050$. Then, as β increases, it is optimal to follow the (R,S) policies found for the non-perishable model.

B.5.4 Sensitivity Analysis for Maximum Proportion of Demand Not Satisfied (γ)

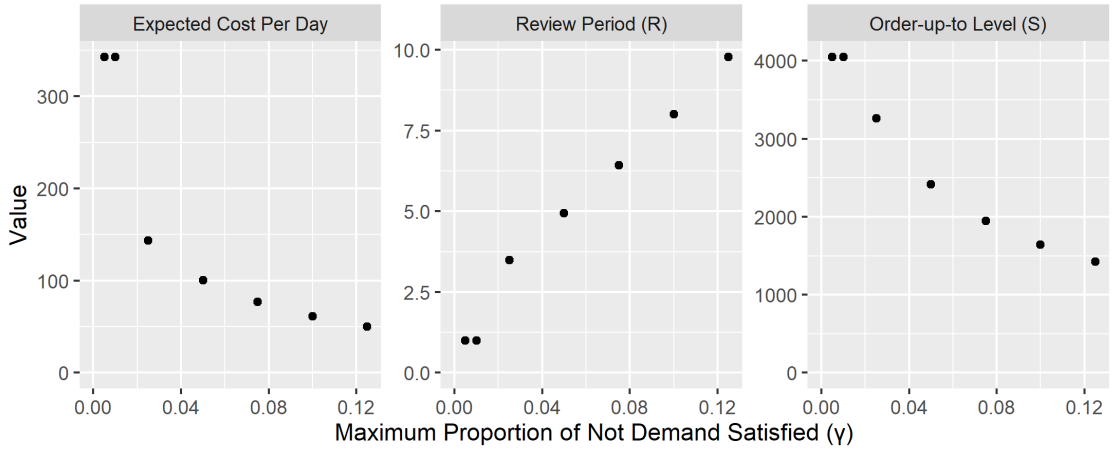

Figure B5: Sensitivity analysis for maximum proportion of demand not satisfied (γ).

Figure B5 shows that as the maximum proportion of demand not satisfied (γ) increases with the other inputs and parameters fixed, the expected cost per day decreases. It is important to note that γ drives how much demand is satisfied and there is no penalty in the objective function for not satisfying demand. Also, when γ is very small (e.g., $\gamma \leq 0.02$), the non-perishable model solves for an order-up-to level where $S > eq$. Therefore, enforcing the perishability condition (i.e., $S = eq = 4050$) is optimal in this case.

B.5.5 Sensitivity Analysis for Holding Cost (h)

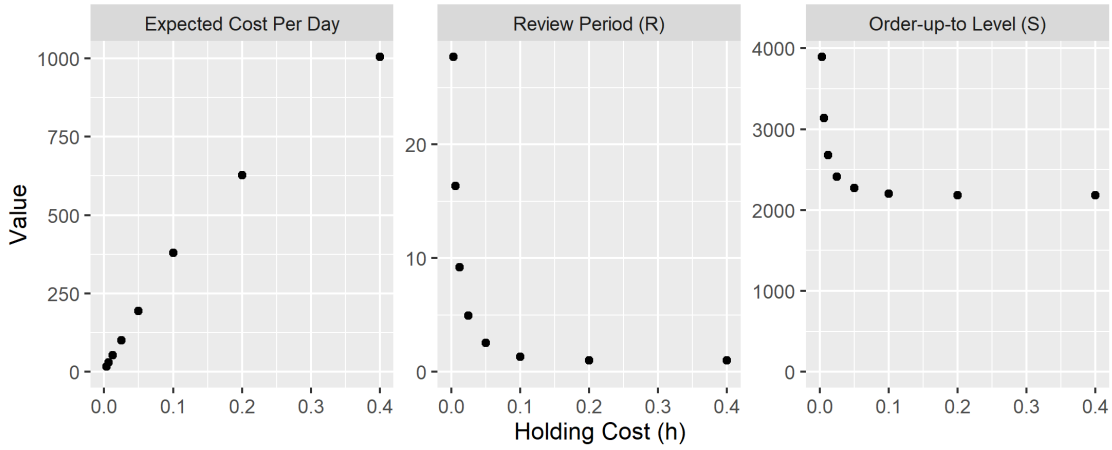

Figure B6: Sensitivity analysis for holding cost (h).

Figure B6 demonstrates that as the holding cost (h) increases with all other inputs and parameters fixed, the expected cost per day increases, but the length of the review period and order-up-to level decrease.

B.5.6 Sensitivity Analysis for Ordering Cost (k)

Figure B7: Sensitivity analysis for ordering cost (k).

Figure B7 illustrates that as the ordering cost (k) increases with all other inputs and parameters fixed, the expected cost per day, the length of the review period, and order-up-to level increase.

APPENDIX C

Chapter 4 Appendix

C.1 The (R,S) Model

The (R,S) model provides the optimal length of the review period R^* and order-up-to level S^* in closed-form for a lost-sales perishable inventory system with supply chain disruptions. We proceed to provide the expression for the ratio of the expected proportion of drug shortages per day given uncertainty in Q_2 in the (R,S) model.

C.1.1 Uncertainty in Q_2 in the $(\mathbf{R},\mathbf{S})$ Model

In Section 4.3.1.1 of Chapter 4, the duration of the supply chain disruption in days (i.e., Q_2) is difficult to quantify in practice. But, it is an important parameter as it defines $\alpha^{(1)}$ and $\beta^{(1)}$ (see Equations (4.2)-(4.3) in Section 4.3.1.1). With this, we provide the expression for the ratio of the expected proportion of drug shortages per day given a value Q_2 and the true value Q_2^*. The true value Q_2^* has a corresponding disruption probability with respect to 1 day (i.e., $\alpha^{*(1)}$), recovery probability with respect to 1 day (i.e., $\beta^{*(1)}$), m^*, R^*, and S^*. The selected value Q_2 has a corresponding disruption probability with respect to 1 day (i.e., $\alpha^{(1)}$), recovery probability with respect to 1 day (i.e., $\beta^{(1)}$), m, R, and S. We denote $P_{short|(q,R,S)}$ as the expected proportion of drug shortages per day given the inventory system has a daily demand q and follows a (R,S) inventory policy.

We first consider the value of $\frac{S^*}{qR^*}$ which corresponds to the true value Q_2^* and $\frac{S}{qR}$ which corresponds to the selected value Q_2. We have $\frac{S^*}{qR^*}$ as shown in Equation (C.1.1) which uses the m^*, R^*, and S^* found using Q_2^*.

$$\frac{S^*}{qR^*} = \frac{-\gamma(\alpha^{*(R^*)}+\beta^{*(R^*)})(1-\beta^{*(R^*)})+\alpha^{*(R^*)}(1-\beta^{*(R^*)})^{m^*}+\alpha^{*(R^*)}\beta^{*(R^*)}(m^*)(1-\beta^{*(R^*)})^{m^*}}{\alpha^{*(R^*)}\beta^{*(R^*)}(1-\beta^{*(R^*)})^{m^*}} \tag{C.1.1}$$

We present $\frac{S}{qR}$ in Equation (C.1.2) which uses the m, R, and S found using Q_2.

$$\frac{S}{qR} = \frac{-\gamma(\alpha^{(R)}+\beta^{(R)})(1-\beta^{(R)}) + \alpha^{(R)}(1-\beta^{(R)})^m + \alpha^{(R)}\beta^{(R)}(m)(1-\beta^{(R)})^m}{\alpha^{(R)}\beta^{(R)}(1-\beta^{(R)})^m} \tag{C.1.2}$$

For Q_2^* which corresponds to $\alpha^{*(1)}$, $\beta^{*(1)}$, m^*, R^*, and S^*, we can use Equation (4.4) in Section 4.3.2 and Equation (C.1.1) to get the expected proportion of drug shortages per day as presented in Equation (C.1.3).

$$P_{short|(q,R^*,S^*)} = \frac{\alpha^{*(R^*)}\beta^{*(R^*)}(1-\beta^{*(R^*)})^{m^*-1}}{(\alpha^{*(R^*)}+\beta^{*(R^*)})}(m^*+1-\frac{S^*}{qR^*}) + \frac{\alpha^{*(R^*)}(1-\beta^{*(R^*)})^{m^*}}{(\alpha^{*(R^*)}+\beta^{*(R^*)})} \tag{C.1.3}$$

$$= \begin{cases} \gamma; \ (R^* \neq 1 \text{ or } S^* \neq eq) \text{ i.e., } \gamma \text{ constraint satisfied (see Chapter 3)} \\ \\ \frac{\alpha^{*(1)}\left(1-\beta^{*(1)}\right)^{e-1}}{\left(\alpha^{*(1)}+\beta^{*(1)}\right)}; \ (R^*=1, S^*=eq) \end{cases}$$

Recall, we select Q_2 which corresponds to $\alpha^{(1)}$, $\beta^{(1)}$, m, R, and S. However, the true value is Q_2^* (i.e., the true disruption and recovery probability values are $\alpha^{*(R)}$ and $\beta^{*(R)}$, respectively, for a review period of length R). Hence, we use Equation (4.4) in Section 4.3.2 with $\alpha^{*(R)}$ and $\beta^{*(R)}$ and Equation (C.1.2) to get the expected proportion of drug shortages per day as presented in Equation (C.1.4).

$$P_{short|(q,R,S)} = \frac{\alpha^{*(R^*)}\beta^{*(R^*)}(1-\beta^{*(R^*)})^{m-1}}{(\alpha^{*(R^*)}+\beta^{*(R^*)})}(m+1-\frac{S}{qR}) + \frac{\alpha^{*(R^*)}(1-\beta^{*(R^*)})^{m}}{(\alpha^{*(R^*)}+\beta^{*(R^*)})} \tag{C.1.4}$$

$$= \begin{cases} \frac{\alpha^{*(R)}\left(1-\beta^{*(R)}\right)^{m-1}}{\left(\alpha^{*(R)}+\beta^{*(R)}\right)}\left(1+\frac{\gamma\beta^{*(R)}\left(\alpha^{(R)}+\beta^{(R)}\right)}{\alpha^{(R)}\beta^{(R)}\left(1-\beta^{(R)}\right)^{m-1}} - \frac{\beta^{*(R)}}{\beta^{(R)}}\right); \ (R \neq 1 \text{ or } S \neq eq) \text{ i.e., } \gamma \text{ constraint satisfied} \\ \\ \frac{\alpha^{*(1)}\left(1-\beta^{*(1)}\right)^{e-1}}{\alpha^{*(1)}+\beta^{*(1)}}; \ (R=1, S=eq) \end{cases}$$

The ratio of the expected proportion of drug shortages per day with Q_2 (i.e., (R,S) policy) and Q_2^* (i.e., (R^*,S^*) policy) is presented as two cases: $Q_2 < Q_2^*$ and $Q_2 > Q_2^*$.

Case 1: $Q_2 < Q_2^*$

When Q_2 selects $(R=1, S=eq)$ as the optimal inventory policy, the ratio of the expected proportion of drug shortages per day with Q_2 (i.e., (R,S) policy) and Q_2^* (i.e., (R^*,S^*) policy) is presented

in Equation (C.1.5).

$$\frac{P_{short|(q,R,S)}}{P_{short|(q,R^*,S^*)}} = \frac{\frac{\alpha^{*(1)}\left(1-\beta^{*(1)}\right)^{e-1}}{\left(\alpha^{*(1)}+\beta^{*(1)}\right)}}{\frac{\alpha^{*(1)}\left(1-\beta^{*(1)}\right)^{e-1}}{\left(\alpha^{*(1)}+\beta^{*(1)}\right)}} = 1 \tag{C.1.5}$$

In Equation (C.1.5), the ratio is 1 because Q_2 selects $(R = 1, S = eq)$ as the optimal inventory policy which corresponds to the smallest R and largest S. In comparison to Q_2, the true value Q_2^* implies the supply chain takes even longer to recover (i.e., $Q_2^* > Q_2$). There will be no change in the inventory policy because the inventory policy is already at the limits for both decision variables.

Given Q_2 selects $(R \neq 1$ or $S \neq eq)$ as the optimal inventory policy, the ratio of the expected proportion of drug shortages per day with Q_2 (i.e., (R,S) policy) and Q_2^* (i.e., (R^*,S^*) policy) is presented in Equation (C.1.6).

$$\frac{P_{short|(q,R,S)}}{P_{short|(q,R^*,S^*)}} = \frac{1}{\max\{\gamma, \frac{\alpha^{*(1)}\left(1-\beta^{*(1)}\right)^{e-1}}{\left(\alpha^{*(1)}+\beta^{*(1)}\right)}\}} \left(\frac{\alpha^{*(R)}\left(1-\beta^{*(R)}\right)^{m-1}}{\left(\alpha^{*(R)}+\beta^{*(R)}\right)} \left(1 + \frac{\gamma\beta^{*(R)}\left(\alpha^{(R)}+\beta^{(R)}\right)}{\alpha^{(R)}\beta^{(R)}\left(1-\beta^{(R)}\right)^{m-1}} - \frac{\beta^{*(R)}}{\beta^{(R)}} \right) \right) \tag{C.1.6}$$

In Equation (C.1.6), the denominator $\max\{\gamma, \frac{\alpha^{*(1)}\left(1-\beta^{*(1)}\right)^{e-1}}{\left(\alpha^{*(1)}+\beta^{*(1)}\right)}\}$ reflects when the maximum proportion of drug shortages per day constraint (i.e., γ) is (1) satisfied or (2) not satisfied, respectively. When $\frac{\alpha^{*(1)}\left(1-\beta^{*(1)}\right)^{e-1}}{\left(\alpha^{*(1)}+\beta^{*(1)}\right)} > \gamma$, this implies that the maximum proportion of drug shortages per day constraint (i.e., γ) cannot be satisfied with Q_2^* (i.e., (2)). The model will select the inventory policy that maximizes the expected inventory on-hand which is $(R^* = 1, S^* = eq)$ and has a corresponding expected proportion of drug shortages per day of $\frac{\alpha^{*(1)}\left(1-\beta^{*(1)}\right)^{e-1}}{\left(\alpha^{*(1)}+\beta^{*(1)}\right)}$. Otherwise (i.e., (1)), the maximum proportion of drug shortages per day constraint is satisfied with Q_2^* and the constraint is tight. Hence, the expected proportion of drug shortages per day is γ.

Case 2: $Q_2 > Q_2^*$

When Q_2 selects $(R = 1, S = eq)$ as the optimal inventory policy, the ratio of the expected proportion of drug shortages per day with Q_2 (i.e., (R,S) policy) and Q_2^* (i.e., (R^*,S^*) policy) is presented in Equation (C.1.7).

$$\frac{P_{short|(q,R,S)}}{P_{short|(q,R^*,S^*)}} = \frac{\frac{\alpha^{*(1)}\left(1-\beta^{*(1)}\right)^{e-1}}{\left(\alpha^{*(1)}+\beta^{*(1)}\right)}}{\max\{\gamma, \frac{\alpha^{*(1)}\left(1-\beta^{*(1)}\right)^{e-1}}{\left(\alpha^{*(1)}+\beta^{*(1)}\right)}\}} \tag{C.1.7}$$

In Equation (C.1.7), the denominator $\max\{\gamma, \frac{\alpha^{*(1)}\left(1-\beta^{*(1)}\right)^{e-1}}{\left(\alpha^{*(1)}+\beta^{*(1)}\right)}\}$ reflects when the maximum proportion of drug shortages per day constraint (i.e., γ) is (1) satisfied or (2) not satisfied, respectively. When $\frac{\alpha^{*(1)}\left(1-\beta^{*(1)}\right)^{e-1}}{\left(\alpha^{*(1)}+\beta^{*(1)}\right)} > \gamma$, this implies that the maximum proportion of drug shortages per day

constraint (i.e., γ) cannot be satisfied with Q_2^* (i.e., (2)). The model will select the inventory policy that maximizes the expected inventory on-hand which is $(R^* = 1, S^* = eq)$ and has a corresponding expected proportion of drug shortages per day of $\frac{\alpha^{*(1)}\left(1-\beta^{*(1)}\right)^{e-1}}{\left(\alpha^{*(1)}+\beta^{*(1)}\right)}$. Otherwise (i.e., (1)), the maximum proportion of drug shortages per day constraint is satisfied with Q_2^* and the constraint is tight. Hence, the expected proportion of drug shortages per day is γ.

Given Q_2 selects ($R \neq 1$ or $S \neq eq$) as the optimal inventory policy, the ratio of the expected proportion of drug shortages per day with Q_2 (i.e., (R,S) policy) and Q_2^* (i.e., (R^*,S^*) policy) is presented in Equation (C.1.8).

$$\frac{P_{short|(q,R,S)}}{P_{short|(q,R^*,S^*)}} = \frac{1}{\gamma}\left(\frac{\alpha^{*(R)}\left(1-\beta^{*(R)}\right)^{m-1}}{\left(\alpha^{*(R)}+\beta^{*(R)}\right)}\left(1+\frac{\gamma\beta^{*(R)}\left(\alpha^{(R)}+\beta^{(R)}\right)}{\alpha^{(R)}\beta^{(R)}\left(1-\beta^{(R)}\right)^{m-1}}-\frac{\beta^{*(R)}}{\beta^{(R)}}\right)\right) \tag{C.1.8}$$

In Equation (C.1.8), we divide by γ because Q_2 selects ($R \neq 1$ or $S \neq eq$) as the optimal inventory policy illustrating the maximum proportion of drug shortages per day constraint (i.e., γ) is satisfied with Q_2. In comparison to Q_2, the true value Q_2^* implies that the supply chain recovers more quickly (i.e., $Q_2^* < Q_2$). If Q_2 satisfies the maximum proportion of drug shortages per day constraint, then Q_2^* will also satisfy the constraint. When the constraint is satisfied, it is tight and hence, the expected proportion of drug shortages per day is γ.

C.1.1.1 Uncertainty in Q_2 Validation

Case 1 and Case 2 provide closed-form expressions for the ratio of the expected proportion of drug shortages per day given a value Q_2 and the expected proportion of drug shortages per day given the true value Q_2^*. We proceed to validate the closed-form expressions presented as Case 1 and Case 2 by comparing the closed-form expression results to a simulated ratio. The simulated ratio is determined using a simulation model of the (R,S) inventory system. Using Q_2 (i.e., the value selected), we first calculate the optimal inventory policy (i.e., (R,S)). For the simulation model, we require R to be integer, so after solving for R, we round R down to the nearest whole number (i.e., conservative rounding; reduces shortages). Then, we use the simulation model to calculate the (a) proportion of drug shortages per day given the inventory system truly has the characteristics of Q_2^*, but operates under the (R,S) inventory policy. Using Q_2^* (i.e., the true value), we calculate the optimal inventory policy (i.e., (R^*,S^*); R^* rounded down to the nearest whole number). We use the simulation model to calculate the (b) proportion of drug shortages per day given the inventory system has the characteristics of Q_2^*, and operates under the (R^*,S^*) inventory policy. We then calculate the ratio by taking $\frac{a}{b}$.

Table C1 provides the input parameters we use for the validation of the ratio. These input parameters are consistent with the numerical analysis in Chapter 4 and the drug case study presented in Chapter 3. Also, consistent with this earlier research, we consider a simulation model of the

(R,S) inventory system with a warm-up period of 360 days (i.e., approx 1 year), a 5-year planning horizon (i.e., 1800 days), and 500 simulation replications.

Table C1 Input parameters for the uncertainty in Q_2 validation.

Notation	Description
$k = 250$	Fixed ordering cost (i.e., for each order attempted; $k > 0$)
$h = 0.025$	Holding cost per day per drug ($h > 0$; Jia & Zhao, 2017)
$q = 45$	Deterministic demand per day ($q > 0$; Chapter 3)
$e = 90$	Expiration lifetime in days ($e \geq 1$)
$\gamma = 0.05$	Maximum proportion of drug shortages per day over the infinite horizon ($0 < \gamma < 1$)
$Q_1^* = 0.25$	True value for the answer to Question 1 (in Section 4.3.1.1)
$Q_2^* \in \{30, 60\}$	True value for the answer to Question 2 (in Section 4.3.1.1)

We validate the ratio using $Q_2^* = 30$ (see Figure C1) and $Q_2^* = 60$ (see Figure C2). It is worth noting that we consider $Q_2^* = 60$ which is not considered in the numerical analysis in Chapter 4 as it provides insights on when the optimal inventory policy is at the limits for both decision variables R^* and S^*. Figures C1 and C2 present the length of the review period (R) and the order-up-to level (S) when using the true value Q_2^* (i.e., Optimal (R^*, S^*)) and selected value Q_2 (i.e., Selected (R,S)). The figures also present the ratio of the expected proportion of drug shortages per day using the simulation model (i.e., Simulated Ratio) and the closed-form expressions (i.e., Closed-form Ratio). Both figures show that the closed-form expressions are valid as the curves for the Simulated Ratio (see blue curves) and Closed-form Ratio (see purple curves) lie fairly close to one another. We want to highlight that the curves are close and not always exactly overlapping because the (R,S) model treats the decision variables as continuous variables. However, a simulation model of a periodic review inventory system must define the length of the review period (R) as an integer. In this analysis, we round R down to the nearest whole number (i.e., conservative rounding; places orders more frequently then necessary). As a result, there are instances where the curves do not exactly coincide.

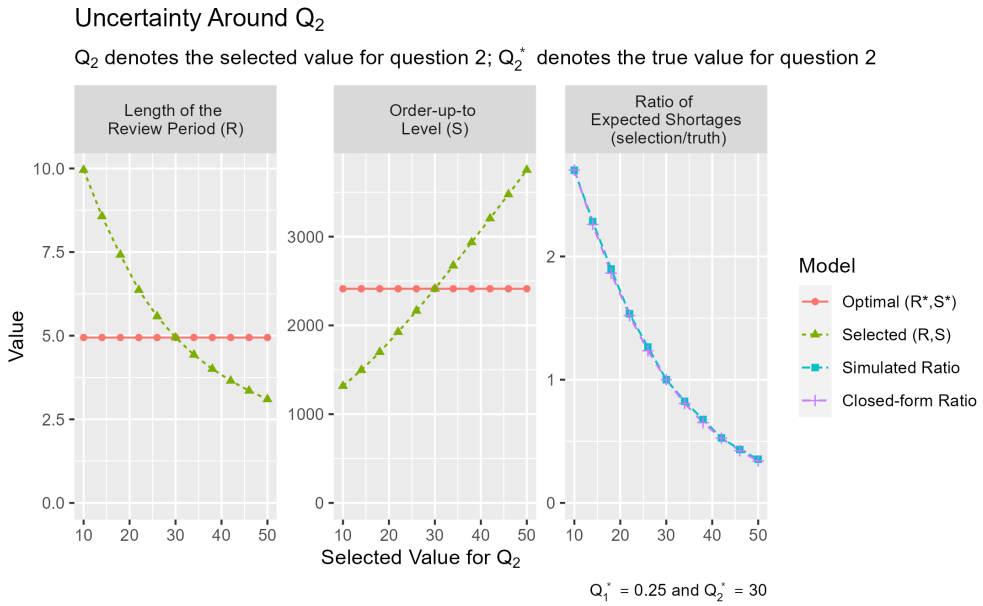

Figure C1: Uncertainty around Q_2 when $Q_2^* = 30$.

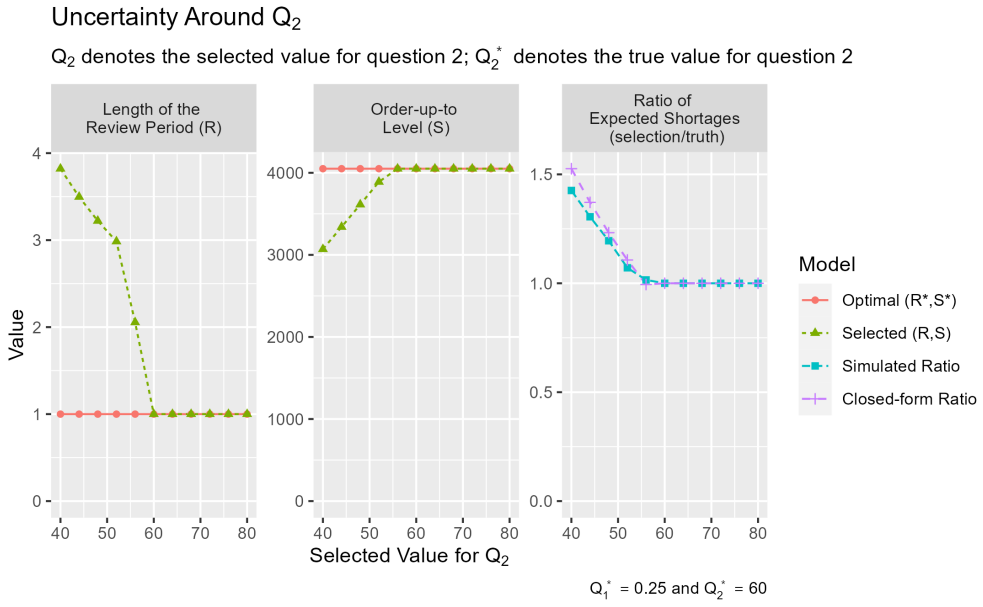

Figure C2: Uncertainty around Q_2 when $Q_2^* = 60$.

In line with our intuition, a key takeaway from Figures C1 and C2 is that when we underestimate Q_2 (i.e., $Q_2 < Q_2^*$; we claim the supply chain recovers more quickly than reality), we expect to have more or no change to shortages (ratio ≥ 1). When we overestimate Q_2 (i.e., $Q_2 > Q_2^*$; we claim the supply chain takes longer to recover than reality), we expect to have less or no change to shortages

(ratio ≤ 1). Specifically, Figure C1 illustrates that underestimating Q_2 by about 20 days keeps the ratio below 2.75, and overestimating Q_2 by about 20 days keeps the ratio above 0.25. Figure C2 illustrates that underestimating Q_2 by about 20 days keeps the ratio below 1.6, and overestimating Q_2 has no impact on the ratio. No impact on the ratio occurs because the optimal (R^*, S^*) policy for $Q_2^* = 60$ is $(R^* = 1, S^* = eq)$, which is already at the limits for both decision variables. Hence, claiming the supply chain takes longer to recover than reality (i.e., $Q_2 > Q_2^*$) will lead to the same optimal inventory policy of $(R = 1, S = eq)$.

C.2 Proportion Metrics

We present how to calculate the expected proportion of drug shortages per day (in Appendix C.2.1) and the expected proportion of drugs wasted per day (in Appendix C.2.2).

C.2.1 Expected Proportion of Drug Shortages

In Section 4.3.2.1 of Chapter 4, we present the expected proportion of drug shortages per day in Equation (4.4). The expected proportion of drug shortages per day depends on the length of the review period (R), the order-up-to level (S), the expected daily demand ($\bar{q}$), the disruption probability ($\alpha^{(R)}$), and the recovery probability ($\beta^{(R)}$). Proportion is measured relative to the expected daily demand $\bar{q}$.

When deriving the closed-form (R,S) inventory policy solutions, Chapter 3 assumes that $m = \lfloor \frac{S}{\bar{q}R} \rfloor \geq 1$ (i.e., S covers at least 1 review period; see Table 4.2 in Section 4.3.1). Therefore, when $m \geq 1$, we directly use the two-state supply process results found in Chapter 3. When $m < 1$, the order-up-to level S does not cover a full review period (i.e., $\bar{q}R > S$). The expected proportion of drug shortages consists of two components: (a) the long-run probability that the supply chain is not disrupted (i.e., $\frac{\beta^{(R)}}{\alpha^{(R)}+\beta^{(R)}}$) multiplied by the expected number of drug shortages in a review period of length R (i.e., $\bar{q}R - S$) divided by the demand in a review period of length R (i.e., $\bar{q}R$) and (b) the long-run probability that the supply chain is disrupted (i.e., $1 - \frac{\beta^{(R)}}{\alpha^{(R)}+\beta^{(R)}}$) multiplied by 1. For (b), m does not cover a full review period, so any time the supply chain is disrupted, the inventory system has zero inventory on-hand resulting in an expected proportion of drug shortages per day of 1. We do not account for stochastic demand as past research illustrates that stochastic demand that is normally distributed has a negligible impact on the expected proportion of drug shortages per day with the (R,S) model.

C.2.2 Expected Proportion of Drugs Wasted

In Section 4.3.2.1 of Chapter 4, we present the expected proportion of drugs wasted per day in Equation (4.5). We calculate the expected proportion of drugs wasted given a (R,S) inventory policy with expected daily demand $\bar{q}$ and standard deviation of daily demand σ (i.e., $P_{waste|(\bar{q},\sigma,R,S)}$). We consider the expiration lifetime (e), the length of the review period (R), the order-up-to level (S), the expected daily demand ($\bar{q}$), the standard deviation of daily demand (σ), the disruption probability ($\alpha^{(R)}$), and the recovery probability ($\beta^{(R)}$). Proportion is measured relative to the expected number of drugs ordered.

To derive Equation (4.5), we first consider the case with deterministic daily demand $\bar{q}$ and no supply chain disruptions. When we follow a (R,S) inventory policy for a perishable drug with expiration lifetime e, the inventory system follows a cyclic pattern where each cycle lasts $\lceil \frac{e}{R} \rceil \cdot R$ days. We note that the (R,S) model assumes $e \geq 1$ and $R \geq 1$. On day e in the cycle, we discard $\max\{0, S - e\bar{q}\}$ drugs. As a specific example, consider $R = 3$ days and $e = 5$ days. In Table C2, we present the cyclic pattern of the inventory system where we vary S in increments of $\bar{q}$. We use the same table convention as presented in Atan & Rousseau (2016). For $S < e\bar{q}$, the inventory system has a cycle of R days which is a factor of $\lceil \frac{e}{R} \rceil \cdot R$ and for $S \geq e\bar{q}$, the inventory system has a cycle of $\lceil \frac{e}{R} \rceil \cdot R$ days, so we simply conclude each cycle lasts $\lceil \frac{e}{R} \rceil \cdot R$ days.

Table C2 Cyclic pattern of the (R,S) inventory system with perishability.

S	Day	Start. Inv.	Waste	Short	End Inv.	Age=1	Age=2	Age=3	Age=4	Order size
$\bar{q}$	1	$\bar{q}$	0	0	0	0	0	0	0	0
	2	0	0	$\bar{q}$	0	0	0	0	0	0
	3	0	0	$\bar{q}$	0	0	0	0	0	$\bar{q}$
$2\bar{q}$	1	$2\bar{q}$	0	0	$\bar{q}$	$\bar{q}$	0	0	0	0
	2	$\bar{q}$	0	0	0	0	0	0	0	0
	3	0	0	0	0	0	0	0	0	$2\bar{q}$
$3\bar{q}$	1	$3\bar{q}$	0	0	$2\bar{q}$	$2\bar{q}$	0	0	0	0
	2	$2\bar{q}$	0	0	$\bar{q}$	0	$\bar{q}$	0	0	0
	3	$\bar{q}$	0	0	0	0	0	0	0	$3\bar{q}$
$4\bar{q}$	1	$4\bar{q}$	0	0	$3\bar{q}$	$3\bar{q}$	0	0	0	0
	2	$3\bar{q}$	0	0	$2\bar{q}$	0	$2\bar{q}$	0	0	0
	3	$2\bar{q}$	0	0	$\bar{q}$	0	0	$\bar{q}$	0	$3\bar{q}$
	4	$4\bar{q}$	0	0	$3\bar{q}$	$3\bar{q}$	0	0	0	0
	5	$3\bar{q}$	0	0	$2\bar{q}$	0	$2\bar{q}$	0	0	0
	6	$2\bar{q}$	0	0	$\bar{q}$	0	0	$\bar{q}$	0	$3\bar{q}$
$S \geq 5\bar{q}$	1	S	0	0	S-$\bar{q}$	S-$\bar{q}$	0	0	0	0
	2	S-$\bar{q}$	0	0	S-$2\bar{q}$	0	S-$2\bar{q}$	0	0	0
	3	S-$2\bar{q}$	0	0	S-$3\bar{q}$	0	0	S-$3\bar{q}$	0	$3\bar{q}$
	4	S	0	0	S-$\bar{q}$	$3\bar{q}$	0	0	S-$4\bar{q}$	0
	5	S-$\bar{q}$	S-$5\bar{q}$	0	$3\bar{q}$	0	$3\bar{q}$	0	0	0
	6	$3\bar{q}$	0	0	$2\bar{q}$	0	0	$2\bar{q}$	0	S-$2\bar{q}$
	7	S	0	0	S-$\bar{q}$	S-$2\bar{q}$	0	0	$\bar{q}$	0
	8	S-$\bar{q}$	0	0	S-$2\bar{q}$	0	S-$2\bar{q}$	0	0	0
	9	S-$2\bar{q}$	0	0	S-$3\bar{q}$	0	0	S-$3\bar{q}$	0	$3\bar{q}$
	10	S	0	0	S-$\bar{q}$	$3\bar{q}$	0	0	S-$4\bar{q}$	0
	11	S-$\bar{q}$	S-$5\bar{q}$	0	$3\bar{q}$	0	$3\bar{q}$	0	0	0
	12	$3\bar{q}$	0	0	$2\bar{q}$	0	0	$2\bar{q}$	0	S-$2\bar{q}$

We now formally define a cycle as a period of time such that on day e in the cycle, $\max\{0, S - e\bar{q}\}$ drugs are wasted. There is zero waste on all other days in the cycle. With deterministic daily demand $\bar{q}$ and no supply chain disruptions, we have the results presented in Equations (C.2.1)-(C.2.2).

$$\mathbb{E}[\text{no. of drugs wasted per cycle}] = \max\{0, S - e\bar{q}\} \tag{C.2.1}$$

$$\mathbb{E}[\text{no. of drugs ordered in a cycle}] = \overbrace{\lceil \frac{e}{R} \rceil R\bar{q}}^{\text{no. used}} + \overbrace{\max\{0, S - e\bar{q}\}}^{\text{no. wasted}} \tag{C.2.2}$$

We next consider that we have deterministic daily demand $\bar{q}$ and supply chain disruptions that follow a two-state supply process. Given we have a supply chain disruption, we need to consider which day in the cycle the supply chain disruption begins to accurately calculate the number of

drugs ordered in the cycle. Recall that a cycle lasts $\lceil \frac{e}{R} \rceil \cdot R$ days when the supply chain is not disrupted. Each cycle has $\lceil \frac{e}{R} \rceil$ days that correspond to a review period day. Given we know that a supply chain disruption occurs in the cycle, out of the $\lceil \frac{e}{R} \rceil$ days in the cycle that correspond to a review period day, there is an equal probability that the supply chain disruption begins on any of these review period days. Hence, the probability that the supply chain disruption starts on any review period day in the cycle is $\frac{1}{\lceil \frac{e}{R} \rceil}$. Taking into account the day the supply chain disruption starts in a cycle, we have the results presented in Equations (C.2.3)-(C.2.4). π_j is the probability that the supply chain is disrupted for exactly j consecutive review periods ($\pi_0 = \frac{\beta^{(R)}}{\alpha^{(R)}+\beta^{(R)}}$; $\pi_j = \frac{\alpha^{(R)}\beta^{(R)}}{(\alpha^{(R)}+\beta^{(R)})(1-\beta^{(R)})}(1-\beta^{(R)})^j,\ j \geq 1$).

$$\mathbb{E}[\text{no. of drugs wasted per cycle}] = \max\{0, S - e\bar{q}\} \tag{C.2.3}$$

$\mathbb{E}[\text{no. of drugs ordered in a cycle}] =$ (C.2.4)

$$\begin{cases} S;\ \lceil \frac{e}{R} \rceil = 1 \\ (\lceil \frac{e}{R} \rceil R\bar{q} + \max\{0, S - e\bar{q}\})\pi_0 + \sum\limits_{j=\lceil \frac{e}{R} \rceil - 1}^{\infty} \pi_j (S + R\bar{q}(\frac{(\lceil \frac{e}{R} \rceil - 1)}{2}));\ \lceil \frac{e}{R} \rceil = 2 \\ (\lceil \frac{e}{R} \rceil R\bar{q} + \max\{0, S - e\bar{q}\})(\pi_0 + \sum\limits_{j=1}^{\lceil \frac{e}{R} \rceil - 2} \pi_j - \frac{1}{\lceil \frac{e}{R} \rceil} \sum\limits_{j=1}^{\lceil \frac{e}{R} \rceil - 2} j\pi_j) + \frac{S}{\lceil \frac{e}{R} \rceil} \sum\limits_{j=1}^{\lceil \frac{e}{R} \rceil - 2} j\pi_j + \frac{R\bar{q}}{\lceil \frac{e}{R} \rceil} \sum\limits_{j=1}^{\lceil \frac{e}{R} \rceil - 2} j^2\pi_j \\ + \sum\limits_{j=\lceil \frac{e}{R} \rceil - 1}^{\infty} \pi_j (S + R\bar{q}(\frac{(\lceil \frac{e}{R} \rceil - 1)}{2}));\ \lceil \frac{e}{R} \rceil \geq 3 \end{cases}$$

where $\pi_0 = \dfrac{\beta^{(R)}}{\alpha^{(R)} + \beta^{(R)}}$

$$\sum_{j=1}^{\lceil \frac{e}{R} \rceil - 2} \pi_j = \frac{\alpha^{(R)}\beta^{(R)}}{\alpha^{(R)} + \beta^{(R)}} \left(\frac{1 - (1 - \beta^{(R)})^{\lceil \frac{e}{R} \rceil - 2}}{\beta^{(R)}} \right)$$

$$\sum_{j=\lceil \frac{e}{R} \rceil - 1}^{\infty} \pi_j = 1 - \sum_{j=0}^{\lceil \frac{e}{R} \rceil - 2} \pi_j = 1 - \pi_0 - \sum_{j=1}^{\lceil \frac{e}{R} \rceil - 2} \pi_j$$

$$\sum_{j=1}^{\lceil \frac{e}{R} \rceil - 2} j\pi_j = \frac{\alpha^{(R)}\beta^{(R)}}{(\alpha^{(R)} + \beta^{(R)})(1 - \beta^{(R)})} \left(\frac{1 - (\lceil \frac{e}{R} \rceil - 1)(1 - \beta^{(R)})^{\lceil \frac{e}{R} \rceil - 2} + (\lceil \frac{e}{R} \rceil - 2)(1 - \beta^{(R)})^{\lceil \frac{e}{R} \rceil - 1}}{(\beta^{(R)})^2} \right)$$

$$\sum_{j=1}^{\lceil \frac{e}{R} \rceil - 2} j^2\pi_j = \frac{\alpha^{(R)}\beta^{(R)}}{(\alpha^{(R)} + \beta^{(R)})(1 - \beta^{(R)})} (\frac{1}{\beta^{(R)}}) \Big(-(1 - \beta^{(R)}) \Big(\frac{1 - (1 - \beta^{(R)})^{\lceil \frac{e}{R} \rceil - 2}}{\beta^{(R)}} \Big) - (\lceil \frac{e}{R} \rceil - 2)^2 (1 - \beta^{(R)})^{\lceil \frac{e}{R} \rceil - 1}$$
$$+ 2 \Big(\frac{(1 - \beta^{(R)}) - (\lceil \frac{e}{R} \rceil - 1)(1 - \beta^{(R)})^{\lceil \frac{e}{R} \rceil - 1} + (\lceil \frac{e}{R} \rceil - 2)(1 - \beta^{(R)})^{\lceil \frac{e}{R} \rceil}}{(\beta^{(R)})^2} \Big) \Big)$$

Next, we incorporate stochastic demand where we assume that the daily demand is normally distributed with mean $\bar{q}$ and standard deviation σ. For the deterministic demand case, we have a

demand of exactly $e\bar{q}$ in a period of e days. Thus, we waste exactly $\max\{0, S - e\bar{q}\}$ drugs in each cycle. When we have stochastic demand, we may not have a demand of exactly $e\bar{q}$ in a period of e days. Letting Q_i denote a normal random variable with mean $\bar{q}$ and standard deviation σ for the demand on day i, we consider the expectation of $\max\{0, S - e\bar{q}\}$ presented in Equation (C.2.5) and simplified in Equation (C.2.6). In Equation (C.2.6), we assume that $\sum_{i=1}^{e} Q_i \geq 0$ as demand is always non-negative. It is worth noting that Equation (C.2.6) can easily be modified to accommodate other probability distributions (e.g., Poisson) by using the appropriate expectation and cumulative distribution function.

$$\mathbb{E}[\max\{0, S - \sum_{i=1}^{e} Q_i\}] = E_w = \Pr(\sum_{i=1}^{e} Q_i < S) \cdot (S - \mathbb{E}[\sum_{i=1}^{e} Q_i | \sum_{i=1}^{e} Q_i < S]) + \Pr(\sum_{i=1}^{e} Q_i \geq S) \cdot 0 \tag{C.2.5}$$

$$\mathbb{E}[\max\{0, S - \sum_{i=1}^{e} Q_i\}] = E_w = S \cdot \Pr(Z < \frac{S - e\bar{q}}{\sqrt{e\sigma^2}}) - e\bar{q} - \frac{1}{2\pi}\Big(- e^{-\frac{(S-e\bar{q})^2}{2e\sigma^2}} + e^{-\frac{(-e\bar{q})^2}{2e\sigma^2}} \Big)\sqrt{e\sigma^2} \tag{C.2.6}$$

Using Equation (C.2.6) in Equations (C.2.3)-(C.2.4) and the relation presented in Equation (C.2.7), we have the expected proportion of drugs wasted per day presented in Equation (4.5) of Chapter 4. Equation (4.5) accounts for a (R, S) inventory system with supply chain disruptions that follow a two-state supply process and stochastic demand that is independent and normally distributed.

$$E_{waste|(\bar{q},\sigma,R,S)} = \frac{\mathbb{E}[\text{no. of drugs wasted per cycle}]}{\mathbb{E}[\text{no. of drugs ordered in a cycle}]} \tag{C.2.7}$$

C.3 Numerical Analysis: Additional Data and Takeaways

We proceed to present the weekly demand data for the four other 503B drugs of interest presented in Section 4.5.4 of the numerical analysis: Avastin, Oxytocin, Cefazolin, Norepinephrine (in Section C.3.1). We also gain additional takeaways by analyzing (T1) the sensitivity of N (i.e., number of past daily demand observations used to estimate the expected daily demand; in Section C.3.2) and (T2) the length of demand disruptions (in Section C.3.3).

C.3.1 Weekly Demand Data for the Other 503B Drugs of Interest

We present the weekly demand data for the four other 503B drugs of interest: Avastin 1.25mg/0.05mL (chemotherapy with several indications; see Figure C3a), Oxytocin 30 units/500 mL (induction of labor; see Figure C3b), Cefazolin 2gm/100mL (antibiotic; see Figure C3c), Norepinephrine 16 mg/250mL (vasopressor used to increase blood pressure; see Figure C3d). The red

lines in the figures denote the corresponding mean weekly demand that minimizes the sum of squared errors for the daily demand data. For Avastin, Oxytocin, and Cefazolin, we consider two mean weekly demand values and for Norepinephrine, we consider three mean weekly demand values.

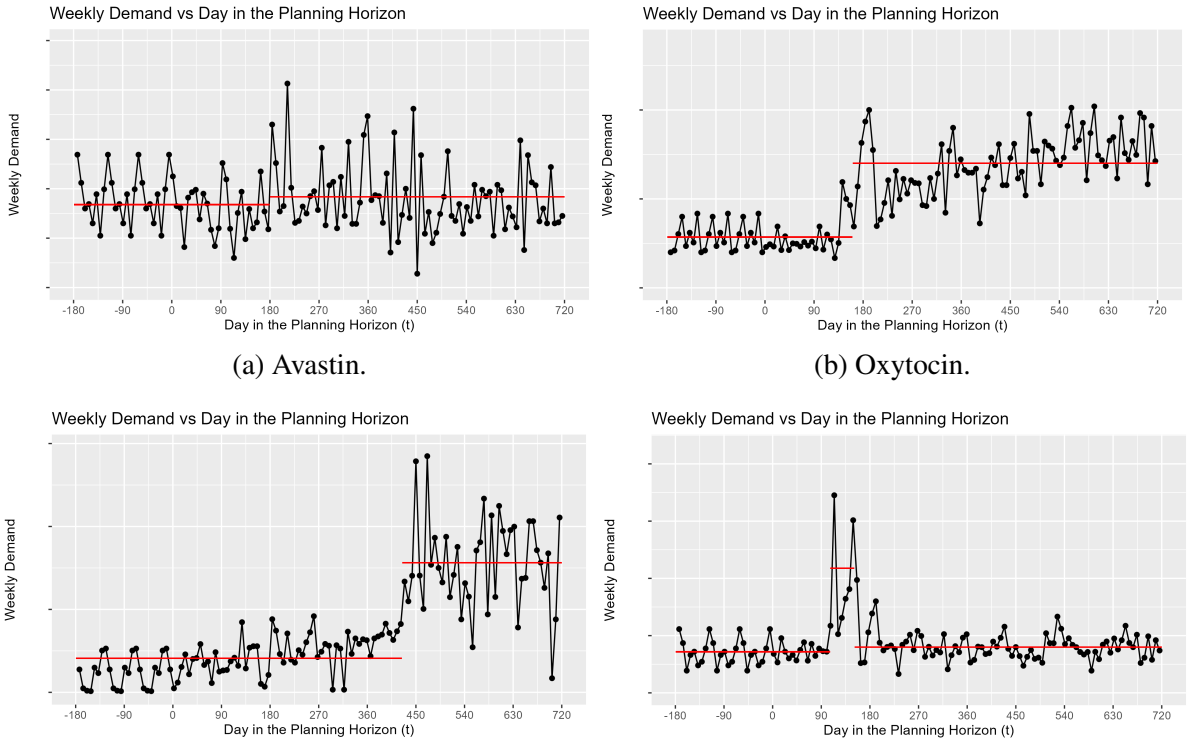

(a) Avastin.

(b) Oxytocin.

(c) Cefazolin.

(d) Norepinephrine.

Figure C3: Weekly demand versus day in the planning horizon. We remove the weekly demand values on the y-axis for data confidentiality.

C.3.2 (T1) Sensitivity to N

In the adaptive inventory system in Chapter 4 (in Section 4.3.2), N denotes the number of past daily demand observations used to estimate the expected daily demand $\bar{q}_{new}$ and standard deviation of daily demand σ_{new}.

We first consider the sensitivity of this parameter N by analyzing the proportion of drug shortages per day, proportion of drugs wasted per day, and number of policy changes during the testing horizon for varying $N \in [7, 105]$ days (i.e., about 1-15 weeks) and varying supply chain disruption durations. We consider the daily demand data for the Rocuronium and Labetalol case study presented in Section 4.5.1. We consider the input parameters defined in Section 4.5.3, $(\delta_s, \delta_w) = (0.05, 0.05)$ for the shortage-waste weighting, and varying supply chain disruption pro-

files: $(\alpha^{(1)},\beta^{(1)}) = \{(\frac{1}{30},\frac{1}{10}),(\frac{1}{90},\frac{1}{30}),(\frac{1}{270},\frac{1}{90}),(\frac{1}{810},\frac{1}{270})\}$. Figures C4 and C5 illustrate the results for Rocuronium (resembles an increasing demand disruption) and Labetalol (resembles a decreasing demand disruption), respectively. We present the (A) Adaptive, (B) Adaptive with Buyback, (C) Benchmark, and (D) Static models. The supply chain disruption profiles are distinguished into four rows where all supply chain disruption profiles have the same long-run probability that the supply chain is disrupted (i.e., $\frac{\alpha^{(1)}}{\alpha^{(1)}+\beta^{(1)}}$). For each supply chain disruption profile, we provide the disruption probability with respect to 1 day (i.e., $\alpha^{(1)}$), recovery probability with respect to 1 day (i.e., $\beta^{(1)}$), and optimal (R^*,S^*) inventory policy when the mean of the first $B = 90$ training demand observations is used for the expected daily demand.

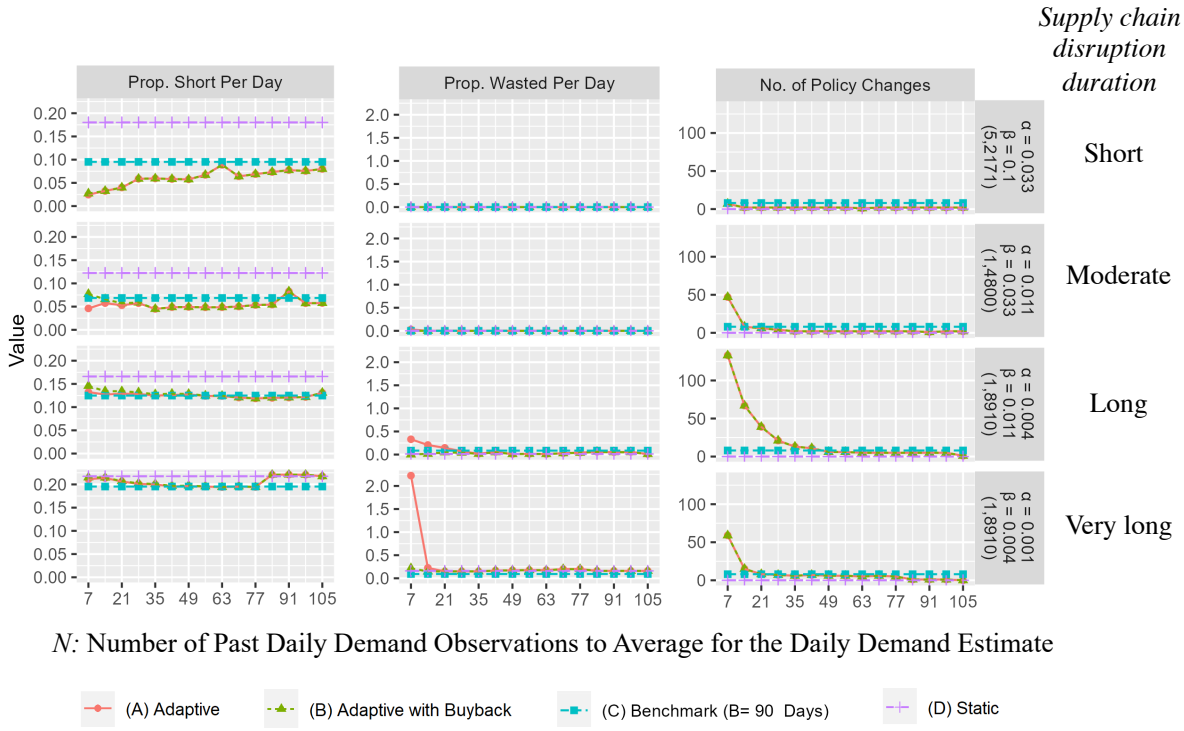

Figure C4: Sensitivity to N for Rocuronium for varying supply chain disruption durations.

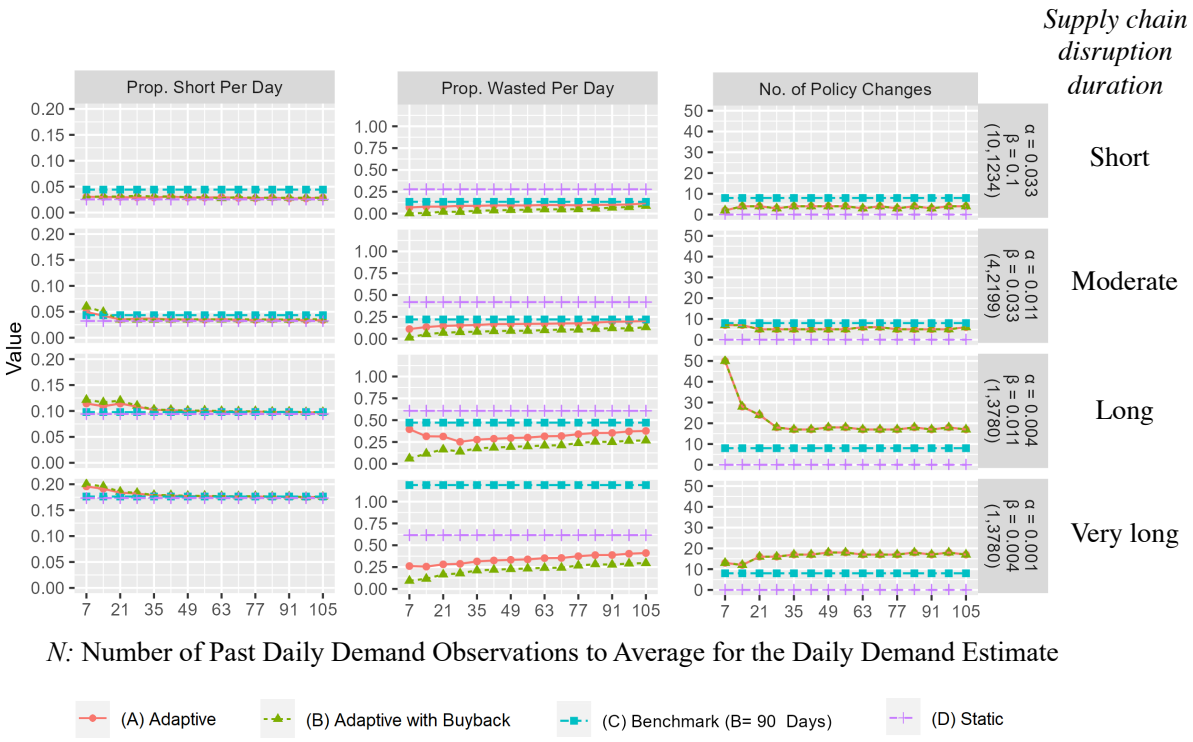

Figure C5: Sensitivity to N for Labetalol for varying supply chain disruption durations.

We second consider the sensitivity of the parameter N by analyzing the proportion of drug shortages per day and proportion of drugs wasted per day during the testing horizon for varying $N \in [7, 105]$ days (i.e., about 1-15 weeks) and multiple variabilities in demand. We consider the daily demand data for the Rocuronium and Labetalol case study presented in Section 4.5.1. However, we modify this real-world demand data. We first take notice to the two red horizontal lines denoting the corresponding mean weekly demand that minimizes the sum of squared errors for the daily demand data in Figure 4.4. We modify the real-world daily demand data using these mean weekly demand values divided by 7 (i.e., mean daily demand values; $\bar{q}_{\text{corresponding red mean daily demand}}$ where *"corresponding red mean daily demand"* implies that we use the red horizontal mean weekly demand divided by 7 that corresponds to the day in the planning horizon for the demand observation) to study multiple variabilities in demand. In particular, for any demand observation $q_{original}$, we obtain a modified demand observation as shown in Equation (C.3.1).

$$q_{modified} = (q_{original} - \bar{q}_{\text{corresponding red mean daily demand}}) + a \cdot \bar{q}_{\text{corresponding red mean daily demand}} \tag{C.3.1}$$

We consider $a \in \{0.5, 1, 2, 3\}$ where $a < 1$ implies that the variability of the modified daily demand data is less than the original daily demand data and $a > 1$ implies that the variability of the

modified daily demand data is greater than the original daily demand data. When $a = 1$, the original and modified daily demand data are equivalent. We consider the input parameters defined in Section 4.5.3, $(\delta_s, \delta_w) = (0.05, 0.05)$ for the shortage-waste weighting, and $(\alpha^{(1)}, \beta^{(1)}) = (\frac{1}{90}, \frac{1}{30})$ for the supply chain disruption profile. Figures C6 and C7 illustrate the results for Rocuronium (resembles an increasing demand disruption) and Labetalol (resembles a decreasing demand disruption), respectively. We present the (A) Adaptive, (B) Adaptive with Buyback, (C) Benchmark, and (D) Static models. Also, $a \in \{0.5, 1, 2, 3\}$ is denoted on the right-hand side of the figures and the vertical arrow denotes that the variance increases when reading the figures from top to bottom.

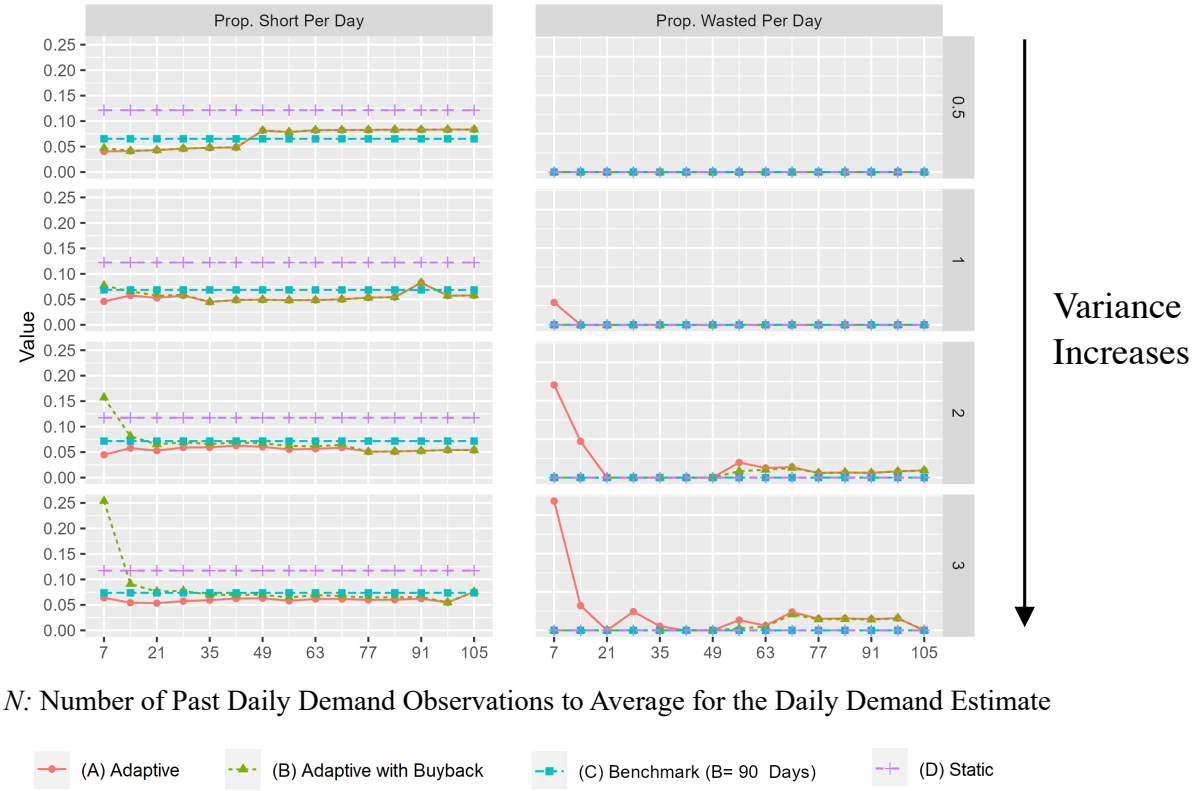

Figure C6: Sensitivity to N for Rocuronium for multiple variabilities in demand.

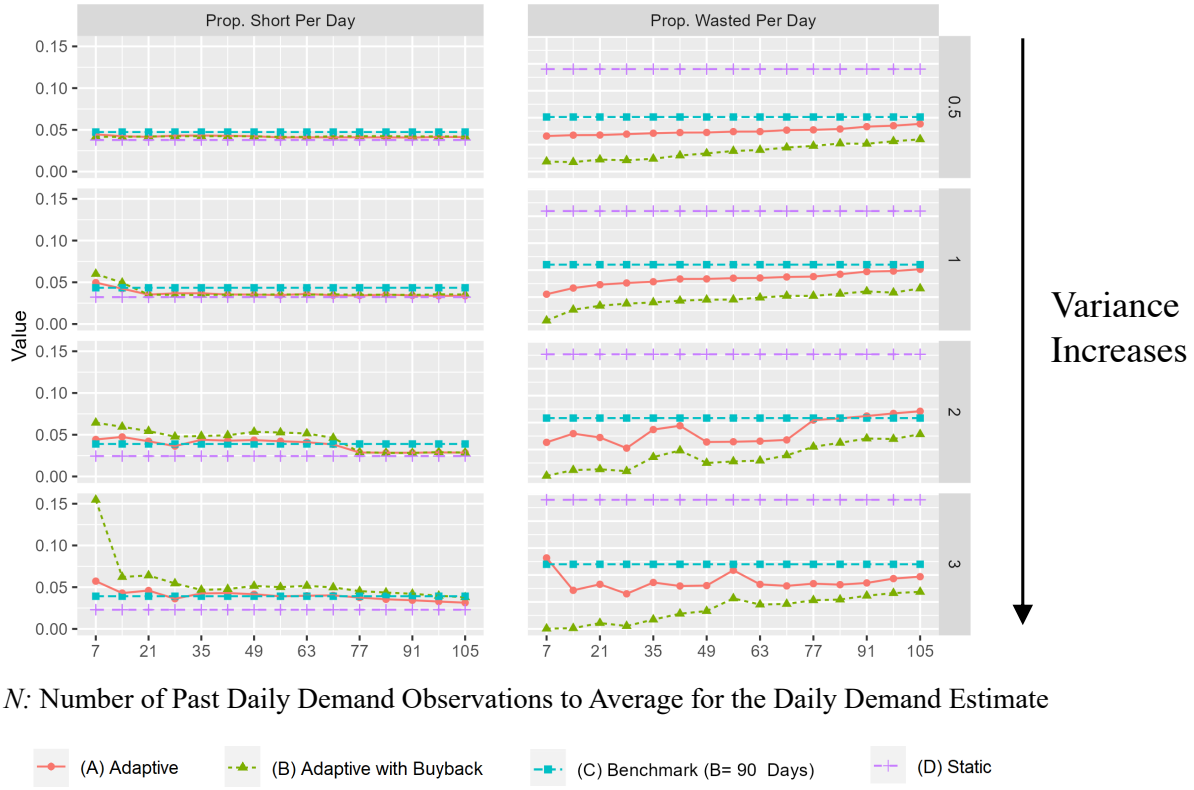

Figure C7: Sensitivity to N for Labetalol for multiple variabilities in demand.

From the results in Figures C4-C7, we find that:

1. As N increases, the metrics often begin to stabilize. Furthermore, a drug with a higher variability in demand often requires a larger N to stabilize.

2. Selecting too small of an N can cause high shortage and/or waste detriments, especially for drugs with a high variability in demand.

3. Selecting too large of an N can start to reduce the benefits of an adaptive model.

4. Selecting a value of $N = 56$ is appropriate for the numerical analysis.

C.3.3 (T2) Length of the Demand Disruption

Another viewpoint of interest is the length of the demand disruption. However, we do not have data that represents every possible length of the disruption. Therefore, we consider a transformed version of the Labetalol data. For the analysis, we first replicate and stack the Labetalol real-world daily demand data from October 1, 2019-January 31, 2020 (i.e., pre-Covid; not disrupted/stable

demand) to have 1,080 daily demand observations (i.e., 180 training and 900 testing). Starting at test day $t = 181$ (i.e., about 6 months into the testing horizon), we transform this "not disrupted" demand based on the length of disruption $l \in [60, 360]$ days. Consistent with the change in the two demand means that minimize the sum of squared errors (see red horizontal lines on Figure 4.4), we assume an increasing demand disruption adds 1.5 times the mean training data and a decreasing demand disruption subtracts 0.9 times the mean training data (non-positive values forced to 1). We assume that the first 30 days and last 30 days of the demand disruption have a linear change. Figure C8 depicts the transformed data for (a) increasing demand disruptions and (b) decreasing demand disruptions of lengths $l \in \{180, 360\}$ days.

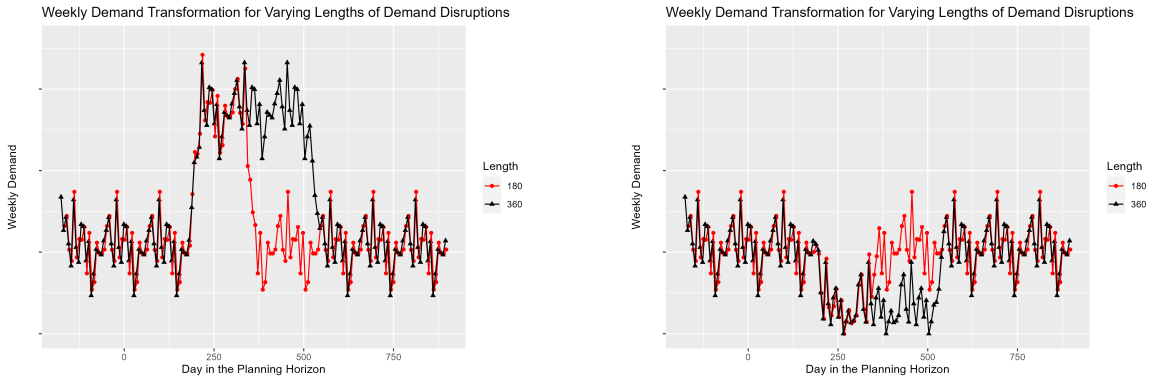

(a) Increasing demand disruptions.

(b) Decreasing demand disruptions.

Figure C8: Labetalol daily demand data for varying lengths of demand disruptions. We remove the numerical values on the y-axis for data confidentiality.

Increasing and Decreasing Demand Disruption Results

Using the transformed data in Figure C8 for the respective type of demand disruption (i.e., increasing or decreasing), we consider the model input parameters defined in Section 4.5.3 with the Labetalol wholesale price of about \$7. We simulate the performance of the (A) Adaptive, (B) Adaptive with Buyback, (C) Benchmark, and (D) Static models. We present the overall expected proportion of drug shortages per day and proportion of drugs wasted per day from test day $t = 181$ (i.e., start of the demand disruption) to test day $t = 900$ (i.e., end of the testing horizon). We consider demand disruptions lengths $l \in [60, 360]$ days, $(\delta_s, \delta_w) = (0.05, 0.05)$ for the shortage-waste weighting, and $(\alpha^{(1)}, \beta^{(1)}) = \{(\frac{1}{30}, \frac{1}{10}), (\frac{1}{90}, \frac{1}{30}), (\frac{1}{270}, \frac{1}{90}), (\frac{1}{810}, \frac{1}{270})\}$ for the varying supply chain disruption profiles.

For increasing demand disruptions (see Figure C9), we find that (a) the length of the demand disruption has the greatest influence on the (D) Static model when we consider the proportion of drug shortages per day and the (C) Benchmark model when we consider the proportion of drugs wasted per day. In particular, the (C) Benchmark model performs very poorly when we consider

the proportion of drugs wasted per day for long and very long supply chain disruption durations. For decreasing demand disruptions (see Figure C10), we find that (b) the length of the demand disruption has little to no impact on the proportion of drug shortages per day for all four models when considering moderate, long, and very long supply chain disruption durations. For short supply chain disruption durations, the length of the demand disruption influences the proportion of drug shortages when considering the (A) Adaptive, (B) Adaptive with Buyback, and (C) Benchmark models. For decreasing demand disruptions, we also find that the superior performance of the (A) Adaptive and (B) Adaptive with Buyback models when it comes to the proportion of drugs wasted per day becomes more apparent as the supply chain disruption duration increases. Focusing on the (C) Benchmark model for both increasing and decreasing demand disruptions, we overall find that (c) the model performs poorly from a shortage perspective when the supply chain disruption duration is short and the model performs poorly from a waste perspective when the supply chain disruption duration is long. Hence, we do not recommend implementing a benchmark policy that updates the inventory policies on long fixed intervals. Furthermore, we find that for both increasing and decreasing demand disruptions, (d) the (B) Adaptive with Buyback model often has a smaller proportion of drugs wasted per day in comparison to the (A) Adaptive model with zero to negligible difference in the proportion of drug shortages. This finding encourages the implementation of such buyback programs.

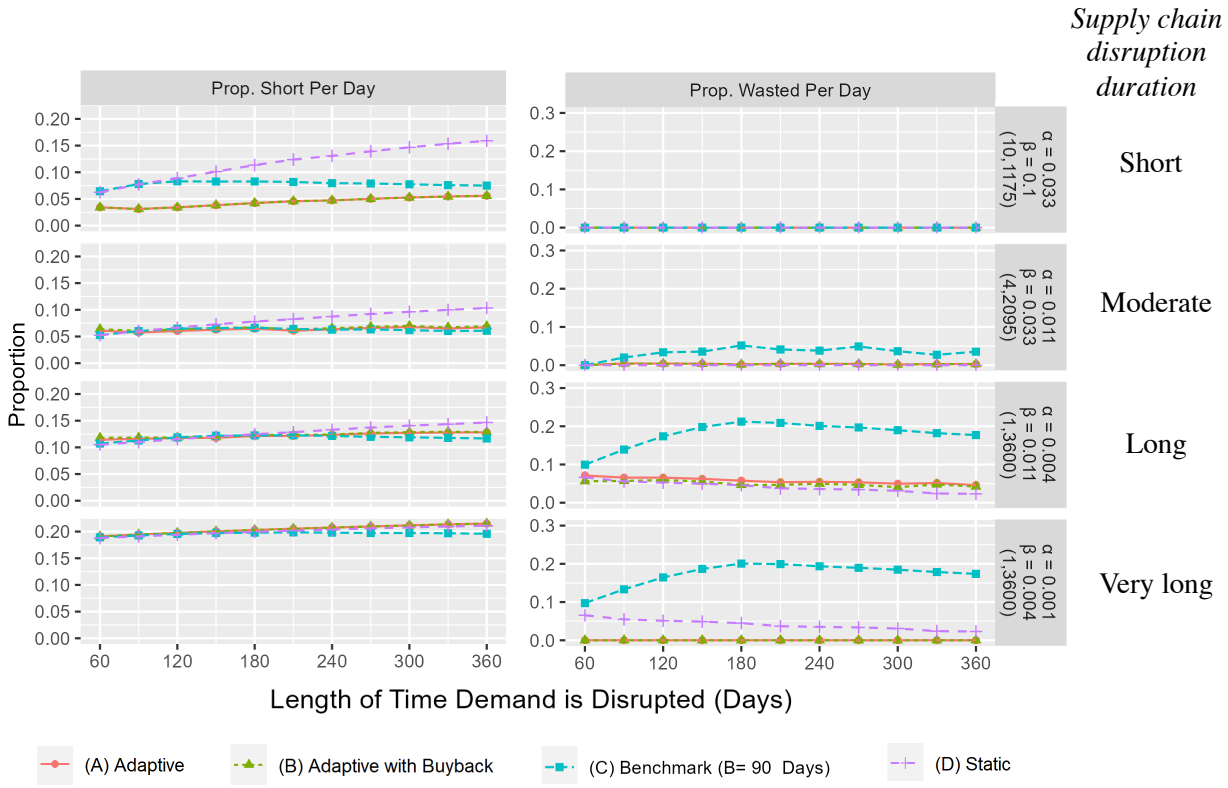

Figure C9: Results for varying lengths of increasing demand disruptions.

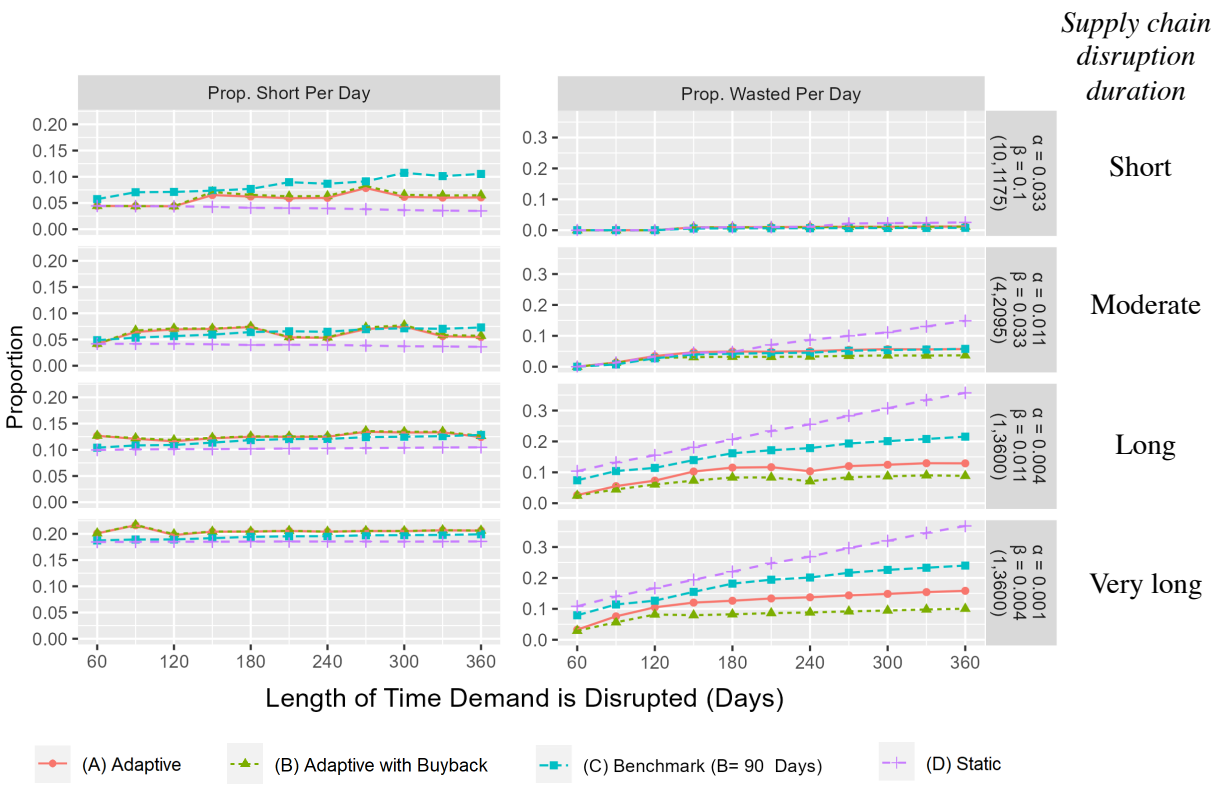

Figure C10: Results for varying lengths of decreasing demand disruptions.

APPENDIX D

Chapter 5 Appendix

D.1 Inventory System Descriptions and Derivations

We present the descriptions and/or derivations for the Two Supplier Integrated Non-perishable (in Section D.1.1), Non-Integrated Non-perishable (in Section D.1.2), Non-Integrated Perishable (in Section D.1.3), One Supplier Integrated Non-perishable (in Section D.1.4), and One Supplier Integrated Perishable (in Section D.1.5) inventory systems.

D.1.1 Non-Perishable Integrated Inventory System with Two Independent Suppliers

D.1.1.1 Exact Model: Long-run Probabilities (Two Supplier Integrated)

Figure 5.2 in Chapter 5 illustrates the two-HNP integrated inventory system continuous time Markov chain when the suppliers are independent from one another. We find the long-run probability of being in each state of the Markov chain by solving the system of equations presented in Equations (D.1.1)-(D.1.14).

$$(\lambda_1+\lambda_2)P_{(S_1+1)(S_2+1)} = \mu_1 \sum_{i=0}^{S_1} P_{i(S_2+1)} + \mu_2 \sum_{j=0}^{S_2} P_{(S_1+1)j} \tag{D.1.1}$$

$$(\lambda_1+q_2+\mu_2)P_{(S_1+1)S_2} = \lambda_2 P_{(S_1+1)(S_2+1)} + \mu_1 \sum_{i=0}^{S_1} P_{iS_2} \tag{D.1.2}$$

$$(\lambda_1+q_2+\mu_2)P_{(S_1+1)j} = q_2 P_{(S_1+1)(j+1)} + \mu_1 \sum_{i=0}^{S_1} P_{ij};\ \ j=1,...,S_2-1 \tag{D.1.3}$$

$$(\lambda_1+\mu_2)P_{(S_1+1)0} = q_2 P_{(S_1+1)1} + \mu_1 \sum_{i=0}^{S_1} P_{i0} \tag{D.1.4}$$

$$(\lambda_2+q_1+\mu_1)P_{S_1(S_2+1)} = \lambda_1 P_{(S_1+1)(S_2+1)} + \mu_2 \sum_{j=0}^{S_2} P_{S_1 j} \tag{D.1.5}$$

$$(\lambda_2+q_1+\mu_1)P_{i(S_2+1)} = q_1 P_{(i+1)(S_2+1)} + \mu_2 \sum_{j=0}^{S_2} P_{ij};\ \ i=1,...,S_1-1 \tag{D.1.6}$$

$$(\lambda_2+\mu_1)P_{0(S_2+1)} = q_1 P_{1(S_2+1)} + \mu_2 \sum_{j=0}^{S_2} P_{0j} \tag{D.1.7}$$

$$(\mu_1+\mu_2+q_1+q_2)P_{S_1S_2} = \lambda_1 P_{(S_1+1)S_2} + \lambda_2 P_{S_1(S_2+1)} \tag{D.1.8}$$

$$(\mu_1+\mu_2+q_1+q_2)P_{S_1 j} = \lambda_1 P_{(S_1+1)j} + q_2 P_{S_1(j+1)};\ \ j=0,1,...,S_2-1 \tag{D.1.9}$$

$$(\mu_1+\mu_2+q_1+q_2)P_{iS_2} = \lambda_2 P_{i(S_2+1)} + q_1 P_{(i+1)S_2};\ \ i=0,1,...,S_1-1 \tag{D.1.10}$$

$$(\mu_1+\mu_2+q_1+q_2)P_{ij} = q_1 P_{(i+1)j} + q_2 P_{i(j+1)};\ \ i=1,...,S_1-1,\ \ j=1,...,S_2-1 \tag{D.1.11}$$

$$(\mu_1+\mu_2+q_1+q_2)P_{i0} = (q_1+q_2)P_{(i+1)0} + q_2 P_{i1};\ \ i=1,...,S_1-1 \tag{D.1.12}$$

$$(\mu_1+\mu_2+q_1+q_2)P_{0j} = (q_1+q_2)P_{0(j+1)} + q_1 P_{1j};\ \ j=1,...,S_2-1 \tag{D.1.13}$$

$$\sum_{i=0}^{S_1+1}\sum_{j=0}^{S_2+1} P_{ij} = 1 \tag{D.1.14}$$

Having two independent suppliers guarantees that $P_{(S_1+1)(S_2+1)} = (\frac{\mu_1}{\mu_1+\lambda_1})(\frac{\mu_2}{\mu_2+\lambda_2})$ (i.e., long-run probability that both HNPs are not disrupted), $\sum_{i=0}^{S_1} P_{i(S_2+1)} = (\frac{\lambda_1}{\mu_1+\lambda_1})(\frac{\mu_2}{\mu_2+\lambda_2})$ (i.e., long-run probability that HNP 1 is disrupted and HNP 2 is not disrupted), and $\sum_{j=0}^{S_2} P_{(S_1+1)j} = (\frac{\mu_1}{\mu_1+\lambda_1})(\frac{\lambda_2}{\mu_2+\lambda_2})$ (i.e., long-run probability that HNP 1 is not disrupted and HNP 2 is disrupted). We verify that these results hold in Equation (D.1.1).

D.1.2 Non-Perishable Non-Integrated Inventory System

In the numerical analysis in Chapter 5 (in Section 5.5), we are interested in comparing a two independent supplier integrated inventory system to a non-integrated inventory system (i.e., no lateral transshipments between HNPs). This implies that each HNP acts independently. Thus, we can solve for the optimal order-up-to level at each HNP by modeling a one HNP one supplier system

as a continuous time Markov chain (see Figure D1; we simply use q to indicate the daily demand rate for the single HNP, λ to indicate the supply chain disruption rate for the single supplier, and μ to indicate the supply chain recovery rate for the single supplier). In this continuous time Markov chain, the state at any time t ($X_t; t \geq 0$) is defined as I. For $0 \leq I \leq S; I \in \mathbb{Z}_0^+$, the supply chain is disrupted and I represents the inventory on-hand at the HNP. For $I = S+1$, the supply chain is not disrupted and the inventory on-hand is S. The structure of the Markov chain assumes that the order-up-to level (i.e., S) is at least 1. Also, the Markov chain is irreducible and it has a finite number of states where all states are positive recurrent. Thus, the limiting distribution exists and is unique (Sericola, 2013).

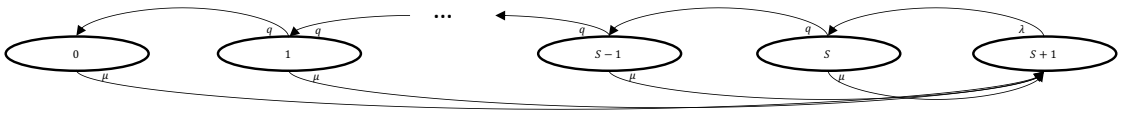

Figure D1: One HNP with one supplier Markov chain.

D.1.2.1 One HNP with One Supplier Long-run Probabilities (Non-Integrated)

Using Figure D1, we write the balance equations for the Markov chain and solve for the long-run probability that the Markov chain is in state $X = (I)$ (i.e., $P_i = \Pr(I = i)$) as shown in Equations (D.1.15)-(D.1.17).

$$P_{S+1} = \frac{\mu}{\mu + \lambda} \tag{D.1.15}$$

$$P_i = \left(\frac{\lambda}{\mu+q}\right)\left(\frac{\mu}{\mu+\lambda}\right)\left(\frac{q}{\mu+q}\right)^{S-i};\ i = 1,2,...,S \tag{D.1.16}$$

$$P_0 = \left(\frac{q}{\mu}\right)\left(\frac{\lambda}{\mu+q}\right)\left(\frac{\mu}{\mu+\lambda}\right)\left(\frac{q}{\mu+q}\right)^{S-1} \tag{D.1.17}$$

D.1.2.2 One HNP with One Supplier Expected Cost Per Day (Non-Integrated)

In a one HNP inventory system, we have no lateral transshipment costs since we are focusing on one HNP. Therefore, the non-perishable expected cost per day only consists of holding costs (h) and shortage costs (b). Using the cost parameters (see Table 5.2) and the long-run probability that the inventory system is in state $X = (I)$ (i.e., P_i; see Equations (D.1.15)-(D.1.17)), we have the expected cost per day as presented in Equation (D.1.18). We define $c_{(1H1S)}$ (corresponds to the constant term) and $e_{(1H1S)}$ (corresponds to the exponential term) which are fixed values that do not depend on the order-up-to level S for the one HNP system with one supplier (i.e., $1H1S$).

$$
\begin{aligned}
\mathbb{E}[\text{cost per day}]_{1H1S} &= h(S)P_{S+1} + \sum_{i=1}^{S} h(i)P_i + b(q)P_0 \\
&= \underbrace{-h\left(\frac{\lambda q}{\mu(\mu+\lambda)}\right)}_{c_{(1H1S)}} \\
&+ \left(\underbrace{h\left(\frac{\mu}{\mu+\lambda}\right) + h\left(\frac{\lambda}{\mu+\lambda}\right)}_{h}\right)S \\
&+ \left(\underbrace{\left(\frac{\lambda q}{\mu+\lambda}\right)\left(\frac{h}{\mu}+b\right)}_{e_{(1H1S)}}\right)\left(\frac{q}{q+\mu}\right)^S \\
&= c_{(1H1S)} + hS + e_{(1H1S)}\left(\frac{q}{q+\mu}\right)^S
\end{aligned}
\tag{D.1.18}
$$

D.1.2.3 One HNP with One Supplier Optimal Order-up-to Level (Non-Integrated)

We consider Equation (D.1.18) to solve for the optimal order-up-to level S for the one HNP inventory system with one supplier. Using the same approach and arguments presented in Section 5.4.1.4 of Chapter 5 along with observing that $e_{(1H1S)} = \left(\frac{\lambda q}{\mu+\lambda}\right)\left(\frac{h}{\mu}+b\right) > 0$, we present the optimal order-up-to level S in Equation (D.1.19).

$$
S = \max\{1, \frac{\ln\left(\frac{-h}{e_{(1H1S)}\ln(\frac{q}{q+\mu})}\right)}{\ln(\frac{q}{q+\mu})}\}
\tag{D.1.19}
$$

D.1.2.4 Solving for the Non-Integrated Inventory Policies (Non-Integrated)

Using Equation (D.1.19), we can independently solve for the optimal non-integrated order-up-to level for HNP 1 (i.e., S_1^*) and HNP 2 (i.e., S_2^*) by using the respective input parameters for each HNP. It is important to consider that when the two-HNP system operates as a non-integrated inventory system, it is not necessary to distinguish two independent suppliers and one supplier given the same disruption rates (i.e., $\lambda_1 = \lambda_2$) and recovery rates (i.e., $\mu_1 = \mu_2$). The reason is that the HNPs act independently and thus, in the long-run, the expected costs and waste at each HNP will be the same regardless of distinguishing two independent suppliers and one supplier. However, if the disruption rates (i.e., $\lambda_1 \neq \lambda_2$) and/or recovery rates (i.e., $\mu_1 \neq \mu_2$) are different from one another, it must be assumed that the system consists of two independent suppliers.

D.1.2.5 Using the Non-Integrated Inventory System to Determine $S_{k(min)}$ and $S_{k(max)}$ (Non-Integrated)

In Section 5.4.1.3 of Chapter 5, we define the feasible integer order-up-to levels as $S_1 \in [S_{1(min)}, S_{1(max)}]$ and $S_2 \in [S_{2(min)}, S_{2(max)}]$. An important observation is that we can use the one HNP with one supplier optimal order-up-to level derived for the non-integrated inventory system to strategically select these values. For $S_{k(min)}$ $(k = 1,2)$, we use Equation (D.1.19) with $q = \min\{q_1, q_2\}$, $h = h_1 + h_2$, $\lambda = \frac{\lambda_1 \lambda_2}{\lambda_1 + \lambda_2}$, and $\mu = \mu_1 + \mu_2$. For $S_{k(max)}$ $(k = 1,2)$, we use Equation (D.1.19) with $q = q_1 + q_2$, $h = \min\{h_1, h_2\}$, $\lambda = \max\{\lambda_1, \lambda_2\}$, and $\mu = \min\{\mu_1, \mu_2\}$.

D.1.3 Perishable Non-Integrated Inventory System

In Section 5.4.2 of Chapter 5, we describe how to extend the non-perishable integrated inventory system with two independent suppliers to a perishable inventory system using a perishability condition. We apply the same approach to the non-integrated inventory system and we proceed to formally describe the procedure.

D.1.3.1 Defining the Perishability Condition (Non-Integrated)

After solving for the optimal non-perishable (S_1^*, S_2^*) inventory policies for a non-integrated inventory system (in Appendix D.1.2), we consider that the drug has a finite and deterministic lifetime of x days. To capture this, we define a perishability condition which can be seen as a chance constraint at each HNP. Like the two independent supplier inventory system, when following the (S_1^*, S_2^*) non-integrated inventory policy, we want the probability that drugs are wasted at each HNP to be at most δ. To calculate the corresponding probability for each HNP, we assume that the supplier for each HNP is never disrupted (i.e., no supply chain disruptions). We define the random variable Q_{x,q_k} as the number of demand arrivals in a time period of x days when the Poisson demand rate is q_k for $k = 1,2$.

Using these definitions, we present the probability that drugs are wasted at HNP 1 given supply at HNP 1 is never disrupted and an expiration lifetime x (i.e., $P_{\text{waste HNP 1}|(\text{ND}, x)}$) in Equation (D.1.20). We present the probability that drugs are wasted at HNP 2 given supply at HNP 2 is never disrupted and an expiration lifetime x (i.e., $P_{\text{waste HNP 2}|(\text{ND}, x)}$) in Equation (D.1.21).

$$P_{\text{waste HNP 1}|(\text{ND}, x)} = \Pr(Q_{x,q_1} \leq S_1 - 1) = \left(e^{-q_1 x} \sum_{i=0}^{S_1 - 1} \frac{(q_1 x)^i}{i!} \right) \tag{D.1.20}$$

$$P_{\text{waste HNP 2}|(\text{ND}, x)} = \Pr(Q_{x,q_2} \leq S_2 - 1) = \left(e^{-q_2 x} \sum_{j=0}^{S_2 - 1} \frac{(q_2 x)^j}{j!} \right) \tag{D.1.21}$$

D.1.3.2 Perishability Condition as a Chance Constraint (Non-Integrated)

When following the (S_1^*, S_2^*) non-integrated inventory policy, we want the probability that drugs are wasted at HNP k ($k = 1, 2$) given supply is never disrupted at HNP k to be at most δ. With this, we formally define the perishability condition in Equation (D.1.22).

$$P_{\text{waste HNP 1}|(\text{ND}, x)} = e^{-q_1 x} \sum_{i=0}^{S_1 - 1} \frac{(q_1 x)^i}{i!} \leq \delta$$

$$\text{and} \tag{D.1.22}$$

$$P_{\text{waste HNP 2}|(\text{ND}, x)} = e^{-q_2 x} \sum_{j=0}^{S_2 - 1} \frac{(q_2 x)^j}{j!} \leq \delta$$

When the perishability condition is satisfied, (S_1^*, S_2^*) provides the optimal non-integrated inventory policy for the perishable system. When the perishability condition is violated (i.e., $P_{\text{waste HNP 1}|(\text{ND}, x)} > \delta$ or $P_{\text{waste HNP 2}|(\text{ND}, x)} > \delta$), we enforce the perishability condition as described in Appendix D.1.3.3.

D.1.3.3 Enforcing the Perishability Condition (Non-Integrated)

When the perishability condition is violated (i.e., $P_{\text{waste HNP 1}|(\text{ND}, x)} > \delta$ or $P_{\text{waste HNP 2}|(\text{ND}, x)} > \delta$), we enforce the perishability condition to ensure that the order-up-to levels (S_1^*, S_2^*) are appropriate for the expiration lifetime of x days. From the perishability condition presented in Equation (D.1.22), we notice that $P_{\text{waste HNP 1}|(\text{ND}, x)}$ only depends on S_1^* and $P_{\text{waste HNP 2}|(\text{ND}, x)}$ only depends S_2^*. Therefore, we reduce S_1^* until $P_{\text{waste HNP 1}|(\text{ND}, x)} \leq \delta$ and similarly, we reduce S_2^* until $P_{\text{waste HNP 2}|(\text{ND}, x)} \leq \delta$. We proceed to formally present the algorithm.

Algorithm Enforcing the perishability condition for the non-integrated inventory system.

Input:
Parameters: h_1, h_2, b, t_{12}, t_{21}, q_1, q_2, λ_1, μ_1, λ_2, μ_2 (see Table 5.2)
Optimal non-perishable policy: (S_1^*, S_2^*)

Step A: Decrease S_1^* Until the Waste Constraint is Satisfied at HNP 1
while $P_{\text{waste HNP 1}|(\text{ND}, x)} > \delta$ *% Waste constraint not satisfied at HNP 1; See equation* (D.1.22)
    Step A.1: Decrease S_1^*
    $S_1^* = \max\{1, S_1^* - 1\}$ *% Decrease the order-up-to level at HNP 1 by 1 if the boundary condition is not met*
    Calculate $P_{\text{waste HNP 1}|(\text{ND}, x)}$ with S_1^* *% See equation* (D.1.22)

    Step A.2: Check if the Order-up-to Level Boundary Condition is Met
    if $S_1^* = 1$ *% Boundary condition is met at HNP 1*
        break while loop (note: perishability condition may still be violated)
    end if
end while

Step B: Decrease S_2^* Until the Waste Constraint is Satisfied at HNP 2
while $P_{\text{waste HNP 2}|(\text{ND}, x)} > \delta$ *% Waste constraint not satisfied at HNP 2; See equation* (D.1.22)
    Step B.1: Decrease S_2^*
    $S_2^* = \max\{1, S_2^* - 1\}$ *% Decrease the order-up-to level at HNP 2 by 1 if the boundary condition is not met*
    Calculate $P_{\text{waste HNP 2}|(\text{ND}, x)}$ with S_2^* *% See equation* (D.1.22)

    Step B.2: Check if the Order-up-to Level Boundary Condition is Met
    if $S_2^* = 1$ *% Boundary condition is met at HNP 2*
        break while loop (note: perishability condition may still be violated)
    end if
end while

Output: (S_1^*, S_2^*) denotes the optimal perishable non-integrated inventory policy

D.1.4 Non-Perishable Integrated Inventory System with One Supplier

D.1.4.1 Long-run Probabilities (One Supplier Integrated)

Figure 5.2 in Chapter 5 illustrates the integrated inventory system Markov chain when the two suppliers are independent from one another. In Section 5.5.3 of the numerical analysis, we are interested in illustrating the impact of operating as an integrated inventory system (i.e., lateral transshipments are permitted) when the HNPs have the same supplier.

For one supplier, we have the modified integrated inventory system Markov chain as presented in Figure D2. We let λ and μ represent the disruption rate and recovery rate, respectively, for the single supplier. We let $X_t = (I_1, I_2)$ represent the state at time t. For $0 \leq I_1 \leq S_1$; $0 \leq I_2 \leq S_2$, we let I_1 and I_2 represent the inventory on-hand at HNP 1 and HNP 2, respectively, when the supply chain is disrupted. We introduce the state ($I_1 = S_1 + 1$, $I_2 = S_2 + 1$) to denote the inventory on-hand at HNP 1 and HNP 2 is S_1 and S_2, respectively, plus the supply chain is not disrupted (recall, S_1 and S_2 denote the maximum inventory on-hand at any point in time). The integrated inventory system Markov chain with one supplier is irreducible and it has a finite number of states where all states are positive recurrent. Thus, the limiting distribution exists and is unique (Sericola, 2013).

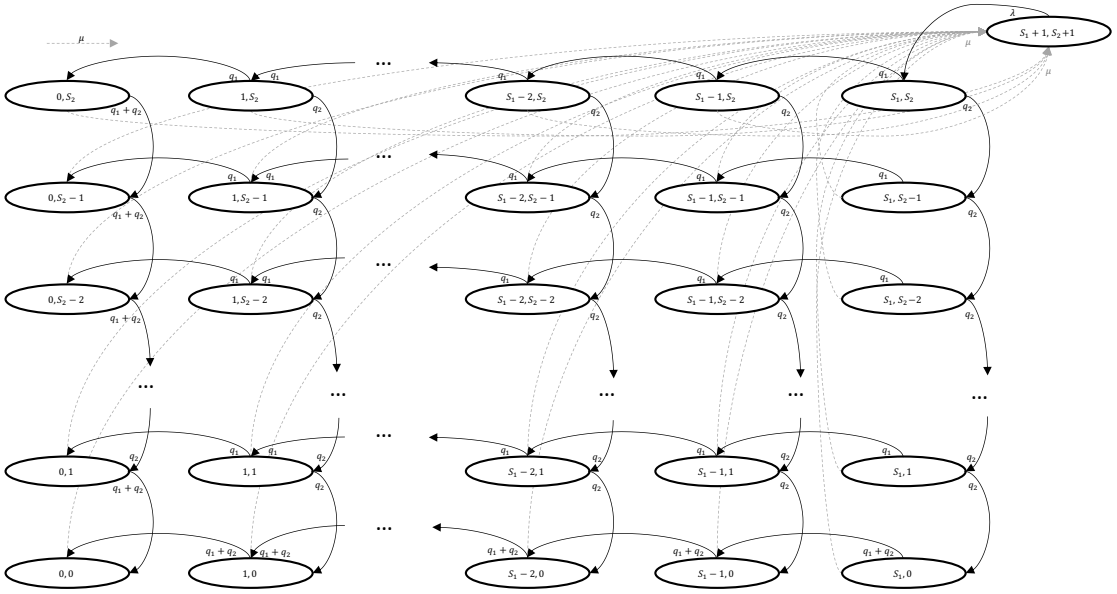

Figure D2: Integrated inventory system Markov chain with one supplier.

Using Figure D2, we write the balance equations for the Markov chain and solve for the long-run probability that the integrated inventory system with one supplier is in state $X = (I_1, I_2)$ (i.e., P_{ij}). We present the long-run probabilities in Equations (D.1.23)-(D.1.30).

$$P_{(S_1+1)(S_2+1)} = \frac{\mu}{\mu+\lambda} \tag{D.1.23}$$

$$P_{S_1 S_2} = \left(\frac{\mu}{\mu+\lambda}\right)\left(\frac{\lambda}{\mu+q_1+q_2}\right) \tag{D.1.24}$$

$$P_{S_1 j} = \left(\frac{q_2}{\mu+q_1+q_2}\right)^{S_2-j}\left(\frac{\mu}{\mu+\lambda}\right)\left(\frac{\lambda}{\mu+q_1+q_2}\right);\ \ j = 0,1,...,S_2-1 \tag{D.1.25}$$

$$P_{iS_2} = \left(\frac{q_1}{\mu+q_1+q_2}\right)^{S_1-i}\left(\frac{\mu}{\mu+\lambda}\right)\left(\frac{\lambda}{\mu+q_1+q_2}\right);\ \ i = 0,1,...,S_1-1 \tag{D.1.26}$$

$$P_{ij} = \binom{S_1-i+S_2-j}{S_1-i}\left(\frac{q_1}{\mu+q_1+q_2}\right)^{S_1-i}\left(\frac{q_2}{\mu+q_1+q_2}\right)^{S_2-j}\left(\frac{\mu}{\mu+\lambda}\right)\left(\frac{\lambda}{\mu+q_1+q_2}\right);$$
$$i = 1,2,...,S_1-1, j = 1,2,...,S_2-1 \tag{D.1.27}$$

$$P_{i0} = \sum_{k=0}^{S_1-i}\left(\binom{S_2+k-1}{k}\left(\frac{q_1+q_2}{\mu+q_1+q_2}\right)^{S_1-i-k}\left(\frac{q_2}{\mu+q_1+q_2}\right)^{S_2}\left(\frac{q_1}{\mu+q_1+q_2}\right)^{k}\left(\frac{\mu}{\mu+\lambda}\right)\left(\frac{\lambda}{\mu+q_1+q_2}\right)\right);$$
$$i = 1,2,...,S_1-1 \tag{D.1.28}$$

$$P_{0j} = \sum_{k=0}^{S_2-j}\left(\binom{S_1+k-1}{k}\left(\frac{q_1+q_2}{\mu+q_1+q_2}\right)^{S_2-j-k}\left(\frac{q_1}{\mu+q_1+q_2}\right)^{S_1}\left(\frac{q_2}{\mu+q_1+q_2}\right)^{k}\left(\frac{\mu}{\mu+\lambda}\right)\left(\frac{\lambda}{\mu+q_1+q_2}\right)\right);$$
$$j = 1,2,...,S_2-1 \tag{D.1.29}$$

$$P_{00} = \left(\frac{q_1+q_2}{\mu}\right)\left(P_{01}+P_{10}\right) \tag{D.1.30}$$

D.1.4.2 Expected Cost Per Day (One Supplier Integrated)

Using the cost parameters (see Table 5.2) and the long-run probability that the inventory system is in state $X = (I_1 = i, I_2 = j)$ (i.e., P_{ij}; see Equations (D.1.23)-(D.1.30)), we have the expected cost per day for the two-HNP integrated inventory system with one supplier (i.e., 2H1S) as presented in Equation (D.1.31).

$$\begin{aligned}\mathbb{E}[\text{cost per day}]_{2H1S} = &\overbrace{(h_1 S_1 + h_2 S_2) P_{(S_1+1)(S_2+1)} + \sum_{i=1}^{S_1}\sum_{j=1}^{S_2}(h_1 i + h_2 j) P_{ij}}^{\text{Holding Cost Only}} \\ &+ \overbrace{\sum_{i=1}^{S_1}\Big(h_1 i + t_{12} q_2\Big) P_{i0}}^{\text{Holding and } t_{12} \text{ Cost}} + \overbrace{\sum_{j=1}^{S_2}\Big(h_2 j + t_{21} q_1\Big) P_{0j}}^{\text{Holding and } t_{21} \text{ Cost}} \\ &+ \overbrace{b(q_1 + q_2) P_{00}}^{\text{Shortage Cost Only}}\end{aligned} \tag{D.1.31}$$

D.1.4.3 Optimal Order-up-to Levels (One Supplier Integrated)

We exhaustively calculate the expected cost per day for all feasible integer order-up-to level combinations (S_1, S_2): $S_1 \in [S_{1(min)}, S_{1(max)}]$ and $S_2 \in [S_{2(min)}, S_{2(max)}]$. $S_{k(min)}$ and $S_{k(max)}$ are bounds that denote the minimum and maximum order-up-to level for HNP k, respectively. In Appendix D.1.2.5, we illustrate how to strategically select $S_{1(min)}$, $S_{1(max)}$, $S_{2(min)}$, and $S_{2(max)}$. We select (S_1^*, S_2^*) corresponding to the solution with the smallest expected cost per day.

D.1.5 Perishable Integrated Inventory System with One Supplier

In Appendix D.1.3, we describe how to extend the non-perishable non-integrated inventory system to a perishable inventory system using a perishability condition. A one supplier system implies that when supply is never disrupted at one HNP (recall we condition on supply never being disrupted when establishing the perishability condition), supply is also never disrupted at the second HNP since they have the same supplier. As a result, the probability that drugs are wasted only depends on the individual HNP's order-up-to level like the non-integrated inventory system. Hence, we can enforce the perishability condition as described in Appendix D.1.3.3.

D.2 Evaluating the Accuracy of the Approximate Model

We proceed to evaluate the accuracy of the Two Supplier Integrated (Approximate Model) inventory system.

D.2.1 Evaluating the Accuracy with the Expected Cost Per Day

We first evaluate the accuracy of the two-HNP integrated approximate model with two independent suppliers using the expected cost per day for the non-perishable integrated inventory system. We consider all input parameter combinations presented in (i) demonstrating when it is beneficial to allow two HNPs with independent suppliers experiencing supply chain disruptions to share inventory (in Section 5.5.1 Figure 5.5) and all input parameter combinations presented in (ii) providing insights on how the supply chain disruption characteristics of the two HNPs influence whether it is beneficial to share inventory (in Section 5.5.2 Figure 5.6). However, to ensure that the system of equations with the two-HNP integrated exact model is computationally possible in an appropriate amount of time, we multiply all costs by $q_1 = q_2 = 45$ (i.e., $h_1, h_2, t_{12}, t_{21}, b$) and simply consider a daily demand rate of $q_1' = q_2' = 1$. Since we convert the input parameters, it is important to keep in mind that the true order-up-to levels are $S^*_{1(true)} = 45 \cdot S^*_1$ and $S^*_{2(true)} = 45 \cdot S^*_2$. However, we highlight that these order-up-to levels are not exact since we are "batching" the demand, and hence, reducing the variance of the demand as it follows a Poisson distribution (i.e., variance equals the demand rate; Kulkarni, 2011).

With these modified input parameter combinations, we first solve for the non-perishable optimal order-up-to levels (S^*_1, S^*_2) using the two-HNP integrated approximate model. We record the expected cost per day for this (S^*_1, S^*_2) integrated inventory policy when using the approximate model. Then, we use this same (S^*_1, S^*_2) policy and calculate the expected cost per day when using the system of equations for the exact model. We find comparable results (difference in expected total cost per day $\leq 2.6\%$) supporting that the approximate model achieves high accuracy.

Above, we consider a daily demand rate of $q_1' = q_2' = 1$ to ensure that the system of equations with the two-HNP integrated exact model is computationally possible in an appropriate amount of time. To further validate the approximate model, we consider the original base case input parameters (see Table 5.4; in Section 5.5) with $\mu_1 = \mu_2 = \frac{1}{30}$ [$\mu_1 = \mu_2 = \frac{1}{90}$]. We solve for the optimal order-up-to levels with the two-HNP perishable integrated approximate model which leads to $S^*_1 = S^*_2 = 2,666$ [$S^*_1 = S^*_2 = 3,952$]. The approximate model calculates that the expected cost per day is \$171.76 [\$603.06]. These optimal order-up-to levels are far too large to calculate the expected cost per day with the exact model in an appropriate amount of time. Instead, we simulate the two-HNP non-perishable integrated inventory system (1,000 simulation replications, warm-up period of 500 days, planning horizon of 10,000 days; since the input parameters for both HNPs are equivalent, these simulation parameters ensure the difference $\leq 2.5\%$ with respect to each cost component). We find an expected cost per day of \$173.23 [\$615.38]. Therefore, the approximate model calculates an expected cost per day that is only 1% different [2% different] than the expected cost per day with the simulation model. This again supports that the approximate model achieves high accuracy. We highlight that we consider a simulation model of the non-perishable

inventory system when evaluating the accuracy of the approximate model because the closed-form expressions for the approximate model are derived for a non-perishable inventory system.

D.2.2 Evaluating the Accuracy with the Order-up-to Levels

We second evaluate the accuracy of the two-HNP integrated approximate model with two independent suppliers using the optimal order-up-to levels for the non-perishable integrated inventory system. We consider the input parameter combinations presented in (i) demonstrating when it is beneficial to allow two HNPs with independent suppliers experiencing supply chain disruptions to share inventory (in Section 5.5.1 Figure 5.5) and the input parameter combinations presented in (ii) providing insights on how the supply chain disruption characteristics of the two HNPs influence whether it is beneficial to share inventory (in Section 5.5.2 Figure 5.6). However, to ensure that the system of equations with the two-HNP integrated exact model is computationally possible in an appropriate amount of time, we multiply all costs by $q_1 = q_2 = 45$ (i.e., $h_1, h_2, t_{12}, t_{21}, b$) and simply consider a daily demand rate of $q_1' = q_2' = 1$. Furthermore, to ensure that the exhaustive search runs in an appropriate amount of time, we only consider cases where the approximate model solves for order-up-to levels ≤ 100 (i.e., $S_1^* \leq 100$ and $S_2^* \leq 100$) and we consider increments of 5 for the exhaustive search.

With the modified input parameter combinations, we first solve for the non-perishable optimal order-up-to levels (S_1^*, S_2^*) using the two-HNP integrated exact model and calculate the expected cost per day when using the exact model. Second, we solve for the optimal order-up-to levels (S_1^*, S_2^*) using the two-HNP integrated approximate model, but calculate the expected cost per day when using the exact model. We find comparable results (difference in expected total cost per day $\leq 0.2\%$) supporting that the approximate model achieves high accuracy.

D.3 Numerical Analysis Explanations

In Section 5.5 of Chapter 5, we complete a numerical analysis where we study a two-HNP inventory system. We proceed to present details regarding the sensitivity of t^* for a broad range of disruption and recovery rates, an instance where a HNP can be “worse-off", and additional takeaways.

D.3.1 Numerical Analysis: Sensitivity of t^* for a Broad Range of Disruption and Recovery Rates

In Section 5.5.1 of Chapter 5, the most interesting finding is that (d) it is not always beneficial to operate as an integrated inventory system given the cost of a lateral transshipment is less than the cost of a shortage for a two-HNP inventory system with independent suppliers. Instead, there exists a lateral transshipment cost $t^* \leq b$ such that it becomes optimal to operate as a non-integrated inventory system given $t_{12} = t_{21}$. This t^* may be strictly less than b. Using the base case input parameters (see Table 5.4; in Section 5.5), we study the sensitivity of t^* for varying disruption and recovery rate combinations. Furthermore, due to the difficulty in selecting the disruption rate and recovery rate, we study an even broader range than those presented in Table 5.4 (in Section 5.5). We consider varying disruption rates such that $\lambda_1 = \lambda_2 \in [0.001, 0.1]$; expected time between supply chain disruptions within 10-1000 days. We consider varying recovery rates such that $\mu_1 = \mu_2 \in [0.001, 0.1]$; expected supply chain disruption durations within 10-1000 days. When finding the value of t^*, we consider t^* in increments of \$2.50.

Figure D3 illustrates the results where for generalizability, we present the dimensionless value $\frac{t^*}{b}$. The value $\frac{t^*}{b}$ represents the ratio of the lateral transshipment cost such that it becomes optimal to operate as a non-integrated inventory system given $t_{12} = t_{21}$ to the shortage cost. We find that (a) the value of t^* is largely influenced by the disruption rate and recovery rate (i.e., supply chain disruption parameters). More specifically, when $\lambda_1 = \lambda_2$ and $\mu_1 = \mu_2$, t^* decreases as the long-run probability that the supply chain is disrupted increases. We also find that (b) regardless of the disruption and recovery rate combination, we reach the same conclusion that the lateral transshipment cost must be sufficiently less than the cost of a shortage for it to be beneficial to operate as an integrated inventory system.

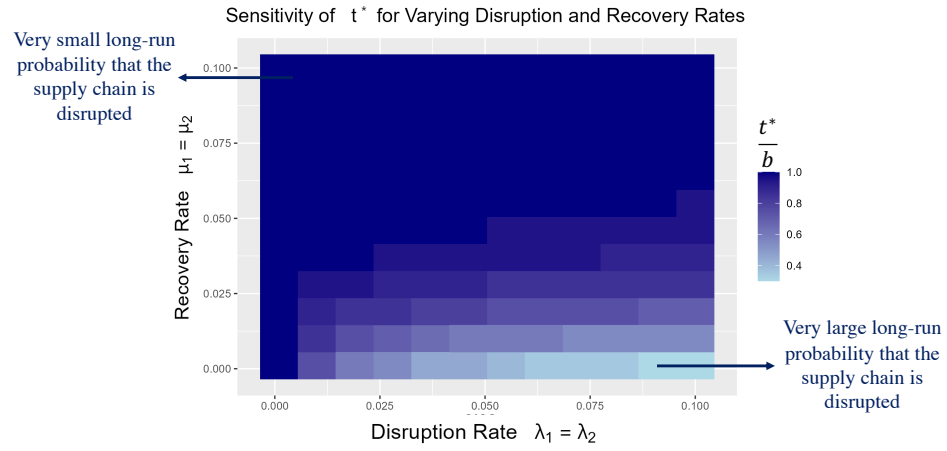

Figure D3: Results for the sensitivity of t^*.

D.3.2 Numerical Analysis: “Worse-Off" Example

In Section 5.5.2 of Chapter 5, the final interesting finding is that (d) when selecting a partner HNP, it is critical to consider the disruption rate and recovery rate of the partner HNP. It is also worth noting that as the long-run probability that the supply chain is disrupted increases, there are instances where failing to consider the disruption rate and recovery rate of the partner HNP can cause one HNP to be “worse-off". In Figure D4, we present the Two Supplier Integrated Perishable (Approximate Model) and Non-Integrated Perishable inventory systems with the base case input parameters (see Table 5.4; in Section 5.5), but varying disruption and recovery rates at HNP 1 and HNP 2. For all instances in Figure D4, the long-run probability that the supply chain is not disrupted is $\frac{\mu_k}{\lambda_k+\mu_k} = \frac{3}{7} = 0.43$ for $k = 1,2$. With the base case input parameters, we want to emphasize that $t_{12} = t_{21} = \$12.50$ is sufficiently less than the cost of a shortage $b = \$50$ for all disruption and recovery rate combinations in this analysis to result in cost savings when using the integrated inventory system (see result (d) in Section 5.5.1). We provide the expected total, lateral transshipment, holding cost at HNP 1, holding cost at HNP 2, shortage cost at HNP 1, and shortage cost at HNP 2 per day. We also present the probability that drugs are wasted given supply is never disrupted at each respective HNP and the optimal perishable order-up-to levels (S_1^*, S_2^*) at each HNP.

In Figure D4, consider when we have long expected disruption durations at one HNP (e.g., $\mu_1 = 0.01$; $\frac{1}{\mu_1} = 100$ days; see x-axis) in comparison to the second HNP (e.g., $\mu_2 = 0.1$; $\frac{1}{\mu_2} = 10$ days; see triangular curves). The results illustrate that the second HNP has a larger order-up-to level for the integrated inventory system in comparison to the non-integrated inventory system (see second row fifth column). Even worse, the second HNP has a larger expected shortage cost (i.e., expected shortage cost almost doubles; second row third column). The first HNP reaps benefits (i.e., expected shortage cost with the integrated inventory system is about 4% the shortage cost of the non-integrated inventory system; see first row third column). But, the results illustrate that HNPs need to take into consideration the disruption rate and recovery rate of the partner HNP’s supplier before deciding if it is beneficial to operate as an integrated inventory system.

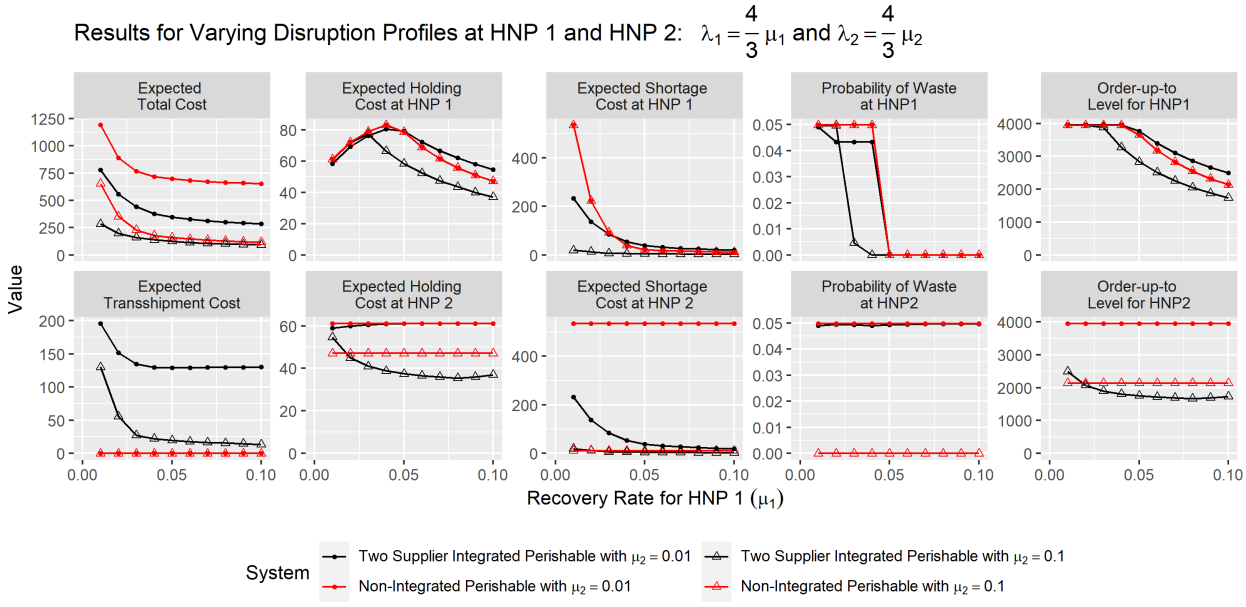

Figure D4: Results for “worse-off" example.

D.3.3 Numerical Analysis: Additional Takeaways

We proceed to present additional takeaways from the numerical analysis. In Appendix D.3.3.1, we provide the expected total, holding, shortage, and lateral transshipment cost per day. We also present the probability that drugs are wasted given supply is never disrupted at each respective HNP and the optimal perishable order-up-to levels (S_1^*, S_2^*) at each HNP. In Appendix D.3.3.2, we summarize the results in tables.

D.3.3.1 (T1) Two Independent Suppliers versus One Supplier

We illustrate the difference between two independent suppliers and one supplier. In Figure D5, we present the Two Supplier Integrated Perishable (Approximate Model), One Supplier Integrated Perishable, and Non-Integrated Perishable inventory systems with the base case input parameters (see Table 5.4; in Section 5.5), but we multiply all costs by $q_1 = q_2 = 45$ (i.e., $h_1, h_2, t_{12}, t_{21}, b$) and simply consider a daily demand rate of $q_1' = q_2' = 1$ to reduce the time of the exhaustive search for the One Supplier Integrated Perishable inventory system. We also consider varying disruption rates $\lambda = \lambda_1 = \lambda_2$ such that $\lambda = \frac{1}{3}\mu$ where $\mu = \mu_1 = \mu_2$. By setting $\lambda = \frac{1}{3}\mu$, all combinations have the same long-run probability that the supply chain is not disrupted (i.e., $\frac{\mu}{\lambda+\mu} = 0.75$). Also, all combinations correspond to expected disruption durations (i.e., down time) that are shorter than the expected time until a disruption given the supply chain is currently not disrupted (i.e., up time).

With the modified base case input parameters, we want to emphasize that $t_{12} = t_{21} = \$12.50(45)$

is sufficiently less than the cost of a shortage $b = \$50(45)$ for all disruption and recovery rate combinations in this analysis to result in cost savings when using the integrated inventory system with two independent suppliers (see result (d) in Section 5.5.1). We find that (a) the greatest benefit of an integrated inventory system is seen with two independent suppliers, as opposed to one supplier. In Figure D5, the Two Supplier Integrated Perishable inventory system (see black circular curves) falls below the One Supplier Integrated Perishable inventory system (see blue triangular curves) for the expected total cost (see row 1 column 1). When comparing the One Supplier Integrated Perishable inventory system (see blue triangular curves) to the Non-Integrated Perishable inventory system (see red square-shaped curves), we find that (b) partnering with a HNP with the same supplier provides negligible to no benefit. This implies that hospital pharmacy managers should focus on drugs that provide the opportunity to partner with a HNP that has a supplier independent of it's own supplier.

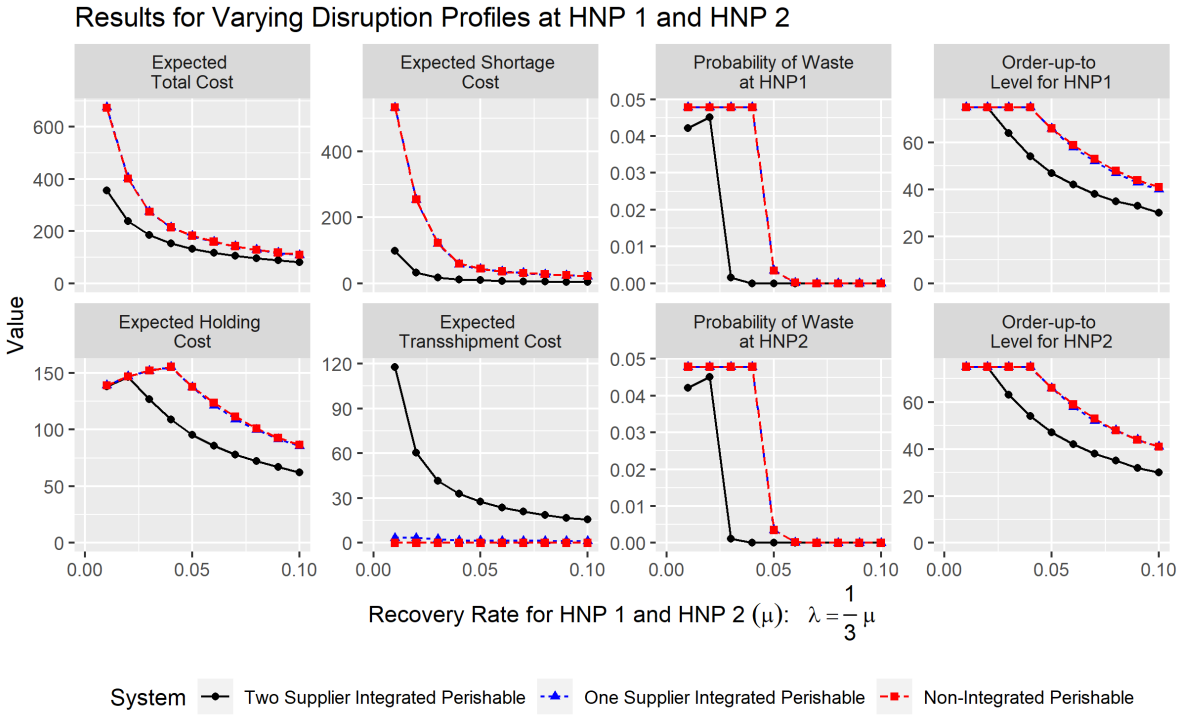

Figure D5: Results for two independent suppliers versus one supplier.

D.3.3.2 (T2) Hoarding Inventory when Both Suppliers are Disrupted

In the Two Supplier Integrated Perishable inventory system, we assume that the HNPs participate in lateral transshipments when one HNP has zero inventory on-hand, and the other HNP has positive inventory on-hand. This implies that these lateral transshipments are even possible when both suppliers are disrupted. We are interested in gaining a better understanding of what happens when

HNPs hoard inventory (i.e., do not participate in lateral transshipments when both suppliers are disrupted). For the analysis, we consider the base case input parameters (see Table 5.4; in Section 5.5), but we multiply all costs by $q_1 = q_2 = 45$ (i.e., $h_1, h_2, t_{12}, t_{21}, b$) and simply consider a daily demand rate of $q_1' = q_2' = 1$ to reduce the run time of the simulation models. Even the modified parameters require a large run time and number of simulation replications (defined later) to ensure that both HNPs have similar holding, lateral transshipment, and shortage costs (difference $\leq 2.5\%$ with respect to each cost component) which is required for model stability given the two HNPs have equivalent input parameters.

With the modified base case parameters, we first solve for the optimal perishable order-up-to levels using the Two Supplier Integrated (Approximate Model). We find $(S_1^*, S_2^*) = (60, 60)$ when $\mu_1 = \mu_2 = \frac{1}{30}$ and $(S_1^*, S_2^*) = (76, 76)$ when $\mu_1 = \mu_2 = \frac{1}{90}$. For both sets of inventory policies, we use the respective recovery rate and we simulate a Two Supplier Integrated Perishable inventory system (i.e., always allow lateral transshipments), Hoarding Two Supplier Integrated Perishable inventory system (i.e., no lateral transshipments when both suppliers are disrupted), and Non-Integrated Perishable inventory system (i.e., no lateral transshipments). For all simulation models, we consider $1,000$ simulation replications, a warm-up period of 500 days, and a planning horizon of $10,000$ days. We present the expected costs ($\mathbb{E}[T]$: expected total cost, $\mathbb{E}[H]$: expected holding cost, $\mathbb{E}[LT]$: expected lateral transshipment cost, $\mathbb{E}[S]$: expected shortage cost) for the three simulation models in Table D1 and Table D2 for $\mu_1 = \mu_2 = \frac{1}{30}$ and $\mu_1 = \mu_2 = \frac{1}{90}$, respectively. We normalize these costs with respect to the Two Supplier Integrated Perishable inventory system.

The results in Table D1 consider $\mu_1 = \mu_2 = \frac{1}{30}$ which is the base case that does not require the perishability condition to be enforced. The results illustrate that (a) the Hoarding Two Supplier Integrated Perishable inventory system (i.e., no lateral transshipments when both suppliers are disrupted) can cause 2.4 times the number of shortages in comparison to the Two Supplier Integrated Perishable inventory system (i.e., always allows lateral transshipments). We also find that (b) the Non-Integrated Perishable inventory system can lead to 9.4 times the number of shortages in comparison to the Two Supplier Integrated Perishable inventory system that always allows lateral transshipments emphasizing the benefit of integration.

Table D1 Hoarding inventory analysis for $\mu_1 = \mu_2 = \frac{1}{30}$.

Model	S_1^*	S_2^*	$\mathbb{E}[T]$	$\mathbb{E}[H]$	$\mathbb{E}[LT]$	$\mathbb{E}[S]$
Two Supplier Integrated Perishable	60	60	1.0	1.0	1.0	1.0
Hoarding Two Supplier Integrated Perishable	60	60	1.1	1.0	0.8	2.4
Non-Integrated Perishable	60	60	1.6	1.0	0	9.4

The results in Table D2 consider $\mu_1 = \mu_2 = \frac{1}{90}$ which is the base case that requires the perishability condition to be enforced. The results illustrate that (a) the Hoarding Two Supplier Integrated Perishable inventory system (i.e., no lateral transshipments when both suppliers are

disrupted) can cause 1.2 times the number of shortages in comparison to the Two Supplier Integrated Perishable inventory system (i.e., always allows lateral transshipments). We also find that (b) the Non-Integrated Perishable inventory system can lead to 2.5 times the number of shortages in comparison to the Two Supplier Integrated Perishable inventory system that always allows lateral transshipments again emphasizing the benefit of integration.

Table D2 Hoarding inventory analysis for $\mu_1 = \mu_2 = \frac{1}{90}$.

Model	S_1^*	S_2^*	$\mathbb{E}[T]$	$\mathbb{E}[H]$	$\mathbb{E}[LT]$	$\mathbb{E}[S]$
Two Supplier Integrated Perishable	76	76	1.0	1.0	1.0	1.0
Hoarding Two Supplier Integrated Perishable	76	76	1.1	1.0	0.7	1.2
Non-Integrated Perishable	76	76	1.6	1.0	0	2.5

BIBLIOGRAPHY

Aldrighetti, R., Zennaro, I., Finco, S., & Battini, D. (2019). Healthcare supply chain simulation with disruption considerations: A case study from Northern Italy. *Global Journal of Flexible Systems Management*, 20 (Suppl 1), 81–102.

Allen, M. (2017). Want to cut health care costs? Start with the obscene amount of waste. *The Washington Post*.

Arani, M., Chan, Y., Liu, X., & Momenitabar, M. (2021). A lateral resupply blood supply chain network design under uncertainties. *Applied Mathematical Modelling*, 93, 165–187.

Arikan, E. & Silbermayr, L. (2018). Risk pooling via unidirectional inventory transshipments in a decentralized supply chain. *International Journal of Production Research*, 56 (17), 5593–5610.

Atan, Z. & Rousseau, M. (2016). Inventory optimization for perishables subject to supply disruptions. *Optimization Letters*, 10, 89–108.

Avci, M. G. (2019). Lateral transshipment and expedited shipping in disruption recovery: A mean-CVaR approach. *Computers & Industrial Engineering*, 130, 35–49.

Avci, M. G. & Selim, H. (2016). A multi-agent system model for supply chains with lateral preventive transshipments: Application in a multi-national automotive supply chain. *Computers in Industry*, 82, 28–39.

Azghandi, R. (2019). Analytic models for effective planning for response to sudden disruptions. *[Doctoral Dissertation] The Department of Mechanical and Industrial Engineering: Northeastern University*, 1–97.

Bashyam, S. & Fu, M. C. (1991). Application of perturbation analysis to (s, S) inventory systems. *Proceedings of the 1991 Winter Simulation Conference*, 994–1003.

Berk, E. & Arreola-Risa, A. (1994). Note on future supply uncertainty in EOQ models. *Naval Research Logistics*, 41 (1), 129–132.

Bischak, D. P., Robb, D. J., Silver, E. A., & Blackburn, J. D. (2014). Analysis and management of periodic review, order-up-to level inventory systems with order crossover. *Production and Operations Management*, 23 (5), 762–772.

Borwein, J. M. & Crandall, R. E. (2013). Closed forms: what they are and why we care. *Notices of the American Mathematical Society*, 60 (1), 50–65.

Bozkir, C. D. C., Kundakcioglu, O. E., & Henry, A. C. (2022). Hospital service levels during drug shortages: Stocking and transshipment policies for pharmaceutical inventory. *Journal of Global Optimization*, 83, 565–584.

Centers for Disease Control and Prevention (2021). *Fentanyl*. URL: https://www.cdc.gov/opioids/basics/fentanyl.html (visited on 05/26/2022).

Czerniak, L. L., Daskin, M. S., Lavieri, M. S., Sweet, B. V., Erley, J., & Tupps, M. A. (2021). Improving simulation optimization run time when solving for periodic review inventory policies in a pharmacy. *Proceedings of the 2021 Winter Simulation Conference*, 1–12.

Czerniak, L. L., Daskin, M. S., Lavieri, M. S., Sweet, B. V., Leja, J., Tupps, M. A., & Renius, K. (2023a). Closed-form (R,S) inventory policies for perishable inventory systems with supply chain disruptions. *INFOR: Information Systems and Operational Research*, 61 (3), 327–367.

Czerniak, L. L., Lavieri, M. S., Daskin, M. S., Byon, E., Renius, K., Sweet, B. V., Leja, J., & Tupps, M. A. (2023b). The benefits (or detriments) of adapting to demand disruptions in a hospital pharmacy with supply chain disruptions. *Submitted paper*.

Czerniak, L. L., Lavieri, M. S., Daskin, M. S., Sweet, B. V., Leja, J., Tupps, M. A., & Renius, K. (2023c). When is it worth it for two hospital network pharmacies to operate as an integrated inventory system in the presence of supply chain disruptions? *IISE Transactions on Healthcare Systems Engineering*.

Dehghani, M. & Abbasi, B. (2018). An age-based lateral-transshipment policy for perishable items. *International Journal of Production Economics*, 198, 93–103.

Dehghani, M., Abbasi, B., & Oliveira, F. (2021). Proactive transshipment in the blood supply chain: A stochastic programming approach. *Omega*, 98, 1–16.

Dillon, M., Oliveira, F., & Abbasi, B. (2017). A two-stage stochastic programming model for inventory management in the blood supply chain. *International Journal of Production Economics*, 187, 27–41.

Dolgui, A., Ivanov, D., & Sokolov, B. (2020). Reconfigurable supply chain: the X-network. *International Journal of Production Research*, 58 (13), 4138–4163.

Dolgui, A., Tiwari, M. K., Sinjana, Y., Kumar, S. K., & Son, Y.-J. (2018). Optimising integrated inventory policy for perishable items in a multi-stage supply chain. *International Journal of Production Research*, 56 (1-2), 902–925.

Duc, N. T. T., Tai, P. D., & Buddhakulsomsiri, J. (2022). A Markovian approach to modeling a periodic order-up-to-level policy under stochastic discrete demand and lead time with lost sales. *International Transactions in Operational Research*, 29 (2), 1132–1158.

Eilon, S. & Elmaleh, J. (1970). Adaptive limits in inventory control. *Management Science*, 16 (8), B533–B548.

Fathi, M., Khakifirooz, M., Diabat, A., & Chen, H. (2021). An integrated queuing-stochastic optimization hybrid Genetic Algorithm for a location-inventory supply chain network. *International Journal of Production Economics*, 237, 1–13.

Firouz, M., Keskin, B. B., & Melouk, S. H. (2017). An integrated supplier selection and inventory problem with multi-sourcing and lateral transshipments. *Omega*, 70, 77–93.

Franco, C. & Alfonso-Lizarazo, E. (2020). Optimization under uncertainty of the pharmaceutical supply chain in hospitals. *Computers and Chemical Engineering*, 135, 1–13.

Fu, M. (2015). *Handbook of Simulation Optimization*. Ed. by M. Fu. New York: Springer.

Fu, M. & Hu, J. Q. (1997). *Conditional Monte Carlo*. Boston, MA: Springer US.

Gebicki, M., Mooney, E., Chen, S.-J., & Mazur, L. M. (2014). Evaluation of hospital medication inventory policies. *Health Care Management Science*, 17, 215–229.

Ghomi-Avili, M., Jalali Naini, S., Tavakkoli-Moghaddam, R., & Jabbarzadeh, A. (2017). A network design model for a resilient closed-loop supply chain with lateral transshipment. *International Journal of Engineering*, 30 (3), 374–383.

Gullu, R., Onol, E., & Erkip, N. (1997). Analysis of a deterministic demand production/inventory system under nonstationary supply uncertainty. *IIE Transactions*, 29 (8), 703–709.

Haijema, R. (2013). A new class of stock-level dependent ordering policies for perishables with a short maximum shelf life. *International Journal of Production Economics*, 143 (2), 434–439.

Hartman, M., Martin, A. B., Espinosa, N., Catlin, A., & The National Health Expenditure Acc (2018). National health care spending in 2016: spending and enrollment growth slow after initial coverage expansions. *Health Affairs*, 37 (1), 150–160.

HDMA (2009). Understanding the drivers of expired pharmaceutical returns. *Healthcare Distribution Management Association*, 1–27.

He, Y. & Wang, S. (2012). Analysis of production-inventory system for deteriorating items with demand disruption. *International Journal of Production Research*, 50 (16), 4580–4592.

Heimann, D. I. & Waage, F. (2007). A closed-form approximation solution for an inventory model with supply disruptions and non-ZIO reorder policy. *Systemics, Cybernetics and Informatics*, 5 (4), 1–12.

Herer, Y. T., Tzur, M., & Yücesan, E. (2006). The multilocation transshipment problem. *IIE Transactions*, 38 (3), 185–200.

Hong, L. J. & Nelson, B. L. (2006). Discrete optimization via simulation using COMPASS. *Operations Research*, 54 (1), 115–129.

Hu, X., Duenyas, I., & Kapuscinski, R. (2007). Existence of coordinating transshipment prices in a two-location inventory model. *Management Science*, 53 (8), 1289–1302.

— (2008). Optimal joint inventory and transshipment control under uncertain capacity. *Operations Research*, 56 (4), 881–897.

Hung, P. D. P. (n.d.). *Manufacturer Image 2216661 from Noun Project CC BY 3.0*. URL: https://thenounproject.com/icon/manufacturer-2216661/.

IBM Watson Health (2022). *Fentanyl: IBM Micromedex Red Book*. URL: https://www.ibm.com/watson-health/about/micromedex (visited on 05/26/2022).

Ivanov, D. (2021). Lean resilience: AURA (Active Usage of Resilience Assets) framework for post-COVID-19 supply chain management. *International Journal of Logistics Management*, 33 (4), 1196–1217.

— (2022). Viable supply chain model: integrating agility, resilience and sustainability perspectives—lessons from and thinking beyond the COVID-19 pandemic. *Annals of Operations Research*, 319 (1), 1411–1431.

Ivanov, D. & Dolgui, A. (2020). Viability of intertwined supply networks: extending the supply chain resilience angles towards survivability. A position paper motivated by COVID-19 outbreak. *International Journal of Production Research*, 58 (10), 2904–2915.

Ivanov, D., Dolgui, A., Sokolov, B., & Ivanova, M. (2017). Literature review on disruption recovery in the supply chain. *International Journal of Production Research*, 55 (20), 6158–6174.

Jabbarzadeh, A., Haughton, M., & Khosrojerdi, A. (2018). Closed-loop supply chain network design under disruption risks: A robust approach with real world application. *Computers & Industrial Engineering*, 116, 178–191.

Janakiraman, G. & Roundy, R. O. (2004). Lost-sales problems with stochastic lead times: Convexity results for base-stock policies. *Operations Research*, 52 (5), 795–803.

Jia, J. & Zhao, H. (2017). Mitigating the US drug shortages through pareto-improving contracts. *Production and Operations Management*, 26 (8), 1463–1480.

Jin, M. & Agirbas, G. (2013). Reducing expiry of slow-moving critical medical items with transshipment. *IIE Transactions on Healthcare Systems Engineering*, 3 (3), 193–206.

Jones, S. (2020). *What are 503B outsourcing facilities?* https://www.fagronsterile.com/newsroom/what-is-a-503b-compounding-pharmacy.

Kara, A. & Dogan, I. (2018). Reinforcement learning approaches for specifying ordering policies of perishable inventory systems. *Expert Systems with Applications*, 91, 150–158.

Karthikeyan, K. & Sudhesh, R. (2016). Recent review article on queueing inventory systems. *Research Journal of Pharmacy and Technology*, 9 (11), 1451–1461.

Khan, S. R., Kona, R., Faustino, P. J., Gupta, A., Taylor, J. S., Porter, D. A., & Khan, M. (2014). United States Food and Drug Administration and Department of Defense shelf-life extension program of pharmaceutical products: Progress and promise. *Journal of Pharmaceutical Sciences*, 103 (5), 1331–1336.

Kim, C., Jun, J., Baek, J., Smith, R., & Kim, Y. (2005). Adaptive inventory control models for supply chain management. *The International Journal of Advanced Manufacturing Technology*, 26, 1184–1192.

Konstantaras, I., Skouri, K., & Lagodimos, A. (2019). EOQ with independent endogenous supply disruptions. *Omega*, 83, 96–106.

Kouki, C., Legros, B., Zied Babai, M., & Jouini, O. (2020). Analysis of base-stock perishable inventory systems with general lifetime and lead-time. *European Journal of Operational Research*, 287 (3), 901–915.

Kulkarni, V. G. (2011). *Introduction to Modeling and Analysis of Stochastic Systems*. Springer Texts in Statistics. New York, NY: Springer New York.

Kumari, A., Wijayanayake, A., & Niwunhella, D. (2022). Lateral transhipment inventory models: A systematic literature review of models and solution approaches. *18th International Conference on Business Management*, 54–80.

Larimi, N. G., Yaghoubi, S., & Hosseini-Motlagh, S.-M. (2019). Itemized platelet supply chain with lateral transshipment under uncertainty evaluating inappropriate output in laboratories. *Socio-Economic Planning Sciences*, 68, 1–17.

Lee, C. & Park, K. S. (2016). Inventory and transshipment decisions in the rationing game under capacity uncertainty. *Omega*, 65, 82–97.

Li, J., Liu, L., Hu, H., Zhao, Q., & Guo, L. (2018). An inventory model for deteriorating drugs with stochastic lead time. *International Journal of Environmental Research and Public Health*, 15 (12), 1–20.

Li, M. & Zhang, J. (2012). Does inventory pooling improve customer service levels? *Operations Research Letters*, 40 (2), 96–98.

Li, Q., Yu, P., & Wu, X. (2016). Managing perishable inventories in retailing: Replenishment, clearance sales, and segregation. *Operations Research*, 64 (6), 1270–1284.

Li, S., He, Y., Huang, H., Lin, J., & Ivanov, D. (2023). Supply chain hoarding and contingent sourcing strategies in anticipation of price hikes and product shortages. *IISE Transactions*, 1–16.

Little, J. & Coughlan, B. (2008). Optimal inventory policy within hospital space constraints. *Health Care Management Science*, 11 (2), 177–183.

Nakandala, D., Lau, H., & Shum, P. K. (2017). A lateral transshipment model for perishable inventory management. *International Journal of Production Research*, 55 (18), 5341–5354.

Nasr, W. W., Salameh, M. K., & Moussawi-Haidar, L. (2012). Transshipment and safety stock under stochastic supply interruption in a production system. *Computers & Industrial Engineering*, 63 (1), 274–284.

National Academies of Sciences, Engineering, & Medicine (NASEM). (2022). *Building resilience into the nation's medical product supply chains*. Washington, DC: The National Academies Press.

Needham, P. M. & Evers, P. T. (1998). The influence of individual cost factors on the use of emergency transshipments. *Transportation Research Part E: Logistics and Transportation Review*, 34 (2), 149–160.

Neve, B. V. & Schmidt, C. P. (2022). Point-of-use hospital inventory management with inaccurate usage capture. *Health Care Management Science*, 25 (1), 126–145.

Nguyen, D. H. & Chen, H. (2019). Optimization of a perishable inventory system with both stochastic demand and supply: comparison of two scenario approaches. *Croatian Operational Research Review*, 10 (1), 175–185.

— (2022). An effective approach for optimization of a perishable inventory system with uncertainty in both demand and supply. *International Transactions in Operational Research*, 29 (4), 2682–2704.

Niziolek, L., Chiam, T. C., & Yih, Y. (2012). A simulation-based study of distribution strategies for pharmaceutical supply chains. *IIE Transactions on Healthcare Systems Engineering*, 2 (3), 181–189.

Noham, R. & Tzur, M. (2014). The single and multi-item transshipment problem with fixed transshipment costs. *Naval Research Logistics (NRL)*, 61 (8), 637–664.

Özdemir, D., Yücesan, E., & Herer, Y. T. (2013). Multi-location transshipment problem with capacitated production. *European Journal of Operational Research*, 226 (3), 425–435.

Parlar, M. & Berkin, D. (1991). Future supply uncertainty in EOQ models. *Naval Research Logistics*, 38 (1), 107–121.

Parlar, M. & Perry, D. (1995). Analysis of a (Q, r, T) inventory policy with deterministic and random yields when future supply is uncertain. *European Journal of Operational Research*, 84 (2), 431–443.

Paterson, C., Kiesmüller, G., Teunter, R., & Glazebrook, K. (2011). Inventory models with lateral transshipments: A review. *European Journal of Operational Research*, 210 (2), 125–136.

Paul, S. K., Sarker, R., & Essam, D. (2016). Managing risk and disruption in production-inventory and supply chain systems: A review. *Journal of Industrial and Management Optimization*, 12 (3), 1009–1029.

Peterson, R., Thomas, J., & Loiseau, A. (1972). Operational inventory control with stochastic, seasonal demand. *INFOR: Information Systems and Operational Research*, 10 (1), 81–93.

Phuong, J. M., Penm, J., Chaar, B., Oldfield, L. D., & Moles, R. (2019). The impacts of medication shortages on patient outcomes: A scoping review. *PLOS ONE*, 14 (5), 1–17.

Poormoaied, S. & Demirci, E. Z. (2021). Analysis of an inventory system with emergency ordering option at the time of supply disruption. *OR Spectrum*, 43 (4), 1007–1045.

Popcomarts (n.d.). *Hospital Building Image 3039718 from Noun Project CC BY 3.0*. URL: https://thenounproject.com/icon/hospital-building-3039718/.

Puranam, K., Novak, D. C., Lucas, M. T., & Fung, M. (2017). Managing blood inventory with multiple independent sources of supply. *European Journal of Operational Research*, 259 (2), 500–511.

Qi, L., Shen, Z.-J. M., & Snyder, L. V. (2009). A continuous-review inventory model with disruptions at both supplier and retailer. *Production and Operations Management*, 18 (5), 516–532.

Rajendran, S. & Ravindran, R. A. (2019). Inventory management of platelets along blood supply chain to minimize wastage and shortage. *Computers & Industrial Engineering*, 130, 714–730.

Rajendran, S. & Srinivas, S. (2020). Hybrid ordering policies for platelet inventory management under demand uncertainty. *IISE Transactions on Healthcare Systems Engineering*, 10 (2), 113–126.

Rana, R. S., Kumar, D., Mor, R. S., & Prasad, K. (2021). Modelling the impact of demand disruptions on two warehouse perishable inventory policy amid COVID-19 lockdown. *International Journal of Logistics Research and Applications*, 1–24.

Ross, A. M., Rong, Y., & Snyder, L. V. (2008). Supply disruptions with time-dependent parameters. *Computers & Operations Research*, 35 (11), 3504–3529.

Rozhkov, M., Ivanov, D., Blackhurst, J., & Nair, A. (2022). Adapting supply chain operations in anticipation of and during the COVID-19 pandemic. *Omega*, 110, 1–17.

Ruel, S., El Baz, J., Ivanov, D., & Das, A. (2021). Supply chain viability: conceptualization, measurement, and nomological validation. *Annals of Operations Research*, 1–30.

Saedi, S., Kundakcioglu, O. E., & Henry, A. C. (2016). Mitigating the impact of drug shortages for a healthcare facility: An inventory management approach. *European Journal of Operational Research*, 251 (1), 107–123.

Saha, E. & Ray, P. K. (2019). Patient condition-based medicine inventory management in healthcare systems. *IISE Transactions on Healthcare Systems Engineering*, 9 (3), 299–312.

Sahin, I. (1982). On the objective function behavior in (s, S) inventory models. *Operations Research*, 30 (4), 709–724.

Saithong, C. & Lekhavat, S. (2020). Derivation of closed-form expression for optimal base stock level considering partial backorder, deterministic demand, and stochastic supply disruption. *Cogent Engineering*, 7 (1), 1–27.

Salehi, H., Taleizadeh, A. A., & Tavakkoli-Moghaddam, R. (2016). An EOQ model with random disruption and partial backordering. *International Journal of Production Research*, 54 (9), 2600–2609.

Schmitt, A. J., Snyder, L. V., & Shen, Z.-J. M. (2010). Inventory systems with stochastic demand and supply: Properties and approximations. *European Journal of Operational Research*, 206 (2), 313–328.

Schmitt, T. G., Kumar, S., Stecke, K. E., Glover, F. W., & Ehlen, M. A. (2017). Mitigating disruptions in a multi-echelon supply chain using adaptive ordering. *Omega*, 68, 185–198.

Sericola, B. (2013). *Markov Chains: Theory, Algorithms and Applications*. Applied stochastic methods series. London, UK & Hoboken, NJ, USA: ISTE; Wiley.

Shen, B. & Li, Q. (2017). Market disruptions in supply chains: a review of operational models. *International Transactions in Operational Research*, 24 (4), 697–711.

Shih, H. & Rajendran, S. (2020). Stochastic inventory model for minimizing blood shortage and outdating in a blood supply chain under supply and demand uncertainty. *Journal of Healthcare Engineering*, 2020, 1–14.

Shokouhifar, M., Sabbaghi, M. M., & Pilevari, N. (2021). Inventory management in blood supply chain considering fuzzy supply/demand uncertainties and lateral transshipment. *Transfusion and Apheresis Science*, 60 (3), 1–8.

Skouri, K., Konstantaras, I., Lagodimos, A., & Papachristos, S. (2014). An EOQ model with backorders and rejection of defective supply batches. *International Journal of Production Economics*, 155, 148–154.

Snyder, L. V. & Shen, Z.-J. M. (2019). *Fundamentals of Supply Chain Theory*. 2nd ed. Hoboken, NJ: Wiley.

Snyder, L. V. (2014). A tight approximation for an EOQ model with supply disruptions. *International Journal of Production Economics*, 155, 91–108.

Snyder, L. V., Atan, Z., Peng, P., Rong, Y., Schmitt, A. J., & Sinsoysal, B. (2016). OR/MS models for supply chain disruptions: a review. *IIE Transactions*, 48 (2), 89–109.

Syawal, H. A. & Alfares, H. K. (2020). Inventory optimization for multiple perishable products with dynamic pricing, dependent stochastic demand, and dynamic reorder policy. *2020 Industrial and Systems Engineering Conference (ISEC)*, 1–5.

Tai, P. D., Huyen, P. P. N., & Buddhakulsomsiri, J. (2021). A novel modeling approach for a capacitated (S,T) inventory system with backlog under stochastic discrete demand and lead time. *International Journal of Industrial Engineering Computations*, 12 (1), 1–14.

Taleizadeh, A. A. (2017). Lot-sizing model with advance payment pricing and disruption in supply under planned partial backordering. *International Transactions in Operational Research*, 24 (4), 783–800.

Taleizadeh, A. A., Tafakkori, K., & Thaichon, P. (2021). Resilience toward supply disruptions: A stochastic inventory control model with partial backordering under the base stock policy. *Journal of Retailing and Consumer Services*, 58, 1–19.

Tucker, E. (2020). Modeling pharmaceutical supply chains to mitigate drug shortages. *[Doctoral Dissertation] The Department of Industrial and Operations Engineering: University of Michigan*, 1–154.

Tucker, E. L., Cao, Y., Fox, E. R., & Sweet, B. V. (2020a). The drug shortage era: a scoping review of the literature 2001–2019. *Clinical Pharmacology & Therapeutics*, 108 (6), 1150–1155.

Tucker, E. L., Daskin, M. S., Sweet, B. V., & Hopp, W. J. (2020b). Incentivizing resilient supply chain design to prevent drug shortages: policy analysis using two- and multi-stage stochastic programs. *IISE Transactions*, 52 (4), 394–412.

Uthayakumar, R. & Karuppasamy, S. K. (2019). An EOQ model for deteriorating items with different types of time-varying demand in healthcare industries. *The Journal of Analysis*, 27, 3–18.

Vicil, O. & Jackson, P. (2016). Computationally efficient optimization of stock pooling and allocation levels for two-demand-classes under general lead time distributions. *IIE Transactions*, 48 (10), 955–974.

Vizient (2019). Drug shortages and labor cost. *Vizient Drug Shortages Impact Report, June 2019*.

Wang, K.-M. & Ma, Z.-J. (2015). Age-based policy for blood transshipment during blood shortage. *Transportation Research Part E: Logistics and Transportation Review*, 80, 166–183.

Wei, Y., Yang, M., Chen, J., Liang, L., & Ding, T. (2022). Dynamic lateral transshipment policy of perishable foods with replenishment and recycling. *Computers & Industrial Engineering*, 172, 1–12.

Xu, Y. & Szmerekovsky, J. (2022). A multi-product multi-period stochastic model for a blood supply chain considering blood substitution and demand uncertainty. *Health Care Management Science*, 25 (3), 441–459.

Yao, M. & Minner, S. (2017). Review of multi-supplier inventory models in supply chain management: An update. *SSRN Electronic Journal*, 1–70.

Zhang, C., Ayer, T., White, C. C., Bodeker, J. N., & Roback, J. D. (2022). Inventory sharing for perishable products: Application to platelet inventory management in hospital blood banks. *Operations Research*, 1–21.

Zhang, X., Meiser, D., Liu, Y., Bonner, B., & Lin, L. (2014). Kroger uses simulation-optimization to improve pharmacy inventory management. *Interfaces*, 44 (1), 70–84.

Zheng, Y.-S. & Federgruen, A. (1991). Finding optimal (s, S) policies is about as simple as evaluating a single policy. *Operations Research*, 39 (4), 654–665.

Zhou, Y., Zou, T., Liu, C., Yu, H., Chen, L., & Su, J. (2021). Blood supply chain operation considering lifetime and transshipment under uncertain environment. *Applied Soft Computing*, 106, 1–16.