

THE UNIVERSITY OF CHICAGO

INVENTORY STRATEGIES AND ONLINE ORDER FULFILLMENT IN A
MULTI-TIER NETWORK

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE UNIVERSITY OF CHICAGO
BOOTH SCHOOL OF BUSINESS
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

BY
YANYANG ZHAO

CHICAGO, ILLINOIS
DECEMBER 2022

Copyright © 2022 by Yanyang Zhao

All Rights Reserved

TABLE OF CONTENTS

LIST OF FIGURES	vi
LIST OF TABLES	vii
ACKNOWLEDGMENTS	viii
ABSTRACT	xii
1 INTRODUCTION	1
1.1 Related Papers	2
2 MANAGING MULTI-TIER INVENTORY NETWORKS WITH EXPEDITING UNDER NORMAL AND DISRUPTED MODES	4
2.1 Introduction	4
2.1.1 Literature review	7
2.1.2 Outline of paper	10
2.2 Distribution Network Model	10
2.2.1 Two-Tier Distribution Network	10
2.2.2 Demand at the Retailers	11
2.2.3 Evolution Equations	11
2.2.4 Cost Parameters	13
2.3 Normal Mode Analysis	13
2.4 Disrupted Mode Analysis	18
2.5 Assessing Policies for Normal and Disrupted Modes	24
2.5.1 Introducing industrial partner's current practice.	24
2.5.2 Normal mode of operations	27
2.5.3 Disrupted mode of operations	35
2.6 Conclusion and Discussion	42
2.7 Technical Proofs	45
2.7.1 Proof of Lemma 1	45
2.7.2 Proof of Lemma 2	49
2.7.3 Proof of Theorem 1	52
2.7.4 Proof of Lemma 3	52
2.7.5 Proof of Lemma 4	55
2.7.6 Proof of Lemma 5	62
2.7.7 Proof of Lemma 6	64
2.7.8 Proof of Lemma 9	65
2.7.9 Proof of Lemma 10	66
2.7.10 Proof of Lemma 11	66
2.7.11 Proof of Lemma 12	67
2.7.12 Proof of Lemma 13	68
2.7.13 Proof of Theorem 3	72

2.7.14	Other proofs	81
2.8	More on numerical experiments	82
2.8.1	Parameter estimations and network descriptions	82
3	MULTI-ITEM ONLINE ORDER FULFILLMENT IN A TWO-LAYER NETWORK	92
3.1	Introduction	92
3.1.1	Literature review	96
3.1.2	Overview of model and contributions	99
3.1.3	Outline of paper	102
3.2	Model	102
3.2.1	Myopic policy	105
3.2.2	Main results	108
3.3	Discussion of the analysis	111
3.3.1	Proof sketch of Theorems 4 and 6	111
3.3.2	High-level proof discussion of Theorems 7 and 8	114
3.4	Implications of the main results & additional results	114
3.4.1	Implications under Assumption (A)	115
3.4.2	Implications under Assumptions (B) and (C)	115
3.4.3	Additional results: a practical upper bound with estimations under Assumption (C)	116
3.4.4	Summary	118
3.5	Extension: When Demand Forecasts Are Available	120
3.6	Numerical Experiments	126
3.6.1	Comparing Myopic with LP Rounding	126
3.6.2	Delay in order fulfillment	131
3.6.3	Managerial insights	132
3.7	Conclusion and discussion	133
3.8	Appendix	134
3.8.1	Fixed-cost estimates for Standard Overnight	134
3.8.2	A brief description of the open dataset	136
3.8.3	Performance of other greedy-type policies	137
3.8.4	Proof of Theorem 4	138
3.8.5	Proof of Theorem 5	144
3.8.6	Proof of Theorem 6	147
3.8.7	Proof of Theorem 7	150
3.8.8	Proof of Theorem 8	154
3.8.9	Proof of Theorem 9	161
3.8.10	Proof of Theorem 10	164
3.8.11	Proof of Theorem 11	165
3.8.12	Other proofs	167
3.8.13	Other numerical experiments	170
3.8.14	Summary of instances	177

REFERENCES 180

LIST OF FIGURES

2.1	Ratio of the improved base-stock-level policy over industrial partner's current policy	34
2.2	Cost ratios between policies with normal distributions	39
2.3	Cost ratios between policies with Weibull distributions	42
3.1	Percentages of items fulfilled by the local designated DC in March 2018	95
3.2	Performance w.r.t. time horizon T , holding $n = 10, \overline{ o } = 5, \underline{ o } = 5, \theta = 0.8, f_F = 50, f_R = 100, c_F = 5, c_R = 0$. Left panel: T from 10 to 200; right panel: T from 356 to 11, 246.	127
3.3	Performance w.r.t. RDC fixed cost f_R , holding $\overline{ o } = 5, \underline{ o } = 5, \theta = 0.8, f_F = 40, c_F = 5, c_R = 0$. Left panels: $T = 100$; right panels: $T = 2,000$	128
3.4	Performance w.r.t. RDC fixed cost f_R , holding $\overline{ o } = 5, \underline{ o } = 5, \theta = 0.8, f_F = 40, c_F = 5, c_R = 0$. Left panels: $T = 100$; right panels: $T = 2,000$	129
3.5	Performance w.r.t. FDC inventory scaling parameter θ , holding $n = 20, \overline{ o } = 5, \underline{ o } = 5, f_F = 50, f_R = 100, c_F = 5, c_R = 0$. Left panel: $T = 100$; right panel: $T = 8,000$	131
3.6	Performance w.r.t. order fulfillment delay T^d , holding $T = 100, n = 10, \overline{ o } = 5, \underline{ o } = 5, f_F = 50, f_R = 100, c_F = 5, c_R = 0$	132
3.7	Performance w.r.t. time horizon T , holding $n = 10, \overline{ o } = 5, \underline{ o } = 5, \theta = 0.8, f_F = 100, f_R = 80, c_F = 0, c_R = 50$. Left panel: T from 10 to 200; right panel: T from 356 to 11, 246.	172
3.8	Performance w.r.t. RDC fixed cost f_R , holding $\overline{ o } = 5, \underline{ o } = 5, \theta = 0.8, f_F = 100, c_F = 0, c_R = 50$. Left panels: $T = 100$; right panels: $T = 2,000$	173
3.9	Performance w.r.t. RDC fixed cost f_R , holding $\overline{ o } = 5, \underline{ o } = 5, \theta = 0.8, f_F = 100, c_F = 0, c_R = 50$. Left panels: $T = 100$; right panels: $T = 2,000$	174
3.10	Performance w.r.t. FDC inventory scaling parameter θ , holding $n = 10, \overline{ o } = 5, \underline{ o } = 5, f_F = 100, f_R = 80, c_F = 0, c_R = 50$. Left panel: $T = 100$; right panel: $T = 2,000$	174
3.11	Performance w.r.t. order fulfillment delay T^d , holding $T = 100, n = 10, \overline{ o } = 5, \underline{ o } = 5, f_F = 100, f_R = 80, c_F = 0, c_R = 50$	175
3.12	Performance w.r.t. FDC fixed cost f_F , holding $T = 100, n = 20, \overline{ o } = 8, \underline{ o } = 5, \theta = 0.5, f_R = 4, c_F = 1, c_R = 15$	176

LIST OF TABLES

2.1	Evaluating $\frac{\lambda-a-1}{2} \frac{\lambda^{a-1}}{(a-1)!e^\lambda}$ for small a	71
2.2	Unit labor cost in different categories	87
3.1	Summary statistics of three highest-demand districts in JD.com’s open dataset .	94
3.2	Alibaba’s data snapshot of a category of fast-moving products in an eastern city of China in 2019	95
3.3	A comparison of fulfillment papers explicitly [◇] dealing with multi-item orders . .	99
3.4	Summary of parameter-dependent bounds. The second and third rows summarize the upper and lower bounds on $\text{RATIO}(\text{Myopic})$, respectively; the last row summarizes the lower bounds on the ratio of any online algorithm (including deterministic and randomized).	109
3.5	Fixed-cost estimates: Ground shipping	118
3.6	Summary of implications/results in Section 3.4.	119
3.7	FedEx Ground Shipping Rates	135
3.8	FedEx Standard Overnight Shipping Rates	135
3.9	Fixed-Cost Estimates: Standard Overnight	136
3.10	A snapshot of the Myopic and Threshold policy parameters under Assumption (A)	176
3.11	Proof of Theorem 7, Assumption (A).	177
3.12	Proof of Theorem 7, Assumption (A).	178
3.13	Proof of Theorem 7, Assumption (B).	178
3.14	Proof of Theorem 7, Assumption (C).	178
3.15	Proof of Theorem 8, Assumption (A).	179
3.16	Proof of Theorem 8, Assumption (A).	179
3.17	Proofs of Theorem 8, Assumption (B) and Theorem 10.	179
3.18	Proof of Theorem 8, Assumption (C).	179

ACKNOWLEDGMENTS

First, I must sincerely thank my advisors Linwei Xin, John Birge, and Levi DeValve. The PhD journey has never been an easy one for me. I had a difficult time in the third year of my PhD, trying to figure out what interested me the most and what I was capable of accomplishing. As I approached the end of my curriculum paper, I was insecure, anxious, and struggling. When I went to Linwei and John for new projects in December 2019, they did not hesitate and offered me several opportunities. I am profoundly grateful for their kindness and openness.

I had a smooth and rewarding experience when working with Linwei. We managed to have the first draft of our paper available on SSRN only six months later. Although we encountered many difficulties in the writing of the paper and in its later revision process, Linwei has always been supportive and provided me with detailed guidance throughout our project. I was never ashamed to ask any types of questions in our weekly meetings, and his patience, enthusiasm, and insightful comments motivated me throughout our research.

I would like to thank John again for introducing me to Levi. I am deeply indebted to them for their constant support and encouragement over the last three years. Discussions with John and Levi have always been intellectually inspiring. In addition to research, I was the TA for Levi's Supply Chain Management class for two quarters. I truly believe Levi has spent more time discussing research, class materials, career plans, and life in general, with me than with any other student he has taught. I not only acquired various mathematical techniques from Levi, but also benefited from his critical thinking, constructive writing, and communication skills. Most notably, he modeled the importance of a well-balanced role between being a professional researcher and a good husband and father.

My advisors are my biggest assets at Chicago Booth. I feel fortunate to work with such intellectuals. Additionally, I would like to thank my collaborators, Xinshang Wang and Robert Inman. They both provided me with different perspectives on research problems.

They offered me tremendous help throughout my projects and demonstrated their expertise in both technical work and industrial knowledge.

Next, I would like to thank all the Operations Management faculty who taught me in the core OM courses during my PhD journey: Baris Ata, John Birge, Rene Caldentey, Levi DeValve, Kipp Martin, Christopher Ryan, Yuan Zhong, and others. Without them, I would not have gained the OR skillsets necessary for research. I sincerely appreciate their time, guidance, and valuable advice during my graduate studies. Additionally, my PhD journey would not have been the same without the friends I met and made at Booth. Special thanks go to Amir Alwan, Cagla Keceli, Antonis Tsiflis, and Yueyang Zhong. I will always remember all the days and nights I spent with them at school and outside of school. Each of them has given me valuable feedback, lessons, and encouragement along the way. I treasure my friendship with them and wish the best for each of them.

I would also like to take this opportunity to thank the professors at Georgia Tech who inspired me and offered their help to get me into the PhD program: Professors Christos Alexopoulos, Craig Tovey, Valerie Thomas, Santanu Dey, George Nemhauser, and Chen Zhou. I greatly appreciate their teaching and mentorship, and I will always remember that their expertise, knowledge, kindness, and openness were what made my undergraduate experience so fruitful and rewarding.

I was never the top dog in school when I was a kid in China. I studied at a really good middle school, but my ranking among the 60-student class was always mediocre, regardless of how hard I was trying. I was losing myself. I did not know how to differentiate myself from other students, and, unfortunately, I did not find my worth in class. However, my life took a 180-degree turn when I met my beloved high school teachers at Saint James School in Maryland. I still remember that after the first marking period, Mrs. Becker, teaching me pre-calculus at that time, wrote to my parents that I was the "brightest student" she had ever taught. Throughout high school, I felt appreciated not only by Mrs. Becker, but

also many other teachers, including Ms. Cerruti, Ms. Clyburn, Ms. Pollock, Mr. Hobbs, Mr. Collin, Mr. Schwab, Mr. Silver, and Mr. Yergey. I would like to give a big thanks to all of them. They helped me through numerous difficulties in academics, sports, and life and helped me regain my confidence. Their words, support, and backing will stay with me forever, and I am constantly reminded of how strong a person I have become because of them.

Last but not least, I would like to extend my thanks to my best friend Yu Zhang and my family. Yu and I met in middle school in China, and we have stayed best friends supporting each other for the last 15 years. Quite frankly, Yu was the person who inspired me to enter a PhD program, while he was deciding between going to a medical school or becoming a biomedical engineering PhD. We share countless memories thanks to the past many years, and the support and love we have provided each other are impossible to describe. I regard him as family and wish him the best of luck in life and work, as he receives his doctorate degree from Johns Hopkins in the near future.

I love my parents; they have unconditionally rooted for me and supported all of my decisions since I was young. Sending a child 8,000 miles away from home for ninth grade is not an ordinary or easy decision for any parent, but my mother and father had no hesitation and fully supported me when I told them one day after school in China that I would like to go to high school in the United States. Now, when I look back, I understand it was a bold decision, but my parents, with endless love, made my dream come true. I am grateful that they have always been proud of me and have always trusted my judgment. I also would like to thank my American parents, Mr. David Levy and Dr. Margaret Levy, who took me in when I first came to the states 12 years ago and have treated me like their own son. I will always remember their sweetness, kindness, wisdom, and humor. I am fortunate enough that I can forever be a child and a student to them and keep learning new life lessons from them. My final and most important thank you goes to my wife, Lisa, who has always been

on my side, offering me constant support in tough times during my PhD studies. Lisa is the most caring and loving person I have ever met, and she offers me another perspective on the meaning of life and shows me the diversity of the world in addition to economics principles. I am grateful for all the time we have spent together supporting each other and all the sacrifices she has made for me in the past five years. I cannot wait to start a new chapter in our lives and start our new journey together in Atlanta next year.

ABSTRACT

Problems originated from multi-tier networks are central to the field of OM/OR. Over the years, multi-tier networks have only gotten more complex. On one hand, companies today are building more local distribution centers and opening up more retail shops to serve regional demand. On the other hand, companies are competing in goods delivery time to take customer service to the next level. In this dissertation, we study inventory and fulfillment policies in a two-tier network, where the upper tier consists of one central warehouse or regional distribution center (RDC), and the bottom tier consists of multiple retailers or front distribution centers (FDCs).

Classic literature on multi-echelon inventory policy assumes a steady state in which a manufacturer is always able to place an order and receive it within a reasonable time frame. In Chapter 2, we first assume the same steady state, and we consider the problem of minimizing the long-run cost of a two-tier network with multiple retailers and with expediting. The features of multi-location and expediting are omnipresent and critical to supply-chain networks in practice. Furthermore, due to the recent pandemic, we later drop the steady-state assumption in Chapter 2 and study the problem of allocating limited inventory across the two-tier network, when the manufacturer is unable to receive external supplies.

Chapter 3 is concerned solely with fulfilling orders in a two-tier network. We assume no inventory replenishment can happen during a fulfillment period, we allow orders to consist of multiple items, and we allow orders to be split into multiple packages for fulfillment. Because an order may contain more than one item, the decision-maker needs to efficiently decide what distribution centers to use to fulfill what part of an order. Chapter 3 studies a widely-implemented myopic policy in such a setting and evaluates the policy's performance in competitive analysis. Chapter 3 also studies a linear program rounding policy and a delay policy, theoretically and numerically.

CHAPTER 1

INTRODUCTION

Two concrete problems are studied in this dissertation.

Chapter 2. In this paper, we design inventory policies for managing a multi-tier/location network with expediting that are effective at minimizing costs and are implementable in practice. The paper also addresses a key practical challenge faced by many supply chains of the possibility of a disruption (from a variety of sources), which typically limits the inventory available in the network. Thus, we consider two “modes” of operation: the *normal mode*, with regular replenishment from an outside supplier, and the *disrupted mode*, where an unexpected disruption limits the available supply in the network. In normal mode, we study the problem of minimizing the long-run holding, backlog, and expediting costs and demonstrate the optimality of a base-stock policy in a *one warehouse multiple retailer* (OWMR) setting with expediting. A key step in our theoretical analysis is the development of a stochastic program lower bound on the optimal cost, which also provides the practical benefit of being adaptable to many additional problem features (which we explore further through simulation). In disrupted mode, we consider the problem of whether to centralize (i.e., keep all inventory at the central warehouse) or decentralize (i.e., keep all inventory at the retailers) the limited inventory, and provide a simple cost criterion to determine when decentralization is preferred. This result is useful practically because it demonstrates that while centralization may appeal intuitively to management’s desire to control limited inventory during a disruption, it may not always be cost effective relative to decentralization; and our criterion provides a practical rule of thumb for comparing the two strategies. Our policy development and analysis are conducted in partnership with an industrial collaborator who operates a nationwide automotive service parts distribution network in the US. We validate our policies by adapting our stochastic programming formulation to accommodate several additional features present in our partner’s setting (including non-stationary demand, multiple demand

classes, and stochastic lead times), and simulating our policies using our partner’s network, cost and demand data. We find that our optimized policies can provide a significant cost savings relative to our collaborator’s current practice.

Chapter 3. In this paper, we study the problem of minimizing fulfillment costs, in which an e-retailer must decide in real-time which warehouse(s) will fulfill each order, subject to warehouses’ inventory constraints. The e-retailer can split an order at an additional cost and fulfill it from different warehouses. We focus on an RDC-FDC distribution network that major e-retailers have implemented in practice. In such a network, the upper layer contains larger regional distribution centers (RDCs) and the lower layer contains smaller front distribution centers (FDCs). We analyze the performance of a simple myopic policy, which chooses the least expensive fulfillment option for each order without considering the impact on future orders. Perhaps surprisingly, myopic policies (and their variations) are standard in the implementation of a typical order-management system. We provide theoretical bounds on the performance ratio of the myopic policy compared with an optimal clairvoyant policy. We also empirically estimate our upper bound on the ratio by using FedEx shipping rates and demonstrate the bound can be as low as 1.13 for reasonable scenarios in practice. Moreover, we extend our study to the setting in which demand forecasting is available and prove the asymptotic optimality of a linear program rounding policy. Finally, we complement our theoretical results with a numerical study.

1.1 Related Papers

The material presented in this dissertation is based on the following papers.

Chapter 2. Yanyang Zhao, John R. Birge, Levi DeValve, Robert R. Inman. *Managing Multi-Tier Inventory Networks With Expediting Under Normal and Disrupted Modes*. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4204008.

Chapter 3. Yanyang Zhao, Xinshang Wang, and Linwei Xin. *Multi-Item Online Order*

Fulfillment in a Two-Layer Network. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3675117.

CHAPTER 2

MANAGING MULTI-TIER INVENTORY NETWORKS WITH EXPEDITING UNDER NORMAL AND DISRUPTED MODES

2.1 Introduction

To effectively make products available to customers, supply chains often operate in multiple tiers (or “echelons”) as well as in multiple locations. This allows the supply chain to take advantage of both the operational efficiency of various tiers, as well as the responsiveness of locations closer to customers. For example, Seven-Eleven Japan supports large clusters of traditional brick and mortar retail stores through a centralized distribution center (Chopra, 2017), while Chinese e-commerce giant JD.com operates a similar network consisting of several front distribution centers supported through a larger regional distribution center (De-Valve et al., 2021); in both cases a central facility efficiently handles large product volumes, while a set of dispersed facilities offer quick response times to customers. This ubiquitous supply chain setting is often called the *one warehouse multiple retailer* (OWMR) problem in the operations management literature (e.g. Federgruen and Zipkin, 1984b; Roundy, 1985), and is the focus of this paper.

In particular, we analyze an OWMR supply chain operated by our industrial collaborator, a large US automotive manufacturer, for distributing service (i.e., repair) parts used in their automobiles. Operating this supply chain presents several practical challenges relative to existing OWMR research, including expediting from the central warehouse, multiple demand classes, stochastic lead times, and stochastic, non-stationary demand. In addition to these challenges, our industrial partner also faces the risk of disruption to its operations from various major identifiable causes. In just the past few years, these have included natural disasters, labor disruptions, and pandemics (among others), which effectively limit the supply available to operate the system. These disruptions highlight a distinction between “common

cause" and "special cause" variation (following the terminology of Deming, 1975) that is important in practice for our industrial partner's operations. Our partner defines common cause variation as the uncertainty due to "common" causes that are in some sense predictable (e.g., forecasting the demand or lead time distribution) vs. special cause variation, which is the uncertainty due to "special" or unforeseen events that are fundamentally difficult to predict (such as the recent Covid-19 pandemic and related supply disruptions).

Within this framework, our industrial collaborator's goal is to design effective inventory management policies that plan primarily for common cause variation, but that can also adapt to special cause disruptions. Therefore, we consider two "modes" of operation: the *normal mode*, which is subject to uncertainty arising from common cause variation, and the *disrupted mode*, where a special cause disruption limits the available supply. In both modes, we propose effective inventory policies, which we validate through both theoretical analysis and simulation with our partner's data. Moreover, while our model and analysis are motivated by a collaboration with a particular automaker, the policies and insights are more broadly applicable to general OWMR settings, and so we present much of the model in more generic terms. We next highlight our key contributions.

Normal Mode. When the system operates in normal mode, our goal is to identify an inventory replenishment and fulfillment policy that minimizes the long-term average holding, backlog, and expediting costs. This leads to a challenging dynamic programming problem in the general case that is intractable due to the curse of dimensionality. However, in the special case with negligible lead times, we are able to establish the optimality of a base-stock policy (i.e., a policy that returns inventory at each location to a fixed level each period) through a comparison with a lower bound on optimal cost defined by a stochastic program (SP). This is interesting theoretically, as it is the first result establishing optimality of a base-stock policy in a OWMR setting with expediting. It is also relevant to our industrial partner's practice for two reasons. First, while the automaker currently stocks inventory primarily at

the retailers during normal mode, our analysis reveals that allowing the central warehouse to hold some inventory is helpful, especially with the expediting option. Second, our stochastic programming formulation offers an improved method for setting base-stock levels across the network, and is adaptable to the challenging practical features of our partner's true problem (including non-negligible and stochastic lead times). We demonstrate this adaptability in our numerical simulations with our partners' data, which we discuss further below.

Disrupted Mode. When supply is disrupted, our goal is to use the limited inventory available in the system to minimize the backlog and expediting costs over the duration of the disruption. Currently, our industrial partner takes centralized control over all inventory at the central warehouse during a disruption, a strategy we call "centralization." We compare this policy with the other extreme of "decentralization" or keeping all inventory at the retailers. We derive a simple criterion for comparing the backlog and expediting costs to determine whether decentralization will perform better than centralization. Intuitively, the condition favors decentralization when the expediting cost is high, but it also provides the insight that as either supply or demand grow large, decentralization also becomes a better option. Thus, our result suggests that centralization is not always preferred in disrupted mode, and also provides conditions for deciding when to keep inventory decentralized, which is useful for our industrial partner. Moreover, in our disrupted-mode analysis, we derive a new concentration bound on the sum of Poisson random variables that requires a novel analysis of the incomplete gamma function, which may be of independent interest.

Validation via Data Driven Simulation. Finally, we test our proposed policies on data from our industrial collaborator's service parts distribution network. As mentioned above, our partner faces several challenges in its operations that go beyond the scope of the theoretical models considered (including non-stationary demand, stochastic lead times, etc.). However, we demonstrate via simulation that our analysis leads to a methodological approach that can adapt to these additional features. First, in normal mode, we adapt

our stochastic program to accommodate non-stationary demand and stochastic lead times, and use it to simultaneously optimize the base-stock levels at all retailers and the central warehouse. Across approximately two dozen parts, this provides about a 5% cost improvement on average over our partner’s current base-stock levels. Second, in disrupted mode, we adapt our analysis comparing the extremes of complete centralization (i.e., all inventory at the central warehouse) vs. complete decentralization (i.e., all inventory at the retailers) to allow for an intermediate balance of some centralization and some decentralization, again via the solution of an appropriate stochastic program. We find that when there are few disruptions, our stochastic programming based policy outperforms our industrial partner’s current practice of complete centralization, and that when disruptions occur more frequently, the two policies perform similarly. Thus, our simulations validate our stochastic programming approach to designing effective inventory policies in both normal and disrupted mode, and suggest that these policies provide cost savings in practice.

2.1.1 Literature review

Aligned with our motivation, our work is generally related to two streams of literature: analysis of replenishment and allocation policies in multi-tier/location systems (normal mode), and analysis of policies for allocating limited resources over a finite time horizon (disrupted mode).

The normal mode we study in our work is first related to traditional periodic-review, multi-tier, multi-location inventory allocation problems. In their seminal paper, Clark and Scarf (1960) use a dynamic programming approximation to solve a two-tier OWMR distribution system problem under periodic review. Numerous papers since then have analyzed the OWMR problem, including Eppen and Schrage (1981), Federgruen and Zipkin (1984a), and Federgruen and Zipkin (1984c). Dođru et al. (2009) survey the papers that studied the OWMR model. Our paper contributes to this literature by incorporating the expediting

feature, which we discuss next.

Studies which explicitly consider expediting in an OWMR model include Moinzadeh and Aggarwal (1997) and Drent and Arts (2021), which both assume use of an augmented base-stock policy with thresholds for expediting, then propose various methods to optimize the policy within this class. While our model differs in some details (e.g., only the retailers can place expedited orders), our work builds on this literature by demonstrating that a base-stock policy can indeed be optimal in an OWMR setting with expediting. Other work studying expediting in various inventory systems includes Moinzadeh and Schmidt (1991), who consider the single location version of Moinzadeh and Aggarwal (1997), and Lawson and Porteus (2000), who consider the serial system of Clark and Scarf (1960) with an expediting option. Lawson and Porteus (2000) show that the optimal solution can be obtained by recursively solving nested single-dimensional convex optimization problems. However, computing such a solution can be tedious in practice. Muharremoglu and Tsitsiklis (2008) consider supermodular expediting costs, and Mamani and Moinzadeh (2014) consider a continuous-review version of Lawson and Porteus (2000). Shen et al. (2022) study a multi-tier model with finite horizon, expediting, and service time target (STT). They assume lead time of zero between stages and each stage consisting of one location and show a base-stock policy and a rationing policy are optimal in their setting. In general, this literature demonstrates that computing the optimal solution for the multi-tier model with non-negligible lead-times, multi-locations, and expediting is nontrivial. Our work identifies an efficient heuristic that can be implemented in practice and can solve large problems with expediting and multiple locations in a single tier.

The problem of allocating limited resources over a finite horizon in our disrupted mode is related to capacity control in the network revenue management (NRM) literature. In particular, our model in the disrupted mode allocates limited resources without replenishment. We refer the interested reader to Talluri and Van Ryzin (2004) for a comprehensive review

of the larger revenue-management literature. Our work differs from the NRM literature due to features such as a cost-driven objective (as opposed to pricing and maximizing revenue), and multi-tier/location model with expediting. Another difference is that we decide where to allocate resources at the beginning of a selling horizon.

There is a growing body of literature on e-commerce fulfillment which is also relevant to our disrupted mode with limited resources (Xu et al. (2009), Acimovic and Graves (2015), Jasin and Sinha (2015)). Chen and Graves (2021) study a problem of choosing fulfillment centers in which to place items, and formulate the problem as a large-scale mixed-integer program modeling thousands of items to be placed in dozens of FCs and shipped to dozens of customer regions. DeValve et al. (2021) also consider the initial inventory allocation in a two-layer e-retail network and evaluate the benefits of flexibility when allowing cross fulfillment among adjacent front distribution centers. Govindarajan et al. (2021) study the problem of an omnichannel retailer who faces both online and in-store demand and has to decide how much inventory to reserve at each store and where to fulfill each online order from. They propose a simple heuristic for the multi-location problem and prove its asymptotic near-optimality for large number of omnichannel stores under certain conditions. For more details on this literature, we refer the interested reader to a recent tutorial on the fulfillment optimization problem by Acimovic and Farias (2019). Our work in the disrupted mode aims to provide useful practical guidance by deriving a simple cost criterion for when to decentralize the limited inventory (i.e., to keep all inventory at the retailers).

The disruption feature in our work is also related to match-up scheduling studies such as Bean et al. (1991) and Akturk and Gorgulu (1999), which consider heuristic approaches to find match-up times. Snyder et al. (2006), Peng et al. (2011), and Mak and Shen (2012) consider the optimal network structure under the possibility of a disruption. Simchi-Levi et al. (2014) and Simchi-Levi et al. (2015) develop a risk-exposure model that defers the need for a company to estimate the probability associated with a disruption risk until after

it has learned the effect such a disruption will have on its operations. The second half of our work is inspired by numerous disruptions that occurred during the recent global pandemic. We are particularly interested in how to allocate inventory from a central location to regional facilities in the event of upstream supply uncertainty caused by the pandemic. We implement a practically-used “designate-for-intervention” policy and discuss its performance under various lead time distributions.

2.1.2 Outline of paper

The rest of the paper is organized as follows. We describe our generic model in Section 2.2. In Section 2.3, we conduct a normal-mode analysis identifying an optimal inventory policy, assuming the system can be replenished regularly. In Section 2.4, the disrupted-mode analysis is conducted, and we derive a simple criterion that identifies when decentralization is preferred to centralization. We assess our policies via simulation in Section 2.5 with additional features such as non-stationary demand and stochastic lead times. Finally, we include a technical appendix in Section 2.7.

2.2 Distribution Network Model

We consider a firm who sells a product through a canonical two-tier distribution network. In this section we describe the general distribution network, product demand, evolution equations, and cost parameters we consider at a high level. In the following sections, we present additional relevant details (such as objective functions) for our specific analyses of the normal and disrupted modes of operation.

2.2.1 Two-Tier Distribution Network

The distribution network consists of two tiers. The upper tier is a warehouse which places orders for the product with an exogenous supplier, while the lower tier consists of n retailers

(indexed by i) which receive shipments from the warehouse and which fulfill exogenous demand from customers. Both the warehouse and the retailers may hold inventory of the product. Each retailer can use its inventory to fulfill demand from only those customers arriving at its location (i.e., no transshipment). The warehouse can use its inventory to either replenish the retailers' inventories, or to expedite a shipment to directly fulfill a customer's demand. We will refer to this latter form of delivery as "expediting." Expedited shipments are differentiated from normal replenishment shipments in the time they take to arrive. We will specify further details of the time dynamics when presenting the evolution equations below.

2.2.2 Demand at the Retailers

Demand for the product occurs only at the retailers (the warehouse receives no exogenous demand stream). Time is discrete and indexed by period t . Demand at retailer i in period t is denoted by the random variable $D_{i,t}$. In general we assume the random variables $D_{i,t}$ are independent across time t , but not necessarily independent across retailers i , nor identical across t or i (specific assumptions on the distribution of demand will be presented alongside the appropriate analyses below). Unmet demand is backlogged at each retailer.

2.2.3 Evolution Equations

We now describe how the system evolves from one period to the next. Let I_t and $X_{i,t}$ denote the warehouse's and retailer i 's inventory level at the end of period t , respectively. Similarly, let $B_{i,t}$ denote retailer i 's backlog level at the end of period t (backlogs are only accumulated at the retailers). We track these state variables at the end of each period because this is when associated holding and backlogging costs will be charged (described below). We assume that the system starts empty, so $I_0 = X_{i,0} = B_{i,0} = 0$ for all i .

To complete the state description, we must specify notation for the firm's replenishment

and fulfillment decisions in each period. Let x_t denote the warehouse's replenishment order to an exogenous supplier in period t , which arrives after a lead time of L periods. The lead time L can be either random or deterministic, and we will specify it in later sections. Let $z_{i,t}$ denote the warehouse's replenishment shipment to retailer i in period t , which arrives after a lead time of l_i periods. Let $y_{i,t}$ denote the warehouse's expedited shipment to retailer i in period t , which arrives after a lead time of l_i^e periods, and is assumed to directly fulfill demand at retailer i which it couldn't fulfill with its local inventory. Naturally, expedited fulfillment has short lead times, and so our analyses focus either on $l_i^e = 0$ or $l_i^e = 1$. Finally, let $w_{i,t}$ denote retailer i 's fulfillment of demand from its local inventory in period t , which occurs without a lead time. To be concrete on the timing of these decisions, we specify the following sequence of events in period t :

1. Receive supplier order x_{t-L} at warehouse.
2. Send shipment $z_{i,t}$ from warehouse to retailer i .
3. Realize demand $D_{i,t}$ at each retailer i .
4. Receive warehouse shipment $z_{i,t-l_i}$ at each retailer i .
5. Fulfill $w_{i,t}$ of demand at each retailer i .
6. Expedite $y_{i,t}$ from warehouse to each retailer i .
7. Receive expedited shipment $y_{i,t-l_i^e}$ at each retailer i .
8. Order x_t from supplier.

With these variables specified, we are now ready to characterize the system's evolution

equations. The warehouse and retailers' inventory levels evolve as follows:

$$I_t = I_{t-1} + x_{t-L} - \sum_i (y_{i,t} + z_{i,t}), \quad (2.1)$$

$$X_{i,t} = X_{i,t-1} + z_{i,t-l_i} - w_{i,t},$$

while retailer i 's backlog evolves as:

$$B_{i,t} = B_{i,t-1} + D_{i,t} - w_{i,t} - y_{i,t-l_i^e}. \quad (2.2)$$

To be feasible, the decisions x_t , $z_{i,t}$, $y_{i,t}$ and $w_{i,t}$ must ensure the state variables remain non-negative, i.e., $I_t \geq 0$, and $X_{i,t}, B_{i,t} \geq 0$ for i and t .

2.2.4 Cost Parameters

The firm incurs unit holding costs h_0 and h_i for inventory held at the warehouse and retailer i , respectively, at the end of each period. Likewise, the firm incurs a unit backlog cost of b_i for backlogged demand at retailer i at the end of each period. We assume that the cost of normal shipments, $z_{i,t}$, and normal fulfillment, $w_{i,t}$, are normalized to zero, while the unit cost of expedited fulfillment, $y_{i,t}$ is f_i . We assume the expediting cost to each retailer is larger than the warehouse's holding cost, i.e., $f_i \geq h_0$ for all i . The cost incurred in period t is then

$$h_0 I_t + \sum_i (h_i X_{i,t} + b_i B_{i,t} + f_i y_{i,t}).$$

2.3 Normal Mode Analysis

In this section, we consider a "normal mode" of operation, where the warehouse can regularly place replenishment orders with the supplier. Here we assume that demand is independent across time periods t , but may be correlated across retailers i within a given time period.

In this setting, since inventory can be replenished steadily over time, it is natural for the firm to consider a long run average cost objective. In particular, in normal mode, the firm minimizes the following objective

$$\limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \mathbb{E} \left[h_0 I_t + \sum_i (h_i X_{i,t} + b_i B_{t,i} + f_i y_{it}) \right]. \quad (2.3)$$

Minimizing 2.3 is well-known to be a challenging dynamic optimization problem, even in the case with no expediting (Federgruen and Zipkin (1984a), Federgruen and Zipkin (1984c), Zipkin (1984)), due to the large state space and the difficulty of allocating inventory across the retailers. Therefore, we do not attempt to minimize (3) exactly for the fully general model. Rather, we set out in this section to solve a simplified version of the problem which will provide two useful insights: i) we gain intuition for an effective class of policies to use as a heuristic, and ii) we gain a framework for how to design these policies in more general settings. For i), we show that a base-stock policy is optimal in a simplified model with negligible lead times. This is useful as base-stock policies are straightforward to implement and generally well understood by practitioners. Perhaps more importantly though, for ii) we show that an appropriately defined stochastic program is key to determining the base-stock levels in the simplified model. This is helpful as it provides intuition on how to generalize the stochastic program to set base-stock levels in more complex environments with non-negligible lead times.

We now consider a model where all lead times are negligible: the external supplier delivers x_t with a lead time of $L = 1$, the warehouse's normal shipments $z_{i,t}$ arrive on the day they are sent (so $l_i = 0$), and the warehouse's expedited shipments arrive immediately (so $l_i^e = 0$). We note that the expedited shipments maintain an advantage over normal shipments because they can be shipped after demand is realized, while the normal shipments are sent before demand is realized. These lead time assumptions suggest the optimality of a base-stock policy: because the supplier's shipment arrives in time to update each retailer's inventory

position at the start of each period, the warehouse is always able to return the system's inventory to the same state, which naturally specify the base-stock levels. Our analysis verifies this intuition in the presence of an expediting option, and also characterizes how to set the base-stock levels.

We now show that a base-stock policy, which keeps the inventory position at each retailer and the warehouse constant, is optimal for this system. Specifically, let $\mathbf{S} = (S_0, S_1, \dots, S_n)$ denote the base-stock levels for the warehouse and the n retailers respectively. Then, the warehouse ships the following quantity to retailer i at the beginning of day t :

$$z'_{i,t} = (S_i - X_{i,t-1} + B_{i,t-1})^+, \quad (2.4)$$

and the warehouse places the following order with the supplier at the end of day $t - 1$:

$$x'_{t-1} = (S_0 - I_{t-1})^+ + \sum_{i=1}^n z'_{i,t}, \quad (2.5)$$

so that the warehouse always orders enough to bring its own inventory position up to S_0 and each retailer's inventory position up to S_i . This completes a specification of the warehouse ordering and retailer shipping policy. We delay specifying the retailer fulfillment and warehouse expediting policy until we have clarified how to choose the base-stock levels, which we do next.

In the following, we use D_i to denote a random variable with the same distribution as $D_{i,t}$. To decide the base-stock levels, which we denote by $\mathbf{S} = (S_0, S_1, \dots, S_n)$, we solve the

following stochastic program:

$$\begin{aligned}
& \min_{\mathbf{S} \geq 0} h_0 S_0 + \mathbb{E}[g(\mathbf{S}; \mathbf{D})] + \sum_{i=1}^n (h_i S_i + b_i \mathbb{E}[D_i]) \\
& \text{where } g(\mathbf{S}; \mathbf{D}) = \min \sum_{i=1}^n (f_i - b_i - h_0) y_i - (h_i + b_i) w_i \\
& \text{s.t. } \sum_{i=1}^n y_i \leq S_0, \\
& w_i \leq S_i, \quad \forall 1 \leq i \leq n, \\
& w_i + y_i \leq D_i, \quad \forall 1 \leq i \leq n \\
& y_i, w_i \geq 0, \quad \forall 1 \leq i \leq n.
\end{aligned} \tag{2.6}$$

The stochastic program (2.6) has a natural interpretation as minimizing the expected cost incurred in a given period, if we were able to begin the period with the inventory vector \mathbf{S} and no backlog. The objective accounts for the warehouse's holding cost, $h_0(S_0 - \sum_i y_i)$, the retailers' holding cost, $\sum_i h_i(S_i - w_i)$, and backlog cost, $\sum_i b_i(D_i - w_i - y_i)$, and the expediting cost $\sum_i f_i y_i$. The constraints enforce that the warehouse's inventory, retailers' inventory, and retailers' backlog must all be non-negative at the end of the period.

Let C^* denote the optimal value of (2.6). The following lemma shows that C^* is a lower bound on the optimal value of (2.3) and follows from a relaxation argument similar to that of Dođru et al. (2010).

Lemma 1. *The optimal value of (2.6), C^* , is less than the value of (2.3) for any feasible policy.*

We are now ready to specify the retailer fulfillment and warehouse expediting policy. In each period t , we perform the fulfillment decisions, w_i and expediting decisions y_i specified by the optimal second stage solution of (2.6), i.e., the minimizing arguments used to compute $g(\mathbf{S}^*; \mathbf{D}_t)$. More specifically, in period t let $w_{i,t}^*$ and $y_{i,t}^*$ for $1 \leq i \leq n$ denote an optimal

solution to $g(\mathbf{S}^*; \mathbf{D}_t)$. Then, in period t fulfill and expedite according to the following:

$$w'_{i,t} = w^*_{i,t} + B_{i,t-1}, \quad (2.7)$$

$$y'_{i,t} = y^*_{i,t}. \quad (2.8)$$

Denote by π the policy with base-stock levels \mathbf{S}^* , ordering policy (2.5), shipping policy (2.4), fulfillment policy (2.7), and expediting policy (2.8). Our goal is to show that the base-stock policy π incurs cost equal to the lower bound C^* , and hence is optimal. To do so, we first show the policy is feasible in the next lemma.

Lemma 2. *The base-stock policy π is feasible in each period t . Furthermore, the shipping and expediting decisions reduce to the following:*

$$z'_{i,t} = S_i^* - X_{i,t-1} + B_{i,t-1},$$

$$x'_{t-1} = S_0^* - I_{t-1} + \sum_{i=1}^n z'_{i,t}.$$

Lemma 2 guarantees that the base-stock policy is always able to return the inventory positions to the base-stock levels in each period. With this in hand, the optimality result follows from a straightforward accounting of the cost incurred in each period.

Theorem 1. *The base-stock policy π incurs expected cost C^* in each period, and is thus optimal.*

Extending Theorem 1 to a model with non-negligible lead times is challenging because the warehouse cannot always return each retailers' inventory position to its base-stock level at the start of each period: some inventory added to retailer i 's pipeline in an earlier period may be better used by a different retailer in a future period based on the intermediate demand realizations, but it cannot be re-allocated at that time. Nevertheless, Theorem 1 provides the insight that using an intuitive base-stock policy is a plausible strategy in this

setting. Moreover, our analysis suggests that a stochastic program can be used to determine the base-stock levels. Extending this approach to obtain a heuristic for the setting with lead times can be done by adjusting the per-period demand in the stochastic program to the lead time demand, as is typically done in the single-location inventory literature. This is the approach we pursue when designing base-stock policies for our industrial partner’s service parts network in Section 2.5.

2.4 Disrupted Mode Analysis

In Section 2.3 we analyzed an inventory system, assuming its inventory level can return to a stable state within a reasonable time frame. However, it is not uncommon for such a system to experience significant delays for its external supply (for example, due to natural disasters or the COVID-19 pandemic), in which case the warehouse cannot expect timely receipt of replenishment orders with the supplier. In this section, we explore effective inventory policies in this setting where only limited inventory is available due to long delays.

A simple policy currently used by our industrial partner is to centralize all inventory at the warehouse when a disruption occurs, then fulfill all orders directly to the customer with expedited shipping. This allows central control of fulfillment decisions, and can reduce the number of customers who have to wait until the disruption is over to have demand fulfilled, i.e., it reduces the backlogging costs. But it also can drastically increase fulfillment costs, since each unit of demand is filled through high cost expedited shipping, rather than by the typical route through the retailers. Thus, in this section, we focus on this tradeoff between backlog and fulfillment costs in order to answer a fundamental question related to our industrial partner’s operations in disrupted mode: under what circumstances is it better to keep the inventory decentralized at the retailers, rather than centralized at the warehouse? We develop a simplified version of our model that captures the key tradeoff between backlog and fulfillment cost which we detail next.

To focus on the tradeoff between backlog and fulfillment costs, in this section we assume the n retailers have identical cost parameters: each retailer i has unit backlog cost $b_i = b$ and unit expedited fulfillment cost $f_i = f$. We also assume the retailers have i.i.d. Poisson demand distributions. Further, to facilitate our analysis of the expected backlog costs, we approximate the periodic arrivals described in the Section 2 model with a continuous Poisson process. In particular, letting T denote the time horizon of the disruption, demand at each retailer i is a Poisson process on the interval $[0, T]$ with rate λ/T , so that the cumulative demand at retailer i over the course of the disruption is a Poisson random variable with rate λ , and we let D_i denote this demand. Thus, the total expected demand across all retailers until replenishment is $n\lambda$, and we let $D^n = \sum_i D_i$ denote this system wide demand. We also assume there are $a \geq 1$ units of inventory available per retailer, so there are na units of inventory available system wide. Since a disruption is only meaningful for our industrial partner when it creates a supply shortage, we focus on the case when $\lambda \geq a$, i.e., expected demand exceeds supply. We also assume holding costs are identical across the retailers and warehouse and so can be effectively ignored, allowing us to further focus on the tradeoff between backlog and fulfillment costs. Finally, to conduct a meaningful analysis we assume $n \geq 2$; otherwise decentralization is always optimal as there is no benefit of aggregation with only one retailer.

Our goal in this model is to identify intuitive conditions on f and b such that holding all inventory at the retailers (i.e., decentralization) is preferred to holding all inventory at the warehouse (i.e., centralization). This “all-or-nothing” comparison of complete decentralization to complete centralization is relevant for our industrial partner’s operations, where the main question in a disruption is whether to centralize all inventory at the warehouse. We are able to provide a simple criterion for this decision in our analysis of this section. Moreover, we also explore the more nuanced strategy of centralizing only a portion of the available inventory during a disruption in our numerical studies of Section 2.5.

We begin our analysis with a characterization of the expected cost of the centralized and decentralized systems. The advantage of using the Poisson process approximation of periodic demand is that we can write these costs in a simple closed form, which is amenable to further analysis. The proofs of all Lemmas in this section are provided in Appendix 2.7.

Lemma 3. *The centralized system's expected total cost is*

$$bT \left(\frac{n(\lambda - a)^2 + a}{2\lambda} \mathbb{P}(D^n \geq na) + \frac{n(\lambda - a) - 1}{2} \mathbb{P}(D^n = na - 1) \right) + fn(\lambda \mathbb{P}(D^n < na) + a \mathbb{P}(D^n > na)). \quad (2.9)$$

The decentralized system's expected total cost is

$$bnT \left(\frac{(\lambda - a)^2 + a}{2\lambda} \mathbb{P}(D_1 \geq a) + \frac{\lambda - a - 1}{2} \mathbb{P}(D_1 = a - 1) \right). \quad (2.10)$$

Lemma 3 presents exact expressions for the expected cost of each policy, and so can be used to evaluate and compare costs directly for any particular problem parameters. However, to gain further insight into the tradeoffs involved, we also derive a simple sufficient condition for decentralization to be preferred. Intuitively, the expedited fulfillment cost, f , should be large relative to the backlog cost, b , in order to prefer decentralization, and our next result establishes such a cutoff.

Theorem 2. *For $n \geq 2$ retailers, $a \geq 1$ inventory per retailer, and Poisson rate $\lambda \geq a$ per retailer, decentralization provides lower cost than centralization if*

$$f \geq bT \left(\frac{0.621}{\lambda} + \frac{0.150}{a} \right).$$

We prove Theorem 2 with two lemmas in two key cases ($\lambda \geq a + 1$ and $a \leq \lambda \leq a + 1$), with lower and upper bounds on the centralized and decentralized costs. The key motivation for our case analyses comes from the following lemma bounding the tail probability of the

sum of Poisson random variables.

Lemma 4. *For D_i , $1 \leq i \leq n$, i.i.d. Poisson random variables with rate $\lambda \geq a + 1$ with $a \in \mathbb{Z}_+$ and $D^n = \sum_i D_i$ we have*

$$\mathbb{P}(D^n \geq na) \geq \mathbb{P}(D_i \geq a).$$

The intuition of Lemma 4 is that when $\lambda \geq a + 1$, the average of n i.i.d. Poisson random variables, $\frac{D^n}{n}$, is more concentrated around the mean λ than that of a single Poisson random variable, and thus there is a higher probability of $\frac{D^n}{n}$ being larger than a value strictly below the mean. The proof of Lemma 4 requires a detailed analysis of the lower incomplete gamma function, which can be used to express the relevant Poisson probabilities in closed form. To the best of our knowledge, this analysis of the incomplete gamma function is novel, as the existing literature effectively proves the inequality in the opposite direction for the case when $\lambda = a$ (Van der Vaart, 1961), and thus may be of independent interest. The main idea of the analysis is to show that when λ is strictly larger than a , scaling both λ and a increases the Poisson tail probability, and the requirement that $\lambda \geq a + 1$ appears tight when a grows large. Lemma 4 then motivates us to divide our analysis below into the two cases of $\lambda \geq a + 1$ and $\lambda \leq a + 1$.

Lemma 5. *For Poisson rate $a \leq \lambda \leq a + 1$, $n \geq 2$ retailers, the centralized system cost is lower bounded by*

$$bnT \frac{(\lambda - a)^2}{4\lambda} + fna(1 - 2e^{-2}), \quad (2.11)$$

while the decentralized system cost is upper bounded by

$$bnT \left(1 - e^{-2}\right) \left(\frac{a + (\lambda - a)^2}{2\lambda}\right). \quad (2.12)$$

Lemma 6. *For Poisson rate $\lambda \geq a + 1$, $n \geq 2$ retailers, the centralized system cost is lower*

bounded by

$$bnT \frac{(\lambda - a)^2}{2\lambda} \mathbb{P}(D_1 \geq a) + fna(1 - \frac{32}{3}e^{-4}), \quad (2.13)$$

while the decentralized system cost is upper bounded by

$$bnT \left(\frac{(\lambda - a)^2}{2\lambda} \mathbb{P}(D_1 \geq a) + \frac{a}{2\lambda} + \frac{1}{2\sqrt{2\pi e}} \right). \quad (2.14)$$

With the bounds of Lemmas 5 and 6 in hand, we now are ready to prove Theorem 2.

Proof of Theorem 2. For $a \leq \lambda \leq a + 1$, setting (2.12) less than (2.11) yields

$$fa(1 - 2e^{-2}) \geq bT \left(\frac{a}{2\lambda}(1 - e^{-2}) + \frac{(\lambda - a)^2}{2\lambda} \left(1 - e^{-2} - \frac{1}{2} \right) \right). \quad (2.15)$$

To derive a sufficient condition, we can further upper bound $\frac{(\lambda - a)^2}{2\lambda}$. Observe that the derivative of $\frac{(\lambda - a)^2}{2\lambda}$ with respect to λ is $\frac{1}{2} - \frac{a^2}{2\lambda^2} \geq 0$, for $a \leq \lambda \leq a + 1$. Therefore, $\frac{(\lambda - a)^2}{2\lambda} \leq \frac{(a+1-a)^2}{2(a+1)} \leq \frac{1}{4}$, where the last inequality follows from $a \geq 1$. Hence, (2.15) is implied by the following sufficient condition,

$$fa(1 - 2e^{-2}) \geq bT \left(\frac{a}{2\lambda}(1 - e^{-2}) + \frac{1}{4} \left(1 - e^{-2} - \frac{1}{2} \right) \right),$$

which simplifies to

$$f \geq bT \left(\frac{1 - e^{-2}}{2(1 - 2e^{-2})\lambda} + \frac{1}{8a} \right) \approx bT \left(\frac{0.593}{\lambda} + \frac{0.125}{a} \right). \quad (2.16)$$

For $\lambda \geq a + 1$, observe that the term $bnT \frac{(\lambda - a)^2}{2\lambda} \mathbb{P}(D_i \geq a)$ cancels in (2.13) and (2.14).

Then, setting (2.14) less than (2.13) yields

$$fa(1 - \frac{32}{3}e^{-4}) \geq bT \left(\frac{a}{2\lambda} + \frac{1}{2\sqrt{2\pi e}} \right),$$

and simplifying we obtain

$$f \geq bT \left(\frac{1}{2(1 - \frac{32}{3}e^{-4})\lambda} + \frac{1}{2\sqrt{2\pi e}(1 - \frac{32}{3}e^{-4})a} \right) \approx bT \left(\frac{0.621}{\lambda} + \frac{0.150}{a} \right). \quad (2.17)$$

Finally, observe that both constants above λ and a in (2.17) are greater than those in (2.16), so the cost condition in (2.17) is a sufficient condition for all $\lambda \geq a \geq 1$. This completes the proof. \square

Theorem 2 has a few important implications. First, it verifies the intuition that decentralization should be preferred when the expediting cost, f , is large relative to the backlog cost, b . Naturally, it also shows that shorter disruption times, T , favor decentralization. More interestingly, however, Theorem 2 also implies that as either demand or supply grow large, decentralization becomes a better option. This is intuitive for large supply, a (since plentiful supply is better sent to the retailers at a lower cost), but perhaps less so for large demand, λ . To explain the preference for decentralization with large λ , we consider two cases. First, when supply and demand are relatively balanced, i.e., $\lambda \approx a$, backlog in the decentralized system occurs toward the end of the time horizon, incurring only a small backlog cost, while the centralized system incurs an expediting cost that grows linearly in λ . Second, when demand far outweighs supply $\lambda \gg a$, most of the retailers would be able to use all their inventory, and so it is best for them to store it locally to avoid the expediting cost. This highlights an important point: decentralization only suffers a disadvantage when inventory at one retailer would have been better used by another retailer, i.e., one retailer stocks out while another has excess. When demand and supply are relatively balanced, this disadvantage is limited by the fact that backlog occurs toward the end of the horizon, and when demand is very large, this disadvantage has a low probability of occurrence. Therefore, Theorem 2 suggests that decentralization can be an effective policy when supply is disrupted. Moreover, our sufficient condition in Theorem 2 provides an easy rule of thumb for deciding

whether decentralization should be preferred, which our industrial partner finds useful for informing such decisions in practice.

2.5 Assessing Policies for Normal and Disrupted Modes

In this section, we adapt our policies for use in the service parts distribution network of a large, U.S. automobile manufacturer under normal and disrupted modes. In the next subsection, we describe relevant details of the automaker’s setting.

2.5.1 Introducing industrial partner’s current practice.

We collected data on 23 of our industrial partner’s service parts, with weights, labor costs, holding costs, shipping costs, and demand information representative of the broad spectrum of parts carried by our partner. Each part is supplied with some lead time by an external supplier in normal mode. In the following sections, we describe the automaker’s current practice for managing inventory of these parts.

2.5.1.1 Three-Tier Distribution Network.

The automaker has a fulfillment network including a central warehouse and 15 distribution centers, or “RDCs”, which operate in the retailer role of our model in Section 2.2 (we thus use the term RDC in place of retailer in this section). For each RDC, its customers are automobile dealers in the region, each of which in turn serve demand from their customers (i.e., end customers who need repairs on their automobile). There are total of more than 4,400 dealers served by our industrial partner. Relative to our model of sections 2.2-2.4, a new important feature here is that the dealers can hold their own inventory and thus represent a third tier in the inventory network. However, since the dealers are independent businesses (i.e., the automaker does not control their inventory policy), our model of Section 2.2 still provides a good approximation because dealers are exogenous to our industrial partner’s

network and are served as customers of the RDCs. Moreover, we validate our approach by verifying a base-stock policy is appropriate for a model with this new feature in Section 2.5.2.1. Finally, in the automaker’s network, some of the RDCs have assigned “partner” RDCs nearby, which act as a backup option for fulfillment (see Appendix 2.8.1 for more detail). We describe the partner RDCs with the fulfillment policy in Section 2.5.1.4 below.

2.5.1.2 Demand Distributions.

The demand that a RDC sees on each day comes in two streams: normal (non-emergent) and customer (emergent) orders. This arises from the fact that dealers hold inventory. A normal order that a dealer places is a replenishment to get up to their base-stock level, whereas a customer order happens when the dealer is stocked out and has a customer waiting. Together with our industrial partner, we estimate daily end-customer demand (who arrive at dealers) as a heterogeneous Poisson process. Moreover, we estimate the dealers base-stock levels based on their frequency of normal versus customer orders. We provide more details on demand and base-stock levels estimation in Appendix Section 2.8.1.

2.5.1.3 Lead Times.

The automaker models the supplier’s lead time to the central warehouse, as well as the warehouse’s lead time to the RDCs as independent normal random variables whose means and variances are estimated from their data. While this generalizes the constant lead times of the model in Section 2, our stochastic programming approach is still applicable by considering the random demand over the lead time (more details on the stochastic program are given in Section 2.5.2.2 below). In accordance with our industrial partner’s practice, we assume immediate fulfillment from local RDCs to dealers (i.e. dealers who place orders on day t receive supply at the beginning of day $t + 1$ to clear backlogs immediately; this reflects our partner’s practice of daily milk run deliveries from the RDC to dealers), and we assume a

one-day lead time for the warehouse’s expedited shipments and a two-day lead time for local partner RDC’s fulfillment.

2.5.1.4 Industrial Partner’s Current Fulfillment Policy.

On each day, customers bring cars to dealers for repair. Dealers, based on their own base-stock policies and on-hand inventories, place normal and customer orders to their local RDCs. The RDCs prioritize customer orders and fulfill normal orders with any remaining inventory. In the event when customer orders exceed an RDC’s inventory, a neighboring partner RDC would help fulfill these orders with its available inventory. If the orders exceed the partner’s inventory, then those orders are escalated and backlogged at the central warehouse. For each service part, the warehouse orders weekly from the external supplier. In the automaker’s current practice, the warehouse is generally merely a pass-through and does not hold inventory. Once the warehouse receives supply, it first expedites shipments for customer-order backlogs that were not fulfilled by the local RDCs previously. Then, the warehouse replenishes the RDCs following their independent base-stock levels.

2.5.1.5 Costs.

In addition to the costs that we considered in our model in Section 2.2 (i.e., holding, backlog, and expediting), our industrial partner faces a few additional costs in their daily operation. As a pass-through, the central warehouse incurs a unit cross-docking fee as part of the labor cost for inventories that pass through it to local RDCs. In Section 2.5.2.3 and 2.5.3.2, we discuss a few cases when the central warehouse needs to hold inventory. Therefore, whenever inventories are stored in the warehouse and are later used to replenish the RDCs or fulfill expedited shipments, we incur an additional labor cost per unit to store, pick, and pack those inventory parts. Moreover, for shipments sent from the warehouse to the RDCs by truck, we incur a unit replenishment shipping cost. We provide more details on all costs in

Appendix Section 2.8.1.

If everything runs smoothly, the automaker is able to receive external supply within a reasonable range of lead time and replenish the RDCs and fulfill demand by normal operations. We simulate this scenario in Section 2.5.2. However, occasionally, our industrial partner encounters an unforeseen event that disrupts its external supply, which we explore in more detail in Section 2.5.3.

2.5.2 Normal mode of operations

As discussed in Section 2.3, optimizing the inventory policy in normal mode with multiple tiers, multiple locations, expediting, and positive lead times is a challenging problem. However, we are interested in finding a policy that can be easily implemented in practice and performs well. We are able to demonstrate that a base-stock policy can be plausible in the following two sections. First in Section 2.5.2.1, we analyze a one-RDC-one-dealer model to get intuition on the structure of a good inventory policy. This model captures the key feature that a third-party dealer follows its own, separate base-stock policy, and we are able to show a base-stock policy is optimal for the manufacturer. In Section 2.5.2.2, we extend the model to a one-warehouse-n-RDCs network with two demand classes and expedited shipping. We consider both customer orders as well as normal orders. We adapt our stochastic program for setting base-stock policies to this more general setting, and it performs well in simulation. To see how our heuristic performs, we solve this stochastic program using data from our industrial partner in Section 2.5.2.3 to derive the optimal base-stock levels for the network. The simulation using this optimized base-stock policy has a non-trivial performance improvement over the one without optimized base-stock levels.

2.5.2.1 One RDC, one dealer in normal mode.

In this section we demonstrate in a simple model that even when the manufacturer's end customer is an independent dealer who also holds inventory, it remains optimal for the manufacturer to use a base-stock policy. This further suggests (in addition to our analysis of Section 2.3) that the class of base-stock policies may be effective for the automaker's setting. The model we analyze in this section considers a single manufacturer location (the RDC) serving a single independent dealer. The dealer receives inventory from the RDC and serves external demand.

Let D_t denote the random external demand at the dealer on day t (we assume D_t is i.i.d. across t). Let $X_t^{\mathcal{D}}$ and $B_t^{\mathcal{D}}$ denote dealer's inventory and backlog levels at time t . To minimize its own costs, we assume the dealer orders according to a base-stock policy in every period, and let $S^{\mathcal{D}}$ denote its base-stock level. Let O_t and v_t denote the dealer's order to the RDC and its demand fulfillment at time t . Let w_t denote what the RDC sends to the dealer at the end of day t , and we assume the fulfillment w_t does not show up until the beginning of day $t + 1$, at which point it clears the dealer's backlog. We assume there is no quicker expediting option here for simplicity. Let L denote the manufacturer's lead time for replenishing the RDC. We assume, for simplicity, that the manufacturer can fulfill any order from the RDC within the lead time (i.e., the manufacturer ordering policy from an external supplier is not our concern here). Finally, let X and B represent the RDC's inventory and backlog.

The evolution equations for the RDC and dealer are

$$\begin{aligned}
 X_t &= X_{t-1} + x_{t-L} - w_t, \\
 B_t &= B_{t-1} + O_t - w_t, \\
 X_{t+1}^{\mathcal{D}} &= X_t^{\mathcal{D}} + w_t - v_{t+1}, \\
 B_{t+1}^{\mathcal{D}} &= B_t^{\mathcal{D}} + D_{t+1} - v_{t+1}.
 \end{aligned} \tag{2.18}$$

The objective we consider in this model consists of the RDC holding cost and two levels of backlog cost: unit cost b at the RDC level and $b^{\mathcal{D}}$ at the dealer level, respectively. This is to model the scenario in which the manufacturer often bears additional cost when a customer is actively waiting for a part at the dealer (eg. daily rental car cost). The manufacturer needs to select a policy π to minimize long run average costs

$$C_{\pi} = \limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\sum_{t=1}^T \left(hX_t + bB_t + b^{\mathcal{D}} B_t^{\mathcal{D}} \right) \right]. \quad (2.19)$$

Next we show that the manufacturer can use a base-stock policy to minimize (2.19).

Theorem 3. *There exists a base-stock policy that is long-run optimal for Problem (2.19).*

We present the correct base-stock level X^* at the RDC below, which is an extension of the newsvendor formula.

$$\mathbb{P} \left(D^L \leq X^* \right) = \frac{b + b^{\mathcal{D}} \mathbb{P} \left(D^L + D^1 \geq X^* + S^{\mathcal{D}}, D^L \geq X^* \right)}{h + b},$$

where D^L is the random demand over the RDC's lead time, and D^1 is the one-day demand over the RDC's fulfillment time to the dealer. The left-hand-side, $\mathbb{P} \left(D^L \leq X^* \right)$, represents the RDC's optimal service level for a lead time L , which would typically equal $b/(h + b)$ in a traditional newsvendor formula. The numerator on the right-hand-side, compared to the traditional newsvendor formula, has an extra penalty term with coefficient $b^{\mathcal{D}}$ which increases the service level due to the extra dealer's backlog cost. The coefficient $b^{\mathcal{D}}$ is moderated by the joint probability of two events. The first event, $D^L + D^1 \geq X^* + S^{\mathcal{D}}$, is intuitive (since it represents that the demand exceeds the supply of both the RDC and dealer, so the dealer carries backlog), but perhaps the second event $D^L \geq X^*$ is less intuitive. To interpret this, consider the complementary condition $D^L \leq X^*$, which together with $D^L + D^1 \geq X^* + S^{\mathcal{D}}$ implies $D^1 \geq S^{\mathcal{D}}$. In this case the manufacturer should not take the dealer's backlog cost

into consideration when designing its own base-stock policy, because $D^1 \geq S^D$ implies the one-day demand exceeds the dealer’s base-stock level, and the manufacturer cannot control this through its own policy. Despite this issue, Theorem 3 demonstrates that a base-stock policy can remain effective with an independent dealer, and we therefore focus on this policy class in our simulations.

Next, we provide a high-level outline of the proof of Theorem 3 here (which follows a similar strategy to the proof of Theorem 1) and defer the proof details to Appendix 2.7.13. In Appendix 2.7.13.1, we formulate a stochastic program of this one-RDC-one-dealer problem by relaxing the time constraint, and in Section 2.7.13.2, we show that the optimal solution from this stochastic program provides a lower bound of Problem (2.19). In Section 2.7.13.3, we characterize the optimal solution as well as the optimal objective value in tractable forms. Finally, in Section 2.7.13.4, we prove that, with a correct base-stock level, the cost of a base-stock policy, naturally serving as an upper bound of Problem (2.19), matches the lower bound from the stochastic program. This optimality result demonstrates that a base-stock policy is plausible in our industrial partner’s setting.

2.5.2.2 An approximation of Problem (2.3).

So far, we have gained knowledge and confidence on the optimality of a base stock policy in networks in Section 2.3 (with expediting) and Section 2.5.2.1 (with dealer’s independent base-stock level). In this section, we incorporate more realistic features of our industrial partner’s setting to the previous models. First, for each RDC, it receives separate customer and normal orders from dealers in its region. Each quantity implicitly depends on dealers’ own base-stock levels which is the key feature we analyzed in the Section 2.5.2.1. Here, we model each as its own stochastic process. Let $D_{i,t}^N$ and $D_{i,t}^C$ denote normal demand and customer demand seen at RDC i at time t . We also include the expedited fulfillment feature considered in Section 2.3. Let $y_{i,t}$ denote the amount of expedited shipments from the

central warehouse to RDC i at time t , which is not received until day $t + 1$. Last, we include positive, non-trivial external supplier lead time and RDC lead times. Let L denote the lead time from the external supplier to the warehouse and let l_i denote the warehouse's lead time for replenishing RDC i . Both lead times are stochastic. The manufacturer's long-run objective is to minimize the cost in (2.3).

Given the intuition from the previous sections in which a base-stock policy is optimal in an n -RDC system (Section 2.3) and in a one-RDC-one-dealer network (Section 2.5.2.1), we design a 3-stage stochastic program to solve daily base-stock levels for the central warehouse as well as the RDCs. This stochastic program is to approximate the dynamic inventory control problem in (2.3).

Let t be the beginning of the time horizon considered by the stochastic program. In the first stage of our stochastic program, we define I and X_i to be the ideal inventory levels that the manufacturer should have at the central warehouse and the RDCs at time t to balance all the costs. More importantly, I and X_i are what the manufacturer will use to set the base-stock levels. In the second stage, the RDCs first see cumulative customer and normal demand, $D_i^{2,C} = \sum_{s=t+1}^{s+L} D_{i,s}^C$ and $D_i^{2,N} = \sum_{s=t+1}^{s+L} D_{i,s}^N$, over the warehouse's lead time L from the external supplier and then make fulfillment decisions, $w_i^{2,C}$ and $w_i^{2,N}$, only using their on-hand inventory X_i . Meanwhile, in the second stage, the central warehouse receives supply x from the external supplier and sends replenishment z_i to RDC i . In the final stage, each RDC first sees the third stage cumulative demand, $D_i^{3,C} = \sum_{s=t+L+1}^{t+L+l_i} D_{i,s}^C$ and $D_i^{3,N} = \sum_{s=t+L+1}^{t+L+l_i} D_{i,s}^N$, over the lead time l_i from the warehouse to the RDCs and then makes fulfillment decisions $w_i^{3,C}$ and $w_i^{3,N}$ using the new supply replenished from the warehouse and leftover inventory from the previous stage. It is worth noting that because the warehouse's lead times from the supplier L and to the RDCs l_i are stochastic, the stochastic program actually considers the demand over random length lead times in both the second stage and the third stage.

In the following 3-stage stochastic program, denote $\mathbf{X} = (X_1, \dots, X_n)$, $\mathbf{D}^2 = (D_1^{2,C}, D_2^{2,C}, \dots, D_n^{2,C}, D_1^{2,N}, D_2^{2,N}, \dots, D_n^{2,N})$ and $\mathbf{D}^3 = (D_1^{3,C}, D_2^{3,C}, \dots, D_n^{3,C}, D_1^{3,N}, D_2^{3,N}, \dots, D_n^{3,N})$.

We want to minimize

$$C_{\text{SP}} = \mathbb{E} \left[g_2(I, \mathbf{X}, \mathbf{D}^2, \mathbf{D}^3) \right] + h_0 I + \sum_i h_i X_i \quad (2.20)$$

where g_2 is the objective function of the second stage after $D_i^{2,C}$ and $D_i^{2,N}$ are realized

$$\begin{aligned} g_2(I, \mathbf{X}, \mathbf{D}^2, \mathbf{D}^3) = & h_0 x + \sum_i \left((h_i - h_0) z_i + (h_i + \frac{b_i^C}{2}) B_i^{2,C} + (h_i + \frac{b_i^N}{2}) B_i^{2,N} \right) \\ & + \mathbb{E} \left[g_3(I, \mathbf{X}, \mathbf{D}^2, \mathbf{D}^3) \right] \end{aligned}$$

where g_3 is the objective function of the third stage after $D_i^{3,C}$ and $D_i^{3,N}$ are realized

$$g_3(I, \mathbf{X}, \mathbf{D}^2, \mathbf{D}^3) = \sum_i \left[\left(h_i - h_0 + \frac{f_i}{l_i} \right) y_i + (h_i + \frac{b_i^C}{2}) B_i^{3,C} + (h_i + \frac{b_i^N}{2}) B_i^{3,N} \right]$$

subject to the following second stage constraints:

$$\sum_i z_i \leq I, \quad (2.21)$$

$$w_i^{2,C} + w_i^{2,N} \leq X_i, \quad \forall i, \quad (2.22)$$

$$w_i^{2,C} + B_i^{2,C} \geq D_i^{2,C}, \quad \forall i, \quad (2.23)$$

$$w_i^{2,N} + B_i^{2,N} \geq D_i^{2,N}, \quad \forall i, \quad (2.24)$$

and the following third stage constraints:

$$\sum_i (z_i + y_i) \leq I + x, \quad (2.25)$$

$$w_i^{2,C} + w_i^{2,N} + w_i^{3,C} + w_i^{3,N} \leq X_i + z_i, \quad \forall i, \quad (2.26)$$

$$w_i^{2,C} + w_i^{3,C} + y_i + B_i^{3,C} \geq D_i^{2,C} + D_i^{3,C}, \quad \forall i, \quad (2.27)$$

$$w_i^{2,N} + w_i^{3,N} + B_i^{3,N} \geq D_i^{2,N} + D_i^{3,N}, \quad \forall i, \quad (2.28)$$

$$I, x, X_i, z_i, y_i, w_i^{2,C}, w_i^{2,N}, w_i^{3,C}, w_i^{3,N}, B_i^2, B_i^3 \geq 0, \quad \forall i. \quad (2.29)$$

Constraints (2.21) and (2.25) ensure the warehouse inventory level is non-negative. Constraints (2.22) and (2.26) ensure RDCs' inventory levels are non-negative. Constraints (2.23), (2.24), (2.27), (2.28) ensure all demand is either fulfilled or backlogged, and the last constraint is for non-negativity. We refer interested readers to Appendix 2.8.1.1 for more details and intuition on how we constructed the stochastic program.

We charge a backlog cost in both the second and third stage of the SP to avoid pushing all fulfillment to the third stage. This double counts the backlog cost though, so we account for this by normalizing the backlog cost by a factor of 2 to keep the cost of the stochastic program on the order of the cost of one period in the dynamic problem. Similarly, the expediting cost f_i is divided by the random lead time l_i to account for the fact that the decision variable y_i represents the cumulative expediting decisions over the lead time l_i , and so the magnitude of the cost should be normalized to be on the order of the cost of one period (to match the holding and backlog costs).

Next in Section 2.5.2.3, we solve the above SP using the automaker's data, use the implied optimal base-stock levels as a heuristic in simulation, and show this improves over the automaker's current base-stock levels.

2.5.2.3 Experiments in normal mode.

Consistent with our normal mode analysis, we first simulate our industrial partner’s current normal mode policy, denoted as N-I (for “normal-independent base-stock levels”). This N-I policy provides us with a baseline cost representing our industrial partner’s status quo. The supplier-to-warehouse lead times are normally distributed with means ranging from 3 weeks to 3 months and with average coefficient of variation being 0.043. In this first set of simulations, the automaker’s current time-based base-stock levels are estimated from the targeted service levels at each RDC (more details in Appendix Section 2.8.1). We run the simulation for 365 days and for 1,000 sample paths.




Figure 2.1: Ratio of the improved base-stock-level policy over industrial partner’s current policy

Further, we propose an improved base-stock policy allowing the warehouse to hold some inventory, so that it can replenish the RDCs more readily, and have some inventory on hand for expediting. We use the SP constructed in Section 2.5.2.2 to set the base stock levels of the RDCs and warehouse together, which is expected to perform better than the current

independent approach.

We solve the 3-stage SP with 10,000 demand sample paths to obtain the optimal base-stock levels I and X_i for the central warehouse and for each RDC on each day. We do this for 52 days of a year when the warehouse orders from the outside supplier for each part-RDC pair. Then, we use these optimized base-stock levels to run the same simulation as we have done for the N-I policy, and denote this new set of simulations as N-SP. We ran N-SP for the same 1,000 sample paths as in N-I and record each sample path cost for each part. Finally, we compute the ratio of each sample path between N-SP and N-I and report the average ratios and 95% confidence intervals (CIs) in Figure 2.1. Each letter (A through W) in Figure 2.1 represents an individual service part in our simulation.

Our heuristic with improved base-stock levels performs better in cost across parts and achieved a 4.80% average savings in system cost compared to the baseline N-I policy. Figure 2.1 shows, with 95% CI, the cost ratio of the optimized base-stock-level model to our industrial partner's current network is below 1 across all 23 parts. The N-SP policy is robust in that it correlates the warehouse base-stock level with the RDCs. Recall that the automaker currently uses the central warehouse predominantly as a cross-docking site and does not use it to hold inventory for its parts. Our heuristic in Section 2.5.2.2, on the other hand, shows that allowing some inventory to be held at the warehouse can lead to substantial cost reductions. Intuitively, this is possible because the SP provides robust interaction between the centrally held inventory and demand.

2.5.3 Disrupted mode of operations

As discussed in Section 2.4, it is not uncommon for the automaker to encounter unexpected events that prohibit them from receiving external supply within a reasonable lead time. In order to cope with this, our industrial partner adopts what is called a "designate-for-intervention" (DFI) policy: it prioritizes fulfilling emergent customer orders via expedited

shipping, and holds future supply at the central warehouse until the disruptive event terminates. A part is designated for intervention when there exists backlogs at the warehouse, and a part is released from centralization when one day a future order arrives on time and there is enough inventory to clear all customer-order backlogs. It is worth noting that in current practice by our industrial partner, DFI is a dynamic policy, triggered whenever there are enough backlogs. Our industrial partner designs DFI as a protection measure for its customers, in case a future disruption happens.

An interesting question with regard to the DFI policy is what to do during a DFI period. Currently, the automaker holds all arriving supplies at the central warehouse during a DFI period. The model and result in Section 2.4 suggest that it is often better to keep some inventory decentralized. Therefore, in Section 2.5.3.1, we explore this option and model the DFI period inventory allocation problem by a 2-stage stochastic program with two demand streams and expedited shipments, in a similar fashion to the normal mode. In Section 2.5.3.2, we first simulate our industrial partner’s current network with their DFI policy. We then compare it to a modified DFI policy where we split some of the supplies to the RDCs and show that doing this can reduce the system cost.

2.5.3.1 One warehouse, n RDCs in disrupted mode.

In this section, we assume that once DFI is initiated, the warehouse receives no replenishment until T time periods later, and that the warehouse does not replenish the RDCs (i.e. $z_i = 0, \forall i$) during an active DFI period. Denote S as the total inventory on hand across the entire system. In the stochastic program, we allow the warehouse and the RDCs to hold any portion of the available inventory S , representing the desired initial inventory levels. Similar to Section 2.4, we assume holding costs are effectively zero at the central warehouse and the RDCs because holding costs associated with the available inventory S are sunk.

The goal here is to allocate limited inventory S to the central warehouse and the RDCs

at the beginning of the planning horizon. Our focus is on the impact of the initial inventory placement at the warehouse and RDCs, and not on optimizing the dynamic fulfillment policy, because our industrial partner prioritizes filling customer orders, and so will do so as much as possible at all points in time. The goal is to minimize the cost over finite time horizon T .

We introduce a 2-stage stochastic program below, which we use to modify the automaker's DFI policy and send supply to local RDCs at the beginning of a DFI period. It is worth noting that the stochastic program is a relaxation of the dynamic problem, where in the first stage, we define I and X_i to be the ideal inventory the warehouse and the RDCs should have on hand, and in the second stage, we can make the fulfillment decisions after all demand has been realized (i.e., at the end of the time horizon T). The real problem has demand arriving throughout the time horizon, but this relaxation provides an approximation.

In the following, let $\mathbf{D} = (D_1^C, D_2^C, \dots, D_n^C, D_1^N, D_2^N, \dots, D_n^N)$. The stochastic program solves the desired on-hand inventory for the central warehouse as well as the RDCs as

$$\min C_{\text{SP}} = \mathbb{E} [g_2(\mathbf{D})], \quad (2.30)$$

where g_2 is the objective function of the 2nd stage after cumulative normal and customer orders $D^N = \{D_i^N\}_{i=1}^{i=n}$, $D^C = \{D_i^C\}_{i=1}^{i=n}$ over the time horizon T are realized, given as

$$g_2(\mathbf{D}) = \sum_i \left[f_i y_i + \frac{T b_i^C}{2} B_i^C + \frac{T b_i^N}{2} B_i^N \right]$$

subject to

$$\begin{aligned}
I + \sum_i X_i &= S, \\
\sum_i y_i &\leq I, \\
w_i^N + w_i^C &\leq X_i, \quad \forall i, \\
w_i^N + B_i^N &\geq D_i^N, \quad \forall i, \\
w_i^C + y_i + B_i^C &\geq D_i^C, \quad \forall i, \\
I, X_i, w_i^N, w_i^C, y_i, B_i^N, B_i^C &\geq 0, \quad \forall i,
\end{aligned}$$

where $\frac{T}{2}$ can be interpreted as expected number of days for orders to be backlogged at RDCs. The first inventory constraint is binding due to the fact that in practice when a supply disruption occurs, a manufacturer's limited inventory S is most likely not enough to satisfy future demand, so the manufacturer needs to use every unit of inventory on hand. The second and third constraints make sure the warehouse and RDCs' on-hand inventory is non-negative. The fourth and fifth constraints make sure orders are either fulfilled or backlogged. The last constraint is for non-negativity.

2.5.3.2 Experiments in disrupted mode.

In conjunction with our industrial partner, we adopt the following approach to model their current practice for the DFI policy in our simulation. If a part is out-of-stock at a local retailer and its partner RDC (if any), then the part is backlogged at the central warehouse, and it is flagged as DFI. Once the part is flagged, future supplies are held at the warehouse, and only customer orders are fulfilled and expedited by the warehouse. The next step is to determine when to unflag the part from DFI. Because the part is flagged due to long delays on its supply, we consider that it has sufficient supply once a future order arrives on




Figure 2.2: Cost ratios between policies with normal distributions

time (i.e., its warehouse’s realized lead time from the supplier is less than the expected lead time) and once all customer-order backlogs are cleared by the current inventory held at the warehouse. If these two conditions are met, then the part is unflagged from DFI, and the warehouse starts to distribute inventory to local RDCs again. It is worth noting that how the DFI policy performs depends on the distribution of the lead times from the supplier to the warehouse.

We test this policy in a few interesting cases with supplier’s normal lead time distribution and also with a higher lead time variance.

First, we assume the supplier-to-warehouse lead times are still normally distributed with small coefficients of variation. We are interested to see how the DFI policy alone affects the automaker’s network. We run our partner’s current policy with the DFI feature added (we call this new policy D-I for “disrupted” and “independent base-stock policies”) for 365 days for the same 1,000 sample paths. The cost ratios compared to the base-line model N-I without the DFI policy are the blue bars in Figure 2.2.

Observe that the D-I policy in general incurs more cost than the N-I policy when the supplier's lead times have small coefficients of variation. The D-I policy is more expensive because the DFI feature holds all inventory at the central warehouse during a DFI period. This behavior increases the system expedited shipping cost significantly as each item now needs to be expedited to the customer during a DFI period. Further, the D-I policy could also increase the mean wait time of customer orders as now every customer needs to wait at least one day for the expedited shipments from the warehouse. The simulation reflects that the average wait time across all 23 parts increases by 1.75% in the D-I policy compared with the N-I policy. Previously, the N-I policy sends every unit of inventory downstream by default, and hence many customer orders could be fulfilled locally and incur zero wait time if the dealer has supply at the time of ordering. Therefore, under the D-I policy, it is possible for the D-I system backlog cost to increase as customers on average could wait longer.

Although the system cost is higher in D-I, our industrial partner uses it as its current practice because it allows the company to have control of every available unit of limited inventory and provides fairness to customers. Our simulation also reflects that, conditional on the wait time being longer than a day (i.e., the customer needed an expediting shipment), the D-I policy shows positive wait time improvements across all 23 parts, and the average conditional wait time is reduced by 0.52%. Guaranteeing to meet customers' demand with one day expediting shipping, and reducing mean wait time conditional on customers have waited more than a day, is valuable to our industrial partner, especially during a possible event of supply disruption.

Now, using the same normal lead-time distribution assumption, we modify the D-I policy a bit, referring back to the two-stage stochastic program we developed in Section 2.5.3.1. One main change in this policy is that we allow the central warehouse to ship out some inventory to RDCs during a DFI time. Specifically, the optimal inventory level solution of the stochastic program in Section 2.5.3.1 enables the warehouse to allocate inventory

proportional to those inventory levels. As we mentioned at the beginning of Section 2.5.3.2, once a part is designated for intervention, the automaker may still receive supplies ordered before the DFI period, and they choose to centralize all the supplies at the warehouse only for expedited shipments as their current practice. Our SP from Section 2.5.3.1 serves as a heuristic to modify this policy. In our industrial partner’s network, we solve the SP right after a part is flagged as DFI, and we are able to allocate its future arriving orders to the warehouse and RDCs according to the solved inventory levels until the the part is unflagged.

We run this modified DFI policy (denoted as D-SP) for the same 1,000 sample paths and present the cost of each part in red in Figure 2.2. Figure 2.2 shows that D-SP performs better than the D-I policy in cost across parts. By design, the D-SP policy tries to bridge the gap between the D-I and N-I policies. It allows the central warehouse to send some inventory downstream. Because of this flexibility, some customers do not need to wait for one day at the time of ordering and also can be fulfilled more cheaply by not incurring the long-distance expedited shipping fee from the warehouse, which reduces the backlog cost and the shipping cost.

It is reasonable to infer that how the DFI policy performs depends on the supplier’s lead time distributions to the warehouse. In the interest of exploring high lead-time variance, we modified the supplier’s lead times to follow a Weibull distribution with the same means as our industrial partner’s data, but with a higher variance. We achieve this by setting the Weibull shape parameter to 0.8, which induces a great variance in each lead time distribution, as the average of coefficient of variation jumped from the previous 0.043 to 1.43.

Figure 2.3 shows the cost ratios between the D-I and N-I policies in blue. As the figure reveals, D-I with the DFI policy performs better across all parts. This is because, with great lead time variance, each unit of inventory of a part is precious and cannot afford to sit idle locally. Without the DFI policy, N-I would by default allocate all inventory downstream, and it is possible inventory is sitting idle in one RDC but needed in a different RDC. Previously,




Figure 2.3: Cost ratios between policies with Weibull distributions

when the lead time distribution has small variance, this was not an issue as the out-of-stock RDC would soon be replenished. Now, with greatly increased lead time uncertainty, the system cannot afford to mis-allocate any part. Further, the lead time from the warehouse to the RDCs is also important. Many customers could wait significantly longer than one day when a replenishment is transported in this pipeline. The N-I policy’s benefit from shipping everything downstream diminishes greatly as external supplies become unreliable. The D-SP policy, colored in red, outperforms the the N-I policy, and performs similarly to the D-I policy. This provides additional insight that, when the supplier’s lead time is unreliable, including some DFI measure can significantly benefit the whole system.

2.6 Conclusion and Discussion

In this paper, we consider a replenishment and fulfillment problem in a multi-tier, multi-retailer distribution network with expediting. In particular, we study two “modes” of operation in practice. In the normal mode, in which the warehouse can regularly place replen-

ishment orders with the supplier, we demonstrate that a base-stock policy is optimal in a simplified model with expediting and negligible lead times. This provides us the insight that using an intuitive base-stock policy is a plausible strategy in a more general setting with stochastic lead times. In fact, we design an appropriate stochastic program that is the key to determine the base-stock levels in more complex environments and validate that the optimized base-stock policy provides a non-trivial performance improvement in the automaker's inventory network. Therefore, we believe that allowing some inventory to be held at the central warehouse (compared to only utilizing the warehouse as a pass-through) and allowing for interaction between the centrally held inventory and demand is a good strategy for our industrial partner. Moreover, in the disrupted mode, where the system needs to work with limited inventory, we explore the effectiveness of a decentralized policy which keeps all inventory at the retailers. Although the automaker currently takes centralized control over all inventory at the central warehouse during a disruption, we show that such a strategy is not always cost effective. We demonstrate that when the expediting cost is high, or when either supply and/or demand is large, decentralization becomes a better strategy, and we provide a simple, intuitive criterion to determine such a threshold. In the disrupted-mode simulation, we extend our analysis and design an appropriate stochastic program to allow for splitting some of the supplies between the central warehouse and the retailers. We demonstrate that our approach outperforms our industrial partner's current practice of complete centralization when there are few disruptions and the two policies perform similarly when disruptions happen more often. Thus, by providing cost savings in the numerical experiments in both the normal mode and disrupted mode, we demonstrate that our stochastic programming method can be an effective approach for designing inventory policies in practice. We also discover a novel concentration bound on the sum of Poisson random variables in our disrupted-mode analysis by bounding the incomplete gamma function.

Lastly, we provide a few thoughts on directions for future research. As we mentioned

in Section 2.3, minimizing (2.3) is a well-known challenging problem. Our normal mode analysis derives the base-stock policy optimality result for a multi-tier/location inventory system with expediting, assuming the lead times are negligible. However, a theoretical analysis of a model with non-negligible lead times is more challenging. Fortunately, our simulation shows a base-stock policy performs well in more complex settings with stochastic lead times, and thus it may be that a bound on performance can be shown theoretically in such settings. Further, in disrupted mode it would be interesting to extend our cost criterion for decentralization (Theorem 2) to other demand distributions besides Poisson, which would likely require extending the concentration bound in Lemma 4 (which we conjecture may hold for other distributions like the Poisson which are “close” to symmetric). Finally, while we identify the effectiveness of a decentralized policy in our disrupted mode analysis and show in simulation that allowing some inventory to be decentralized can provide cost savings, it would also be interesting to prove a performance bound for such a policy balanced between complete centralization and complete decentralization.

2.7 Technical Proofs

2.7.1 Proof of Lemma 1

Proof. We first introduce a slightly different stochastic program that also optimizes backlog levels at each retailer (this is useful to directly compare to the state of the system in a given period t , but we will show this problem is equivalent to (2.6)):

$$\min_{\mathbf{S}, \mathbf{B} \geq 0} h_0 S_0 + \mathbb{E}[g(\mathbf{S}, \mathbf{B}; \mathbf{D})] + \sum_{i=1}^n (h_i S_i + b_i B_i + b_i \mathbb{E}[D_i])$$

where $g(\mathbf{S}, \mathbf{B}; \mathbf{D}) = \min \sum_{i=1}^n (f_i - b_i - h_0) y_i - (h_i + b_i) w_i$

$$\text{s.t. } \sum_{i=1}^n y_i \leq S_0, \tag{2.31}$$

$$w_i \leq S_i, \quad \forall 1 \leq i \leq n,$$

$$w_i + y_i \leq B_i + D_i, \quad \forall 1 \leq i \leq n$$

$$y_i, w_i \geq 0, \quad \forall 1 \leq i \leq n.$$

Let C' denote the optimal value of (2.31), then we first claim that $C' = C^*$. To see this, first note that $C' \leq C^*$, since the optimal solution of (2.6) is feasible to (2.31) with $\mathbf{B} = 0$, and gives the same objective value as C^* . Next we show that we also have $C^* \leq C'$. To see this, consider an optimal solution to (2.31), denoted $\mathbf{S}^*, \mathbf{B}^*, \mathbf{w}^*, \mathbf{y}^*$, and construct the

following different feasible solution for (2.31)

$$\begin{aligned}
S'_0 &= S_0^*, \\
S'_i &= (S_i^* - B_i^*)^+, \quad 1 \leq i \leq n, \\
B'_i &= 0, \quad 1 \leq i \leq n, \\
w'_i &= w_i^* - \min(S_i^*, B_i^*), \quad 1 \leq i \leq n, \\
y'_i &= (y_i^* - (B_i^* - S_i^*)^+)^+, \quad 1 \leq i \leq n.
\end{aligned}$$

To see that this is feasible, we first observe that since $f_i - b_i - h_0 \geq -(b_i + h_i)$ (since $f_i \geq 0$ and $h_i \geq h_0$), the second stage problem in (2.31) always set w_i as large as it can first, so that the optimal solution always satisfies $w_i^* = \min(S_i^*, B_i^* + D_i)$ (since if this is not true we could decrease y_i by ϵ and increase w_i by ϵ to reduce the cost). From this we can deduce that $w'_i \geq 0$, since if $S_i^* < B_i^*$ then we have $w'_i = 0$, and otherwise we have $w'_i = \min(S_i^* - B_i^*, D_i) \geq 0$. The other variables are non-negative by definition. The first constraint of (2.31) is maintained because we have $y'_i \leq y_i^*$. For the second constraint, by our construction of w'_i and S'_i , it is equivalent to

$$w_i^* \leq (S_i^* - B_i^*)^+ + \min(S_i^*, B_i^*) = S_i^*,$$

which holds because the optimal solution was feasible for this constraint. Similarly, the third constraint holds because by construction of w'_i and y'_i , if $y'_i > 0$ then this constraint is equivalent to

$$w_i^* + y_i^* \leq (B_i^* - S_i^*)^+ + \min(S_i^*, B_i^*) + D_i = B_i^* + D_i,$$

which holds from the feasibility of the optimal solution, while if $y'_i = 0$ then the third

constraint simply becomes

$$w'_i \leq D_i,$$

which holds because if $S_i^* < B_i^*$ then we have $w'_i = 0 \leq D_i$, and otherwise we have $w'_i = \min(S_i^* - B_i^*, D_i) \leq D_i$. This completes the proof of feasibility of the constructed solution for (2.31).

Next we show that the cost of the constructed solution is no larger than the original optimal solution, and hence is also optimal. To see this, consider the difference in cost for retailer i comparing the optimal solution (with superscript $*$) to the constructed solution (with superscript $'$):

$$\begin{aligned} & h_i \min(S_i^*, B_i^*) + b_i B_i^* + \mathbb{E}[(f_i - b_i - h_0) \min((B_i^* - S_i^*)^+, y_i^*) - (h_i + b_i) \min(S_i^*, B_i^*)], \\ &= (f_i - h_0) \mathbb{E}[\min((B_i^* - S_i^*)^+, y_i^*)] + b_i \mathbb{E}[B_i^* - \min((B_i^* - S_i^*)^+, y_i^*) - \min(S_i^*, B_i^*)], \\ &\geq b_i \mathbb{E}[B_i^* - \min((B_i^* - S_i^*)^+, y_i^*) - \min(S_i^*, B_i^*)], \\ &\geq 0, \end{aligned}$$

where the first inequality follows from $f_i - h_0 \geq 0$, and the second inequality follows from observing that

$$\min((B_i^* - S_i^*)^+, y_i^*) + \min(S_i^*, B_i^*) \leq (B_i^* - S_i^*)^+ + \min(S_i^*, B_i^*) = B_i^*.$$

Thus, the cost of the constructed solution is no larger than the optimal solution, and hence the constructed solution is also optimal. But since the constructed solution has $B'_i = 0$ for all $1 \leq i \leq n$, it is also feasible for (2.6) and yields the same objective. Therefore we must have $C^* \leq C'$, and this completes the proof that $C^* = C'$.

Next, to establish the lower bound in terms of C' , let the following σ -algebra denote the

information available at the beginning of period t , after all $z_{i,t}$ allocation decisions have been made, but before any demand has been realized

$$\mathcal{S}_t = \sigma \{x_s, z_{i,s}y_{i,s}, w_{i,s}, 1 \leq i \leq n, s \leq t-1; z_{i,t}, 1 \leq i \leq n; I_0, X_{i,0}, B_{i,0}, 1 \leq i \leq n\}$$

Conditional on the information in \mathcal{S}_t , consider the expected costs incurred in (2.3) during period t , and rewrite them using the evolution equations as follows:

$$\begin{aligned} & \mathbb{E} \left[h_0 I_t + \sum_i (h_i X_{i,t} + b_i B_{t,i} + f_i y_{i,t}) | \mathcal{S}_t \right] \\ &= \mathbb{E} \left[h_0 \left(I_{t-1} + x_{t-1} - \sum_i (y_{i,t} + z_{i,t}) \right) | \mathcal{S}_t \right] \\ & \quad + \mathbb{E} \left[\sum_i h_i (X_{i,t-1} + z_{i,t} - w_{i,t}) + b_i (B_{i,t-1} + D_{i,t} - w_{i,t} - y_{i,t}) + f_i y_{i,t} | \mathcal{S}_t \right], \\ &= h_0 \left(I_{t-1} + x_{t-1} - \sum_i z_{i,t} \right) + \sum_i h_i (X_{i,t-1} + z_{i,t}) + b_i B_{i,t-1} + b_i \mathbb{E}[D_i] \\ & \quad + \mathbb{E} \left[\sum_i (f_i - b_i - h_0) y_{i,t} - (h_i + b_i) w_{i,t} | \mathcal{S}_t \right], \\ &\geq C' \\ &= C^* \end{aligned}$$

where the second equality follows because $I_{t-1}, x_{t-1}, z_{i,t}, X_{i,t-1}$, and $B_{i,t-1}$ for all i , are already determined by the information in \mathcal{S}_t , and $D_{i,t}$ is independent of \mathcal{S}_t and has the same distribution as D_i , and the first inequality follows because letting $S_0 = I_{t-1} + x_{t-1} - \sum_i z_{i,t}$, $S_i = X_{i,t} + z_{i,t}$ and $B_i = B_{i,t}$ for $1 \leq i \leq n$ and noting the constraints that $I_t \geq 0$, $X_{i,t} \geq 0$, and $B_{i,t} \geq 0$ for $1 \leq i \leq n$, we have a feasible solution for problem (2.31). Next, taking a

final expectation, we have

$$\begin{aligned} & \mathbb{E} \left[h_0 I_t + \sum_i (h_i X_{i,t} + b_i B_{t,i} + f_i y_{i,t}) \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[h_0 I_t + \sum_i (h_i X_{i,t} + b_i B_{t,i} + f_i y_{i,t}) \mid \mathcal{S}_t \right] \right] \geq C^*, \end{aligned}$$

implying that the expected cost in each period is larger than C^* , so the long run average cost must also be larger than C^* . \square

2.7.2 Proof of Lemma 2

Proof. We prove the result by induction on the period t . A feasible policy must, in each period t , keep all state variables non-negative, i.e., $I_t, X_{i,t}, B_{i,t} \geq 0$ for all $1 \leq i \leq n$, and also must make non-negative decisions, i.e., $x_t, z_{i,t}, w_{i,t}, y_{i,t} \geq 0$ for all $1 \leq i \leq n$. It is clear from (2.5), (2.4), (2.7), and (2.8) that all the decisions are non-negative, so we focus on the non-negativity of the state variables, as well as the claim of the lemma. First, recall that $I_0 = X_{i,0} = B_{i,0} = 0$ for all $1 \leq i \leq n$. Then, for the base case of period 1, observe that

$$z'_{i,1} = (S_i^* - X_{i,0} + B_{i,0})^+ = (S_i^*)^+ = S_i^* = S_i^* - X_{i,0} + B_{i,0},$$

since $X_{i,0} = B_{i,0} = 0$ and $S_i^* \geq 0$. Similarly, then, we have

$$x'_0 = (S_0^* - I_0)^+ + \sum_{i=1}^n z'_{i,1} = (S_0^*)^+ + \sum_{i=1}^n z'_{i,t} = S_0^* - I_0 + \sum_{i=1}^n z'_{i,1},$$

since $I_0 = 0$ and $S_0^* \geq 0$, so that the lemma's claim holds for $t = 1$. Further, by the evolution equations we have

$$\begin{aligned} I_1 &= I_0 + x'_0 - \sum_i (y'_{i,1} + z'_{i,1}) = S_0^* - \sum_i y_{i,1}^* \geq 0, \\ X_{i,1} &= X_{i,0} + z'_{i,1} - w'_{i,1} = S_i^* - w_{i,1}^* \geq 0, \\ B_{i,1} &= B_{i,0} + D_{i,1} - w'_{i,1} - y'_{i,1} = D_{i,1} - w_{i,1}^* - y_{i,1}^* \geq 0, \end{aligned}$$

where the last inequality in each line follows from feasibility of $w_{i,1}^*$ and $y_{i,1}^*$ for the first, second, and third constraint in (2.6), respectively. Thus, all state variables are non-negative for $t = 1$.

For the induction step, in period $t \geq 1$, assume that all state variables are non-negative and that the claim of the lemma holds. Then consider period $t - 1$. First observe that

$$\begin{aligned} S_i^* - X_{i,t} + B_{i,t} &= S_i^* - X_{i,t-1} - z'_{i,t} - w'_{i,t} + B_{i,t-1} + D_{i,t} - w'_{i,t} - y'_{i,t}, \\ &= S_i^* - X_{i,t-1} - S_i^* + X_{i,t-1} - B_{i,t-1} + B_{i,t-1} + D_{i,t} - y'_{i,t}, \\ &= D_{i,t} - y_{i,t}^* \geq 0, \end{aligned}$$

where the first equality follows from the evolution equations, the second equality follows from the induction hypothesis of $z'_{i,t} = S_i^* - X_{i,t-1} + B_{i,t-1}$, and the final inequality follows from the third constraint in (2.6). Therefore, we have $S_i^* - X_{i,t} + B_{i,t} \geq 0$, which implies

$$z'_{i,t+1} = (S_i^* - X_{i,t} + B_{i,t})^+ = S_i^* - X_{i,t} + B_{i,t}.$$

Similarly, we have

$$\begin{aligned}
S_0^* - I_t &= S_0^* - I_{t-1} - x'_{t-1} + \sum_{i=1}^n (y'_{i,t} + z'_{i,t}), \\
&= S_0^* - I_{t-1} - (S_0^* - I_{t-1} + \sum_{i=1}^n z'_{i,t}) + \sum_{i=1}^n (y'_{i,t} + z'_{i,t}), \\
&= \sum_{i=1}^n y'_{i,t} \geq 0,
\end{aligned}$$

where the first equality follows from the evolution equation for I_t , and the second equality follows from the induction hypothesis of $x'_{t-1} = S_0^* - I_{t-1} + \sum_{i=1}^n z'_{i,t}$. Therefore, we have $S_0^* - I_t \geq 0$, which implies

$$x'_t = (S_0^* - I_t)^+ + \sum_{i=1}^n z'_{i,t+1} = S_0^* - I_t + \sum_{i=1}^n z'_{i,t+1},$$

so that the claim of the lemma holds for period $t + 1$. Now it remains to show the state variables are non-negative. Again, the evolution equations give

$$\begin{aligned}
I_{t+1} &= I_t + x'_t - \sum_i (y'_{i,t+1} + z'_{i,t+1}) = S_0^* - \sum_i y_{i,t}^* \geq 0, \\
X_{i,t+1} &= X_{i,t} + z'_{i,t+1} - w'_{i,t+1} = S_i^* - w_{i,t+1}^* \geq 0, \\
B_{i,t+1} &= B_{i,t} + D_{i,t} - w'_{i,t+1} - y'_{i,t+1} = D_{i,t+1} - w_{i,t+1}^* - y_{i,t+1}^* \geq 0,
\end{aligned}$$

where the last inequality in each line follows from feasibility of $w_{i,t+1}^*$ and $y_{i,t+1}^*$ for the first, second, and third constraint in (2.6), respectively. Thus, all state variables are non-negative for $t + 1$ and the induction is complete. \square

2.7.3 Proof of Theorem 1

Proof. In period t , the expected cost incurred is

$$\begin{aligned}
& \mathbb{E} \left[h_0 I_t + \sum_i (h_i X_{i,t} + b_i B_{t,i} + f_i y_{i,t}) \right] \\
&= \mathbb{E} \left[h_0 \left(I_{t-1} + x'_{t-1} - \sum_i (y'_{i,t} + z'_{i,t}) \right) \right] \\
&\quad + \mathbb{E} \left[\sum_i h_i \left(X_{i,t-1} + z'_{i,t} - w'_{i,t} \right) + b_i \left(B_{i,t-1} + D_{i,t} - w'_{i,t} - y'_{i,t} \right) + f_i y'_{i,t} \right], \\
&= \mathbb{E} \left[h_0 \left(I_{t-1} + S_0^* - I_{t-1} + \sum_i z'_{i,t} - \sum_i (y_{i,t}^* + z'_{i,t}) \right) \right] \\
&\quad + \mathbb{E} \left[\sum_i h_i \left(X_{i,t-1} + S_i^* - X_{i,t-1} + B_{i,t-1} - w_{i,t}^* - B_{i,t-1} \right) \right] \\
&\quad + \mathbb{E} \left[\sum_i b_i \left(B_{i,t-1} + D_{i,t} - w_{i,t}^* - B_{i,t-1} - y_{i,t}^* \right) + f_i y_{i,t}^* \right], \\
&= \mathbb{E} \left[h_0 \left(S_0^* - \sum_i y_{i,t}^* \right) + \sum_i h_i \left(S_i^* - w_{i,t}^* \right) + b_i \left(D_{i,t} - w_{i,t}^* - y_{i,t}^* \right) + f_i y_{i,t}^* \right], \\
&= C^*
\end{aligned}$$

□

2.7.4 Proof of Lemma 3

Proof. We first find the expected backlog cost of an individual retailer i for Poisson D with rate λ and inventory a using the law of total expectation. Conditioning on there being $D = d$ arrivals of the Poisson process at a retailer, the arrivals are uniformly distributed over the time horizon T (Durrett, 1999). Thus, using order statistics of uniform sample of size d and

assuming that a is integral, the expected backlog cost is

$$\begin{aligned}\mathbb{E}[B|D = d] &= bT \sum_{k=a+1}^d \left(1 - \frac{k}{d+1}\right) = \frac{bT}{d+1} \sum_{j=1}^{d-a} j \\ &= \frac{bT(d-a)^+(d-a+1)}{2(d+1)} \\ &= bT \left(\frac{1}{2}d + \frac{a(a+1)}{2(d+1)} - a\right) \mathbb{1}_{\{d \geq a\}},\end{aligned}$$

where the first equality follows from a re-indexing with $j = d + 1 - k$ and reversing the order of the sum, and the second equality is a standard formula for the sum of consecutive integers, with the positive part coming from noting that the initial sum is 0 if $d \leq a$ (i.e., if there are fewer arrivals than a , there is no backlog).

The law of total expectation dictates the expected backlog cost of retailer i for Poisson D with rate λ and inventory a is

$$\mathbb{E}[B] = \mathbb{E}[\mathbb{E}[B|D = d]].$$

We take expectations of the individual parts of the backlog expression,

$$\begin{aligned}\frac{1}{2}\mathbb{E}D\mathbb{1}_{\{D \geq a\}} &= \frac{1}{2} \sum_{k=a}^{\infty} \frac{k\lambda^k e^{-\lambda}}{k!} = \frac{\lambda}{2} \sum_{k=a}^{\infty} \frac{\lambda^{k-1} e^{-\lambda}}{(k-1)!} = \frac{\lambda}{2} \sum_{j=a-1}^{\infty} \frac{\lambda^j e^{-\lambda}}{j!} = \frac{\lambda}{2} \mathbb{P}(D \geq a-1), \\ \frac{a(a+1)}{2}\mathbb{E}\frac{1}{D+1}\mathbb{1}_{\{D \geq a\}} &= \frac{a(a+1)}{2} \sum_{k=a}^{\infty} \frac{\lambda^k e^{-\lambda}}{(k+1)k!} = \frac{a(a+1)}{2\lambda} \sum_{k=a}^{\infty} \frac{\lambda^{k+1} e^{-\lambda}}{(k+1)!} = \frac{a(a+1)}{2\lambda} \mathbb{P}(D \geq a+1), \\ a\mathbb{E}\mathbb{1}_{\{D \geq a\}} &= a\mathbb{P}(D \geq a).\end{aligned}$$

Therefore the expected backlog cost for Poisson D with rate λ and inventory a is

$$\begin{aligned}
& bT \left(\frac{\lambda}{2} \mathbb{P}(D \geq a-1) + \frac{a(a+1)}{2\lambda} \mathbb{P}(D \geq a+1) - a \mathbb{P}(D \geq a) \right), \\
&= bT \left(\left(\frac{\lambda^2 + a(a+1) - 2a\lambda}{2\lambda} \right) \mathbb{P}(D \geq a) + \frac{\lambda}{2} \mathbb{P}(D = a-1) - \frac{a(a+1)}{2\lambda} \mathbb{P}(D = a) \right), \\
&= bT \left(\left(\frac{(\lambda-a)^2 + a}{2\lambda} \right) \mathbb{P}(D \geq a) + \frac{\lambda-a-1}{2} \mathbb{P}(D = a-1) \right).
\end{aligned} \tag{2.32}$$

For the decentralized system, the total system cost is simply the sum of n retailers backlog cost in (2.32), which completes the proof of the second part of Lemma 3. Next, for the centralized system, it sees aggregate Poisson demand $D^n = \sum_i D_i$ with rate $n\lambda$ and has total inventory of na , so its expected backlog cost is

$$bT \left(\left(\frac{n(\lambda-a)^2 + a}{2\lambda} \right) \mathbb{P}(D^n \geq na) + \frac{n(\lambda-a) - 1}{2} \mathbb{P}(D^n = na-1) \right).$$

Then, the expected fulfillment cost (with unit cost f) for the centralized system is

$$\begin{aligned}
f \mathbb{E} \min(D^n, na) &= f \left(\mathbb{E} D \mathbb{1}_{\{D^n \leq na\}} + \mathbb{E} a \mathbb{1}_{\{D^n > na\}} \right), \\
&= f \left(\sum_{k=0}^{na} \frac{k(n\lambda)^k e^{-(n\lambda)}}{k!} + a \mathbb{P}(D^n > na) \right), \\
&= f \left((n\lambda) \sum_{k=1}^{na} \frac{(n\lambda)^{k-1} e^{-(n\lambda)}}{(k-1)!} + a \mathbb{P}(D^n > na) \right), \\
&= f \left((n\lambda) \sum_{j=0}^{na-1} \frac{(n\lambda)^j e^{-(n\lambda)}}{j!} + a \mathbb{P}(D^n > na) \right), \\
&= fn (\lambda \mathbb{P}(D^n < na) + a \mathbb{P}(D^n > na)).
\end{aligned}$$

Adding the expedited fulfillment cost to the backlog cost completes the proof of the first part of Lemma 3. □

2.7.5 Proof of Lemma 4

To prove Lemma 4, we first define a few gamma functions and relate them to the Poisson distribution. The gamma function and lower incomplete gamma function are defined for real α and β as

$$\begin{aligned}\Gamma(\alpha) &= \int_0^{\infty} t^{\alpha-1} e^{-t} dt, \\ \gamma(\alpha, x) &= \int_0^x t^{\alpha-1} e^{-t} dt.\end{aligned}$$

For an integer $\alpha \geq 1$ the gamma function satisfies $\Gamma(\alpha) = (\alpha - 1)!$ (Press et al., 1992). Following Press et al. (1992), the reverse cumulative distribution function of a Poisson with rate λ can be written in terms of the gamma functions as follows

$$\mathbb{P}(X_i \geq a) = \frac{\gamma(a, \lambda)}{\Gamma(a)}.$$

Therefore, noting that Lemma 4 is trivially true when $a = 0$ (since both probabilities equal 1), we can prove the lemma by proving the following for $a \geq 1$ and $\lambda \geq a + 1$:

$$\frac{\gamma(na, n\lambda)}{\Gamma(na)} \geq \frac{\gamma(a, \lambda)}{\Gamma(a)}. \quad (2.33)$$

The key building block to prove this result is the following bound comparing the integrand to a partial integral involved in the gamma function.

Lemma 7. *For positive integers $\alpha, \beta \in \mathbb{Z}_+$ such that $\alpha \geq \beta \geq 1$, and a positive real $x \in \mathbb{R}_+$ such that $x \geq \beta + 1$, we have*

$$\left(\frac{x\alpha}{\beta}\right)^{\alpha} e^{-\frac{x\alpha}{\beta}} \leq \int_{\frac{x\alpha}{\beta}}^{\frac{x(\alpha+1)}{\beta}} t^{\alpha} e^{-t} dt.$$

Proof. First, for generic limits of integration $y < z$, integration by parts gives

$$\int_y^z t^\alpha e^{-t} dt = y^\alpha e^{-y} - z^\alpha e^{-z} + \alpha \int_y^z t^{\alpha-1} e^{-t} dt,$$

from which it is clear that $y^\alpha e^{-y} \leq \int_y^z t^\alpha e^{-t} dt$ is equivalent to $z^\alpha e^{-z} \leq \alpha \int_y^z t^{\alpha-1} e^{-t} dt$.

Thus, to prove the lemma we prove the equivalent claim that

$$\left(\frac{x(\alpha+1)}{\beta} \right)^\alpha e^{-\frac{x(\alpha+1)}{\beta}} \leq a \int_{\frac{x\alpha}{\beta}}^{\frac{x(\alpha+1)}{\beta}} t^{\alpha-1} e^{-t} dt. \quad (2.34)$$

To prove this, we first establish a few facts about the function $f(t) = t^{\alpha-1} e^{-t}$. First, we show that $f(t)$ is decreasing in t for $t \geq \alpha - 1$. To see this consider the derivative $f'(t) = t^{\alpha-1} e^{-t} \left(\frac{\alpha-1}{t} - 1 \right)$, which is negative when $t \geq \alpha - 1$.

Next, we show for $\alpha - 1 \leq t \leq \alpha - 1 + \sqrt{\alpha - 1}$ that $f(t)$ is concave, and for $t \geq \alpha - 1 + \sqrt{\alpha - 1}$, $f(t)$ is convex. To see this, consider the second derivative

$$f''(t) = t^{\alpha-3} e^{-t} (t^2 - 2(\alpha-1)t + (\alpha-1)^2 - (\alpha-1)),$$

whose sign is determined by the polynomial $t^2 - 2(\alpha-1)t + (\alpha-1)^2 - (\alpha-1)$ since $t^{\alpha-3} e^{-t}$ is always non-negative for $t \geq 0$. The polynomial $t^2 - 2(\alpha-1)t + (\alpha-1)^2 - (\alpha-1)$ is a convex quadratic (since the leading coefficient is positive) and hence is negative between its roots and positive otherwise. The two roots are given by the quadratic formula $t = \alpha - 1 \pm \sqrt{\alpha - 1}$, so for $\alpha - 1 \leq t \leq \alpha - 1 + \sqrt{\alpha - 1}$ we have $f''(t) \leq 0$ and hence $f(t)$ is concave, while for $t \geq \alpha - 1 + \sqrt{\alpha - 1}$ we have $f''(t) \geq 0$ and hence $f(t)$ is convex.

Next, we progress toward proving (2.34) by proving a lower bound on the integrand $f(t) = t^{\alpha-1} e^{-t}$ between generic limits of integration $y < z$ with $y \geq \alpha - 1$. To do this we

define two auxiliary functions

$$g(t) = \frac{(f(z) - f(y))t + f(y)z - yf(z)}{z - y},$$

$$h(t) = f(z) + f'(z)(t - z),$$

where $g(t)$ represents the secant line to f through y and z , and $h(t)$ represents the tangent line to f at z . Then we claim that for $t \in [y, z]$ we have

$$f(t) \geq \min(g(t), h(t)). \tag{2.35}$$

First, we note that if $z \leq \alpha - 1 + \sqrt{\alpha - 1}$ then $f(t)$ is concave for all $t \in [y, z]$ and so $f(t) \geq g(t)$ follows directly, while if $y \geq \alpha - 1 + \sqrt{\alpha - 1}$ then $f(t)$ is convex for all $t \in [y, z]$ and so $f(t) \geq h(t)$ follows directly.

Thus, it remains to consider the case $y \leq \alpha - 1 + \sqrt{\alpha - 1} \leq z$. For $t \in [\alpha - 1 + \sqrt{\alpha - 1}, z]$, the convexity of f implies $f(t) \geq h(t)$, so the claim holds in this region. Next consider the region $[y, \alpha - 1 + \sqrt{\alpha - 1}]$ and consider two cases. First, consider if $h(y) \leq f(y)$, and recall that the convexity of f in the region $[\alpha - 1 + \sqrt{\alpha - 1}, z]$ means that $h(\alpha - 1 + \sqrt{\alpha - 1}) \leq f(\alpha - 1 + \sqrt{\alpha - 1})$. Thus, the concavity of f in the region $[y, \alpha - 1 + \sqrt{\alpha - 1}]$ implies that $h(t) < f(t)$ for all t in this region, and the claim holds. Otherwise, if $h(y) \geq f(y)$, recall that $g(y) = f(y)$ by definition, so $h(y) \geq g(y)$. Further, since $h(z) = g(z)$ by definition, we must have $h(t) \geq g(t)$ for all $t \leq z$ (since h and g are lines in the plane, which can only cross once). Therefore, we have $g(\alpha - 1 + \sqrt{\alpha - 1}) \leq h(\alpha - 1 + \sqrt{\alpha - 1}) \leq f(\alpha - 1 + \sqrt{\alpha - 1})$, and thus, by the concavity of f we have $g(t) \leq f(t)$ for all $t \in [y, \alpha - 1 + \sqrt{\alpha - 1}]$. This completes the proof of the claim in (2.35).

From (2.35) we have that

$$\int_y^z f(t)dt \geq \int_y^z \min(g(t), h(t))dt = \min\left(\int_y^z g(t)dt, \int_y^z h(t)dt\right),$$

where the equality follows since we have $h(z) = g(z)$ by definition, so we either have $h(t) \geq g(t)$ for all $t \in [y, z]$ or $h(t) \leq g(t)$ for all $t \in [y, z]$, since h and g are lines in the plane and can only cross once. Thus, we will establish (2.34) by proving it is true when we replace the integrand on the right hand side with either lower bound of h or g .

For the lower bound h , after evaluating the integral (2.34) becomes

$$\left(\frac{x(\alpha+1)}{\beta}\right)^\alpha e^{-\frac{x(\alpha+1)}{\beta}} \leq \frac{\alpha x}{2\beta} \left(1 + \frac{x}{\beta} + \frac{2}{\alpha+1}\right) \left(\frac{x(\alpha+1)}{\beta}\right)^{\alpha-1} e^{-\frac{x(\alpha+1)}{\beta}},$$

from which the exponential terms can be canceled on each side, as well as an $(x/\beta)^\alpha(\alpha+1)^{\alpha-1}$ term, to give the equivalent expression

$$\alpha + 1 \leq \frac{\alpha}{2} \left(1 + \frac{x}{\beta} + \frac{2}{\alpha+1}\right).$$

Subtracting α from both sides and combining the fractions on the right gives the equivalent expression

$$1 \leq \frac{\alpha((x-\beta)(\alpha+1) + 2\beta)}{2\beta(\alpha+1)},$$

and multiplying both sides by $2\beta(\alpha+1)$ and canceling common terms gives the equivalent expression

$$2\beta \leq \alpha(x-\beta)(\alpha+1),$$

which follows from $x \geq \beta + 1$, $\alpha \geq \beta$, and $\alpha + 1 \geq \beta + 1 \geq 2$, and thus completes the

verification of (2.34) for the lower bound h .

For the lower bound g , after evaluating the integral (2.34) becomes

$$\left(\frac{x(\alpha+1)}{\beta}\right)^\alpha e^{-\frac{x(\alpha+1)}{\beta}} \leq \frac{\alpha x}{2\beta} \left(\left(\frac{xa}{\beta}\right)^{\alpha-1} e^{-\frac{xa}{\beta}} + \left(\frac{x(\alpha+1)}{\beta}\right)^{\alpha-1} e^{-\frac{x(\alpha+1)}{\beta}} \right),$$

from which the $(x/\beta)^\alpha$ terms cancel on each side, and multiplying each side by the exponential term $e^{\frac{x(\alpha+1)}{\beta}}$ gives the equivalent expression

$$(\alpha+1)^\alpha \leq \frac{\alpha}{2} \left(\alpha^{\alpha-1} e^{\frac{x}{\beta}} + (\alpha+1)^{\alpha-1} \right).$$

Multiplying by 2 on each side and collecting the $\alpha+1$ terms on the left hand side gives the equivalent expression

$$\left(2 - \frac{\alpha}{\alpha+1}\right) (\alpha+1)^\alpha \leq \alpha^\alpha e^{\frac{x}{\beta}},$$

and rearranging the leading fraction and dividing by α^α gives the equivalent expression

$$\left(\frac{\alpha+2}{\alpha+1}\right) \left(1 + \frac{1}{\alpha}\right)^\alpha \leq e^{\frac{x}{\beta}}.$$

Here we note that $x \geq \beta + 1$ and $\alpha \geq \beta$ imply that $1 + \frac{1}{\alpha} \leq \frac{x}{\beta}$, so that the inequality is satisfied if we have

$$\left(\frac{\alpha+2}{\alpha+1}\right) \left(1 + \frac{1}{\alpha}\right)^\alpha \leq e^{1+\frac{1}{\alpha}},$$

which follows from the standard inequality $1 + \frac{1}{\alpha} \leq e^{\frac{1}{\alpha}}$ (i.e., the elementary inequality $1 + t \leq e^t$ for $t = 1/\alpha$). To see this note that $1 + \frac{1}{\alpha} \leq e^{\frac{1}{\alpha}}$ implies $\left(1 + \frac{1}{\alpha}\right)^\alpha \leq e$, and also

implies

$$e^{\frac{1}{\alpha}} \geq 1 + \frac{1}{\alpha} = \frac{\alpha + 1}{\alpha} \geq \frac{\alpha + 2}{\alpha + 1}.$$

Putting these two inequalities together gives

$$\left(\frac{\alpha + 2}{\alpha + 1}\right) \left(1 + \frac{1}{\alpha}\right)^\alpha \leq e \left(e^{\frac{1}{\alpha}}\right) = e^{1 + \frac{1}{\alpha}},$$

which completes the verification of (2.34) for the lower bound g , and thus also the proof. \square

Using Lemma 7, we can then prove the following recursive bound for the incomplete gamma function.

Lemma 8. *For positive integers $\alpha, b \in \mathbb{Z}_+$ such that $\alpha \geq b \geq 1$, and a positive real $x \in \mathbb{R}_+$ such that $x \geq \beta + 1$, we have*

$$a\gamma\left(\alpha, x\frac{\alpha}{\beta}\right) \leq \gamma\left(\alpha + 1, x\left(\frac{\alpha + 1}{\beta}\right)\right).$$

Proof. Integration by parts on the integral representation of the incomplete gamma function gives the following

$$\begin{aligned} \alpha\gamma\left(\alpha, x\frac{\alpha}{\beta}\right) &= \left(\frac{x\alpha}{\beta}\right)^\alpha e^{-\frac{x\alpha}{\beta}} + \int_0^{\frac{x\alpha}{\beta}} t^\alpha e^{-t} dt, \\ &\leq \int_{\frac{x\alpha}{\beta}}^{\frac{x(\alpha+1)}{\beta}} t^\alpha e^{-t} dt + \int_0^{\frac{x\alpha}{\beta}} t^\alpha e^{-t} dt, \\ &= \gamma\left(\alpha + 1, x\left(\frac{\alpha + 1}{\beta}\right)\right), \end{aligned}$$

where the second line follows from Lemma 7. \square

With Lemma 8, the proof of Lemma 4 now follows with a simple recursion.

Proof of Lemma 4. We prove (2.33) by establishing the following claim through induction on the integers $j \geq 0$

$$\frac{\Gamma(a+j)}{\Gamma\alpha(a)}\gamma(a, \lambda) \leq \gamma\left(a+j, \lambda\left(\frac{a+j}{a}\right)\right), \quad (2.36)$$

from which (2.33) follows by letting $j = (n-1)a$. For $j = 0$, (2.36) trivially holds with equality, while for $j = 1$ by Lemma 8, letting $\alpha = \beta = a$ and $x = \lambda$ we have

$$a\gamma(a, \lambda) \leq \gamma\left(a+1, \lambda\left(\frac{a+1}{a}\right)\right),$$

from which (2.36) follows since $a = \Gamma(a+1)/\Gamma(a)$. Now assume (2.36) holds for $j \geq 1$ and consider $j+1$, for which we have:

$$\begin{aligned} \frac{\Gamma(a+j+1)}{\Gamma\alpha(a)}\gamma(a, \lambda) &= \frac{(a+j)\Gamma(a+j)}{\Gamma(a)}\gamma(a, \lambda), \\ &\leq (a+j)\gamma\left(a+j, \lambda\left(\frac{a+j}{a}\right)\right), \\ &\leq \gamma\left(a+j+1, \lambda\left(\frac{a+j+1}{a}\right)\right), \end{aligned}$$

where the first line follows since $\Gamma(a+j+1) = (a+j)\Gamma(a+j)$, the second line follows from the induction hypothesis, and the third line follows from Lemma 8. \square

2.7.6 Proof of Lemma 5

Proof. For $a \leq \lambda \leq a + 1$, the centralized system cost in (2.9) can be lower bounded by

$$\begin{aligned}
& bT \left(\frac{n(\lambda - a)^2 + a}{2\lambda} \mathbb{P}(D^n \geq na) + \frac{n(\lambda - a) - 1}{2} \mathbb{P}(D^n = na - 1) \right) \\
& + fn (\lambda \mathbb{P}(D^n < na) + a \mathbb{P}(D^n > na)) \\
\geq & bT \left(\frac{n(\lambda - a)^2 + a}{2\lambda} \mathbb{P}(D^n \geq na) - \frac{1}{2} \mathbb{P}(D^n = na - 1) \right) \\
& + fn (\lambda \mathbb{P}(D^n < na) + a \mathbb{P}(D^n > na)),
\end{aligned}$$

where the inequality follows because $\lambda \geq a$. Observe that

$$\mathbb{P}(D^n = na - 1) = \frac{(n\lambda)^{na-1}}{(na-1)!e^{n\lambda}} = \frac{na}{n\lambda} \frac{(n\lambda)^{na}}{(na)!e^{n\lambda}} = \frac{a}{\lambda} \mathbb{P}(D^n = na).$$

Then,

$$bT \left(\frac{a}{2\lambda} \mathbb{P}(D^n \geq na) - \frac{1}{2} \mathbb{P}(D^n = na - 1) \right) = bT \left(\frac{a}{2\lambda} \mathbb{P}(D^n > na) \right) \geq 0.$$

Therefore, the centralized system cost can be lower bounded by

$$\begin{aligned}
& bnT \frac{(\lambda - a)^2}{2\lambda} \mathbb{P}(D^n \geq na) + fn (\lambda \mathbb{P}(D^n < na) + a \mathbb{P}(D^n > na)) \\
\geq & bnT \frac{(\lambda - a)^2}{2\lambda} \mathbb{P}(D^n \geq na) + fna (\mathbb{P}(D^n < na) + \mathbb{P}(D^n > na)) \\
= & bnT \frac{(\lambda - a)^2}{2\lambda} \mathbb{P}(D^n \geq na) + fna (1 - \mathbb{P}(D^n = na)),
\end{aligned}$$

where the first inequality follows from $\lambda \geq a$. We introduce the following lemmas to lower bound the remaining terms.

Lemma 9. For all $a \geq 1$ and for any $\lambda \geq a$,

$$\mathbb{P}(D_1 \geq a) \geq \frac{1}{2}.$$

As an immediate consequence of Lemma 9, $\mathbb{P}(D^n \geq na) \geq \frac{1}{2}$. Next, we want to upper bound the term $\mathbb{P}(D^n = na)$.

Lemma 10. $\mathbb{P}(D^n = na)$ with rate $n\lambda$, where $\lambda \geq a \geq 1$, is decreasing in λ .

Lemma 11. $\mathbb{P}(D^n = na)$ with rate na is decreasing in na , where $n \geq 2$, $a \geq 1$, and n and a are integers.

Lemma 10 shows that for fixed n, a , the probability mass function $\mathbb{P}(D^n = na)$ is greater when λ is closer to a . Therefore, if $a \leq \lambda \leq a+1$, we set $\lambda = a$ to upper bound $\mathbb{P}(D^n = na)$. Lemma 11 further implies that we want to set na to its minimum value which happens at $n = 2$ and $a = 1$. Hence, for $a \leq \lambda \leq a+1$,

$$\begin{aligned} & bnT \frac{(\lambda - a)^2}{2\lambda} \mathbb{P}(D^n \geq na) + fna(1 - \mathbb{P}(D^n = na)) \\ & \geq bnT \frac{(\lambda - a)^2}{2\lambda} \frac{1}{2} + fna \left(1 - \frac{2^2}{2!e^2}\right) \\ & = bnT \frac{(\lambda - a)^2}{2\lambda} \frac{1}{2} + fna(1 - 2e^{-2}). \end{aligned}$$

Next, for $a \leq \lambda \leq a+1$ the decentralized system cost in (2.10) can be upper bounded by

$$\begin{aligned} & bnT \left(\frac{(\lambda - a)^2 + a}{2\lambda} \mathbb{P}(D_1 \geq a) + \frac{\lambda - a - 1}{2} \mathbb{P}(D_1 = a - 1) \right) \\ & \leq bnT \left(\frac{(\lambda - a)^2 + a}{2\lambda} \mathbb{P}(D_1 \geq a) \right), \end{aligned}$$

where the inequality follows from the fact that $\lambda - a - 1 \leq 0$. Next, for $a < \lambda \leq a+1$, observe that $\mathbb{P}(D_1 \geq a)$ is maximized at $\lambda = a+1$. The following lemma helps upper bound $\mathbb{P}(D_1 \geq a)$.

Lemma 12. $\mathbb{P}(D_1 \geq a)$ with rate $a + 1$ is decreasing in a for $a \geq 1$.

By Lemma 12, $\mathbb{P}(D_1 \geq a)$ with rate $\lambda = a + 1$ is decreasing in a , so $\mathbb{P}(D_1 \geq a)$ is maximized at $a = 1$. Therefore, $\mathbb{P}(D_1 \geq a) \leq 1 - e^{-2}$, and the decentralized system cost is upper bounded by

$$bnT \left(\frac{(\lambda - a)^2 + a}{2\lambda} (1 - e^{-2}) \right),$$

which completes the proof of Lemma 5. □

2.7.7 Proof of Lemma 6

Proof. For $\lambda \geq a + 1$, by following the same steps in **Proof of Lemma 5**, the centralized system cost in (2.9) can be lower bounded by

$$bnT \frac{(\lambda - a)^2}{2\lambda} \mathbb{P}(D^n \geq na) + fna(1 - \mathbb{P}(D^n = na)).$$

By Lemma 4, $\mathbb{P}(D^n \geq na) \geq \mathbb{P}(D_1 \geq a)$ for $\lambda \geq a + 1$. Next, we upper bound $\mathbb{P}(D^n = na)$. By Lemma 10, if $a + 1 \leq \lambda$, we set $\lambda = a + 1$ to upper bound $\mathbb{P}(D^n = na)$. Then, observe that with a Poisson distribution with rate $n(a + 1)$, $\mathbb{P}(D^n = na) \leq \mathbb{P}(D^n = n(a + 1)) \leq \max_{n(a+1)} \mathbb{P}(D^n = n(a + 1))$. Lemma 11 shows that $\mathbb{P}(D^n = n(a + 1))$ is maximized at the minimal value of $n(a + 1)$, which is at $n = 2$ and $a + 1 = 2$. Hence, for $\lambda \geq a + 1$, the centralized system cost is lower bounded by

$$\begin{aligned} & bnT \frac{(\lambda - a)^2}{2\lambda} \mathbb{P}(D^n \geq na) + fna(1 - \mathbb{P}(D^n = na)) \\ & \geq bnT \frac{(\lambda - a)^2}{2\lambda} \mathbb{P}(D_1 \geq a) + fna \left(1 - \frac{4^4}{4!e^4}\right) \\ & = bnT \frac{(\lambda - a)^2}{2\lambda} \mathbb{P}(D_1 \geq a) + fna \left(1 - \frac{32}{3}e^{-4}\right). \end{aligned}$$

Next, for $\lambda \geq a + 1$, the decentralized cost in (2.10) can be upper bounded by

$$\begin{aligned} & bnT \left(\frac{(\lambda - a)^2 + a}{2\lambda} \mathbb{P}(D_1 \geq a) + \frac{\lambda - a - 1}{2} \mathbb{P}(D_1 = a - 1) \right) \\ & \leq bnT \left(\frac{(\lambda - a)^2}{2\lambda} \mathbb{P}(D_1 \geq a) + \frac{a}{2\lambda} + \frac{\lambda - a - 1}{2} \frac{\lambda^{a-1}}{(a-1)!e^\lambda} \right) \end{aligned}$$

where $\mathbb{P}(D_1 \geq a)$ in the second term is upper bounded by 1, and $\mathbb{P}(D_1 = a - 1) = \frac{\lambda^{a-1}}{(a-1)!e^\lambda}$ is the Poisson probability mass function in the last term. The lemma below is helpful, showing $\frac{\lambda - a - 1}{2} \frac{\lambda^{a-1}}{(a-1)!e^\lambda}$ can be bounded by a constant.

Lemma 13. For $\lambda \geq a \geq 1$,

$$\frac{\lambda - a - 1}{2} \frac{\lambda^{a-1}}{(a-1)!e^\lambda} \leq \frac{1}{2\sqrt{2\pi e}}.$$

Lemma 13 implies that, for $a + 1 \leq \lambda$, the decentralized system cost is upper bounded by

$$bnT \left(\frac{(\lambda - a)^2}{2\lambda} \mathbb{P}(D_1 \geq a) + \frac{a}{2\lambda} + \frac{1}{2\sqrt{2\pi e}} \right),$$

which completes the proof of Lemma 6. □

2.7.8 Proof of Lemma 9

Proof. For a fixed rate $\hat{\lambda}$ and for all $a \leq \hat{\lambda}$,

$$\mathbb{P}(D_1 \geq a) \geq \mathbb{P}(D_1 \geq \hat{\lambda})$$

by the definition of reverse cumulative distribution function. Therefore, it suffices to show that for $\lambda = a$,

$$\mathbb{P}(D_1 \geq \lambda) \geq \frac{1}{2},$$

which is true, shown in (5.5) in Van der Vaart (1961). □

2.7.9 Proof of Lemma 10

Proof. With rate $n\lambda$

$$\mathbb{P}(D^n = na) = \frac{(n\lambda)^{na} e^{-n\lambda}}{(na)!} = \frac{n^{na}}{(na)!} \lambda^{na} e^{-n\lambda}.$$

Taking the derivative with respect to λ yields

$$\left[\frac{n^{na}}{(na)!} \lambda^{na} e^{-n\lambda} \right]' = \frac{n^{na}}{(na)!} \left[n e^{-n\lambda} \lambda^{na-1} (a - \lambda) \right] \leq 0$$

because $\lambda \geq a$. □

2.7.10 Proof of Lemma 11

Proof. Let $x = na$. We claim that $\mathbb{P}(X = x)$ where X is a Poisson random variable with rate x is decreasing in x , where x is an integer. By the Poisson probability mass function, $\mathbb{P}(X = x) = \frac{x^x}{x!e^x}$, this is equivalent to showing,

$$\frac{(x+1)^{x+1}}{(x+1)!e^{x+1}} - \frac{x^x}{x!e^x} < 0.$$

The left-hand-side is equivalent to

$$\frac{1}{x!e^x} \left[\frac{(x+1)^x}{e} - x^x \right].$$

Following Taylor series expansion of $e^y > 1 + y$, we have

$$\begin{aligned}
 1 + \frac{1}{x} &< e^{\frac{1}{x}} \\
 \implies x + 1 &< x e^{\frac{1}{x}} \\
 \implies (x + 1)^x &< x^x e \\
 \implies \frac{(x + 1)^x}{e} &< x^x.
 \end{aligned}$$

Hence, $\frac{(x+1)^x}{e} - x^x < 0$, completing the proof. □

2.7.11 Proof of Lemma 12

Proof. Using the gamma and lower incomplete gamma functions, the reverse cumulative distribution function of a Poisson with rate $\lambda = a + 1$ can be written as

$$\mathbb{P}(D_1 \geq a) = \frac{\gamma(a, a + 1)}{\Gamma(a)},$$

where $\gamma(a, a + 1) = \int_0^{a+1} t^{a-1} e^{-t} dt$. We show that for all $a \geq 1$,

$$\frac{\gamma(a, a + 1)}{\Gamma(a)} > \frac{\gamma(a + 1, a + 2)}{\Gamma(a + 1)}.$$

We first show that for $t > a$, function $e^{-t} t^a$ is decreasing in t . Observe that

$$\frac{a}{at} \left[e^{-t} t^a \right] = e^{-t} t^{a-1} (a - t) < 0. \tag{2.37}$$

Then, (2.37) implies

$$e^{-(a+1)} (a + 1)^a > e^{-t} t^a,$$

for all $t \in [a + 1, a + 2]$. Therefore, integrating both sides of the equation with respect to t , we obtain

$$e^{-(a+1)}(a + 1)^a > \int_{a+1}^{a+2} e^{-t}t^a dt.$$

Next, integration by parts gives

$$\begin{aligned} a \int_0^{a+1} e^{-t}t^{a-1} dt &= e^{-(a+1)}(a + 1)^a + \int_0^{a+1} e^{-t}t^a dt \\ &> \int_{a+1}^{a+2} e^{-t}t^a dt + \int_0^{a+1} e^{-t}t^a dt \\ &= \int_0^{a+2} e^{-t}t^a dt. \end{aligned}$$

Therefore, by definition of the incomplete gamma function,

$$a\gamma(a, a + 1) > \gamma(a + 1, a + 2),$$

which implies

$$\frac{\gamma(a, a + 1)}{\Gamma(a)} > \frac{\gamma(a + 1, a + 2)}{\Gamma(a + 1)}.$$

□

2.7.12 Proof of Lemma 13

Proof. Using the standard Stirling's approximation (Dutkay et al., 2013), we lower bound the factorial,

$$(a - 1)! \geq \sqrt{2\pi(a - 1)}\left(\frac{a - 1}{e}\right)^{a-1},$$

and as a consequence, $\frac{\lambda-a-1}{2} \frac{\lambda^{a-1}}{(a-1)!e^\lambda}$ can be upper bounded by

$$\frac{1}{2\sqrt{2\pi}} \frac{\lambda-a-1}{\sqrt{(a-1)}} \left(\frac{\lambda}{a-1}\right)^{a-1} e^{-(\lambda-a+1)}. \quad (2.38)$$

Since the function in (2.38) is divergent for $a = 1$. We divide this proof into two cases. We first show for large a , (2.38) can be upper bounded by a constant. Then, we show for small a , the same constant remains as a valid upper bound. The following two lemmas help show that (2.38) is upper bounded by a constant for all a .

Lemma 14. *The function in (2.38) is maximized at $\lambda^* = a + \frac{1}{2} + \sqrt{a + \frac{5}{4}}$, for $\lambda \geq a \geq 1$.*

Proof. The partial derivative with respect to λ is

$$\begin{aligned} & \frac{\partial}{\partial \lambda} \left[\frac{1}{2\sqrt{2\pi}} \frac{\lambda-a-1}{\sqrt{(a-1)}} \left(\frac{\lambda}{a-1}\right)^{a-1} e^{-(\lambda-a+1)} \right] \\ &= \frac{1}{2\sqrt{2\pi}(a-1)^{d-\frac{1}{2}} e^{\lambda-a+1}} \lambda^{d-2} \left[-\lambda^2 + (2a+1)\lambda + 1 - a^2 \right], \end{aligned}$$

which setting equal to zero yields critical points

$$\lambda_1 = a + \frac{1}{2} - \sqrt{a + \frac{5}{4}}, \quad \lambda_2 = a + \frac{1}{2} + \sqrt{a + \frac{5}{4}}.$$

Observe that λ_1 is not in the domain $\lambda \geq a$ because $\lambda_1 < a + \frac{1}{2} - \sqrt{1 + \frac{5}{4}} = d - 1 < d$. λ_2 , therefore, is the only candidate. Next, observe that the coefficient $\frac{1}{2\sqrt{2\pi}(a-1)^{d-\frac{1}{2}} e^{\lambda-a+1}} \lambda^{d-2}$ is non-negative, and the leading coefficient of the quadratic term λ^2 is negative. This implies that the derivative is positive between λ_1 and λ_2 and negative before λ_1 and after λ_2 . Therefore, λ_2 is a maximum of the function over the range $\lambda \geq a$. \square

Next, We substitute $\lambda^* = a + \frac{1}{2} + \sqrt{a + \frac{5}{4}}$ in (2.38).

Lemma 15. For $a \geq 5$,

$$\frac{1}{2\sqrt{2\pi}} \frac{\sqrt{a + \frac{5}{4}} - \frac{1}{2}}{\sqrt{(a-1)}} \left(\frac{a + \frac{1}{2} + \sqrt{a + \frac{5}{4}}}{a-1} \right)^{a-1} e^{-(\frac{3}{2} + \sqrt{a + \frac{5}{4}})} \leq \frac{1}{2\sqrt{2\pi}e}. \quad (2.39)$$

Proof. We first show $\frac{\sqrt{a + \frac{5}{4}} - \frac{1}{2}}{\sqrt{(a-1)}} \leq 1$ for $a \geq 5$. Setting $\frac{\sqrt{a + \frac{5}{4}} - \frac{1}{2}}{\sqrt{(a-1)}} \leq 1$, we obtain

$$\begin{aligned} \sqrt{a + \frac{5}{4}} - \frac{1}{2} &\leq \sqrt{a-1} \\ a + \frac{6}{4} - \sqrt{a + \frac{5}{4}} &\leq a-1 \\ \implies 5 &\leq a. \end{aligned}$$

Because $\frac{\sqrt{a + \frac{5}{4}} - \frac{1}{2}}{\sqrt{(a-1)}} \leq 1$ for large a , what is left is to show

$$\left(\frac{a + \frac{1}{2} + \sqrt{a + \frac{5}{4}}}{a-1} \right)^{a-1} e^{-(\frac{3}{2} + \sqrt{a + \frac{5}{4}})} \leq \frac{1}{\sqrt{e}}.$$

Equivalently, we want to show

$$\frac{a + \frac{1}{2} + \sqrt{a + \frac{5}{4}}}{a-1} \leq e^{\frac{1 + \sqrt{a + \frac{5}{4}}}{a-1}}. \quad (2.40)$$

Taylor series expansion of e^x gives that $e^x \geq 1 + x + \frac{x^2}{2}$ for $x \geq 0$. Substituting $x = \frac{1 + \sqrt{a + \frac{5}{4}}}{a-1}$,

we obtain

$$\begin{aligned}
e^{\frac{1+\sqrt{a+\frac{5}{4}}}{a-1}} &\geq 1 + \frac{1 + \sqrt{a + \frac{5}{4}}}{a - 1} + \frac{\left(\frac{1+\sqrt{a+\frac{5}{4}}}{a-1}\right)^2}{2} \\
&= \frac{2d^2 + 2a\sqrt{a + \frac{5}{4}} - a + \frac{9}{4}}{2(a - 1)^2} \\
&\geq \frac{2a^2 + 2(a - 2)\sqrt{a + \frac{5}{4}} - a - 1}{2(a - 1)^2} \\
&= \frac{a + \frac{1}{2} + \sqrt{a + \frac{5}{4}}}{a - 1}
\end{aligned}$$

where the last equality follows from dividing both the numerator and denominator by $2(a - 1)$. □

Lemma 15 implies that for $a \geq 5$,

$$\frac{\lambda - a - 1}{2} \frac{\lambda^{a-1}}{(a - 1)!e^\lambda} \leq \frac{1}{2\sqrt{2\pi e}} \approx 0.121.$$

We next numerically validate that this upper bound still holds for $4 \geq a \geq 1$. We set $\lambda = a + \frac{1}{2} + \sqrt{a + \frac{5}{4}}$ and evaluate expression $\frac{\lambda - a - 1}{2} \frac{\lambda^{a-1}}{(a-1)!e^\lambda}$ for small a .

Table 2.1: Evaluating $\frac{\lambda - a - 1}{2} \frac{\lambda^{a-1}}{(a-1)!e^\lambda}$ for small a

a	λ	Expression Value
1	3.0	0.0249
2	4.3	0.0379
3	5.6	0.0464
4	6.8	0.0525

The maximum ratio from the table is $0.0525 < \frac{1}{2\sqrt{2\pi e}}$. This completes the proof of

Lemma 13. □

2.7.13 Proof of Theorem 3

2.7.13.1 A 3-stage stochastic program.

Observe that because the dealer follows a base-stock policy, $O_t = D_t$, the dealer's base-stock policy means that they keep their inventory position at the constant level $S^{\mathcal{D}}$, i.e., they order O_t at the end of period t to satisfy:

$$S^{\mathcal{D}} = X_t^{\mathcal{D}} - B_t^{\mathcal{D}} + B_{t-1} + O_t,$$

where $X_t^{\mathcal{D}} - B_t^{\mathcal{D}}$ denotes their current inventory level at the end of period t when they place their order, B_{t-1} denotes their current unfulfilled orders at period t after receiving the orders the retailer shipped yesterday, and O_t is the order they place at the end of day t to bring the inventory position up to $S^{\mathcal{D}}$. We can re-write it as

$$\begin{aligned} S^{\mathcal{D}} &= X_t^{\mathcal{D}} - B_t^{\mathcal{D}} + B_{t-1} + O_t \\ &= X_t^{\mathcal{D}} - B_t^{\mathcal{D}} + B_t + w_t, \end{aligned}$$

where the last equality follows from the second equation in (2.18). The long run problem is therefore

$$C_{\pi} = \limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\sum_{t=1}^T \left(hX_t + bB_t + b^{\mathcal{D}}B_t^{\mathcal{D}} \right) \right]. \quad (2.41)$$

subject to the following constraints

$$\begin{aligned}
X_{t-1} + x_{t-L} - w_t &\geq 0, \quad \forall t \\
B_{t-1} + O_t - w_t &\geq 0, \quad \forall t \\
X_t^{\mathcal{D}} + w_t - v_{t+1} &\geq 0, \quad \forall t \\
B_t^{\mathcal{D}} + D_{t+1} - v_{t+1} &\geq 0, \quad \forall t \\
S^{\mathcal{D}} &= X_t^{\mathcal{D}} - B_t^{\mathcal{D}} + B_t + w_t, \quad \forall t.
\end{aligned}$$

We relax the time constraint, aggregate information in a lead time L , and allow the manufacturer to make decisions after a lead time. As a result, we define the following. Let

- $D_{t-L+1}^t = \sum_{s=t-L+1}^t D_s$ denote the dealer's lead time demand.
- $x_{t-L+1}^t = \sum_{s=t-L+1}^t x_s$ denote the retailer's lead time order from the supplier.
- $w_{t-L+1}^t = \sum_{s=t-L+1}^{t-1} w_s$ denote lead time dealer orders fulfilled by the retailer.
- $v_{t-L+1}^t = \sum_{s=t-L+1}^t v_s$ denote lead time external customers served by the dealer.

We derive the lead-time evolution equations as follows

$$\begin{aligned}
X_t &= X_{t-L} + x_{t-2L+1}^{t-L} - w_{t-L+1}^t, \\
B_t &= B_{t-L} + D_{t-L+1}^t - w_{t-L+1}^t, \\
X_{t+1}^{\mathcal{D}} &= X_{t-L}^{\mathcal{D}} + (w_{t-L} + w_{t-L+1}^t) - (v_{t-L+1}^t + v_{t+1}), \\
B_{t+1}^{\mathcal{D}} &= B_{t-L}^{\mathcal{D}} + D_{t-L+1}^t + D_{t+1} - (v_{t-L+1}^t + v_{t+1}).
\end{aligned} \tag{2.42}$$

We develop a 3-stage SP by considering lead-time-plus-one-day demand (i.e., $L + 1$ days of demand). In the first stage, initial inventory positions ($X_{t-L} + x_{t-2L+1}^{t-L}, X_{t-L}^{\mathcal{D}} + w_{t-L}$) and backlog levels ($B_{t-L}, B_{t-L}^{\mathcal{D}}$) are chosen. In the second stage, the lead-time demand D_{t-L+1}^t is observed. Then, allocation quantities w_{t-L+1}^t and v_{t-L+1}^t are chosen over the lead time.

In the third stage, an extra day of demand, D_{t+1} , is observed. Then, the last-day allocation quantity v_{t+1} is chosen. Therefore,

$$C_{\text{SP}} = \mathbb{E}[g_1(\cdot)] + h(X_{t-L} + x_{t-2L+1}^{t-L}) + bB_{t-L} + b^{\mathcal{D}}B_{t-L}^{\mathcal{D}},$$

where g_1 is the objective function of the second-stage problem

$$g_1 = (b + b^{\mathcal{D}}) D_{t-L+1}^t + \min \left(- (h + b) w_{t-L+1}^t - b^{\mathcal{D}} v_{t-L+1}^t + \mathbb{E}_D [g_2(\cdot)] \right),$$

where g_2 is the objective function of the third-stage problem

$$g_2(\cdot) = b^{\mathcal{D}} D_{t+1} - \max \left(b^{\mathcal{D}} v_{t+1} \right),$$

subject to

$$\begin{aligned} X_{t-L} + x_{t-2L+1}^{t-L} - w_{t-L+1}^t &\geq 0, \\ B_{t-L} + D_{t-L+1}^t - w_{t-L+1}^t &\geq 0, \\ I^{\mathcal{D}} + w_{t-L+1}^t - (v_{t-L+1}^t + v_{t+1}) &\geq 0, \\ B_{t-L}^{\mathcal{D}} + D_{t-L+1}^t + D_{t+1} - (v_{t-L+1}^t + v_{t+1}) &\geq 0, \\ S^{\mathcal{D}} = X_{t-L}^{\mathcal{D}} + w_{t-L} + B_{t-L} - B_{t-L}^{\mathcal{D}}. \end{aligned}$$

2.7.13.2 Lower bound.

We show that the optimal solution to the 3-stage SP above yields a lower bound for Problem (2.19) .

Lemma 16. $C_{\text{SP}}^* \leq C_{\pi}$.

Proof. We want to show that for $t \geq L$,

$$C_{\text{SP}}^* \leq C_\pi = \mathbb{E}[hX_t + bB_t + b^{\mathcal{D}}B_{t+1}^{\mathcal{D}}].$$

We introduce the σ -algebra, representing the information available at the end of time t

$$A_t = \sigma\{x_s, -L + 1 \leq s \leq t; w_s, v_s, D_s, 1 \leq s \leq t; X_0, X_0^{\mathcal{D}}, B_0, B_0^{\mathcal{D}}\}.$$

It follows that

$$\begin{aligned} \mathbb{E}[X_{t-L}|A_{t-L}] &= X_{t-L}; \\ \mathbb{E}[X_{t-L}^{\mathcal{D}}|A_{t-L}] &= X_{t-L}^{\mathcal{D}}; \quad \mathbb{E}[w_{t-L}|A_{t-L}] = w_{t-L}; \\ \mathbb{E}[B_{t-L}|A_{t-L}] &= B_{t-L}; \quad \mathbb{E}[B_{t-L}^{\mathcal{D}}|A_{t-L}] = B_{t-L}^{\mathcal{D}}; \\ \mathbb{E}[x_s|A_{t-L}] &= x_s, \quad \text{for all } s \leq t - L. \end{aligned}$$

Taking the conditional expectation of $\mathbb{E}[hX_t + bB_t + b^{\mathcal{D}}B_{t+1}^{\mathcal{D}}]$ with respect to A_{t-L} , we

obtain

$$\begin{aligned}
& \mathbb{E} \left[hX_t + bB_t + b^{\mathcal{D}} B_{t+1}^{\mathcal{D}} \middle| A_{t-L} \right] \\
&= \mathbb{E} \left[h \left(X_{t-L} + x_{t-2L+1}^{t-L} - w_{t-L+1}^t \right) \right. \\
&\quad + b \left(B_{t-L} + D_{t-L+1}^t - w_{t-L+1}^t \right) \\
&\quad \left. + b^{\mathcal{D}} \left(B_{t-L}^{\mathcal{D}} + D_{t-L+1}^t + D_{t+1} - \left(v_{t-L+1}^t + v_{t+1} \right) \right) \middle| A_{t-L} \right] \\
&= \mathbb{E} \left[h \left(X_{t-L} + x_{t-2L+1}^{t-L} \right) + bB_{t-L} + b^{\mathcal{D}} B_{t-L}^{\mathcal{D}} \right. \\
&\quad + \left(b + b^{\mathcal{D}} \right) D_{t-L+1}^t - (h + b) w_{t-L+1}^t - b^{\mathcal{D}} v_{t-L+1}^t \\
&\quad \left. + b^{\mathcal{D}} D_{t+1} - b^{\mathcal{D}} v_{t+1} \middle| A_{t-L} \right] \\
&= h \left(X_{t-L} + x_{t-2L+1}^{t-L} \right) + bB_{t-L} + b^{\mathcal{D}} B_{t-L}^{\mathcal{D}} \\
&\quad + \mathbb{E} \left[\left(b + b^{\mathcal{D}} \right) D_{t-L+1}^t - (h + b) w_{t-L+1}^t - b^{\mathcal{D}} v_{t-L+1}^t \right. \\
&\quad \left. + b^{\mathcal{D}} D_{t+1} - b^{\mathcal{D}} v_{t+1} \middle| A_{t-L} \right] \\
&\geq C_{\text{SP}}^*,
\end{aligned}$$

where the first equality follows from (2.42).

Applying the law of total expectation,

$$\mathbb{E} \left[hI_t + bB_t + b^{\mathcal{D}} B_{t+1}^{\mathcal{D}} \right] = \mathbb{E} \left[\mathbb{E} \left[hI_t + bB_t + b^{\mathcal{D}} B_{t+1}^{\mathcal{D}} \middle| A_{t-L} \right] \right] \geq C_{\text{SP}}^*.$$

Finally, for the long-run average cost,

$$\begin{aligned}
C_\pi &= \limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\sum_{t=1}^{T-1} \left(hX_t + bB_t + b^{\mathcal{D}} B_{t+1}^{\mathcal{D}} \right) + b^{\mathcal{D}} B_1 + hX_T + bB_T \right] \\
&\geq \limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\sum_{t=1}^{T-1} \left(hX_t + bB_t + b^{\mathcal{D}} B_{t+1}^{\mathcal{D}} \right) \right] \\
&\geq \limsup_{T \rightarrow \infty} \frac{1}{T} (T-1) C_{\text{SP}}^* \\
&= C_{\text{SP}}^*,
\end{aligned}$$

where the first inequality follows from the fact that $B_1, I_T, B_T^{\mathcal{D}} \geq 0$. □

2.7.13.3 SP solution.

Next, we analyze the 3-stage SP and characterize its optimal solution and the corresponding optimal cost.

To simplify the notation a bit, we use X to denote $X_{t-L} + x_{t-2L+1}^{t-L}$, $X^{\mathcal{D}}$ to denote $X_{t-L}^{\mathcal{D}} + w_{t-L}$, B to denote B_{t-L} , $B^{\mathcal{D}}$ to denote $B_{t-L}^{\mathcal{D}}$, w to denote w_{t-L+1}^t , D^L and D^1 to denote the second-stage lead time demand D_{t-L+1}^t and third-stage one-day fulfillment demand D_{t+1} , and v^L and v^1 to denote the second and third stage customer fulfillment decisions.

The SP provides explicit solutions for the second and third stage variables. In particular, end-customers and the dealer should always be served as much as possible, up to the demand. This leads to the following explicit optimal allocation solutions:

$$w = X \wedge (B + D^L),$$

and

$$\begin{aligned}
v^L + v^1 &= (X^{\mathcal{D}} + w) \wedge (B^{\mathcal{D}} + D^L + D^1) \\
v^1 &= (X^{\mathcal{D}} + X \wedge (B + D^L)) \wedge (B^{\mathcal{D}} + D^L + D^1) - v^L \\
&= (X + X^{\mathcal{D}}) \wedge (B + X^{\mathcal{D}} + D^L) \wedge (B^{\mathcal{D}} + D^L + D^1) - v^L \\
&= (X + X^{\mathcal{D}}) \wedge (B + X^{\mathcal{D}} + D^L) \wedge (B + X^{\mathcal{D}} - S^{\mathcal{D}} + D^L + D^1) - v^L \\
&= (X - B + S^{\mathcal{D}} - D^L) \wedge S^{\mathcal{D}} \wedge D^1 + B + X^{\mathcal{D}} - S^{\mathcal{D}} + D^L - v^L,
\end{aligned}$$

where the second to the last equality follows from the last SP constraint (i.e. $S^{\mathcal{D}} = X^{\mathcal{D}} + B - B^{\mathcal{D}}$).

Inserting w, v^1 into the SP objective and substituting $B^{\mathcal{D}} = B + X^{\mathcal{D}} - S^{\mathcal{D}}$, we obtain

$$C_{\text{SP}} = \mathbb{E}[g_1(\cdot)] + hX + bB + b^{\mathcal{D}}(B + X^{\mathcal{D}} - S^{\mathcal{D}})$$

where

$$g_1 = (b + b^{\mathcal{D}}) D^L - (h + b)[(X - B) \wedge D^L + B] - b^{\mathcal{D}} v^L + \mathbb{E}_D[g_2(\cdot)],$$

where

$$g_2(\cdot) = b^{\mathcal{D}} D^1 - b^{\mathcal{D}} [(X - B + S^{\mathcal{D}} - D^L) \wedge S^{\mathcal{D}} \wedge D^1 + B + X^{\mathcal{D}} - S^{\mathcal{D}} + D^L - v^L].$$

After some simplification, the objective becomes

$$\begin{aligned}
C_{\text{SP}} &= h(X - B) + \mathbb{E} \left[bD^L - (h + b) \left((X - B) \wedge D^L \right) \right. \\
&\quad \left. + \mathbb{E} \left[b^{\mathcal{D}} D^1 - b^{\mathcal{D}} \left((X - B + S^{\mathcal{D}} - D^L) \wedge S^{\mathcal{D}} \wedge D^1 \right) \right] \right]. \tag{2.43}
\end{aligned}$$

Without loss of optimality, we let $B = 0$, since $X - B$ appears together in the objective.

Take the derivative of C_{SP} with respect to X in (2.43), we have

and let F_{D^L} and F_{D^1} denote the second and third stage demand CDF's, respectively, we obtain

$$C'_{\text{SP}} = h - (h + b)(1 - \mathbb{P}(D^L \leq X^*)) - b^D(1 - \mathbb{P}(D^L + D^1 \geq X^* + S^D, D^L \geq X^*)).$$

We want optimal X^* such that

$$0 = h - (h + b)(1 - \mathbb{P}(D^L \leq X^*)) - b^D(1 - \mathbb{P}(D^L + D^1 \geq X^* + S^D, D^L \geq X^*))$$

or

$$\mathbb{P}(D^L \leq X^*) = \frac{b + b^D \mathbb{P}(D^L + D^1 \geq X^* + S^D, D^L \geq X^*)}{h + b}.$$

Finally, the corresponding optimal SP objective value is

$$\begin{aligned} C_{\text{SP}}^* &= hX^* + b\mathbb{E}[D^L] + b^D\mathbb{E}[D^1] - (h + b)\mathbb{E}[X^* \wedge D^L] \\ &\quad - b^D\mathbb{E}[(X^* + S^D - D^L) \wedge S^D \wedge D^1]. \end{aligned} \tag{2.44}$$

2.7.13.4 Upper bound.

Last, we show that base stock policy is optimal for the manufacturer by matching the upper bound with the lower bound. Let S and S^D denote the base-stock levels of the manufacturer and the dealer, respectively. Let $C_{\text{BS}}(S)$ denote the manufacturer's cost given base-stock level S .

Lemma 17. *There exists optimal base-stock level S^* such that $C_{\text{BS}}(S^*) = C_{\text{SP}}^*$.*

Proof. For the manufacturer, its base-stock level at time $t \geq 0$ equals its on-hand inventory,

X_t , plus its pipeline inventory, x_{t-L+1}^t , minus its backlog level, B_t . Additionally, if following a base-stock policy, the manufacturer's pipeline inventory is exact the lead-time demand (i.e., $x_{t-L+1}^t = D_{t-L+1}^t$). Therefore, the manufacturer's base-stock level is

$$\begin{aligned} S &= X_t + x_{t-L+1}^t - B_t \\ &= X_t + D_{t-L+1}^t - B_t. \end{aligned}$$

This implies $X_t - B_t = S - D_{t-L+1}^t$ and since only one of X_t and B_t is positive,

$$X_t = \left(S - D_{t-L+1}^t \right)^+, \quad B_t = \left(D_{t-L+1}^t - S \right)^+. \quad (2.45)$$

Similarly, for the dealer, its base-stock level at time $t+1 \geq 1$ equals its on-hand inventory, X_{t+1}^D , plus its pipeline inventory, $O_{t+1} + B_t$, minus its backlog level, B_{t+1}^D . Then, the dealer's base-stock level is

$$\begin{aligned} S^D &= I_{t+1}^D + O_{t+1} + B_t - B_{t+1}^D \\ &= I_{t+1}^D + D_{t+1} + B_t - B_{t+1}^D. \end{aligned}$$

Using the definition of B_t in Eqn (2.45), we obtain

$$\begin{aligned} B_{t+1}^D - I_{t+1}^D &= -S^D + D_{t+1} + B_t \\ &= -S^D + D_{t+1} + (D_{t-L+1}^t - S)^+ \\ &= -S^D + D_{t+1} + D_{t-L+1}^t - D_{t-L+1}^t \wedge S \\ &= D_{t+1} - S^D \wedge (S + S^D - D_{t-L+1}^t). \end{aligned}$$

This implies

$$B_{t+1}^D = \left(D_{t+1} - S^D \wedge (S + S^D - D_{t-L+1}^t) \right)^+. \quad (2.46)$$

Combining (2.45) and (2.46), the cost of manufacturer's base stock policy with level S is

$$\begin{aligned}
C_{\text{BS}} &= \mathbb{E}[hX_t + bB_t + b^D B_{t+1}^D] \\
&= h\mathbb{E}[(S - D_{t-L+1}^t)^+] + b\mathbb{E}[(D_{t-L+1}^t - S)^+] \\
&\quad + b^D \mathbb{E}[D_{t+1} - S^D \wedge (S + S^D - D_{t-L+1}^t)]^+ \\
&= h\mathbb{E}[S - S \wedge D_{t-L+1}^t] + b\mathbb{E}[D_{t-L+1}^t - S \wedge D_{t-L+1}^t] \\
&\quad + b^D \mathbb{E}[D_{t+1} - (S + S^D - D_{t-L+1}^t) \wedge S^D \wedge D_{t+1}] \\
&= hS + b\mathbb{E}[D_{t-L+1}^t] + b^D \mathbb{E}[D_{t+1}] - (h+b)\mathbb{E}[S \wedge D_{t-L+1}^t] \\
&\quad - b^D \mathbb{E}[(S + S^D - D_{t-L+1}^t) \wedge S^D \wedge D_{t+1}] \\
&= C_{\text{SP}}^*,
\end{aligned}$$

where the last step follows by setting $S = X^*$ in (2.44). □

Combining Lemma 16 and Lemma 17, we obtain Theorem 3.

2.7.14 Other proofs

Proposition 1. *For a decentralized system, where the whole system has S (divisible by n) units of inventory, distributing the inventory evenly across all the retailers is an optimal solution.*

Proof. First, we show that the expected backlog function $f(a)$ is convex in a . By definition,

$$f(a) = \int_0^T (D_t - a)^+ dt.$$

Because the positive function $(D_t - a)^+$ is convex, the integral of the function is also convex

(Boyd and Vandenberghe, 2009). Next, consider the following optimization problem

$$\begin{aligned} \min \quad & \sum_{i=1}^n f(a_i) \\ \text{s.t.} \quad & \sum_{i=1}^n a_i = S \\ & a_i \geq 0, \forall i. \end{aligned}$$

Assume in the optimal solution, there exists two retailers j and k such that $a_j \neq a_k$. Because each cost function f is convex, we know that

$$2f\left(\frac{a_j + a_k}{2}\right) \leq f(a_j) + f(a_k).$$

Then, for these two retailers, we can construct a better solution that yields a lower cost. We can keep doing this for any two retailers that do not hold the same amount of inventory until inventory across all the retailers are symmetric. This completes the proof. \square

2.8 More on numerical experiments

2.8.1 Parameter estimations and network descriptions

In this section, we provide more details on our parameter estimation procedure. It is worth noting that all data provided are masked but represent the realistic setting.

Poisson demand. There are two challenges when it comes to estimating dealer's daily demand. First, the automaker does not directly observe when and how many end customers arrive at dealers. The data only contains information when a dealer orders parts, which could be misleading on the demand arrival rate since many dealers follow a base-stock policy unknown to our industrial partner. The second difficulty comes from limited data. We only have one year-long data and wish to estimate demand rate for each day. In conjunction with

our industrial partner, we agree on the data representing their normal mode of operations as the data year is pre-COVID and the system did not experience major supply or technology disruption.

To estimate weekly demand for each part, we first select a partial demand not affected by dealers' base-stock policies and then apply James-Stein estimator (Stein, 1981) to estimate each week's overall demand.

For each part and for each retailer, we are able to identify a small group of dealers who, during the test year, only placed customer orders. This unique ordering behavior implies these dealers do not carry inventory (equivalently zero base-stock level), and we are able to observe the uncensored demand trend from them. We name these dealers ZBSL. We first summarize weekly retailer customer-order demand from these dealers. We then deseasonalize the 52 data points by what we call a "dynamic moving average" method. For retailer i in week t with demand $D_{t,i}$, its demand is detrended to be the average of

$$D'_{t,i} = \frac{\sum_{t-r}^{t-1} D_{t,i} + D_{t,i} + \sum_{t+1}^{t+r} D_{t,i}}{2r + 1}$$

where $r = 6$ in our detrend method. To apply this method for all 52 weeks, we need to *wrap around* the year and make the whole year a cycle. In addition to the retailer-level data, we can adopt the same detrend method to the national-level data and obtain the deseasonalized national trend (D'_t) for each part.

We then apply the James-Stein estimator to estimate aggregated demand (customer orders and normal orders) for each week t . For a part, we define

$$N_{t,i} = D'_t \frac{\sum_t D'_{t,i}}{\sum_t D'_t}$$

to be the scaled national estimate for retailer i . The intuition behind the James-Stein estimation is that it takes into account both the local trend $D'_{t,i}$ as well as the scaled

national trend $N_{t,i}$ to provide a more accurate and reasonable estimate for weekly demand. By calculating the Stein's parameter, we obtain a convex combination of the two trends, denoted by $J_{t,i}$. Observe that $J_{t,i}$ is determined only by those uncensored dealers having effectively zero base-stock levels. Finally, we need to apply this trend to all dealers. We split the raw aggregated demand (customer and normal orders at the national level) proportionally each week by $J_{t,i}$'s weight. Denote this scaled retailer demand as $\tilde{J}_{t,i}$.

So far, we have shown how to derive a weekly, deseasonalized, part-retailer demand by the James-Stein estimator. We next split $\tilde{J}_{t,i}$ proportionally to each dealer's yearly demand compared to the belonged retailer's yearly demand to obtain dealer's weekly demand. Finally, we divide the weekly demand by seven to obtain customer's daily arrival rate λ_t at a dealer.

Dealers' review policies. In conjunction with our industrial partner, we model four major types of dealers for each part in our simulation. The first type is those who do not hold inventory. These dealers are easily to be identified as in data they only place customer orders, implying they do not carry any inventory and hence their base-stock level is zero. The second type of dealers are those who carry inventory and conduct daily review. The third and fourth types are dealers who review their inventory every seven days on average. We categorize the third type as those following a review interval of Binomial(n, p) where $np = 7$ and the fourth type as those following a review length of Poisson (7). From the raw data and estimated demand distribution, we need to determine for each dealer, whether it belongs to the second, the third, or the fourth group by finding the review policy that gives the maximum likelihood (MLE).

We need to introduce a few parameters before detailing our MLE approach. From data, for each part-dealer pair, we first calculate inter-order times. Let T_k denote the number of days for the k th inter-order time. Let λ_{T_k} denote the cumulative Poisson rate for the k th inter-order interval. Last, let $\bar{\lambda}_{T_k} = \frac{\lambda_{T_k}}{T_k}$ denote the average daily arrival rate during the k th interval.

For daily review, the log likelihood is given as follows

$$LL_{\text{daily}} = \sum_k \log \left(e^{-\bar{\lambda}T_k(T_k-1)}(1 - e^{-\bar{\lambda}T_k}) \right).$$

For dealers with Binomial(n, p) where $np = 7$, we adopt a grid search and a Fibonacci search to search for n^* that gives the highest Binomial likelihood. We refer to Nocedal and Wright (1999) for implementing search algorithms. The log likelihood of Binomial(n, p) is

$$LL_{\text{Bin}}(n, p) = \sum_k \log \left(\sum_{i=1}^{T_k} \sum_{z=0}^{T_k} \Pr^{\text{B}}(z; (i-1)n, p) \Pr^{\text{B}}(T_k - z; n, p) e^{-\bar{\lambda}z} (1 - e^{-\bar{\lambda}(T_k-z)}) \right),$$

where $\Pr^{\text{B}}(m; n, p) = \binom{n}{m} p^m (1-p)^{n-m}$ represents the Binomial probability mass function.

Last, the log likelihood of Poisson(7) is similar to the Binomial case because of the Poisson approximation of Binomial for large n :

$$LL_{\text{Pois}}(7) = \sum_k \log \left(\sum_{i=1}^{T_k} \sum_{z=0}^{T_k} \Pr^{\text{P}}(z; 7(i-1)) \Pr^{\text{P}}(T_k - z; 7) e^{-\bar{\lambda}z} (1 - e^{-\bar{\lambda}(T_k-z)}) \right),$$

where $\Pr^{\text{P}}(m; \lambda) = \frac{\lambda^m e^{-\lambda}}{m!}$ represents the Poisson probability mass function.

For each part-dealer pair, we calculate the log likelihood of each review policy (daily, Binomial(n, p), and Poisson(7)) and assign the policy with the maximum likelihood to that dealer.

Dealers' base-stock levels. For dealers who hold inventory, we also need to infer their base-stock levels. First, we determine each part-dealer pair's service level β by setting it equal to

$$\beta = \frac{\text{number of normal orders}}{\text{number of normal orders} + \text{customer orders}}.$$

Recall that each dealer with a review policy either reviews daily or reviews weekly on average. In conjunction with our industrial partner, we add 5 days to each dealer’s review period length for a “lead time” estimate. This approach is standard in the automaker’s inventory calculation. So each dealer for each part will have a period length of either 6 or 12 days. Next, for each part, we take the dealer’s average daily demand over the year, $\bar{\lambda}$ (i.e. their annual demand divided by 365). Using a Poisson cumulative distribution with rate $\bar{\lambda}$, we are able to calculate the smallest Poisson random variable that guarantees the dealer’s service level β . We treat this Poisson variable as the dealer’s base-stock level for this part. Most of dealers for most parts have base-stock level of 1 in our simulation.

Part-retailer service levels. For each part-retailer pair, we define m_1 to be the total number of customer and normal orders that are originated in the local retailer and fulfilled by the retailer. Define m_2 to be total number of customer and normal orders originated in the local retailer and fulfilled by the retailer, by the partner retailer, and by the warehouse. We define service level r to be

$$r = \frac{m_1}{m_2}.$$

We then map these service levels to the automaker’s service level categories shown below (floored at 80% and capped at 99.5%): 80%, 94%, 95%, 96%, 97%, 98%, 99.5%. For example, 100% would become 99.5%, 96.7% would become 97%, and anything lower than 80% would be raised to 80%.

Initial part-retailer inventory. From raw data, we have information of month-end inventory position for each part-retailer pair. For simplicity, we average 12 month-end inventories and use it as the initial part-retailer inventory position when simulation starts. Observe that the initial inventory can be any arbitrary number, and its initial effect diminishes as the simulation runs for a sufficient length of time. We choose the average method to shorten the time for the system to reach its steady state.

Lead Times. We estimate both supplier-to-warehouse lead time and warehouse-to-

retailer lead times for each part. For each origin-destination pair, we estimate mean as well as standard deviation of the lead times from a year-long shipments data. In all simulations, we assume warehouse-to-retailer lead times to be normally distributed with the estimated parameters. We also assume supplier-to-warehouse lead times normally distributed in most cases. We, however, do explore other scenarios in which lead times follow Weibull distributions in Section 2.5.3.2.

No-order-crossing policy. In conjunction with our industrial partner, we assume no orders can cross. That means a newly placed order has to arrive after all the previously placed orders. This is inspired by what our industrial partners sees in practice. They witness that if an old order experience an delay, future orders will most likely experience more delays.

Partner RDCs. The automaker has one warehouse and 15 RDCs across the contiguous United States and Canada. 10 of the 15 RDCs have one or two partner RDCs based on geographic proximity. These partner RDCs act as backup options for fulfillment. We provide the following masked RDC IDs along with their partner RDC pairs for demonstration.

RDC List: 1,2,3,4,5,6,7,8,9,10,11,12,13,14,15

Partner RDC sets: {1 ,3,7}, {2, 11, 13}, {9,12}, {14,15}.

Costs. For each unit of inventory that passes through the central warehouse, we charge a unit cross-docking fee. For each unit of inventory that is stored at the warehouse, we charge a unit storing fee. For each unit of those stored inventory used for RDC replenishment or expedited shipments, we charge a unit picking and packing fee. These disguised but realistic unit costs depend on repair parts' weights, which is shown in the table below.

Weight (lb)	Picking & Packing	Storing	Cross-docking
≤ 1	\$0.42	\$0.83	\$0.83
10	\$1.25	\$2.50	\$2.50
100	\$5.00	\$10.00	\$10.00

Table 2.2: Unit labor cost in different categories

Estimating the expected delay in Section 2.5.3.1. One challenge in the implementation of the stochastic program from Section 2.5.3.1 is to estimate the number of days of a DFI event (parameter T) right when it occurs. According to the automaker, they often have some information on the lead time when they place an order. For example, if there is a great production delay due to external supplier's labor shortage, they would know it at the time when an order is placed. We, therefore, approximate T by the following approach. We define an order arrives "on time" if it arrives within the expected lead time. If a previously placed order would arrive on time, then set T to be the length of days until that day comes. Otherwise, set T to be the maximum of (1) the length from today to the date of the latest arriving order and (2) one week plus the expected lead time. Because no late order can arrive before an early order by our "no crossing" assumption, then the earliest time for the next order to arrive is the date of the latest arriving order. In addition, because the central warehouse reviews weekly, we should not expect the next order to arrive sooner than one week plus the expected lead time during the DFI period. Hence, the maximizer of the two gives an optimistic approximation on parameter T .

2.8.1.1 Constructing the stochastic program in Section 2.5.2.2.

In this section, we provide some intuition on how we constructed the stochastic program in Section 2.5.2.2.

For each retailer i , we can think about the decisions affecting its fulfillment at time $t + L + l_i$ from the perspective of starting at time t , i.e., the earliest decisions that will impact its fulfillment. The central warehouse and retailers' inventory levels evolve as follows

$$\begin{aligned}
 I_t &= I_{t-1} + x_{t-L} - \sum_i (y_{i,t} + z_{i,t}), \\
 X_{i,t} &= X_{i,t-1} + z_{i,t-l_i} - w_{i,t}^C - w_{i,t}^N,
 \end{aligned} \tag{2.47}$$

where $X_{i,t-1} + z_{i,t-l_i} \geq w_{i,t}^C + w_{i,t}^N$, and $w_{i,t}^C$ ($w_{i,t}^N$) is the inventory retailer i uses for customer

orders (normal orders) at time t . Each retailer is allowed to hold inventory and not fulfill the current orders deliberately for possible future emergent orders. As a consequence, retailer i 's customer-order and normal-order backlog levels evolve as follows

$$\begin{aligned} B_{i,t}^C &= B_{i,t-1}^C + D_{i,t}^C - w_{i,t}^C - y_{i,t-1}, \\ B_{i,t}^N &= B_{i,t-1}^N + D_{i,t}^N - w_{i,t}^N, \end{aligned} \tag{2.48}$$

where $B_{i,t-1}^C + D_{i,t}^C \geq w_{i,t}^C + y_{i,t-1}$ and $B_{i,t-1}^N + D_{i,t}^N \geq w_{i,t}^N$.

From the evolution equations in (2.47) and (2.48), we have

$$\begin{aligned} I_{t+L+l_i-1} &= \left(I_t + \sum_{s=t-L+1}^t x_s \right) + \left(\sum_{s=t+1}^{t+l_i-1} x_s - \sum_{s=t+1}^{t+L} \sum_i (y_{i,s} + z_{i,s}) \right) \\ &\quad - \left(\sum_{s=t+L+1}^{t+L+l_i-1} \sum_i (y_{i,s} + z_{i,s}) \right), \\ X_{i,t+L+l_i} &= \left(X_{it} + y_{it} + \sum_{s=t-l_i+1}^t z_{is} \right) + \left(\sum_{s=t+1}^{t+L} (y_{is} + z_{is}) - \sum_{s=t+1}^{t+L} w_{i,s}^C - \sum_{s=t+1}^{t+L} w_{i,s}^N \right) \\ &\quad + \left(\sum_{s=t+L+1}^{t+L+l_i-1} y_{is} - \sum_{s=t+L+1}^{t+L+l_i} w_{i,s}^C - \sum_{s=t+L+1}^{t+L+l_i} w_{i,s}^N \right), \\ B_{i,t+L+l_i}^C &= B_{i,t}^C + \sum_{s=t+1}^{t+L} (D_{i,t}^C - w_{i,s}^C) + \sum_{s=t+L+1}^{t+L+l_i} (D_{i,t}^C - w_{i,s}^C), \\ B_{i,t+L+l_i}^N &= B_{i,t}^N + \sum_{s=t+1}^{t+L} (D_{i,t}^N - w_{i,s}^N) + \sum_{s=t+L+1}^{t+L+l_i} (D_{i,t}^N - w_{i,s}^N). \end{aligned}$$

Then, the SP works as follows. In the first stage, the warehouse and retailers' initial inventory are determined. In the second stage, the retailers see the second-stage cumulative customer and normal demand and fulfill orders only using their on-hand inventory. Meanwhile, the warehouse receives supply from the external supplier and decides the quantity to send to each retailer. In the final stage, each retailer sees the third-stage cumulative demand and

is able to fulfill orders using the new supply from the warehouse and the leftover inventory from the previous stage. We define the following variables that are later used in this SP.

Finally, we let

- $D_i^{2,C} = \sum_{s=t+1}^{t+L} D_{i,s}^C$ ($D_i^{2,N} = \sum_{s=t+1}^{t+L} D_{i,s}^N$) denote the warehouse lead time customer-order (normal-order) demand, which will be part of the second stage demand in the SP.
- $D_i^{3,C} = \sum_{s=t+L+1}^{t+L+l_i} D_{i,s}^C$ ($D_i^{3,N} = \sum_{s=t+L+1}^{t+L+l_i} D_{i,s}^N$) denote the retailer i lead time customer-order (normal-order) demand, which will be part of the third stage demand in the SP.
- $I = I_t + \sum_{s=t-L+1}^t x_s$ ($X_i = X_{i,t} + y_{i,t} + \sum_{s=t-l_i+1}^t z_{i,s}$) denote the warehouse's (retailer i 's) starting inventory.
- $x = \sum_{s=t+1}^{t+l_i-1} x_s$ denote the second stage order from supplier of inventory for the warehouse, assumed to occur after seeing the realization of second stage demand. $D_{i,t}^{2,C}$ and $D_{i,t}^{2,N}$;
- $z_i = \sum_{s=t+1}^{t+L} \sum_i (y_{i,s} + z_{i,s})$ denote the second stage allocation of the warehouse's inventory to retailer i , assumed to occur after seeing the realization of second stage demand $D_{i,t}^{2,C}$ and $D_{i,t}^{2,N}$.
- $w_i^{2,C} = \sum_{s=t+1}^{t+L} w_{i,s}^C$ ($w_i^{2,N} = \sum_{s=t+1}^{t+L} w_{i,s}^N$) denote the second stage customer-order (normal-order) fulfillment at retailer i .
- $w_i^{3,C} = \sum_{s=t+L+1}^{t+L+l_i} w_{i,s}^C$ ($w_i^{3,N} = \sum_{s=t+L+1}^{t+L+l_i} w_{i,s}^N$) denote the third stage customer-order (normal-order) fulfillment at retailer i .
- $y_i = \sum_{s=t+L+1}^{t+L+l_i-1} \sum_i (y_{i,s} + z_{i,s})$ denote the third stage expedited shipment to retailer i , assumed to occur after seeing the realization of third stage demand $D_{i,t}^{3,C}$ and $D_{i,t}^{3,N}$.

- $B_i^{2,C} = D_i^{2,C} - w_i^{2,C}$ ($B_i^{2,N} = D_i^{2,N} - w_i^{2,N}$) denote the customer-order (normal-order) backlog at retailer i at the end of the second stage.
- $B_i^{3,C} = D_i^{2,C} + D_i^{3,C} - w_i^{2,C} - w_i^{3,C} - y_i$ ($B_i^{3,N} = D_i^{2,N} + D_i^{3,N} - w_i^{2,N} - w_i^{3,N}$) denotes the customer-order (normal-order) backlog at retailer i at the end of the third stage.

CHAPTER 3

MULTI-ITEM ONLINE ORDER FULFILLMENT IN A TWO-LAYER NETWORK

3.1 Introduction

We are in a golden age of e-commerce: US online retail sales of physical goods amounted to \$365.2 billion in 2019, contributing to 11.1% of all retail sales in the US, and were projected to reach close to \$600 billion in 2024 (Statista, 2020a). The boom of e-commerce is not unique to the US market and is witnessed globally. Emerging economies have become especially essential for e-commerce growth because internet adoption is increasing amid growing middle-class populations and fosters influential e-retailers such as Alibaba and JD.com in China, Flipkart in India, Lazada in Southeast Asia, MercadoLibre in Latin America, and Jumia in Africa. The ongoing COVID-19 pandemic has only made e-commerce growth even faster.

To meet growing demands and provide high-quality logistics services, e-retailers have aggressively expanded their fulfillment networks to shorten the distance to end consumers and offer faster delivery service than ever. For instance, in North America, Amazon has been rapidly increasing its number of fulfillment centers (FCs)¹ since 2005 when it only had six FCs (Wulfraat, 2020), and now it has more than 110 active FCs in the US and more than 175 facilities globally (Amazon, 2020). JD.com, a leading e-retailer in China, is currently operating over 900 warehouses, a significant step forward from its 123 warehouses in 2014, and 90% of its first-party retail orders were delivered on the same day or the day after the order was placed in 2020 (Qin et al., 2022). As a result of expansions of fulfillment networks, e-retailers face more complex fulfillment operations. For each online order, they must decide which warehouse(s) will fulfill the order to minimize the fulfillment cost, subject to many

1. In this paper, we use the terms “fulfillment center,” “distribution center,” and “warehouse” interchangeably, although they may have different meanings in practice.

operational constraints such as inventory constraints and guaranteed delivery times (e.g., Acimovic and Graves, 2015, Chen and Graves, 2021). The associated logistics costs are increasing rapidly as well. According to Statista (2020b), as the largest e-retailer, Amazon's outbound shipping costs and operating costs within FCs have skyrocketed over the past decade: its total shipping and fulfillment costs amounted to \$3.8 billion in 2009, and those costs had risen to \$78.1 billion by 2019, more than a 20-fold increase. E-retailers have been implementing various initiatives to mitigate the rising fulfillment costs, including encouraging online customers to do self-pickups (e.g., Walmart's order-online-store-pickup), reducing delivery frequency (e.g., Amazon Day), and exploring omnichannel strategies (e.g., Walmart.com and Amazon effectively utilize existing Walmart stores and Whole Foods stores, respectively, to directly fulfill online orders). Despite all these efforts, the challenge remains.

In practice, many e-retailers such as JD.com and Alibaba are implementing a two-layer (or even multi-layer) distribution network. This network has two types of distribution centers: front distribution centers (FDCs) and regional distribution centers (RDCs). FDCs are lower-layer warehouses that are strategically located closer to end customers. They have limited capacities, usually only stocking fast-moving items. FDCs are responsible for fulfilling online orders from the associated areas in a timely manner. The importance of building FDCs for reducing shipping distances can be reflected in e-retailers' strategic expansions: in the warehouse network of JD.com, the number of cities with FDC coverage leaped from 12 in 2014 to 28 in 2018 (Colliers, 2020). By contrast, RDCs are upper-layer warehouses with larger capacities for storing both fast-moving and slow-moving items and are responsible for replenishing several lower-layer FDCs. When an FDC runs out of inventory, the upper-layer RDC provides the "back-up fulfillment" option and can fulfill orders directly, subject to delivery delay or higher overnight shipping costs. We refer to Ge et al. (2019) and Shen et al. (2020) for more discussions of JD.com's two-layer RDC-FDC network. The concept of FDCs is not uncommon and not unique to the China e-commerce market. Perakis et al.

(2020) document an e-retailer in India that owns 12 major warehouses, including three that have large storage capacity (referred to as “Big Boxes”) and nine FDCs. In the US, Amazon has recently opened a series of small warehouses closer to big cities to speed up same-day delivery (Dastin, 2020) and reportedly plans to open more (Soper, 2020). Compared with typical FCs that are farther from urban cores and hold more inventory, the newly established warehouses are smaller and stock fewer products.

We further illustrate the importance of FDCs in terms of fulfilling local demands by using real data. In Table 3.1, we present the summary statistics of three highest-demand districts, according to an open dataset provided by JD.com (Shen et al., 2020). This dataset is based on transactions within one anonymized consumable category during March 2018, and we refer interested readers to Section 3.8.2 for more details. In Figure 3.1, we present percentages of items fulfilled by the local designated distribution centers (DCs) throughout March 2018. On average, more than 90% of items were fulfilled by local DCs, demonstrating the strategic significance of building and utilizing FDCs in JD.com’s fulfillment network. This phenomenon is not unique. Table 3.2 shows another data snapshot obtained from Alibaba. The table summarizes the fulfillment data of a category of fast-moving products in an eastern city of China. We observe similar statistics: the fulfillment percentages by local FDCs were more than 95% during off-peak season and around 87% during peak season.

Table 3.1: Summary statistics of three highest-demand districts in JD.com’s open dataset

District #	5	9	2
Total Orders	47,676	41,864	38,566
Multi-item Orders (%)	10.97	9.27	10.26
Total Items	54,541	46,889	43,772

Table 3.2: Alibaba’s data snapshot of a category of fast-moving products in an eastern city of China in 2019

	Off-Peak Season				Peak Season
	Day 1	Day 2	Day 3	Day 4	Day 5
Total Orders	7,963	6,872	6,582	6,795	25,502
Multi-item Orders (%)	13.95%	13.22%	13.39%	13.05%	14.96%
Entirely Fulfilled from the FDC (%)	96.72%	97.42%	96.85%	97.98%	87.13%

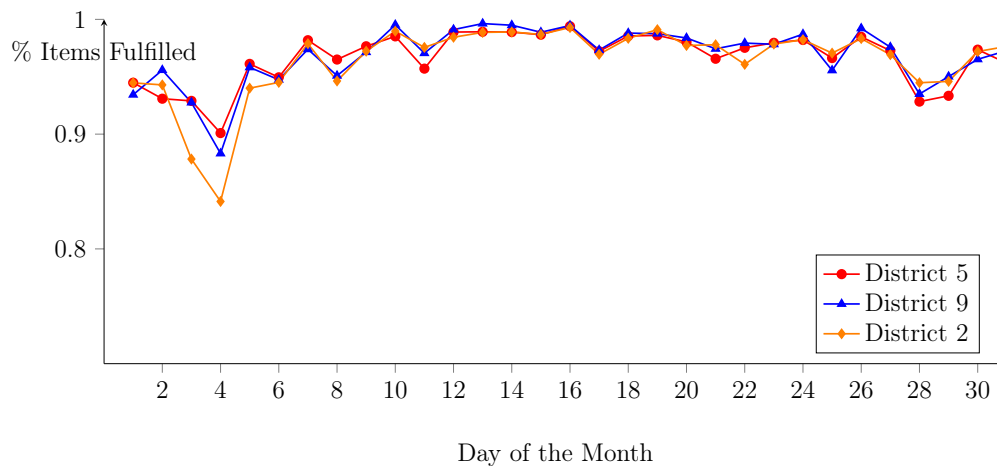


Figure 3.1: Percentages of items fulfilled by the local designated DC in March 2018

Despite the capability to provide faster deliveries in the presence of FDCs, the fulfillment decision itself becomes more intricate in an RDC-FDC network, due to the existence of multi-item orders, namely, orders consisting of more than one item. A multi-item order may lead to *order split*, which is unique to e-commerce fulfillment. Specifically, some items are fulfilled by more distant FCs because the nearest FC does not have all the inventory, resulting in multiple shipments and potentially inconsistent delivery times. Making effective real-time fulfillment decisions at the occurrence of order split is notoriously challenging, which has become a major problem for e-retailers. According to Xu et al. (2009), approximately 35% of Amazon’s online orders in a typical off-peak-season day (44% in a peak-season day) in 2004 were multi-item orders. In addition, Amazon reported that approximately 10% of all

the multi-item orders in a typical off-peak-season day and 15% in a peak-season day were split. Although one may expect that the ratios of multi-item orders have decreased in recent years due to the success of paid-subscription programs (e.g., Amazon Prime) so that free shipping is available to subscribing customers with no minimum spend, the ratios are still not negligible (e.g., see the ratios of JD.com and Alibaba in Tables 3.1 and 3.2). In fact, order split is so common that e-retailers such as Macy’s provide customer service websites explaining it (Macy’s, 2020). Given the high cost associated with order split, coming up with effective fulfillment policies is important for e-retailers.

Last, those fulfillment decisions must be made at the instant the customer orders, as explained in Xu et al. (2009). First, the e-retailer wants to provide an estimated to-ship date as well as the expected number of shipments to the customer, which depends on the fulfillment plan. This instant feedback allows the customer to change her shipping preferences accordingly in real time. Second, the e-retailer wants to virtually secure the inventory in the fulfillment plan. Third, the e-retailer wants to reliably make order-assignment decisions in a 24/7 operating environment, due to variability in order arrivals over a day. Delaying fulfillment decisions may cause warehouse congestion, resulting in delivery delays. Thus, after receiving an order, the e-retailer quickly searches for a feasible fulfillment plan and then virtually reserves inventory. In this paper, we analyze real-time fulfillment decisions in an RDC-FDC network in the presence of multi-item orders.

3.1.1 Literature review

To the best of our knowledge, Xu et al. (2009) were the first to explicitly model and propose a solution strategy for the multi-item online order fulfillment problem. The authors focus on the benefits of delaying fulfillment decisions and batching orders. They propose a myopic-type heuristic that periodically reevaluates all orders that have been assigned to warehouses but have not yet been picked, without considering the impact on future orders.

Wei et al. (2021) investigate the optimal timing to delay fulfillment decisions and how to balance the tradeoff between consolidating shipments and “last-minute” expedited shipping costs. Acimovic and Graves (2015) use an elegant “CD - Textbook” example to demonstrate the challenge of the problem especially due to order split and how fulfillment decisions on a given order affect the e-retailer’s ability to efficiently fulfill future orders. They argue the importance of adopting a forward-looking fulfillment policy. Under the assumption that reliable forecasts on location-specific demand for a product are available, which is typically true for high-volume products at large e-retailers, Acimovic and Graves (2015) propose a forward-looking heuristic that assigns orders based on the dual variables computed through the solution of an offline approximation. When location-specific demand forecasts are unavailable, Andrews et al. (2019) propose a primal-dual algorithm and prove upper and lower bounds on the competitive ratio assuming adversarial demand. This approach is particularly valuable when the fulfillment network consists of hundreds or even thousands of locations that hold inventory, where locations can be either FCs or brick-and-mortar stores. Acimovic and Graves (2017) explore how to use inventory replenishment as a lever to mitigate demand spillover, when the closest FC cannot serve a demand because of a local stockout and that demand is served from a more distant FC at a higher shipping cost. Chen and Graves (2021) study a problem of choosing FCs in which to place items, and formulate the problem as a large-scale mixed-integer program modeling thousands of items to be placed in dozens of FCs and shipped to dozens of customer regions. Li et al. (2019) and Lim et al. (2021) consider joint inventory replenishment/allocation and online order fulfillment problems. DeValve et al. (2021) evaluate the benefits of flexibility when allowing cross fulfillments among adjacent FDCs in a two-layer RDC-FDC distribution network. Joint-pricing and order-fulfillment problems are examined in Lei et al. (2018) and Harsha et al. (2019). Sun et al. (2020) develop a risk-adjusted fulfillment model to analyze two fulfillment choices: using the inventory stored in third-party distribution centers (e.g., fulfillment by Amazon) or

in-house inventory (e.g., fulfillment by seller). We refer to a recent tutorial on the fulfillment optimization problem by Acimovic and Farias (2019).

The paper most relevant to ours is Jasin and Sinha (2015), who explicitly deal with multi-item orders and provide heuristics with performance guarantees. Compared with the earlier paper by Xu et al. (2009), Jasin and Sinha (2015) add two additional layers of complexity. First, they consider the total shipping costs, including a fixed-cost component and a variable-cost component, instead of the number of shipments. Second, they incorporate demand forecast into the model, develop a correlated rounding scheme using the solution of a deterministic linear program (LP) to construct a probabilistic fulfillment policy, and provide an upper bound on its asymptotic competitive ratio. In Table 3.3, we summarize a comparison of the three fulfillment papers (including ours) explicitly dealing with multi-item orders.

We conduct a competitive analysis and use a so-called competitive ratio as the performance measure. It allows for arbitrary non-stationary and even adversarial demand arrivals and does not require any prior knowledge about the arrival patterns. This type of competitive analysis without any knowledge of future inputs in the context of online decision-making has become a popular tool especially in theoretical computer science, and we refer the interested reader to an excellent book on this topic by Borodin and El-Yaniv (2005). It has also been previously applied to various operations management problems, for example, in a two-fare single-leg booking problem (Ball and Queyranne, 2009), in a joint-replenishment problem (Buchbinder et al., 2013), in personalized recommendation/assortment problems (Golrezaei et al., 2014; Chen et al., 2020), in a general class of customer-selection problems (Elmachtoub and Levi, 2016), in a scheduling problem (Wang and Truong, 2018), in an online order fulfillment problem (Andrews et al., 2019), in a general problem of resource allocation (Ma and Simchi-Levi, 2020), and in single-leg revenue management (Ma et al., 2021). We also want to emphasize that analyzing a minimization problem is usually harder than analyzing

Table 3.3: A comparison of fulfillment papers explicitly[◇] dealing with multi-item orders

	Xu et al. (2009)	Jasin and Sinha (2015)	This Paper	This Paper
<i>Forecast Required</i>	no	yes	no	yes
<i>Policy</i>	myopic	forward-looking	myopic	forward-looking
<i># of DCs</i>	multiple	multiple	two	two
<i># of Customer Regions</i>	one	multiple	one	one
<i>Cost Structure</i>	fixed	fixed plus variable	fixed plus variable	fixed plus variable
<i>Competitive Ratio</i>	none	average-case parameter-dependent [♡] asymptotic	worse-case parameter-dependent [♣] non-asymptotic	average-case asymptotic optimality

◇ Following the summary in Table 1 of Acimovic and Farias (2019), we compare our paper only with the two papers by Xu et al. (2009), and Jasin and Sinha (2015), who explicitly incorporate the feature of multi-item orders. Two other papers, by Acimovic and Graves (2015) and Andrews et al. (2019), approximately account for this feature according to Acimovic and Farias (2019), and we do not include the comparison here.

♡ Jasin and Sinha (2015) establish an asymptotic parameter-dependent bound on the competitive ratio as both the demand and inventory scale up. This bound can be roughly interpreted as the weighted fixed costs incurred by an average order, which grows with the expected number of items per order.

♣ We establish non-asymptotic parameter-dependent bounds that depend on the four cost parameters f_R, f_F, c_R, c_F . Note our bounds can become parameter-independent under certain conditions (see Table 3.6).

the maximization counterpart, especially when the “no action” option is available (e.g., allowing a customer order to be intentionally rejected in our fulfillment problem). Although classical problems in competitive analysis (e.g., ski rental, paging/caching) are minimization problems, they are fundamentally different from our fulfillment one.

Finally, our paper is also related to the flexible resource pooling literature such as Asadpour et al. (2020) and Xu et al. (2020), who study online-resource-allocation problems in which flexible resources can satisfy different types of demands and investigate the performance of structures with limited flexibility.

3.1.2 Overview of model and contributions

In this paper, we consider a multi-item online order fulfillment problem motivated by the need to understand how e-retailers can make better real-time fulfillment decisions. The main features of our model are as follows. First, we allow multi-item orders, which is a key feature of online retailing. Second, we consider a two-layer RDC-FDC distribution network consisting of only one RDC and one FDC. According to DeValve et al. (2021), the current fulfillment practice of JD.com only allows spillover fulfillment from the RDC to the FDC

and does not allow cross fulfillment among different FDCs in the same region; that is, no interaction occurs among FDCs. In addition, since in practice RDCs usually have much larger inventory capacities than FDCs, we assume in our model that the RDC has unlimited inventory. As a result, how the RDC inventory is shared among different FDCs becomes irrelevant. In other words, even when the fulfillment network has multiple FDCs, it can reduce to separate networks each consisting of one RDC and one FDC. The one-RDC-one-FDC assumption also aligns with the practice that “over 90% of split orders consist of two shipments” (Xu et al., 2009). Note this model considers fewer warehouses and customer regions than Jasin and Sinha (2015). Third, following Andrews et al. (2019), we assume no demand forecast exists. This assumption is motivated by the fact that predicting product demands at the level of order patterns is generally challenging; additionally, the prediction has to be location-specific. This assumption ensures the robustness of the algorithm, which is an important practical advantage especially in settings where demand is highly volatile (e.g., in the apparel category). Finally, we focus on analyzing a myopic policy. Perhaps surprisingly, myopic policies (and their variations) are standard in the implementation of a typical order-management system (OMS), including IBM Sterling, the most popular OMS on the market (Andrews et al., 2019). This observation also coincides with the claim that many e-retailers simply follow myopic policies in practice (Xu et al., 2009).

Our main theoretical results include parameter-dependent upper and lower bounds on the competitive ratio of our myopic policy, as well as parameter-dependent lower bounds of any online algorithm (including deterministic and randomized). We then discuss their implications. In particular, under certain conditions, we prove the myopic policy is 2-competitive. Under other conditions, we show obtaining parameter-independent bounds is no longer possible, and the lower bound of any online algorithm (including deterministic and randomized) grows with the fixed cost of the RDC. We also prove our bounds on the performance of the myopic policy is tight. Moreover, in the special case in which no variable cost exists, we prove

our upper bound on the competitive ratio of the myopic policy matches the lower bound of any online deterministic policy. Finally, we empirically estimate our upper bound on the competitive ratio of the myopic policy by using FedEx shipping rates, and demonstrate the bound can be as low as 1.13 for reasonable scenarios in practice. Overall, our theoretical results suggest the myopic policy performs well in the considered two-layer fulfillment network. We also extend our study to the setting in which demand forecasting is available and prove the asymptotic optimality of an LP rounding policy. Last, we complement our theoretical results by conducting a numerical study and demonstrating the good performance of the myopic policy even compared with forward-looking policies.

Within the order-fulfillment literature, Jasin and Sinha (2015) is the only extant paper dealing with multi-item orders and having performance guarantees. Our paper differs from theirs in several key ways. First, Jasin and Sinha (2015) consider a generic distribution network allowing any FC to fulfill any customer region, whereas we consider a specialized RDC-FDC network. Second, Jasin and Sinha (2015) assume demand forecasting is available and demands are stationary. By contrast, we assume demand forecasting is unavailable and demands can be non-stationary. Third, Jasin and Sinha (2015) prove asymptotic parameter-dependent bounds on the competitive ratio of forward-looking policies in the average case. By contrast, we focus on the performance of a myopic policy, and our bounds on the competitive ratio are non-asymptotic and worst case. Under certain conditions, we prove the myopic policy has a constant competitive ratio. Furthermore, we empirically estimate our upper bound by using FedEx shipping rates and demonstrate the bound can be close to one for reasonable scenarios in practice. Finally, we also extend our analysis to the setting with demand forecasting (essentially the same setting as Jasin and Sinha (2015) but with only two DCs) and improve the performance bound of LP rounding policies. See Table 3.3 for a detailed comparison.

3.1.3 Outline of paper

The rest of the paper is organized as follows. We describe our model in Section 3.2, introduce the myopic policy in Section 3.2.1 and state the main results in Section 3.2.2. We provide a proof sketch of Theorems 4 and 6 and a high-level picture of other proofs in Section 3.3. We discuss the implications of the main results in Sections 3.4.1 and 3.4.2, and provide additional results in Section 3.4.3. In Section 3.5, we extend our study to the setting in which demand forecasting is available. We conduct a numerical study in Section 3.6. We summarize our contributions and provide a discussion in Section 3.7. We include a technical appendix in Section 3.8.

3.2 Model

In this section, we introduce our fulfillment model. The considered region has one RDC and one FDC. There are n items stocked in the FDC with initial inventory $I_{0,i}$ for item i ($I_{0,i} \geq 0$). The RDC stocks the same assortment of items but with an infinite amount of inventory. Due to this assumption, every customer order can be satisfied within the region, although the cost varies depending on the fulfillment strategy. The FDC inventory cannot be replenished amid the selling horizon, as the horizon can be interpreted as the time between two consecutive replenishments. Customer orders arrive sequentially, and each order is non-empty and may contain one or multiple items. Each item in the order is requested for only one unit of inventory. In other words, $2^n - 1$ types of orders exist in total. Each order can be fulfilled by using either (i) only the FDC, (ii) only the RDC, or (iii) both (i.e., an order split occurs in this case). Note an order that requests k items can be split up to 2^k possible ways. The fulfillment decision is made right after receiving each order, and orders cannot be batched. We assume all the orders have the same delivery-time guarantee (e.g., two-day shipping for Amazon Prime members), and we do not consider orders with urgent deadlines (e.g., orders with two-hour delivery windows), because this type of rush order is

usually fulfilled exclusively by a local dedicated facility. The objective is to minimize the sum of two types of fulfillment costs: fixed costs and variable costs. A fixed fulfillment cost f_R (f_F) is incurred every time the RDC (FDC) is used to fulfill an order, regardless of how many items are shipped for that order. In addition, a variable fulfillment cost c_R (c_F) is incurred for each unit of item fulfilled by the RDC (FDC). We assume the fixed costs are the same for all the orders, and the variable costs are the same for all the items. The assumption of homogeneous variable costs is reasonable because we primarily focus on small-sized items that can be consolidated into one package. Order split is not an issue for median or large-sized items because they are usually packed and shipped in a separate box (e.g., laptops, boxes of diapers).

In an adversarial setting, the e-retailer has no knowledge of future demands, such as the total number of orders or order sequence. The adversary is allowed to choose an arbitrary sequence of orders and present these orders to the e-retailer over time. We conduct a competitive analysis and use *competitive ratio* as our performance measure, which is defined as the worst-case ratio of the cost incurred by an online algorithm² (denoted as ALG) to the cost generated by an optimal offline algorithm (denoted as OPT) that knows a priori the total number of orders T and the exact order sequence (o_1, o_2, \dots, o_T) . Here, each o_i is the set of items in the i -th order. More specifically, the competitive ratio is defined as follows:

$$\text{RATIO}(\text{ALG}) \triangleq \max_{n, I_{0,i}} \max_{T \geq 1} \max_{o_1, o_2, \dots, o_T} \frac{\text{ALG}(o_1, o_2, \dots, o_T)}{\text{OPT}(o_1, o_2, \dots, o_T)}. \quad (3.1)$$

In the above definition, the online algorithm cannot see any future orders, whereas the optimal offline algorithm knows exactly (o_1, o_2, \dots, o_T) in advance. They both have the same knowledge of the costs (f_R, f_F, c_R, c_F) , the number of items n , and the initial inventory levels $\{I_{0,i}\}_{i=1}^n$. In addition, the way we define $\text{RATIO}(\text{ALG})$ is such that it depends on the cost parameters (f_R, f_F, c_R, c_F) because we want to see how it changes with respect to

2. In this paper, we use the words "algorithm," "policy," and "heuristic" interchangeably.

(w.r.t.) (f_R, f_F, c_R, c_F) . When no ambiguity exists, we sometimes suppress the dependence on the parameters. We call an online algorithm a -competitive if $\text{RATIO}(\text{ALG}) \leq a$. Our objective is to develop an efficient online algorithm that minimizes the competitive ratio, which is at least one by definition.

Because we do not impose any conditions on f_R, f_F, c_R, c_F , we consider the following four cases with different cost assumptions separately:

A $f_R \leq f_F, c_R > c_F$;

B $f_R > f_F, c_R > c_F$;

C $f_R > f_F, c_R \leq c_F$;

D $f_R \leq f_F, c_R \leq c_F$.

Note the last case above is trivial, because it is optimal to exclusively fulfill from the RDC such that the competitive ratio of a myopic policy is one. The rest of the paper focuses on the cases under Assumptions (A), (B), and (C).

In practice, different implementations of picking and delivery systems may fit different assumptions. In particular, we argue the implementation of Amazon’s FCs may fit Assumption (C). First, fixed costs primarily come from outbound shipping, that is, transporting packages from distribution centers to customers. In addition, outbound shipping costs are mainly determined by the distances between distribution centers and customers. FDCs are closer to the end consumers, and thus, the assumption that $f_R > f_F$ for a given shipping mode is reasonable. We also estimate f_R and f_F in Section 3.4.3 by using FedEx shipping rates. Second, a major portion of variable costs comes from the picking and packing steps within DCs. RDCs are usually equipped with advanced picking/sorting systems that lead to a lower variable cost of processing and packing products compared with FDCs, which usually rely on the intensive use of human pickers (Berg and Knights, 2021). In addition, the labor costs of workers in RDCs are lower than the ones in FDCs, because RDCs are regional and

often located where labor costs are low, which is not the case for FDCs that are closer to the customer in high-density areas where real estate is high and labor costs of workers are high. Hence, the assumption that $c_R \leq c_F$ is reasonable.

Assumptions (A) and (B) can be relevant as well. For example, in some businesses, outbound shipping costs dominate warehousing costs (or warehousing costs are simply not a concern), and the business goal is to minimize order split. This goal translates to reducing the total fixed cost (and setting the variable costs to almost zero) in our model. These business scenarios would fit Assumption (B). Note that in the special case in which both variable costs are zero, Assumptions (B) and (C) coincide and are relevant at the same time. Finally, let us talk about Assumptions (A). It is not uncommon for companies to interpret fixed and variable costs differently. For example, for many warehouses in Alibaba’s logistic network, one of the most important business constraints is the daily upper limit on the number of packages that each warehouse can process. FDCs’ processing capacities are often binding, especially during shopping festivals such as June 18 and November 11. The fixed cost in this context can be interpreted as the shadow price associated with the constraint on the maximum daily number of orders a distribution center can process and ship. Because FDCs are often more constrained than RDCs, it leads to $f_F \geq f_R$. Regarding the variable costs, since the picking/packing systems in many warehouses in China are similar, the cost components from the warehousing part are similar between RDCs and FDCs. Thus, $c_F \leq c_R$ because FDCs are closer to the end consumers and have lower shipping costs.

3.2.1 Myopic policy

We propose the following myopic policy (denoted as Myopic) that minimizes the fulfillment cost of each arriving order. It is a greedy policy that does not take into consideration future fulfillment costs. For convenience of analysis, we also specify a fulfillment rule when there is a tie among different fulfillment options that incur the same cost.

The following definition is critical to our main theorems as well as the proofs: for any $f_R, f_F, c_R, c_F > 0$,

$$\alpha \triangleq \begin{cases} \left\lfloor \frac{f_F - f_R}{c_R - c_F} \right\rfloor + 1 & \text{if } \frac{f_F - f_R}{c_R - c_F} > 0; \\ 1 & \text{if } \frac{f_F - f_R}{c_R - c_F} \leq 0. \end{cases} \quad (3.2)$$

Here, we allow $\alpha = +\infty$ when $c_R = c_F$. Under Assumption (A) or (B), we have

$$\begin{aligned} f_F + \alpha' c_F &\geq f_R + \alpha' c_R && \text{for all } 1 \leq \alpha' \leq \alpha - 1, \\ f_F + \alpha' c_F &< f_R + \alpha' c_R && \text{for all } \alpha' \geq \alpha. \end{aligned} \quad (3.3)$$

Here, α can be interpreted as an order-size threshold such that Myopic fulfills the order exclusively from the FDC (RDC) if the number of items in that order is greater (smaller) than this threshold, assuming all the requested items are available at the FDC. In other words, under Assumption (A) or (B), orders requesting more items are more likely to be fulfilled by the FDC and orders requesting fewer items are more likely to be fulfilled by the RDC under Myopic. When there is a tie between the two fulfillment options that incur the same cost, namely, $f_F + (\alpha - 1)c_F = f_R + (\alpha - 1)c_R$, Myopic chooses the one that uses the RDC.

Similarly, under Assumption (C), we have

$$\begin{aligned} f_F + \alpha' c_F &\leq f_R + \alpha' c_R && \text{for all } 1 \leq \alpha' \leq \alpha - 1, \\ f_F + \alpha' c_F &> f_R + \alpha' c_R && \text{for all } \alpha' \geq \alpha. \end{aligned} \quad (3.4)$$

Note both inequalities have opposite directions compared with equation (3.3), due to the opposite sign of $c_F - c_R$. Similarly, α can be interpreted as an order-size threshold such that Myopic fulfills the order exclusively from the FDC (RDC) if the number of items in that order is smaller (greater) than this threshold, assuming all the requested items are available at the FDC. When there is a tie between the two fulfillment options that incur the same

cost, namely, $f_F + (\alpha - 1)c_F = f_R + (\alpha - 1)c_R$, Myopic chooses the one that uses the FDC. Myopic under Assumption (C) also aligns with the setting in which demands are sporadic and the order sizes are small (e.g., orders from Amazon Prime members); those small-sized orders should be mostly fulfilled by the FDC.

Note that when the FDC has sufficient inventory, namely, all the requested items are available at the FDC, it is always not optimal to split the order. This is formally stated below.

Observation 1. *Under Assumption (A) or (B), Myopic uses only one of the two DCs when the inventory is sufficient at the FDC.*

We additionally make a similar observation under Assumption (C).

Observation 2. *Under Assumption (C), both Myopic and OPT always use only one of the two DCs for any order.*

The intuition behind Observation 2 is that if we must use the RDC, then because it has infinite inventory and $c_R \leq c_F$, just fulfilling the entire order from the RDC is better. Otherwise, the RDC is not used and everything is fulfilled from the FDC. Compared with Observation 1, Observation 2 does not require the inventory be sufficient at the FDC and even holds when stockouts occur. In addition, Observation 2 holds for both Myopic and OPT. Observation 2 holds for OPT because splitting an order is more costly than exclusively using the RDC (which does not consume the FDC inventory either).

When $c_R > c_F$, namely, under Assumption (A) or (B), we additionally define

$$\beta \triangleq \left\lceil \frac{f_F}{c_R - c_F} \right\rceil + 1. \quad (3.5)$$

It immediately follows that

$$\begin{aligned} f_F + \beta' c_F &\geq \beta' c_R \quad \text{for all } 1 \leq \beta' \leq \beta - 1, \\ f_F + \beta' c_F &< \beta' c_R \quad \text{for all } \beta' \geq \beta. \end{aligned} \tag{3.6}$$

Similar to the order-size threshold α that is designed for when the FDC inventory is sufficient, β is another order-size threshold for when FDC stockouts occur. More specifically, when a subset of the requested items in an order are out of stock at the FDC, the RDC must be used. Myopic's fulfillment decision depends on the number of remaining items. If the number of remaining items is less than $\beta - 1$, Myopic fulfills the order exclusively from the RDC; otherwise, Myopic splits the order: out-of-stock items are fulfilled by the RDC and the rest are fulfilled by the FDC. When there is a tie between the two fulfillment options that incur the same cost, namely, $f_F + (\beta - 1)c_F = (\beta - 1)c_R$, Myopic chooses the one that uses only the RDC. Moreover, one can easily see

$$\beta \geq \alpha. \tag{3.7}$$

This relation is intuitive because the RDC is always utilized and a fixed cost f_R is always incurred when stockouts occur, and thus, the threshold β is no smaller than the non-stockout counterpart α .

3.2.2 Main results

In this section, we present our main results. We refer to Table 3.4 for a summary of all the main theoretical results in this paper. Define $x^+ \triangleq \max\{x, 0\}$, $x \vee y \triangleq \max\{x, y\}$, $x \wedge y \triangleq \min\{x, y\}$, $\mathbb{I}(\cdot)$ as the indicator function, \mathbb{Z} as the set of integers, and $\Theta(g)$ as the big theta of g (meaning it is asymptotically bounded both above and below by g). We first present an upper bound on the competitive ratio of Myopic under Assumption (A) or (B).

Table 3.4: Summary of parameter-dependent bounds. The second and third rows summarize the upper and lower bounds on $\text{RATIO}(\text{Myopic})$, respectively; the last row summarizes the lower bounds on the ratio of any online algorithm (including deterministic and randomized).

	Assumption (A) $f_R \leq f_F, c_R > c_F$	Assumption (B) $f_R > f_F, c_R > c_F$	Assumption (C) $f_R > f_F, c_R \leq c_F$
UB on $\text{RATIO}(\text{Myopic})$	$1 + \left(\frac{f_R}{f_F + \alpha c_F} \vee \frac{f_F}{f_R + \alpha c_R} \right)^\diamond$ (Theorem 4)	$1 + \frac{f_R}{f_F + c_F}$ (if $\beta \geq 2$) $1 + \frac{f_R + f_F}{f_F + c_F + c_R}$ (if $\beta = 1$) (Theorems 4, 5)	$1 + \left(\frac{(f_R - f_F) - (c_F - c_R)}{f_F + c_F} \vee 0 \right)$ (Theorem 6)
LB on $\text{RATIO}(\text{Myopic})$	$1 + \left(\frac{f_R}{f_F + (\beta + 1)c_F + c_R} \vee \frac{f_F - f_R}{f_R + \alpha(c_F + c_R)} \right)^\heartsuit$ (Theorem 7)	$1 + \left(\frac{f_R - f_F}{f_F + c_F + c_R} \vee \frac{f_F}{f_R + f_F + c_R + (\beta + 1)c_F} \right)$ (Theorem 7)	$1 + \left(\frac{(\alpha - 2)(f_R - f_F)}{f_R + (\alpha - 1)(f_F + c_R + c_F)} \vee 0 \right)$ (Theorem 7)
LB on $\text{RATIO}(\text{rand. alg.})$	$1 + \left(\frac{f_F + (\beta + 1)c_F + c_R}{f_R} + \frac{c_F}{c_R - c_F} \right)^{-1}$ (Theorem 8)	$1 + \sup_{n \geq 2, n \in \mathbb{Z}} \left(\frac{f_F + nc_F}{f_R - f_F + n(c_R - c_F)} + \frac{f_R + n(f_F + c_R + c_F)}{(n - 1)(f_R - f_F)} \right)^{-1}$ (Theorem 8)	$1 + \max_{\alpha' \in \{1, \dots, \alpha - 1\}} \left(\frac{f_R + \alpha'(f_F + c_F + c_R)}{(\alpha' - 1)(f_R - f_F)} + \frac{f_F + \alpha' c_F}{f_R - f_F - \alpha'(c_F - c_R)} \right)^{-1}$ (Theorem 8)

\diamond This upper bound is always no greater than 2 (Corollary 1).

\heartsuit This lower bound converges to 2 under a certain sequence of cost parameters (Corollary 2).

\clubsuit The optimal n is explicitly characterized in the proof of Corollary 3.

\spadesuit The optimal α' is explicitly characterized in the proof of Corollary 3.

Theorem 4. Under Assumption (A) or (B),

$$\text{RATIO}(\text{Myopic}) \leq 1 + \left(\frac{f_R}{f_F + \alpha c_F} \vee \frac{f_F}{f_R + \alpha c_R} \right). \quad (3.8)$$

In addition, under Assumption (B), the above upper bound is simplified to $1 + f_R/(f_F + c_F)$.

Under Assumption (B), we can further improve the bound for the case in which $\beta = 1$ (equivalent to $f_F + c_f < c_R$), which is useful for proving the tightness of our bounds on $\text{RATIO}(\text{Myopic})$ in Theorem 9. More specifically, the upper bound in Theorem 4 is too loose to derive the results in Theorem 9 when $\beta = 1$ and we need the following improved upper bound.

Theorem 5. Under Assumption (B), when $\beta = 1$,

$$\text{RATIO}(\text{Myopic}) \leq 1 + \frac{f_R + f_F}{f_F + c_F + c_R}.$$

We next present the upper bound under Assumption (C).

Theorem 6. *Under Assumption (C),*

$$RATIO(Myopic) \leq 1 + \left(\frac{(f_R - f_F) - (c_F - c_R)}{f_F + c_F} \vee 0 \right).$$

Note that the above competitive ratios are parameter-dependent bounds. This follows a growing trend of literature in operations management (e.g., Ball and Queyranne, 2009, Jasin and Sinha, 2015, Andrews et al., 2019, Ma and Simchi-Levi, 2020, Ma et al., 2021) deriving parameter-dependent bounds that provide more managerial insights.

We also provide lower bounds on $RATIO(Myopic)$.

Theorem 7.

$$RATIO(Myopic) \geq \begin{cases} 1 + \left(\frac{f_R}{f_F + (\beta+1)c_F + c_R} \vee \frac{f_F - f_R}{f_R + \alpha(c_F + c_R)} \right) & \text{under Assumption (A);} \\ 1 + \left(\frac{f_R - f_F}{f_F + c_F + c_R} \vee \frac{f_F}{f_R + f_F + c_R + (\beta+1)c_F} \right) & \text{under Assumption (B);} \\ 1 + \left(\frac{(\alpha-2)(f_R - f_F)}{f_R + (\alpha-1)(f_F + c_R + c_F)} \vee 0 \right) & \text{under Assumption (C).} \end{cases}$$

Finally, we provide lower bounds on competitive ratios of any online randomized algorithm.

Theorem 8. *The competitive ratio of any online randomized algorithm cannot be less than*

$$\begin{cases} 1 + \left(\frac{f_F + (\beta+1)c_F + c_R}{f_R} + \frac{c_F}{c_R - c_F} \right)^{-1} & \text{under Assumption (A);} \\ 1 + \sup_{n \geq 2, n \in \mathbb{Z}} \left(\frac{f_F + nc_F}{f_R - f_F + n(c_R - c_F)} + \frac{f_R + n(f_F + c_R + c_F)}{(n-1)(f_R - f_F)} \right)^{-1} & \text{under Assumption (B);} \\ 1 + \max_{\alpha' \in \{1, \dots, \alpha-1\}} \left(\frac{f_R + \alpha'(f_F + c_F + c_R)}{(\alpha'-1)(f_R - f_F)} + \frac{f_F + \alpha'c_F}{f_R - f_F - \alpha'(c_F - c_R)} \right)^{-1} & \text{under Assumption (C).} \end{cases} \quad (3.9)$$

3.3 Discussion of the analysis

In this section, we provide a proof sketch of Theorems 4 and 6 in Section 3.3.1, and a high-level discussion of the proofs for Theorems 7 and 8 in Section 3.3.2.

3.3.1 Proof sketch of Theorems 4 and 6

We use a *potential function* argument to prove Theorems 4 and 6. The idea of such an argument is to (judiciously) design a potential function that captures the cumulative difference between two policies (i.e., our online policy and the optimal offline policy).

We first introduce some notations. We use F and R to denote parameters associated with the FDC and RDC, respectively, and m and * for parameters associated with Myopic and OPT, respectively. In our analysis, we use the words "period" and "order" interchangeably.

- We use $I_{t,i}^m$ and $I_{t,i}^*$ to denote the FDC inventory levels of item i in the end of period t under Myopic and OPT, respectively, for $i = 1, \dots, n, t \geq 0$. In particular, the initial FDC inventory levels satisfy $I_{0,i}^m = I_{0,i}^* = I_{0,i}$.
- We use $w_t^{F,m}(w_t^{R,m})$ to denote the number of times that Myopic uses the FDC (RDC) to fulfill orders from period 1 to period t . Similarly, $w_t^{F,*}(w_t^{R,*})$ denotes the number of times that OPT uses the FDC (RDC) to fulfill orders from period 1 to period t .
- We use o_t to denote the set of requested items in period t . We additionally define $N_t \triangleq |o_t|$, meaning the total number of requested items in the order t .
- We use L_t^m and L_t^* to denote the numbers of items fulfilled from the FDC in period t under Myopic and OPT, respectively.
- We use $V_t^m = V_t^m(o_1, \dots, o_T)$ and $V_t^* = V_t^*(o_1, \dots, o_T)$ to denote the total costs incurred from period 1 to period t given the order sequence (o_1, \dots, o_T) under Myopic

and OPT, respectively. From the definitions,

$$\begin{aligned}
V_t^m &= \left(\sum_{i=1}^n (I_{0,i}^m - I_{t,i}^m) \right) c_F + w_t^{F,m} f_F \\
&\quad + \left(\sum_{j=1}^t N_j - \sum_{i=1}^n (I_{0,i}^m - I_{t,i}^m) \right) c_R + w_t^{R,m} f_R, \\
V_t^* &= \left(\sum_{i=1}^n (I_{0,i}^* - I_{t,i}^*) \right) c_F + w_t^{F,*} f_F \\
&\quad + \left(\sum_{j=1}^t N_j - \sum_{i=1}^n (I_{0,i}^* - I_{t,i}^*) \right) c_R + w_t^{R,*} f_R,
\end{aligned} \tag{3.10}$$

where the terms $\sum_{i=1}^n (I_{0,i}^* - I_{t,i}^*)$ and $\sum_{j=1}^t N_j - \sum_{i=1}^n (I_{0,i}^* - I_{t,i}^*)$ represent the total number of items shipped from period 1 to period t from the FDC and RDC, respectively. In addition, the difference $V_t^* - V_{t-1}^*$ is simply the added fulfillment cost of OPT in period t :

$$V_t^* - V_{t-1}^* = L_t^* c_F + (w_t^{F,*} - w_{t-1}^{F,*}) f_F + (N_t - L_t^*) c_R + (w_t^{R,*} - w_{t-1}^{R,*}) f_R. \tag{3.11}$$

We begin to prove Theorem 4. The proof contains two key components in a high level:

(1) We construct a potential function that captures the difference between V_t and V_t^* ; (2) we prove this potential function is bounded from above by $\theta \cdot V_t^*$ (θ to be specified), such that $1 + \theta$ provides an upper bound on $\text{RATIO}(\text{Myopic})$. We present the proof in several steps.

Step 1: potential function. We construct the following potential function F_t :

$$F_t \triangleq \left(\sum_{i=1}^n (I_{t,i}^m - I_{t,i}^*)^+ \right) (c_R - c_F) + (w_t^{F,m} - w_t^{F,*}) f_F + (w_t^{R,m} - w_t^{R,*}) f_R \quad \text{for } t \geq 1. \tag{3.12}$$

The first term of F_t captures the incremental variable cost when Myopic uses the RDC

compared with OPT, and the second (third) term is the cumulative FDC (RDC) fixed-cost difference between Myopic and OPT up to period t . In particular, the first term $\left(I_{t,i}^m - I_{t,i}^*\right)^+$ takes the positive part of the inventory difference and eliminates potentially large negative values when Myopic has a relatively lower FDC inventory level than OPT (i.e., $I_{t,i}^m < I_{t,i}^*$). Because we eventually bound the difference $F_t - F_{t-1}$ (see (3.15) below), this special trick decreases its variability in each period against the worst-case instance, so that we are able to achieve the desired upper bound. Without this trick, the upper bound would depend on N_t .

To prove Theorem 4, it suffices to show that for any $T \geq 1$,

$$V_T^m \leq V_T^* + F_T, \quad (3.13)$$

$$F_T \leq \left(\frac{f_R}{f_F + \alpha c_F} \vee \frac{f_F}{f_R + \alpha c_R} \right) V_T^*. \quad (3.14)$$

From (3.10),

$$\begin{aligned} V_T^m - V_T^* &= \left(\sum_{i=1}^n (I_{T,i}^m - I_{T,i}^*) \right) (c_R - c_F) + (w_T^{F,m} - w_T^{F,*}) f_F + (w_T^{R,m} - w_T^{R,*}) f_R \\ &\leq \left(\sum_{i=1}^n \left(I_{T,i}^m - I_{T,i}^* \right)^+ \right) (c_R - c_F) + (w_T^{F,m} - w_T^{F,*}) f_F + (w_T^{R,m} - w_T^{R,*}) f_R \\ &= F_T, \end{aligned}$$

completing the proof of (3.13). It remains to prove (3.14), the key component of the proof. We additionally define $V_0^m = V_0^* = 0$ and $F_0 = 0$. It suffices to show that for all $1 \leq t \leq T$, we have

$$F_t - F_{t-1} \leq \left(\frac{f_R}{f_F + \alpha c_F} \vee \frac{f_F}{f_R + \alpha c_R} \right) (V_t^* - V_{t-1}^*). \quad (3.15)$$

We defer the remainder of the proof to the appendix Section 3.8.4. The proof of Theorem

6 follows a similar procedure, by constructing a potential function without $(\cdot)^+$ on the first term. It is detailed in Section 3.8.6.

3.3.2 High-level proof discussion of Theorems 7 and 8

In the proof of Theorem 7, we essentially construct several instances (more specifically, order sequences) and calculate the corresponding ratios of the costs under Myopic to the costs under OPT. A typical instance has a large order arriving at first, followed by a sequence of small orders. Myopic uses the FDC to fulfill the first order, resulting in FDC stockouts and order split for the remaining orders.

In the proof of Theorem 8, we use Yao's minimax principle (Yao, 1977) to derive the desired lower bounds. It essentially says the competitive ratio of any randomized algorithm that has a lack of knowledge about the arrival sequence is at least the ratio of the best deterministic algorithm that knows the probability distribution over the order sequence for any distribution. As a consequence, instead of bounding the competitive ratio for any randomized algorithm, it suffices to construct a probability distribution over the order sequence and then bound the expected competitive ratio for any deterministic algorithm. Similar to the proof of Theorem 7, we also need to construct several instances. The difference is that here, the instances are probabilistic, namely, assigning probabilities to deterministic instances. A typical probabilistic instance looks as follows: a large order arrives at first with probability one; with a certain probability, a sequence of small orders may appear later.

3.4 Implications of the main results & additional results

In this section, we discuss the implications of our parameter-dependent bounds in Sections 3.4.1 and 3.4.2, and provide additional results in Section 3.4.3.

3.4.1 Implications under Assumption (A)

First, as a consequence of Theorem 4, we obtain the following constant upper bound under Assumption (A), which says the competitive ratio of the myopic policy is at most 2.

Corollary 1. *Under Assumption (A), $RATIO(Myopic) \leq 2$.*

Second, as a consequence of Theorem 8, we obtain the following constant lower bound under Assumption (A), which says the competitive ratio of any online algorithm (deterministic or randomized) is at least 2.

Corollary 2. *Under Assumption (A), a sequence of $\{f_R, f_F, c_R, c_F\}$ exists such that the lower bound in equation (3.9) converges to 2.*

Combining Corollaries 1 and 2, we conclude our bound on the competitive ratio of the myopic policy is tight under Assumption (A).

3.4.2 Implications under Assumptions (B) and (C)

Next, we look at cases under Assumptions (B) and (C). From Theorems 4-7, we see the upper and lower bounds on $RATIO(Myopic)$ both grow in f_R under both Assumptions (B) and (C). In other words, Myopic performs poorly compared with OPT for large f_R . This is not unique to Myopic. In fact, we show that no algorithm can achieve a constant competitive ratio in this case, as stated in the next result.

Corollary 3. *Under Assumptions (B) and (C), the lower bound in Theorem 8 grows on the order of $\sqrt{f_R}$ as $f_R \rightarrow \infty$.*

We also investigate the tightness of our bounds on $RATIO(Myopic)$ and show the ratio of the upper bound to the lower bound is always no greater than 2.

Theorem 9. *Under Assumptions (B) and (C), the ratio of the upper bound on $RATIO(Myopic)$ to the lower bound on $RATIO(Myopic)$ is no greater than 2. In addition, it is tight; that is, a sequence of $\{f_R, f_F, c_R, c_F\}$ exists such that the ratio converges to 2.*

3.4.3 Additional results: a practical upper bound with estimations under Assumption (C)

So far, we have proved that the upper bound on $\text{RATIO}(\text{Myopic})$ is no greater than 2 under Assumption (A), and the ratio of the upper to lower bounds on $\text{RATIO}(\text{Myopic})$ is no greater than 2 under Assumptions (B) and (C), demonstrating the tightness of our bounds on $\text{RATIO}(\text{Myopic})$. However, it does not really answer the question of how well Myopic performs compared with other online algorithms under Assumptions (B) and (C). We answer this question for a special case in which $c_R = c_F = 0$, namely, when the objective is to minimize the weighted fixed costs (or equivalently weighted number of split orders). Note that in this case, Assumptions (B) and (C) are equivalent (both requiring $f_R > f_F$).

When $c_R = c_F = 0$, Myopic entirely uses the FDC when all the requested items are available, and entirely uses the RDC otherwise. This observation coincides with the case under Assumption (C) such that we can simply apply Theorem 6 to obtain an upper bound on $\text{RATIO}(\text{Myopic})$: $\frac{f_R}{f_F}$. This bound is also intuitive: Myopic uses only one of the DCs and never splits an order, and thus, its cost is at most f_R for each order. Meanwhile, OPT incurs at least f_F for each order, leading to the bound $\frac{f_R}{f_F}$. Note it is trivial to prove an upper bound of $1 + \frac{f_R}{f_F}$ for any online algorithm, because any algorithm incurs at most $f_R + f_F$ for any order. We next prove the upper bound $\frac{f_R}{f_F}$ matches a lower bound of any online deterministic algorithm, and this upper bound is also no greater than 2 times a lower bound on the ratio of any online randomized algorithm. Note the tightness of the bound $\frac{f_R}{f_F}$ is less obvious, because the lower bound could be 1.

Theorem 10. *Suppose $c_R = c_F = 0$ and $f_R > f_F$. The competitive ratio of any online deterministic algorithm cannot be less than $\frac{f_R}{f_F}$, and the competitive ratio of any online randomized algorithm cannot be less than $\frac{1}{2} \left(\frac{f_R}{f_F} + 1 \right)$. As an immediate consequence, the ratio of the upper bound on $\text{RATIO}(\text{Myopic})$ matches the lower bound on the ratio of any online deterministic algorithm, and the ratio of the upper bound on $\text{RATIO}(\text{Myopic})$ to the*

lower bound on the ratio of any online randomized algorithm is no greater than 2.

For general cases with $c_R, c_F > 0$, we suspect similar results still hold and leave it as an interesting future research question. Note that the variable costs c_F, c_R in the bound of Theorem 6 are usually hard to estimate because they involve many components, including labor costs. By contrast, the fixed costs f_F, f_R are much easier to estimate because they primarily come from outbound shippings. The following observation provides an upper bound on the ratio that is independent of the variable costs. We also use practical data to numerically evaluate this upper bound.

Observation 3. *Under Assumption (C), the upper bound $1 + \left(\frac{(f_R - f_F) - (c_F - c_R)}{f_F + c_F} \vee 0 \right)$ on $RATIO(Myopic)$ in Theorem 6 is no greater than $\frac{f_R}{f_F}$. Namely, the ratio $\frac{f_R}{f_F}$ is a universal upper bound on $RATIO(Myopic)$ under Assumption (C).*

Observation 3 has an important implication: we can focus on the ratio $\frac{f_R}{f_F}$ and ignore the variable costs. To have a sense of how large the ratio is, we use the FedEx Ground shipping rates (FedEx, 2021). Note that ground shipping is a cost-effective option heavily utilized by e-retailers. The shipping rates are zone based. For instance, Zone 2 includes shipments moving within 150 miles from origin to destination, Zone 3 includes shipments moving 151-300 miles, and Zone 8 includes shipments moving 1,801 miles or more. For each zone, we have the shipping rates for a package of different weights and fit the following linear model: $\text{cost}(n) = \beta_0 + \beta_1 n$, where n represents the package weight. Because the shipping cost is usually piecewise linear and concave in n and is not necessarily linear for a large range of n , we restrict the range to weights up to 10 pounds (due to the fact that most online retailing packages weigh less than 10 pounds). Hence, we use $n \in \{1, 2, \dots, 10\}$ to fit a linear model, and β_0 is interpreted as the fixed cost. All the rates are presented in Table 3.7 in the appendix. To match the practice, we choose the closest zone (Zone 2) as the location of the FDC, leading to the value of f_F , and treat the other zones (Zones 3-8) as possible RDC locations, leading to a range of values of f_R . For instance, a shipment to Chicago

Parameter	Zone	Parameter Estimate (\$)	R^2	$\frac{f_R}{f_F}$
f_F	2	8.88	98.41%	—
f_R	3	10.05	98.01%	1.13
	4	10.92	97.50%	1.23
	5	11.17	95.63%	1.26
	6	11.61	95.92%	1.31
	7	11.66	97.86%	1.31
	8	11.65	97.42%	1.31

Table 3.5: Fixed-cost estimates: Ground shipping

Downtown is in Zone 2 when from Channahon (Illinois), Zone 3 when from Indianapolis, and Zone 5 when from New York. Table 3.5 summarizes the estimated f_F from Zone 2 and f_R 's from Zones 3-8 along with the corresponding R^2 values. The last column demonstrates the values of the ratio $\frac{f_R}{f_F}$. In particular, this ratio is 1.13 when the RDC is located in Zone 3, and is 1.31 when the RDC is located in Zone 8 (the farthest possible in the contiguous US). Given that the RDC is most likely not too far way from customers in the associated region (e.g., Zone 3 is within 300 miles and Zone 4 is within 600 miles), we conclude the ratio $\frac{f_R}{f_F}$ is small, which implies that Myopic can be very close to the optimal algorithm in practice. This analysis provides a more intuitive picture and sheds light on the values of the competitive ratio in practice. We also analyze the practical ratio based on rates for FedEx Standard Overnight (one-day shipping) and defer the results to the appendix Section 3.8.1.

3.4.4 Summary

In conclusion, we summarize the implications/results in Table 3.6. Under Assumption (A), Myopic is 2-competitive, and we also prove it is tight in the sense that no online algorithm exists that can do strictly better than 2. Under Assumption (B) or (C), we prove that obtaining a parameter-independent bound is no longer possible and the competitive ratio of any online algorithm grows with f_R . We also investigate the tightness of our bounds and prove the ratio of the upper bound on $\text{RATIO}(\text{Myopic})$ to the lower bound on $\text{RATIO}(\text{Myopic})$ is

Table 3.6: Summary of implications/results in Section 3.4.

$c_R, c_F > 0$	Assumption (A) $f_R \leq f_F, c_R > c_F$	Assumption (B) $f_R > f_F, c_R > c_F$	Assumption (C) $f_R > f_F, c_R \leq c_F$
<i>UB on RATIO(Myopic)</i>	2	$\Theta(f_R)^\diamond$	$\Theta(f_R)$
<i>LB on RATIO(Myopic)</i>	2	$\Theta(f_R)$	$\Theta(f_R)$
<i>LB on RATIO(any rand. alg.)</i>	2	$\Theta(\sqrt{f_R})$	$\Theta(\sqrt{f_R})$
$\frac{UB \text{ on } RATIO(Myopic)}{LB \text{ on } RATIO(Myopic)}$	–	≤ 2	≤ 2
$c_R, c_F = 0$	Assumptions (B) and (C): $f_R > f_F$		
<i>UB on RATIO(Myopic)</i>	f_R/f_F		
<i>LB on RATIO(any det. alg.)</i>	f_R/f_F		
<i>LB on RATIO(any rand. alg.)</i>	$\frac{1}{2} \left(\frac{f_R}{f_F} + 1 \right)$		

\diamond Here, we emphasize the dependence on f_R and suppress its dependence on other cost parameters.

always no greater than 2. In the special case in which $c_R = c_F = 0$, the ratio of the upper bound on RATIO(Myopic) matches the lower bound on the ratio of any online deterministic algorithm, and the ratio of the upper bound on RATIO(Myopic) to the lower bound on the ratio of any online randomized algorithm is no greater than 2. Last, under Assumption (C), which fits Amazon’s business arguably, the competitive ratio RATIO(Myopic) can be upper-bounded by $\frac{f_R}{f_F}$ (Observation 3) and we are able to estimate this ratio by using FedEx Ground Shipping rates (Table 3.5).

The good performance of Myopic leads to the question of whether other myopic- (greedy-) type policies have similar performance in this RDC-FDC network, for example, the policy that fulfills as many items from the FDC as possible (even at the cost of order split) and the policy that minimizes order split as much as possible (even at the cost of fulfilling more items from the RDC). Unfortunately, the performance of both policies can be arbitrarily poor. The intuition is that Myopic has the flexibility of adjusting its behavior and evolving into a different policy under different cost parameters. For example, Myopic always uses the FDC inventory first whenever possible under Assumption (B) and $\beta = 1$, whereas Myopic minimizes order split under Assumption (C). Such flexibility does not exist in other greedy-type policies aforementioned. We defer more details to the appendix Section 3.8.3.

3.5 Extension: When Demand Forecasts Are Available

In this section, we extend our study to the setting in which demand forecasts are available and investigate a stochastic counterpart of Problem (3.1). Specifically, assume there are N order types indexed by \mathcal{J} . Each order of type \mathcal{J} consists of a unique (non-empty) combination of items. Here, $N \leq 2^n - 1$. We write $i \in \mathcal{J}$ if order type \mathcal{J} contains item i . There are T time periods. In each period, at most one order arrives. Let $D^{\mathcal{J},t} \in \{0, 1\}$ be the demand of order type \mathcal{J} in period t , namely, $D^{\mathcal{J},t} = 1$ if an order of type \mathcal{J} arrives in period t and 0 otherwise. We assume demands are independent and identically distributed (i.i.d.) and let $\lambda^{\mathcal{J}} \triangleq \mathbb{E} [D^{\mathcal{J},t}]$, where $\sum_{\mathcal{J}} \lambda^{\mathcal{J}} = 1$. We also assume the variable costs can be item-dependent: $c_{R,i}$ and $c_{F,i}$. Then, the stochastic version of the fulfillment problem can be formulated as below, which largely follows the setting studied in Jasin and Sinha (2015):

$$C^*(T) = \min \sum_{t=1}^T \sum_{\mathcal{J}} \mathbb{E} \left[D^{\mathcal{J},t} \left(\sum_{i \in \mathcal{J}} c_{R,i} X_{R,i}^{\mathcal{J},t} + f_R \max_{i \in \mathcal{J}} \{ X_{R,i}^{\mathcal{J},t} \} \right. \right. \\ \left. \left. + \sum_{i \in \mathcal{J}} c_{F,i} X_{F,i}^{\mathcal{J},t} + f_F \max_{i \in \mathcal{J}} \{ X_{F,i}^{\mathcal{J},t} \} \right) \right] \\ \text{s.t. } \sum_{t=1}^T \sum_{\mathcal{J} \ni i} D^{\mathcal{J},t} X_{F,i}^{\mathcal{J},t} \leq I_{0,i} \quad \forall i \quad (3.16)$$

$$X_{F,i}^{\mathcal{J},t} + X_{R,i}^{\mathcal{J},t} = 1 \quad \forall \mathcal{J}, i, t \quad (3.17)$$

$$X_{F,i}^{\mathcal{J},t}, X_{R,i}^{\mathcal{J},t} \in \{0, 1\} \quad \forall \mathcal{J}, i, t, \quad (3.18)$$

where the decision variable $X_{F,i}^{\mathcal{J},t}$ ($X_{R,i}^{\mathcal{J},t}$) indicates if the e-retailer fulfills item i in order type \mathcal{J} from the FDC (RDC) at time t . Constraint (3.16) is a set of inventory constraints and (3.17) ensures all requested items are always fulfilled. Constraints (3.16) and (3.17) must hold with probability one. Note the values of $X_{F,i}^{\mathcal{J},t}$ and $X_{R,i}^{\mathcal{J},t}$ are irrelevant if order type \mathcal{J} does not appear in period t . We next discuss an LP relaxation and propose an LP Rounding scheme.

Recall that $\lambda^{\mathcal{J}}$ is the demand rate for order type \mathcal{J} . By replacing the random variables by their means and removing the integrality constraint (3.18) in the aforementioned formulation, we obtain the following lower bound on $C^*(T)$:

$$\begin{aligned}
\min \quad & \sum_{t=1}^T \sum_{\mathcal{J}} \lambda^{\mathcal{J}} \left[\sum_{i \in \mathcal{J}} c_{R,i} x_{R,i}^{\mathcal{J},t} + f_R \max_{i \in \mathcal{J}} \{x_{R,i}^{\mathcal{J},t}\} + \sum_{i \in \mathcal{J}} c_{F,i} x_{F,i}^{\mathcal{J},t} + f_F \max_{i \in \mathcal{J}} \{x_{F,i}^{\mathcal{J},t}\} \right] \\
\text{s.t.} \quad & \sum_{t=1}^T \sum_{\mathcal{J} \ni i} \lambda^{\mathcal{J}} x_{F,i}^{\mathcal{J},t} \leq I_{0,i} \quad \forall i \\
& x_{F,i}^{\mathcal{J},t} + x_{R,i}^{\mathcal{J},t} = 1 \quad \forall \mathcal{J}, i, t \\
& x_{F,i}^{\mathcal{J},t}, x_{R,i}^{\mathcal{J},t} \geq 0 \quad \forall \mathcal{J}, i, t.
\end{aligned} \tag{3.19}$$

We further linearize the $\max\{\cdot\}$ terms in the objective function (3.19) to obtain the following equivalent LP:

$$\begin{aligned}
C_{LP}^*(T) = \min \quad & \sum_{t=1}^T \sum_{\mathcal{J}} \lambda^{\mathcal{J}} \left[\sum_{i \in \mathcal{J}} c_{R,i} x_{R,i}^{\mathcal{J},t} + f_R y_R^{\mathcal{J},t} + \sum_{i \in \mathcal{J}} c_{F,i} x_{F,i}^{\mathcal{J},t} + f_F y_F^{\mathcal{J},t} \right] \\
\text{s.t.} \quad & \sum_{t=1}^T \sum_{\mathcal{J} \ni i} \lambda^{\mathcal{J}} x_{F,i}^{\mathcal{J},t} \leq I_{0,i} \quad \forall i \\
& x_{F,i}^{\mathcal{J},t} + x_{R,i}^{\mathcal{J},t} = 1 \quad \forall \mathcal{J}, i, t \\
& y_F^{\mathcal{J},t} \geq x_{F,i}^{\mathcal{J},t}, y_R^{\mathcal{J},t} \geq x_{R,i}^{\mathcal{J},t} \quad \forall \mathcal{J}, i, t \\
& x_{F,i}^{\mathcal{J},t}, x_{R,i}^{\mathcal{J},t} \geq 0 \quad \forall \mathcal{J}, i, t.
\end{aligned}$$

Let $x_{F,i}^{\mathcal{J}}$ ($x_{R,i}^{\mathcal{J}}$) denote the average number of times item i in order type \mathcal{J} is fulfilled from the FDC (RDC) during the selling horizon and let $y_F^{\mathcal{J}}$ ($y_R^{\mathcal{J}}$) denote the average number of times order type \mathcal{J} is (partially) fulfilled from the FDC (RDC) during the selling horizon.

The time-aggregate formulation is of $C_{LP}^*(T)$ is given by

$$\tilde{C}_{LP}^*(T) = \min T \cdot \sum_{\mathcal{J}} \lambda^{\mathcal{J}} \left[\sum_{i \in \mathcal{J}} c_{R,i} x_{R,i}^{\mathcal{J}} + f_R y_R^{\mathcal{J}} + \sum_{i \in \mathcal{J}} c_{F,i} x_{F,i}^{\mathcal{J}} + f_F y_F^{\mathcal{J}} \right] \quad (3.20)$$

$$\text{s.t. } T \cdot \sum_{\mathcal{J} \ni i} \lambda^{\mathcal{J}} x_{F,i}^{\mathcal{J}} \leq I_{0,i} \quad \forall i \quad (3.21)$$

$$x_{F,i}^{\mathcal{J}} + x_{R,i}^{\mathcal{J}} = 1 \quad \forall \mathcal{J}, i \quad (3.22)$$

$$y_F^{\mathcal{J}} \geq x_{F,i}^{\mathcal{J}}, y_R^{\mathcal{J}} \geq x_{R,i}^{\mathcal{J}} \quad \forall \mathcal{J}, i$$

$$x_{F,i}^{\mathcal{J}}, x_{R,i}^{\mathcal{J}} \geq 0 \quad \forall \mathcal{J}, i.$$

It is not difficult to see that $\tilde{C}_{LP}^*(T) = C_{LP}^*(T) \leq C^*(T)$ (we refer the interested reader to Jasin and Sinha (2015), for more details). We next propose an LP rounding scheme of how to round the optimal fractional solution of problem (3.20), which leads to a randomized fulfillment policy. The randomized policy aims to split orders as little as possible to match the optimal fixed costs in the lower bound (3.20).

For order type \mathcal{J} , let $m^{\mathcal{J}}$ be its order size and let $x_R^{\mathcal{J},*} = \{x_{R,i}^{\mathcal{J},*}\}_i$, $x_F^{\mathcal{J},*} = \{x_{F,i}^{\mathcal{J},*}\}_i$ be an optimal solution of problem (3.20). As a notational convenience, we drop the index \mathcal{J} for the remaining description. From constraint (3.22), $x_{F,i}^* + x_{R,i}^* = 1$ for all i . Without loss of generality, we assume $\{x_{R,i}^*\}_i$ are sorted in ascending order, namely, $x_{R,1}^* \leq \dots \leq x_{R,m}^*$. As

an immediate consequence, x_R^* can be decomposed as follows:

$$\begin{aligned}
x_R^* &= \begin{pmatrix} x_{R,1}^* \\ x_{R,2}^* \\ \dots \\ x_{R,m-1}^* \\ x_{R,m}^* \end{pmatrix} = x_{R,1}^* \cdot \begin{pmatrix} 1 \\ 1 \\ \dots \\ 1 \\ 1 \end{pmatrix} \\
&+ (x_{R,2}^* - x_{R,1}^*) \cdot \begin{pmatrix} 0 \\ 1 \\ \dots \\ 1 \\ 1 \end{pmatrix} + \dots + (x_{R,m}^* - x_{R,m-1}^*) \cdot \begin{pmatrix} 0 \\ 0 \\ \dots \\ 0 \\ 1 \end{pmatrix}.
\end{aligned} \tag{3.23}$$

Similarly, because $\{x_{F,i}^*\}_i$ are sorted in descending order, namely, $x_{F,1}^* \geq \dots \geq x_{F,m}^*$, x_F^* can also be decomposed:

$$\begin{aligned}
x_F^* &= \begin{pmatrix} x_{F,1}^* \\ x_{F,2}^* \\ \dots \\ x_{F,m-1}^* \\ x_{F,m}^* \end{pmatrix} = \begin{pmatrix} 1 - x_{R,1}^* \\ 1 - x_{R,2}^* \\ \dots \\ 1 - x_{R,m-1}^* \\ 1 - x_{R,m}^* \end{pmatrix} \\
&= (x_{R,2}^* - x_{R,1}^*) \cdot \begin{pmatrix} 1 \\ 0 \\ \dots \\ 0 \\ 0 \end{pmatrix} + \dots + (x_{R,m}^* - x_{R,m-1}^*) \cdot \begin{pmatrix} 1 \\ 1 \\ \dots \\ 1 \\ 0 \end{pmatrix} + (1 - x_{R,m}^*) \cdot \begin{pmatrix} 1 \\ 1 \\ \dots \\ 1 \\ 1 \end{pmatrix}.
\end{aligned} \tag{3.24}$$

Based on the above decompositions, we propose the following randomized policy (called LP Rounding): it fulfills all the items from the RDC with probability $x_{R,1}^*$, fulfills the first k items from the FDC and the remaining $m - k$ items from the RDC with probability $(x_{R,k+1}^* - x_{R,k}^*)$, $k = 1, \dots, m - 1$, and fulfills all the items from the FDC with probability $(1 - x_{R,m}^*)$. If an item is out of stock at the FDC, fulfill that item from the RDC.

To illustrate the above randomized fulfillment policy, consider the following toy example. Suppose the order type \mathcal{J} has four items and the associated optimal LP solutions are

$$x_R^* = \begin{pmatrix} 0 \\ 0.3 \\ 0.4 \\ 0.8 \end{pmatrix}, \quad x_F^* = \begin{pmatrix} 1 \\ 0.7 \\ 0.6 \\ 0.2 \end{pmatrix}.$$

The solutions say the first item is always fulfilled from the FDC with probability one, the second item is fulfilled from the RDC and FDC with probabilities 0.3 and 0.7, respectively, and so on. Note the two vectors can be decomposed as follows:

$$x_R^* = 0 \cdot \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \end{pmatrix} + 0.3 \cdot \begin{pmatrix} 0 \\ 1 \\ 1 \\ 1 \end{pmatrix} + 0.1 \cdot \begin{pmatrix} 0 \\ 0 \\ 1 \\ 1 \end{pmatrix} + 0.4 \cdot \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix},$$

$$x_F^* = 0.3 \cdot \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} + 0.1 \cdot \begin{pmatrix} 1 \\ 1 \\ 0 \\ 0 \end{pmatrix} + 0.4 \cdot \begin{pmatrix} 1 \\ 1 \\ 1 \\ 0 \end{pmatrix} + 0.2 \cdot \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \end{pmatrix}.$$

Then, the LP Rounding policy works as follows: with probability 0.2, fulfill all the items from the FDC; with probability 0.4, fulfill items 1, 2, 3 from the FDC and item 4 from the

RDC; with probability 0.1, fulfill items 1, 2 from the FDC and items 3, 4 from the RDC; with probability 0.3, fulfill item 1 from the FDC and items 2, 3, 4 from the RDC. If an item is out of stock at the FDC, fulfill that item from the RDC. It is not difficult to see that the above rounding scheme minimizes the order split. Because the two vectors x_R^* and x_F^* are complementary to each other in the RDC-FDC network, our rounding scheme is constructed in a way that each item is “well assigned,” as opposed to potentially “unassigned” in the rounding schemes of Jasin and Sinha (2015) that study a network consisting of multiple warehouses.

For each i , assume $I_{0,i} = \theta_i T$ for some $\theta_i > 0$. Namely, the initial inventory scales linearly in the number of orders T . We next prove the proposed LP Rounding policy is asymptotically optimal as $T \rightarrow \infty$. We use $C^{LPR}(T)$ to denote its average cost over T periods.

Theorem 11. *Assume $I_{0,i} = \theta_i T$ for each i . Then,*

$$\lim_{T \rightarrow \infty} \frac{C^{LPR}(T)}{C^*(T)} = 1.$$

The intuition behind this asymptotic optimality result is the following. By our LP rounding scheme, when no FDC stockout exists, the variable and fixed costs of LP Rounding match the ones in the LP lower bound (3.20). However, an error term appears when the total demand exceeds the supply, and this error term is on the order of $O(\sqrt{T})$ because of the strong law of large numbers. Hence, as $T \rightarrow \infty$, the error term becomes negligible compared with $C^*(T)$, which grows linearly in T .

Note Jasin and Sinha (2015) study two LP rounding schemes in the same asymptotic regime but with multiple DCs and derived parameter-dependent bounds. Their bounds can diverge as the expected order size grows. By contrast, we demonstrate the performance bound can be improved in the RDC-FDC setting, and our LP rounding policy is always asymptotically optimal regardless of the expected order size. The reason leading to this

improvement is that our “sparse” RDC-FDC network has fewer decision variables in the LP formulation such that the optimal LP solutions can be better rounded.

3.6 Numerical Experiments

In this section, we complement our theoretical results with numerical experiments. Recall that LP Rounding is designed only for settings where demand forecasting is available, whereas Myopic is applicable even when e-retailers have no access to demand information. In practice, e-retailers may have access to demand forecasts to some extent. As such, we incorporate demand forecasting into our simulations and conduct two sets of numerical experiments in Sections 3.6.1 and 3.6.2. We only present the results under Assumptions (C) and (D) here, and defer the results under Assumptions (A) and (B) to the appendix Section 3.8.13 (the insights are similar). Additionally, Section 3.8.13.3 explores a family of threshold policies, which can further improve Myopic’s performance.

3.6.1 Comparing Myopic with LP Rounding

In our first set of numerical experiments, we compare Myopic with LP Rounding. We consider the following numerical setting with i.i.d. demands. Given a set of items $\{1, 2, \dots, n\}$, we select $N = \overline{|o|} \cdot \underline{|o|}$ different order types in total, where $\overline{|o|}$ represents the largest size of any order and $\underline{|o|}$ represents the number of order types for each order size. After $\overline{|o|}$ and $\underline{|o|}$ are specified, a set of order types are randomly generated from the n items before each simulation run. In addition, in each sample path, the realization of each order type is chosen uniformly at random.

We model the initial FDC inventory level $I_{0,i}$ of item i by letting $I_{0,i} = \theta (p_i T)$, where $p_i T$ is the mean demand for item i over a time horizon T summed over the mean demand of all order types containing item i , and θ is an inventory scaling parameter that allows us to vary the supply-demand ratio. In addition, we assume $c_F = c_{F,i}$ and $c_R = c_{R,i}$ for

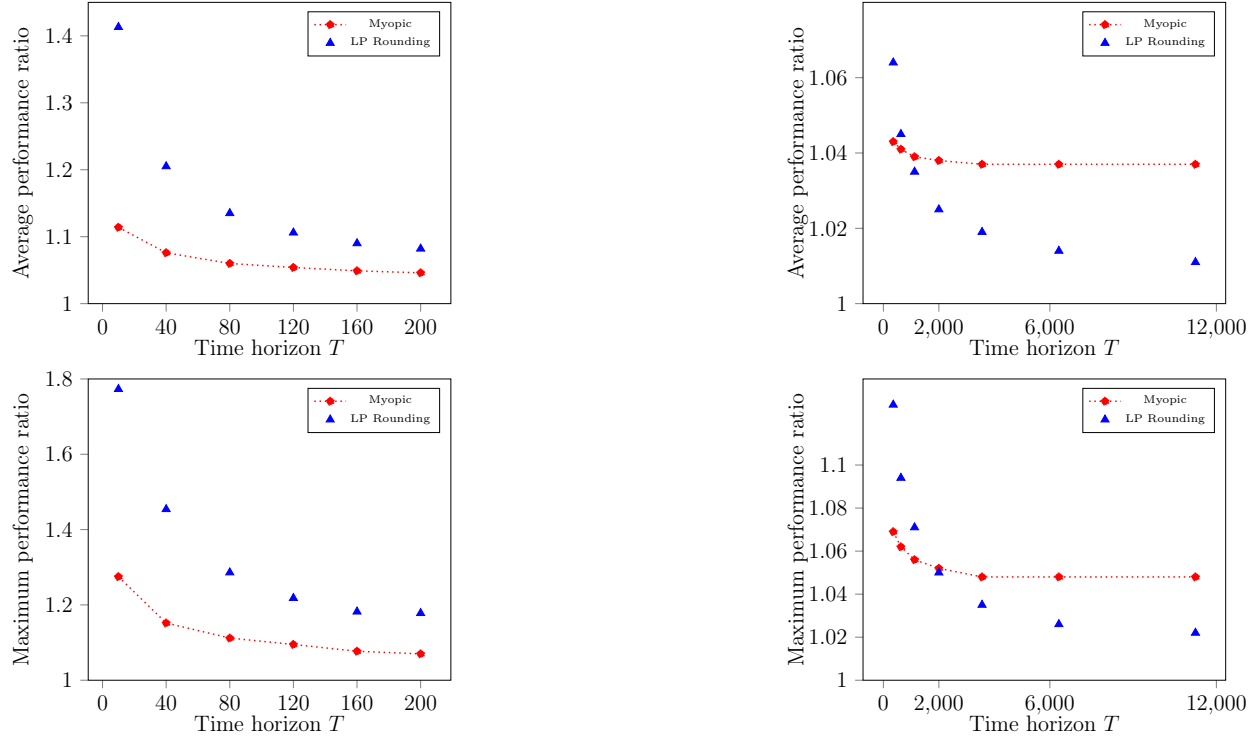


Figure 3.2: Performance w.r.t. time horizon T , holding $n = 10$, $\overline{|o|} = 5$, $|o| = 5$, $\theta = 0.8$, $f_F = 50$, $f_R = 100$, $c_F = 5$, $c_R = 0$. Left panel: T from 10 to 200; right panel: T from 356 to 11,246.

all i for simplicity. For each set of parameters, we first randomly select N order types 30 times. For each of these selections, we randomly generate $M = 300$ independent demand sequences. Finally, for each demand sequence, we define the performance ratio as the ratio of the cost under a given heuristic to OPT. We use the following two metrics to compare different heuristics: (1) the maximum performance ratio among the M simulation trials, as a proxy for the competitive ratio; and (2) the average performance ratio of the same simulation trials, motivated by the fact that in practice firms may be more concerned about the average performance instead of the (theoretical) worst-case scenario. Unless specified otherwise, the model parameters n , $\overline{|o|}$, $|o|$, θ , f_F , f_R , c_F , c_R are fixed in the rest of the section to provide a clear graphic demonstration. The trends, however, hold generally.

Figure 3.2 demonstrates the observed performance ratios under Assumption (C), as the time horizon T increases from 10 to more than 10,000, holding other parameters constant.

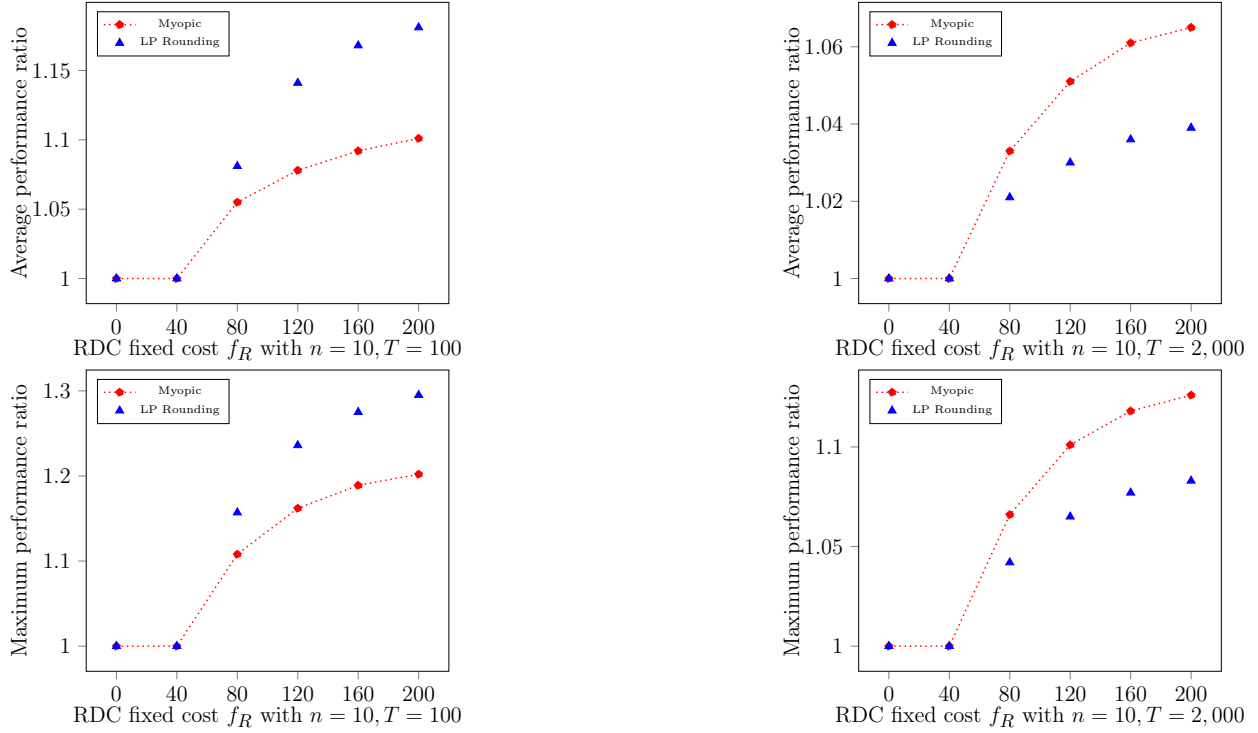


Figure 3.3: Performance w.r.t. RDC fixed cost f_R , holding $\overline{|o|} = 5, \underline{|o|} = 5, \theta = 0.8, f_F = 40, c_F = 5, c_R = 0$. Left panels: $T = 100$; right panels: $T = 2,000$.

Three observations are key here. First, although the upper bound on Myopic’s competitive ratio is 1.8 under given parameters (from Theorem 6), the empirical maximum performance ratio is less than 1.3, suggesting the actual gap between Myopic and OPT could be (much) smaller than the one indicated from our theoretical results. Second, both maximum and average performance ratios of Myopic and LP Rounding are low, and their performance ratios decrease and stabilize when T increases. Third, LP Rounding is outperformed by Myopic when T is small and outperforms Myopic when T is sufficiently large, consistent with its asymptotic optimality. Because of this phenomenon, in the remainder of this subsection, we show our simulation results based on two different T ’s, representing small and large time horizons, respectively.

Figures 3.3 and 3.4 compare Myopic with LP Rounding under Assumptions (C) and (D) by varying the RDC fixed cost f_R . The number of different items is small ($n = 10$) in

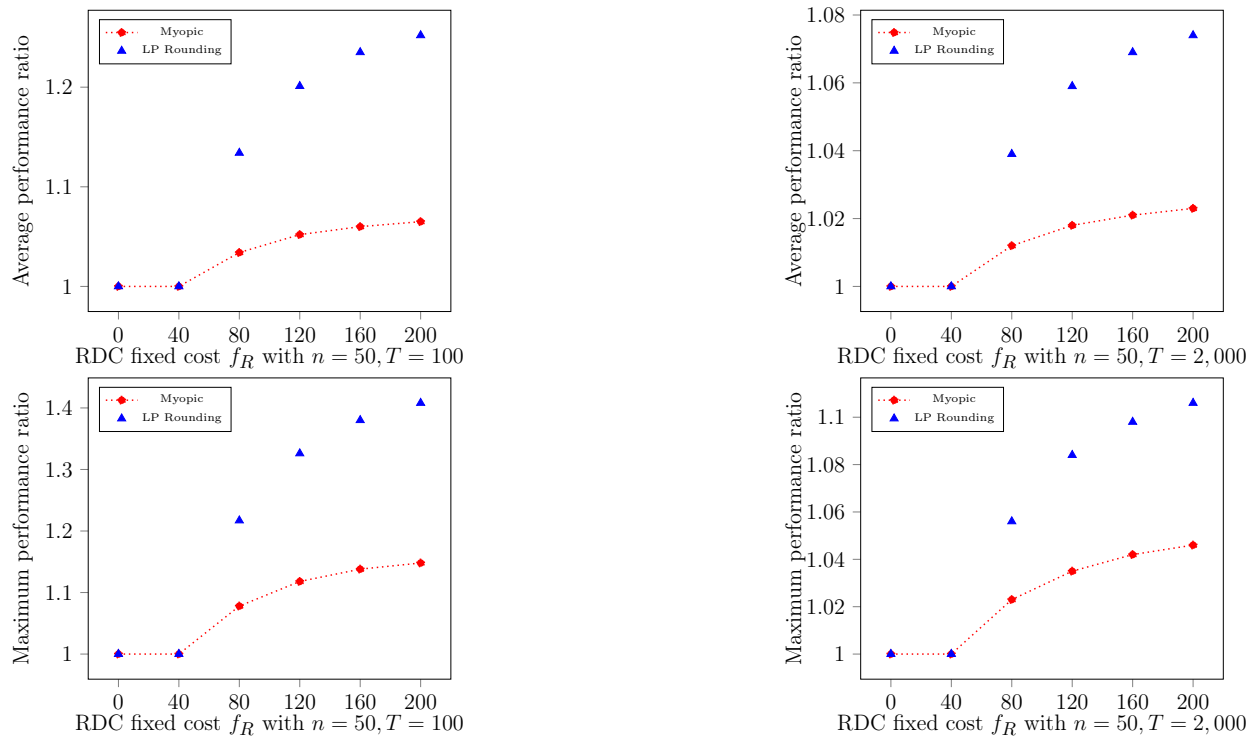


Figure 3.4: Performance w.r.t. RDC fixed cost f_R , holding $\overline{|o|} = 5$, $\underline{|o|} = 5$, $\theta = 0.8$, $f_F = 40$, $c_F = 5$, $c_R = 0$. Left panels: $T = 100$; right panels: $T = 2,000$.

Figure 3.3 and is large ($n = 50$) in Figure 3.4. Recall that under Assumption (D), because $f_R \leq f_F$, it is optimal to exclusively fulfill from the RDC such that $RATIO(Myopic) = 1$. Three key findings emerge. First, under Assumption (C), both ratios grow with f_R , which is consistent with our theoretical results. Second, when comparing the left columns with the right columns, unsurprisingly, both policies perform better for larger T (with LP Rounding having larger improvement), restating what we already observed in Figure 3.2. Third, let us look at the impact of the number of items n by comparing Figures 3.3 and 3.4. While fixing the order-type characteristics $(\overline{|o|}, \underline{|o|})$, the gap between LP Rounding and Myopic increases when the total number of items n increases, and LP Rounding is not guaranteed to outperform Myopic even for large T . The reason is that as the number of items n becomes relatively large (compared with the number of order types), order split becomes less frequent under Myopic, and thus, Myopic performs better. Hence, even when demand forecasting is available, Myopic is still valuable and competitive for such situations.

Next, we explore the impact of the initial FDC inventory under Assumption (C). Note that decreasing the parameter θ increases the FDC stockout probability while all other parameters are fixed. In one extreme case, when θ is close to 0, the FDC has very limited inventory and this limitation applies to all policies, including OPT. In such a case, one would expect all policies to perform close to OPT. By contrast, if we increase θ to the other extreme—the FDC has a sufficiently large amount of inventory—one would expect both Myopic and LP Rounding to perform close to OPT because current decisions are no longer penalized by future FDC stockout events. When θ is in the middle range, Myopic and LP Rounding have the largest performance gap. These phenomena are what we observe in Figure 3.5. In addition, Figure 3.5 again confirms Myopic’s average performance has a competitive advantage over LP Rounding when the time horizon T is small.



Figure 3.5: Performance w.r.t. FDC inventory scaling parameter θ , holding $n = 20$, $|\overline{o}| = 5$, $|\underline{o}| = 5$, $f_F = 50$, $f_R = 100$, $c_F = 5$, $c_R = 0$. Left panel: $T = 100$; right panel: $T = 8,000$.

3.6.2 Delay in order fulfillment

One common practice in online retailing is to impose a time delay between when an order is received and when the order is fulfilled, as explained in Xu et al. (2009). By delaying the decisions on order fulfillment, online retailers can gather more demand information to make better fulfillment decisions. In our second set of numerical experiments, we investigate the benefits of such (intentional) fulfillment-delay policies.

We implement this practice by considering various delay intervals. Let us fix a delay policy (denoted as Delay) with delay interval T^d . For each time t , we wait until orders $o_{t+1}, \dots, o_{t+T^d}$ have arrived, solve the corresponding T^d -period rolling horizon integer program, and fulfill only o_{t+1} at time $t + T^d$. We apply this practice from time T^d to T and benchmark it against Myopic (under Assumption (C)) facing the same order sequence. Note Delay performs no worse than Myopic because Myopic is a special case in which the delay interval equals exactly one period. In addition, Delay is expected to perform better with longer delay intervals, because it aggregates more order information and better coordinates between the FDC inventory and demand. In the extreme case in which $T^d = T$, Delay becomes the offline OPT. The results are summarized in Figure 3.6. One interesting observation is that the marginal benefit decreases as T^d increases, and most of the gap between Myopic and OPT are closed by implementing a delay policy of less than 40 periods.




Figure 3.6: Performance w.r.t. order fulfillment delay T^d , holding $T = 100, n = 10, \overline{|o|} = 5, \underline{|o|} = 5, f_F = 50, f_R = 100, c_F = 5, c_R = 0$.

3.6.3 Managerial insights

Finally, we summarize the managerial insights obtained from the numerical results in this section. First, when demands are i.i.d., the myopic policy performs reasonably well on average, even compared with the LP Rounding policy having forecasting information (Figures 3.2, 3.3, and 3.4). In particular, Myopic outperforms LP Rounding when demand becomes more uncertain (e.g., smaller T) or/and the number of items becomes relatively large. Second, we find little difference among different heuristics when the FDC inventory is too much or little, and the power of a good fulfillment policy only becomes evident when the FDC inventory is in the middle range (Figure 3.5). Third, delaying fulfillment decisions (even by a little) is helpful (Figure 3.6). Finally, we also propose a family of threshold policies can further improve Myopic’s performance and effectively hedge against the sensitivity of the cost parameters (see Figure 3.12 and Section 3.8.13.3 in the appendix for more details).

Combining our theoretical and numerical results, a key message of our paper is that Myopic actually performs well in the two-layer RDC-FDC network. This result is particularly valuable to e-retailers with unreliable demand forecasts. We want to emphasize this insight

departs from the one demonstrated earlier in Acimovic and Graves (2015). They study a different lateral fulfillment network in which any DC can fulfill orders from any location (such as Amazon's), and demonstrate that myopic policies are undesirable and e-retailers must use more complex forward-looking policies. This is because there are many possible options to fulfill an order in a lateral network given any DC in the network can be used, whereas there are fewer options in the more "sparse" RDC-FDC network in which an order is usually fulfilled by the local FDC/RDC. Although publicized discussion on the discrepancy between JD.com and Amazon's fulfillment networks is scant, we suspect the reason leading to this discrepancy is the difference in their promised delivery times. JD has a short promised delivery time (e.g., usually one day) and uses a tree-shaped two-layer network such that most orders are fulfilled locally. By contrast, Amazon has a (relatively) longer promised delivery time (e.g., usually two day), which allows them to search over a wider range of FCs. Amazon also heavily utilizes reactive transshipment among FCs. Given the trend of shorter and shorter promised delivery times, it is not impossible for Amazon to convert its fulfillment network (partially) to a tree-shaped one in the future.

3.7 Conclusion and discussion

In this paper, we considered a multi-item online order fulfillment problem in a two-layer distribution network that major e-retailers have implemented in practice. We analyzed the performance of a simple myopic policy that does not rely on demand forecasts and has been widely implemented in practice. We provided theoretical bounds on the competitive ratio of the myopic policy and showed our bounds are tight. We also empirically estimated our upper bound on the ratio by using FedEx shipping rates and demonstrated the bound can be as low as 1.13 for reasonable scenarios in practice. Moreover, we extended our study to the setting in which demand forecasting is available and proved the asymptotic optimality of an LP rounding policy. Finally, we complemented our theoretical results with a numerical

study. The main take-away of the paper is that Myopic performs well in the two-layer RDC-FDC network. It is comparable to more complex forward-looking policies that rely on demand forecasts, and the insight is especially meaningful for e-retailers struggling to cope with unreliable demand forecasts.

As a final note, although our framework seems applicable to only pure-play e-retailers, it can be applied to omnichannel retailers under certain circumstances. As more and more brick-and-mortar retailers are moving their business online, how to leverage existing stores to fulfill online orders becomes a critical question. If a store has a dedicated pool of inventory for online demands, for example, Walmart’s market fulfillment centers (Ward, 2020), the store can be treated as an FDC and our framework can be applied. Otherwise, if the store shares the pool of inventory for both online and offline demands (e.g., in-store picking), the situation would become more complex and additional challenges would arise (e.g., how to “protect” inventory for offline customers).

Acknowledgements. The authors thank Dan Adelman, Barış Ata, John R. Birge, René Caldentey, Nicole DeHoratius, Levi DeValve, Donald D. Eisenstein, Stephen C. Graves, David A. Goldberg, Stefanus Jasin, Will Ma, Rad Niazadeh, Amitabh Sinha, Amy R. Ward, Yehua Wei, and Jiawei Zhang for stimulating discussions.

3.8 Appendix

3.8.1 Fixed-cost estimates for Standard Overnight

In this section, we analyze the ratio $\frac{f_R}{f_F}$ by using FedEx Standard Overnight rates, similar to the analysis for ground shipping in Section 3.4. Note Standard Overnight provides next-business-day delivery services by 4:30 p.m. to U.S. businesses and by 8 p.m. to residences. We fit a similar linear model by using the rates presented in Table 3.8. The results are summarized in Table 3.9. Compared with ground shipping, the ratio $\frac{f_R}{f_F}$ is higher for Standard Overnight, ranging from 1.30 to 2.48, because overnight shipping is more expensive and

Weight (lbs)	Zone 2	Zone 3	Zone 4	Zone 5	Zone 6	Zone 7	Zone 8
1	\$8.76	\$9.09	\$9.91	\$10.35	\$10.71	\$10.83	\$11.01
2	\$9.42	\$10.48	\$11.33	\$11.58	\$12.08	\$12.54	\$12.75
3	\$9.90	\$11.04	\$11.94	\$12.43	\$12.94	\$13.38	\$14.03
4	\$10.19	\$11.11	\$12.32	\$13.12	\$13.48	\$14.36	\$15.03
5	\$10.46	\$11.62	\$12.75	\$13.70	\$14.29	\$15.02	\$16.02
6	\$10.65	\$11.77	\$12.87	\$13.90	\$14.41	\$15.19	\$16.09
7	\$11.25	\$12.05	\$13.21	\$14.34	\$14.68	\$15.60	\$16.70
8	\$11.58	\$12.49	\$13.68	\$14.75	\$15.27	\$16.23	\$17.43
9	\$11.76	\$12.66	\$13.86	\$14.88	\$15.59	\$16.90	\$18.36
10	\$11.92	\$12.79	\$14.00	\$15.29	\$15.77	\$17.71	\$19.87

Table 3.7: FedEx Ground Shipping Rates

Weight (lbs)	Zone 2	Zone 3	Zone 4	Zone 5	Zone 6	Zone 7	Zone 8
1	\$29.05	\$39.59	\$53.82	\$60.41	\$65.47	\$71.16	\$75.02
2	\$30.98	\$42.36	\$61.89	\$68.15	\$75.37	\$81.50	\$86.97
3	\$33.67	\$44.80	\$68.64	\$75.50	\$80.55	\$89.69	\$95.07
4	\$36.32	\$47.50	\$74.36	\$82.74	\$91.29	\$97.61	\$104.09
5	\$36.96	\$47.91	\$79.10	\$89.96	\$93.30	\$99.75	\$105.37
6	\$39.01	\$54.57	\$88.00	\$95.91	\$105.50	\$111.44	\$118.56
7	\$40.27	\$56.94	\$94.17	\$104.23	\$113.00	\$119.83	\$126.69
8	\$42.39	\$58.58	\$99.45	\$110.97	\$120.62	\$128.32	\$134.61
9	\$43.84	\$61.11	\$106.00	\$118.07	\$123.87	\$136.81	\$143.32
10	\$44.06	\$62.91	\$107.92	\$119.94	\$125.56	\$138.56	\$145.08

Table 3.8: FedEx Standard Overnight Shipping Rates

more sensitive to origin-destination distance. For instance, the cost of shipping a one-pound package by ground increases from 8.76 to 11.01 (approximately 25% increase) when the destination moves from Zone 2 to Zone 8, whereas the corresponding cost by Standard Overnight increases from 29.05 to 75.02 (more than 150% increase). However, given that Standard Overnight is considered one of the most expensive shipping methods, it is far less popular than ground shipping in our context (the e-retailer is unlikely to frequently use a shipping method whose cost is higher than the value of package). Nonetheless, the ratio is still within a reasonable range and the largest one from the table is 2.48, suggesting reasonably good performance of Myopic even for one of the most premium shipping methods.

Parameter	Zone	Parameter Estimate (\$)	R^2	$\frac{f_R}{f_F}$
f_F	2	28.23	97.94%	—
f_R	3	36.79	98.27%	1.30
	4	49.40	99.31%	1.75
	5	54.82	99.47%	1.94
	6	60.85	98.09%	2.16
	7	65.29	99.08%	2.31
	8	70.02	98.84%	2.48

Table 3.9: Fixed-Cost Estimates: Standard Overnight

3.8.2 A brief description of the open dataset

In this section, we briefly describe the open dataset provided by JD.com, which can be accessed from the 2020 MSOM Data Driven Research Challenge (see the full description in Shen et al. (2020)). This dataset is based on transactions associated with over 2.5 million users (450,000 purchases) and 30,000 SKUs during March 2018. All SKUs are within one anonymized consumable category, for example, beauty care (e.g., face moisturizers) or men’s grooming (e.g., electric shavers). The dataset also provides information about JD.com’s fulfillment network: the country is divided into eight geographical regions, and each region contains a certain number of districts. For each district, a designated warehouse is responsible for fulfilling orders in that district, if its inventory permits.

During that month, 446,142 orders were placed and 614,380 items were purchased. Approximately 90% of orders were single-item orders. We rank districts by their demands and select the top three most demanding districts. Their corresponding designated warehouses are DC 5, DC 9, and DC 2, respectively. For each of the three districts, we calculate the total number of orders, the percentage of single-item orders, and the total number of purchased items among the orders. These summary statistics are presented in Table 3.1. We also calculate the percentage of the number of items fulfilled by the designated DC and present the results in Figure 3.1.

3.8.3 Performance of other greedy-type policies

In this section, we investigate the performance of two other greedy-type policies. Both policies seem reasonable in practice. The first policy (FDC-Max) fulfills as many items from the FDC as possible and only uses the RDC when requested items are out of stock at the FDC. This policy is designed to maximize the FDC utilization and ignores order split. The second policy (Split-Min) is designed to minimize order split: if an order contains an item that is out of stock at the FDC, use the RDC to fulfill the entire order; otherwise, use the FDC to fulfill the entire order.

We next construct two specific instances. Both instances have only one order that requests all n items, and the FDC initial inventory is such that $I_{0,1} = 0$ and $I_{0,i} = 1$ for all $i > 1$. We adjust cost parameters such that one instance is under Assumption (B) and the other is under Assumption (C). We demonstrate that the performance of both FDC-Max and Split-Min can be arbitrarily poor, whereas Myopic performs well in these instances. The intuition is that Myopic can be equivalent to FDC-Max (or Split-Min) under certain cost parameters: when cost parameters change, Myopic has the flexibility to adjust its behavior and evolve into a different policy, whereas neither FDC-Max nor Split-Min make any adjustments.

Under Assumption (C). Let $c_R = 0$ and $c_F > 0$. Because $c_R = 0$, Myopic exclusively uses the RDC to fulfill this order, which is what Split-Min does in this case. In addition, this fulfillment decision by Myopic and Split-Min is optimal. FDC-Max, however, uses the RDC to fulfill the first item and then use the FDC to fulfill the remaining items. Hence,

$$\frac{V^{\text{FDC-Max}}}{V^*} = \frac{f_R + f_F + (n-1)c_F}{f_R},$$

which goes to infinity as $c_F \rightarrow \infty$. This simple example illustrates that the performance of FDC-Max can be arbitrarily poor.

Under Assumption (B). Let $f_F = c_F = 0$. Note Myopic only uses the RDC to fulfill

the out-of-stock item and exclusively uses the FDC for the remaining items, which is what FDC-Max does in this case. Again, this fulfillment decision is optimal. Split-Min, however, uses the RDC to fulfill the entire order to minimize order-split. Hence,

$$\frac{V^{\text{Split-Min}}}{V^*} = \frac{f_R + nc_R}{f_R + c_R},$$

which goes to infinity as the number of items $n \rightarrow \infty$. This simple example illustrates the performance of Split-Min can be arbitrarily poor.

From the two examples above, we see Myopic generally performs better across all assumptions than FDC-Max and Split-Min due to the flexibility of evolving into different policies under different cost parameters.

3.8.4 Proof of Theorem 4

In this section, we provide the remaining proof of Theorem 4.

From (3.12),

$$F_t - F_{t-1} = A_1 + A_2,$$

where

$$\begin{aligned} A_1 &\triangleq \left(\sum_{i=1}^n (I_{t,i}^m - I_{t,i}^*)^+ - \sum_{i=1}^n (I_{t-1,i}^m - I_{t-1,i}^*)^+ \right) (c_R - c_F) \\ &\quad + (w_t^{F,m} - w_{t-1}^{F,m})f_F - (w_t^{F,*} - w_{t-1}^{F,*})f_F, \\ A_2 &\triangleq (w_t^{R,m} - w_{t-1}^{R,m})f_R - (w_t^{R,*} - w_{t-1}^{R,*})f_R. \end{aligned}$$

Step 2: upper bounds on $F_t - F_{t-1}$. We first provide upper bounds on A_1 . Depending on whether Myopic and OPT use the FDC in period t , we consider the following four cases.

Case 1: $w_t^{F,m} = w_{t-1}^{F,m}$ and $w_t^{F,*} = w_{t-1}^{F,*}$. It implies both Myopic and OPT do not use the FDC in period t . Hence, they must both use the RDC, namely, $w_t^{R,m} = w_{t-1}^{R,m} + 1$ and

$w_t^{R,*} = w_{t-1}^{R,*} + 1$. It follows that

$$\begin{aligned}
A_1 &= \left(\sum_{i=1}^n \left(I_{t,i}^m - I_{t,i}^* \right)^+ - \sum_{i=1}^n \left(I_{t-1,i}^m - I_{t-1,i}^* \right)^+ \right) (c_R - c_F) \\
&= \left(\sum_{i=1}^n \left(I_{t-1,i}^m - I_{t-1,i}^* \right)^+ - \sum_{i=1}^n \left(I_{t-1,i}^m - I_{t-1,i}^* \right)^+ \right) (c_R - c_F) \\
&= 0,
\end{aligned}$$

where the second equality comes from the fact that neither Myopic nor OPT uses the FDC such that $I_{t,i}^m = I_{t-1,i}^m$ and $I_{t,i}^* = I_{t-1,i}^*$.

Case 2: $w_t^{F,m} = w_{t-1}^{F,m} + 1$ and $w_t^{F,*} = w_{t-1}^{F,*}$. It implies only Myopic uses the FDC in period t such that $I_{t,i}^m \leq I_{t-1,i}^m$ and $I_{t,i}^* = I_{t-1,i}^*$ for all i . It follows that

$$\begin{aligned}
A_1 &= \left(\sum_{i=1}^n \left(I_{t,i}^m - I_{t,i}^* \right)^+ - \sum_{i=1}^n \left(I_{t-1,i}^m - I_{t-1,i}^* \right)^+ \right) (c_R - c_F) + f_F \\
&\leq \left(\sum_{i=1}^n \left(I_{t-1,i}^m - I_{t-1,i}^* \right)^+ - \sum_{i=1}^n \left(I_{t-1,i}^m - I_{t-1,i}^* \right)^+ \right) (c_R - c_F) + f_F \\
&= f_F.
\end{aligned}$$

Case 3: $w_t^{F,m} = w_{t-1}^{F,m}$ and $w_t^{F,*} = w_{t-1}^{F,*} + 1$. It implies only OPT uses the FDC in period

t such that $I_{t,i}^m = I_{t-1,i}^m$ for all i . It follows that

$$\begin{aligned}
A_1 &= \left(\sum_{i=1}^n \left(I_{t-1,i}^m - I_{t,i}^* \right)^+ - \sum_{i=1}^n \left(I_{t-1,i}^m - I_{t-1,i}^* \right)^+ \right) (c_R - c_F) - f_F \\
&= \left(\sum_{i \in o_t} \left(I_{t-1,i}^m - I_{t,i}^* \right)^+ \mathbb{I} \left(I_{t-1,i}^m \geq 1 \right) \right. \\
&\quad \left. - \sum_{i \in o_t} \left(I_{t-1,i}^m - I_{t-1,i}^* \right)^+ \mathbb{I} \left(I_{t-1,i}^m \geq 1 \right) \right) (c_R - c_F) - f_F \\
&\leq \left(\sum_{i \in o_t} \left(I_{t-1,i}^* - I_{t,i}^* \right)^+ \mathbb{I} \left(I_{t-1,i}^m \geq 1 \right) \right) (c_R - c_F) - f_F \\
&\leq \left(\sum_{i \in o_t} \mathbb{I} \left(I_{t-1,i}^m \geq 1 \right) \right) (c_R - c_F) - f_F,
\end{aligned}$$

where the second equality comes from the facts that the inventory level remains the same for those items not in o_t and $\left(I_{t-1,i}^m - I_{t,i}^* \right)^+ = \left(I_{t-1,i}^m - I_{t,i}^* \right)^+ = 0$ for those items with $I_{t-1,i}^m = 0$, the third inequality is from the fact that $x^+ - y^+ \leq (x - y)^+$, and the last inequality is from the fact that $I_{t-1,i}^* \leq I_{t,i}^* + 1$ for all i and $|o_t| = N_t$. In addition, because Myopic only uses the RDC, due to (3.3), (3.6), and (3.7),

$$\begin{aligned}
\sum_{i \in o_t} \mathbb{I} \left(I_{t-1,i}^m \geq 1 \right) &\leq \begin{cases} \alpha - 1 & \text{if } I_{t-1,i}^m \geq 1 \ \forall i \in o_t; \\ \beta - 1 & \text{if there exists } i \in o_t \text{ s.t. } I_{t-1,i}^m = 0, \end{cases} \\
&\leq \beta - 1.
\end{aligned}$$

Hence,

$$A_1 \leq (\beta - 1)(c_R - c_F) - f_F \leq 0,$$

where the second inequality comes from (3.6).

Case 4: $w_t^{F,m} = w_{t-1}^{F,m} + 1$ and $w_t^{F,*} = w_{t-1}^{F,*} + 1$. Because Myopic uses the FDC, for those requested items that are available at the FDC in the end of period $t - 1$, they must be

fulfilled by FDC, that is, for all $i \in o_t$,

$$I_{t,i}^m = \begin{cases} I_{t-1,i}^m - 1 & \text{if } I_{t-1,i}^m \geq 1; \\ 0 & \text{if } I_{t-1,i}^m = 0. \end{cases}$$

It follows that

$$\begin{aligned} A_1 &= \left(\sum_{i \in o_t} (I_{t,i}^m - I_{t,i}^*)^+ - \sum_{i \in o_t} (I_{t-1,i}^m - I_{t-1,i}^*)^+ \right) (c_R - c_F) \\ &= \left(\sum_{i \in o_t, I_{t-1,i}^m \geq 1} (I_{t-1,i}^m - 1 - I_{t,i}^*)^+ - \sum_{i \in o_t, I_{t-1,i}^m \geq 1} (I_{t-1,i}^m - I_{t-1,i}^*)^+ \right) (c_R - c_F) \\ &\leq \left(\sum_{i \in o_t, I_{t-1,i}^m \geq 1} (I_{t-1,i}^* - 1 - I_{t,i}^*)^+ \right) (c_R - c_F) \\ &\leq 0, \end{aligned}$$

where the first equality comes from the fact that the inventory level remains the same for those items not in o_t , the second equality is from the fact that

$$(I_{t,i}^m - I_{t,i}^*)^+ = (I_{t-1,i}^m - I_{t-1,i}^*)^+ = 0$$

for those $i \in o_t$ with $I_{t-1,i}^m = 0$, the third inequality comes from the fact that $x^+ - y^+ \leq (x - y)^+$, and the last inequality is from the fact that $I_{t-1,i}^* - I_{t,i}^* \leq 1$ for all i .

In summary,

$$A_1 \leq \begin{cases} 0 & \text{if } w_t^{F,m} = w_{t-1}^{F,m} \text{ and } w_t^{F,*} = w_{t-1}^{F,*}; \\ f_F & \text{if } w_t^{F,m} = w_{t-1}^{F,m} + 1 \text{ and } w_t^{F,*} = w_{t-1}^{F,*}; \\ 0 & \text{if } w_t^{F,m} = w_{t-1}^{F,m} \text{ and } w_t^{F,*} = w_{t-1}^{F,*} + 1; \\ 0 & \text{if } w_t^{F,m} = w_{t-1}^{F,m} + 1 \text{ and } w_t^{F,*} = w_{t-1}^{F,*} + 1. \end{cases} \quad (3.25)$$

We then evaluate A_2 . Depending on whether Myopic (OPT) uses the RDC in period t , we have

$$A_2 = \begin{cases} 0 & \text{if } w_t^{R,m} = w_{t-1}^{R,m} \text{ and } w_t^{R,*} = w_{t-1}^{R,*}; \\ f_R & \text{if } w_t^{R,m} = w_{t-1}^{R,m} + 1 \text{ and } w_t^{R,*} = w_{t-1}^{R,*}; \\ -f_R & \text{if } w_t^{R,m} = w_{t-1}^{R,m} \text{ and } w_t^{R,*} = w_{t-1}^{R,*} + 1; \\ 0 & \text{if } w_t^{R,m} = w_{t-1}^{R,m} + 1 \text{ and } w_t^{R,*} = w_{t-1}^{R,*} + 1. \end{cases} \quad (3.26)$$

After combining (3.25) and (3.26), we have $4 \times 4 = 16$ cases in total, but because we assume all orders are non-empty, both Myopic and OPT use at least one of the DCs such that all the cases including $w_t^{F,m} = w_{t-1}^{F,m}$ and $w_t^{R,m} = w_{t-1}^{R,m}$ (i.e., Myopic does not use either of the DCs), and $w_t^{F,*} = w_{t-1}^{F,*}$ and $w_t^{R,*} = w_{t-1}^{R,*}$ (i.e., OPT does not use neither of the DCs) are excluded. We also exclude the cases with upper bounds on $F_t - F_{t-1} = A_1 + A_2$ that are less than or equal to 0. Combining all the above, it remains to consider only four

cases, and their positive upper bounds on $F_t - F_{t-1}$ are summarized as follows:

$$\left\{ \begin{array}{ll} f_R & \text{if } w_t^{F,m} = w_{t-1}^{F,m}, w_t^{F,*} = w_{t-1}^{F,*} + 1, w_t^{R,m} = w_{t-1}^{R,m} + 1, \text{ and } w_t^{R,*} = w_{t-1}^{R,*}; \\ f_F & \text{if } w_t^{F,m} = w_{t-1}^{F,m} + 1, w_t^{F,*} = w_{t-1}^{F,*}, w_t^{R,m} = w_{t-1}^{R,m} + 1, \text{ and } w_t^{R,*} = w_{t-1}^{R,*} + 1; \\ f_F - f_R & \text{if } w_t^{F,m} = w_{t-1}^{F,m} + 1, w_t^{F,*} = w_{t-1}^{F,*}, w_t^{R,m} = w_{t-1}^{R,m}, \text{ and } w_t^{R,*} = w_{t-1}^{R,*} + 1; \\ f_R & \text{if } w_t^{F,m} = w_{t-1}^{F,m} + 1, w_t^{F,*} = w_{t-1}^{F,*} + 1, w_t^{R,m} = w_{t-1}^{R,m} + 1, \text{ and } w_t^{R,*} = w_{t-1}^{R,*}. \end{array} \right. \quad (3.27)$$

Step 3: proving (3.15). The final step is to prove (3.15). It is convenient to make the following assumption on OPT: when multiple optimal offline algorithms exist, we let OPT be the one that uses the RDC as much as possible. The assumption is also intuitive, because the RDC has more inventory than the FDC. Recall from (3.11),

$$V_t^* - V_{t-1}^* = L_t^* c_F + (w_t^{F,*} - w_{t-1}^{F,*}) f_F + (N_t - L_t^*) c_R + (w_t^{R,*} - w_{t-1}^{R,*}) f_R.$$

We proceed by a case analysis.

Case 1: $w_t^{F,m} = w_{t-1}^{F,m}, w_t^{F,*} = w_{t-1}^{F,*} + 1, w_t^{R,m} = w_{t-1}^{R,m} + 1, \text{ and } w_t^{R,*} = w_{t-1}^{R,*}$.

Because OPT only uses the FDC, we claim $N_t \geq \alpha$. Otherwise, we can construct another policy OPT_1 that is identical to OPT, except OPT_1 fulfills the t -th order exclusively from the RDC. From (3.3), we see that either OPT_1 incurs strictly less cost than OPT does, or OPT_1 incurs the same cost but uses RDC more than OPT does, contradicting either the optimality of OPT or the aforementioned assumption on OPT. It follows that

$$V_t^* - V_{t-1}^* = f_F + N_t c_F \geq f_F + \alpha c_F. \quad (3.28)$$

Combining (3.28) with (3.27) proves (3.15) in this case.

Case 2: $w_t^{F,m} = w_{t-1}^{F,m} + 1, w_t^{F,*} = w_{t-1}^{F,*}, w_t^{R,m} = w_{t-1}^{R,m} + 1, \text{ and } w_t^{R,*} = w_{t-1}^{R,*} + 1$.

Because Myopic uses both FDC and RDC, we claim the inventory is insufficient for some

requested items in period t from Observation 1. Hence, $N_t \geq \beta \geq \alpha$ from (3.6) and (3.7). In addition, note OPT only uses the RDC. It follows that

$$V_t^* - V_{t-1}^* = f_R + N_t c_R \geq f_R + \alpha c_R. \quad (3.29)$$

Combining (3.29) with (3.27) proves (3.15) in this case.

Case 3: $w_t^{F,m} = w_{t-1}^{F,m} + 1$, $w_t^{F,*} = w_{t-1}^{F,*}$, $w_t^{R,m} = w_{t-1}^{R,m}$, and $w_t^{R,*} = w_{t-1}^{R,*} + 1$.

The argument is exactly the same as Case 2 because: (1) the lower bound (3.29) still holds due to the fact that OPT only uses the RDC; (2) $F_t - F_{t-1}$ has a smaller upper bound in this case than in Case 2.

Case 4: $w_t^{F,m} = w_{t-1}^{F,m} + 1$, $w_t^{F,*} = w_{t-1}^{F,*} + 1$, $w_t^{R,m} = w_{t-1}^{R,m} + 1$, and $w_t^{R,*} = w_{t-1}^{R,*}$.

Because OPT only uses the FDC, we claim $N_t \geq \alpha$, similar to the reasoning in Case 1. It follows that

$$V_t^* - V_{t-1}^* = f_F + N_t c_F \geq f_F + \alpha c_F. \quad (3.30)$$

Combining (3.30) with (3.27) proves (3.15) in this case. The proof of (3.8) is completed.

Under Assumption (B) (i.e., $f_R > f_F$, $c_R > c_F$), plugging the expression of α into (3.8) leads to

$$\text{RATIO(Myopic)} \leq 1 + \left(\frac{f_R}{f_F + c_F} \vee \frac{f_F}{f_R + c_R} \right),$$

which is equal to $1 + f_R/(f_F + c_F)$. The proof of Theorem 4 is completed.

3.8.5 Proof of Theorem 5

In this section, we present the proof of Theorem 5.

Proof. Proof of Theorem 5.

Suppose that T orders are placed for $Q \in [T, nT]$ items in total. Let L^m , L^* be the total numbers of items shipped from the FDC by Myopic and OPT, respectively, and let $w^{F,m}$, $w^{R,m}$ ($w^{F,*}$, $w^{R,*}$) denote the total numbers of times Myopic (OPT) uses the FDC

and RDC, respectively. Then, the total costs incurred by Myopic and OPT are $w^{R,m}f_R + w^{F,m}f_F + L^m c_F + (Q - L^m)c_R$ and $w^{R,*}f_R + w^{F,*}f_F + L^*c_F + (Q - L^*)c_R$, respectively.

Note $\beta = 1$ is equivalent to $f_F + c_F < c_R$, which implies Myopic always uses the FDC as long as the FDC has inventory for requested items (despite order split). It follows that $L^m \geq L^*$. Next, we claim they are indeed identical: $L^m = L^*$. Suppose by contradiction that under OPT at least one item has leftover FDC inventory at the end and that leftover inventory could be used to (partially) fulfill one of the T orders. Then, we can see the policy identical to OPT but consuming that leftover inventory incurs strictly less cost than OPT, because $f_F + c_F < c_R$. In summary, Myopic and OPT use the same amount of RDC and FDC inventory:

$$L^m = L^* = L \quad \text{and} \quad Q - L^m = Q - L^* = Q - L.$$

It follows that

$$RATIO(Myopic) \leq \max_{T, w^{R,m}, w^{F,m}, w^{R,*}, w^{F,*}, Q, L} \frac{w^{R,m}f_R + w^{F,m}f_F + Lc_F + (Q - L)c_R}{w^{R,*}f_R + w^{F,*}f_F + Lc_F + (Q - L)c_R}, \quad (3.31)$$

where the constraints on the parameters are

$$\begin{aligned} 0 \leq w^{R,m}, w^{F,m}, w^{R,*}, w^{F,*} \leq T, \quad w^{R,m} + w^{F,m} \geq T, \quad w^{R,*} + w^{F,*} \geq T, \\ w^{F,m} \vee w^{F,*} \leq L, \quad w^{R,m} \vee w^{R,*} \leq Q - L. \end{aligned}$$

By using the fact that for all $a \geq c > 0$ and $b \geq 0$,

$$\frac{a+b}{c+b} \leq \frac{a}{c},$$

we obtain the following further upper bound on $RATIO(Myopic)$:

$$\begin{aligned}
& RATIO(Myopic) \tag{3.32} \\
& \leq \max_{\substack{0 \leq w^{R,m}, w^{F,m}, w^{R,*}, w^{F,*} \leq T \\ w^{R,m} + w^{F,m} \geq T \\ w^{R,*} + w^{F,*} \geq T}} \frac{w^{R,m} f_R + w^{F,m} f_F + (w^{F,m} \vee w^{F,*}) c_F + (w^{R,m} \vee w^{R,*}) c_R}{w^{R,*} f_R + w^{F,*} f_F + (w^{F,m} \vee w^{F,*}) c_F + (w^{R,m} \vee w^{R,*}) c_R}, \tag{3.33}
\end{aligned}$$

by using the fact that

$$w^{F,m} \vee w^{F,*} \leq L \quad \text{and} \quad w^{R,m} \vee w^{R,*} \leq Q - L.$$

We next observe that the maximum value of Problem (3.32) is attained when $w^{F,m} \geq w^{F,*}$ and $w^{R,m} \geq w^{R,*}$. Otherwise, if $w^{F,m} < w^{F,*}$ or $w^{R,m} < w^{R,*}$, one can easily increase $w^{F,m}$ or $w^{R,m}$ to obtain a strictly higher value. Therefore,

$$\begin{aligned}
RATIO(Myopic) \leq \max_{\substack{0 \leq w^{R,m}, w^{F,m}, w^{R,*}, w^{F,*} \leq T \\ w^{R,m} + w^{F,m} \geq T \\ w^{R,*} + w^{F,*} \geq T \\ w^{F,m} \geq w^{F,*}, w^{R,m} \geq w^{R,*}}} \frac{w^{R,m} f_R + w^{F,m} f_F + w^{F,m} c_F + w^{R,m} c_R}{w^{R,*} f_R + w^{F,*} f_F + w^{F,m} c_F + w^{R,m} c_R}. \tag{3.34}
\end{aligned}$$

For any $(w^{R,m}, w^{F,m}, w^{R,*}, w^{F,*})$ satisfying the constraints in (3.34), it suffices to prove

$$\frac{w^{R,m} f_R + w^{F,m} f_F + w^{F,m} c_F + w^{R,m} c_R}{w^{R,*} f_R + w^{F,*} f_F + w^{F,m} c_F + w^{R,m} c_R} \leq \frac{f_R + 2f_F + c_F + c_R}{f_F + c_F + c_R},$$

which is equivalent to

$$\begin{aligned}
& (w^{R,m} f_R + w^{F,m} f_F + w^{F,m} c_F + w^{R,m} c_R) (f_F + c_F + c_R) \\
& \leq (f_R + 2f_F + c_F + c_R) (w^{R,*} f_R + w^{F,*} f_F + w^{F,m} c_F + w^{R,m} c_R). \tag{3.35}
\end{aligned}$$

After further simplification, (3.35) reduces to

$$\begin{aligned}
w^{R,m} f_R(f_F + c_F) + w^{F,m} f_F(f_F + c_R) &\leq f_R \left(w^{R,*} f_R + w^{F,*} f_F + w^{F,m} c_F \right) \\
&+ f_F \left(w^{R,*} f_R + w^{F,*} f_F + w^{R,m} c_R \right) + (f_F + c_F + c_R) \left(w^{R,*} f_R + w^{F,*} f_F \right).
\end{aligned} \tag{3.36}$$

Note the right-hand side of (3.36) is no smaller than the following:

$$\begin{aligned}
&f_R \left(w^{R,*} f_R + w^{F,*} f_F \right) + f_F \left(w^{R,*} f_R + w^{F,*} f_F \right) \\
&+ f_R c_F \left(w^{F,m} + w^{R,*} \right) + f_F c_R \left(w^{R,m} + w^{F,*} \right).
\end{aligned}$$

Combining the above with the facts that $w^{R,*} + w^{F,*} \geq T \geq w^{R,m}$, $w^{R,*} + w^{F,*} \geq T \geq w^{F,m}$, $w^{F,m} \geq w^{F,*}$, $w^{R,m} \geq w^{R,*}$, and $f_R \geq f_F$, we have

$$\begin{aligned}
f_R \left(w^{R,*} f_R + w^{F,*} f_F \right) &\geq f_R \left(w^{R,*} f_F + w^{F,*} f_F \right) \geq w^{R,m} f_R f_F, \\
f_F \left(w^{R,*} f_R + w^{F,*} f_F \right) &\geq f_F \left(w^{R,*} f_F + w^{F,*} f_F \right) \geq w^{F,m} f_F f_F, \\
f_R c_F \left(w^{F,m} + w^{R,*} \right) &\geq f_R c_F \left(w^{F,*} + w^{R,*} \right) \geq w^{R,m} f_R c_F, \\
f_F c_R \left(w^{R,m} + w^{F,*} \right) &\geq f_F c_R \left(w^{R,*} + w^{F,*} \right) \geq w^{F,m} f_F c_R,
\end{aligned}$$

completing the proof of (3.36). The proof of Theorem 5 is completed. \square

3.8.6 Proof of Theorem 6

In this section, we present the proof of Theorem 6, which is similar to the proof of Theorem 4.

Proof. Proof of Theorem 6. Without loss of generality, we assume $f_R - f_F \geq c_F - c_R$ in this section; otherwise we would have $\alpha = 1$ from the definition (i.e., Myopic always uses the RDC) such that $\text{RATIO}(\text{Myopic}) = 1$. From Observation 2, we have the following relation:

for all $t \geq 1$,

$$w_t^{F,m} + w_t^{R,m} = w_t^{F,*} + w_t^{R,*} = t. \quad (3.37)$$

Similar to the proof of Theorem 4, we present the proof in several steps.

Step 1: potential function. In this case, we consider the following potential function F_t :

$$\begin{aligned} F_t &\triangleq \left(\sum_{i=1}^n (I_{t,i}^* - I_{t,i}^m) \right) (c_F - c_R) + (w_t^{F,m} - w_t^{F,*})(f_F - f_R) \\ &= V_t^m - V_t^*, \end{aligned} \quad (3.38)$$

where the second equality comes from (3.10) and (3.37). Note this potential function is different from the one used in the proof of Theorem 4 in two aspects. First, this potential function is the exact difference between the cumulative costs incurred by Myopic and OPT. Second, we can simplify the terms involving the usages of FDC and RDC due to (3.37).

Following the same proof technique, it suffices to show that for any $t \geq 1$,

$$F_t - F_{t-1} \leq \frac{(f_R - f_F) - (c_F - c_R)}{f_F + c_F} (V_t^* - V_{t-1}^*). \quad (3.39)$$

From the definition in (3.38),

$$\begin{aligned} F_t - F_{t-1} &= \left(\sum_{i=1}^n (I_{t,i}^* - I_{t-1,i}^*) - \sum_{i=1}^n (I_{t,i}^m - I_{t-1,i}^m) \right) (c_F - c_R) \\ &\quad + (w_t^{F,m} - w_{t-1}^{F,m})(f_F - f_R) - (w_t^{F,*} - w_{t-1}^{F,*})(f_F - f_R). \end{aligned}$$

Step 2: upper bounds on $F_t - F_{t-1}$. Recall that $w_t^{F,m}$ ($w_t^{F,*}$) is perfectly correlated with $w_t^{R,m}$ ($w_t^{R,*}$) from (3.37). Hence, we consider the following four cases depending on whether Myopic and OPT use the FDC.

Case 1: $w_t^{F,m} = w_{t-1}^{F,m}$ and $w_t^{F,*} = w_{t-1}^{F,*}$. It follows that $w_t^{R,m} = w_{t-1}^{R,m} + 1$ and $w_t^{R,*} =$

$w_{t-1}^{R,*} + 1$. In addition, $I_{t,i}^m = I_{t-1,i}^m$ and $I_{t,i}^* = I_{t-1,i}^*$ for all i . Hence,

$$F_t - F_{t-1} = \left(\sum_{i=1}^n (I_{t,i}^* - I_{t-1,i}^*) - \sum_{i=1}^n (I_{t,i}^m - I_{t-1,i}^m) \right) (c_F - c_R) = 0.$$

Case 2: $w_t^{F,m} = w_{t-1}^{F,m} + 1$ and $w_t^{F,*} = w_{t-1}^{F,*}$. It follows that $w_t^{R,m} = w_{t-1}^{R,m}$ and $w_t^{R,*} = w_{t-1}^{R,*} + 1$. In addition, $\sum_{i=1}^n (I_{t-1,i}^m - I_{t,i}^m) = N_t$ and $I_{t,i}^* = I_{t-1,i}^*$ for all i . Hence,

$$\begin{aligned} F_t - F_{t-1} &= \left(\sum_{i=1}^n (I_{t,i}^* - I_{t-1,i}^*) - \sum_{i=1}^n (I_{t,i}^m - I_{t-1,i}^m) \right) (c_F - c_R) + f_F - f_R \\ &= N_t (c_F - c_R) + f_F - f_R. \end{aligned}$$

Because all the N_t items are fulfilled by the FDC under Myopic, we have $N_t \leq \alpha - 1$ due to (3.4). Combining all the above implies

$$F_t - F_{t-1} \leq (\alpha - 1) (c_F - c_R) + f_F - f_R \leq 0,$$

where the second inequality comes from (3.4).

Case 3: $w_t^{F,m} = w_{t-1}^{F,m}$ and $w_t^{F,*} = w_{t-1}^{F,*} + 1$. It follows that $w_t^{R,m} = w_{t-1}^{R,m} + 1$ and $w_t^{R,*} = w_{t-1}^{R,*}$. In addition, $\sum_{i=1}^n (I_{t-1,i}^* - I_{t,i}^*) = N_t$ and $I_{t,i}^m = I_{t-1,i}^m$ for all i . Hence,

$$\begin{aligned} F_t - F_{t-1} &= \left(\sum_{i=1}^n (I_{t,i}^* - I_{t-1,i}^*) - \sum_{i=1}^n (I_{t,i}^m - I_{t-1,i}^m) \right) (c_F - c_R) - f_F + f_R \\ &= -N_t (c_F - c_R) - f_F + f_R \\ &\leq -(c_F - c_R) - f_F + f_R, \end{aligned}$$

where the last inequality is from Assumption (C).

Case 4: $w_t^{F,m} = w_{t-1}^{F,m} + 1$ and $w_t^{F,*} = w_{t-1}^{F,*} + 1$. It follows that $w_t^{R,m} = w_{t-1}^{R,m}$ and

$w_t^{R,*} = w_{t-1}^{R,*}$. In addition, $I_{t-1,i}^m - I_{t,i}^m = I_{t-1,i}^* - I_{t,i}^*$ for all i . Hence,

$$F_t - F_{t-1} = \left(\sum_{i=1}^n (I_{t,i}^* - I_{t-1,i}^*) - \sum_{i=1}^n (I_{t,i}^m - I_{t-1,i}^m) \right) (c_F - c_R) = 0.$$

In summary,

$$F_t - F_{t-1} \tag{3.40}$$

$$\leq \begin{cases} (f_R - f_F) - (c_F - c_R) & \text{if } w_t^{F,m} = w_{t-1}^{F,m}, w_t^{F,*} = w_{t-1}^{F,*} + 1, \\ & w_t^{R,m} = w_{t-1}^{R,m} + 1, \text{ and } w_t^{R,*} = w_{t-1}^{R,*}; \\ 0 & \text{otherwise.} \end{cases} \tag{3.41}$$

We can again exclude the cases with upper bounds less than or equal to zero.

Step 3: proving (3.39). The final step is to prove (3.39). Recall from (3.11),

$$V_t^* - V_{t-1}^* = L_t^* c_F + (w_t^{F,*} - w_{t-1}^{F,*}) f_F + (N_t - L_t^*) c_R + (w_t^{R,*} - w_{t-1}^{R,*}) f_R.$$

From (3.40), the only remaining case is when $w_t^{F,m} = w_{t-1}^{F,m}$, $w_t^{F,*} = w_{t-1}^{F,*} + 1$, $w_t^{R,m} = w_{t-1}^{R,m} + 1$, and $w_t^{R,*} = w_{t-1}^{R,*}$. Because OPT only uses the FDC under this case, we have

$$V_t^* - V_{t-1}^* = f_F + N_t c_F \geq f_F + c_F. \tag{3.42}$$

Combining (3.42) and (3.40) proves (3.39), completing the proof of Theorem 6. \square

3.8.7 Proof of Theorem 7

In this section, we present the proof of Theorem 7. Note all the instances used in this proof have exactly one unit of inventory for all the items, that is, $I_{0,i} = 1$ for all $i = 1, \dots, n$. We specify the number of items and the order sequences under Assumptions (A)-(C) separately.

Section 3.8.14 provides a summary of the instances used in this proof.

Proof. Proof of Theorem 7. **Under Assumption (A).** We proceed by constructing the following two order sequences I_1 and I_2 . For I_1 , fix a sufficiently large integer $m > 0$ and let $n = \alpha m$. There are m α -item orders: $o_i = \{(i-1)\alpha + 1, \dots, i\alpha\}$ for $i = 1, \dots, m$, followed by one n -item order: $o_{m+1} = \{1, \dots, n\}$. Note OPT uses the RDC to fulfill the orders from $t = 1$ to $t = m$ and uses the FDC to fulfill the last order. Hence, the cost under OPT is

$$V^*(I_1) = mf_R + \alpha mc_R + f_F + \alpha mc_F.$$

In addition, from (3.3), Myopic uses the FDC for the first m periods and uses the RDC in the last period, due to FDC stockouts. Hence, the cost under Myopic is

$$V^m(I_1) = mf_F + \alpha mc_F + f_R + \alpha mc_R,$$

which implies

$$\frac{V^m(I_1)}{V^*(I_1)} = \frac{mf_F + \alpha mc_F + f_R + \alpha mc_R}{mf_R + \alpha mc_R + f_F + \alpha mc_F},$$

which converges to

$$1 + \frac{f_F - f_R}{f_R + \alpha(c_F + c_R)} \text{ as } m \rightarrow \infty.$$

For I_2 , fix an integer $m > 0$ and let $n = m(\beta + 1)$. One m -item order is placed in period one: $o_1 = \{1, \dots, m\}$, followed by m $(\beta + 1)$ -item orders:

$$o_{i+1} = \{i, m + (i-1)\beta + 1, \dots, m + i\beta\}$$

for $i = 1, \dots, m$. Note OPT uses the RDC to fulfill the first order and then uses the FDC

to fulfill the remaining m orders. Hence, the cost under OPT is

$$V^*(I_2) = f_R + mc_R + m(f_F + (\beta + 1)c_F).$$

In addition, when m is sufficiently large, Myopic uses the FDC for the first order from (3.3). Then, it uses the RDC to fulfill the FDC out-of-stock item and the FDC to fulfill the remaining β items for each of the remaining orders from (3.6). Hence, the cost under Myopic is

$$V^m(I_2) = f_F + mc_F + m(f_R + c_R + f_F + \beta c_F),$$

which implies

$$\frac{V^m(I_2)}{V^*(I_2)} = \frac{f_F + mc_F + m(f_R + c_R + f_F + \beta c_F)}{f_R + mc_R + m(f_F + (\beta + 1)c_F)},$$

which converges to

$$1 + \frac{f_R}{f_F + (\beta + 1)c_F + c_R} \text{ as } m \rightarrow \infty.$$

Therefore,

$$\text{RATIO(Myopic)} \geq 1 + \left(\frac{f_R}{f_F + (\beta + 1)c_F + c_R} \vee \frac{f_F - f_R}{f_R + \alpha(c_F + c_R)} \right).$$

Under Assumption (B). Similarly, we construct the following order sequences I_1 and I_2 . One n -item order is placed: $o_1 = \{1, 2, \dots, n\}$, followed by n single-item orders, $o_{i+1} = \{i\}$. Note OPT uses the RDC to fulfill the first order and uses the FDC to fulfill the remaining single-item orders. Hence, the cost under OPT is

$$V^*(I_1) = f_R + nc_R + nf_F + nc_F.$$

In addition, Myopic uses the FDC to fulfill the first order from (3.3) and uses the RDC to

fulfill the remaining single-item orders. Hence, the cost under Myopic is

$$V^m(I_1) = f_F + nc_F + nf_R + nc_R,$$

Which implies

$$\frac{V^m(I_1)}{V^*(I_1)} = \frac{f_F + nc_F + nf_R + nc_R}{f_R + nc_R + nf_F + nc_F},$$

which converges to

$$1 + \frac{f_R - f_F}{f_F + c_F + c_R} \text{ as } n \rightarrow \infty.$$

For I_2 , let $n = \beta + 1$. One single-item order is placed: $o_1 = \{1\}$, followed by one $(\beta + 1)$ -item order: $o_2 = \{1, 2, \dots, \beta + 1\}$. Note OPT uses the RDC to fulfill the first single-item order and then uses the FDC to fulfill the multi-item order. Hence, the cost under OPT is

$$V^*(I_2) = f_R + c_R + f_F + (\beta + 1)c_F.$$

In addition, Myopic uses the FDC for the first order from (3.3). Then, it uses the RDC to fulfill the FDC out-of-stock item and the FDC to fulfill the remaining β items for each of the remaining orders from (3.6). Hence, the cost under Myopic is

$$V^m(I_2) = f_F + c_F + f_R + c_R + f_F + \beta c_F,$$

which implies

$$\begin{aligned} \frac{V^m(I_2)}{V^*(I_2)} &= \frac{f_F + c_F + f_R + c_R + f_F + \beta c_F}{f_R + c_R + f_F + (\beta + 1)c_F} \\ &= 1 + \frac{f_F}{f_R + f_F + c_R + (\beta + 1)c_F}. \end{aligned}$$

Therefore,

$$\text{RATIO(Myopic)} \geq 1 + \left(\frac{f_R - f_F}{f_F + c_F + c_R} \vee \frac{f_F}{f_R + f_F + c_R + (\beta + 1)c_F} \right).$$

Under Assumption (C). Without loss of generality, we assume $\alpha \geq 2$; otherwise, the result holds trivially. We construct the following order sequence I_1 . Let $n = \alpha - 1$. One $(\alpha - 1)$ -item order is placed: $o_1 = \{1, \dots, \alpha - 1\}$, followed by $(\alpha - 1)$ single-item orders: $o_{i+1} = \{i\}$ for $i = 1, \dots, \alpha - 1$. Note OPT uses the RDC to fulfill the orders in period $t = 1$ and then uses the FDC to fulfill the remaining single-item orders. Hence, the cost under OPT is

$$V^*(I_1) = f_R + (\alpha - 1)c_R + (\alpha - 1)(f_F + c_F).$$

In addition, from (3.4), Myopic uses the FDC to fulfill the orders in period $t = 1$. Then, it uses the RDC to fulfill the remaining single-item orders because the FDC is out of stock. Hence, the cost under Myopic is

$$V^m(I_1) = f_F + (\alpha - 1)c_F + (\alpha - 1)(f_R + c_R),$$

which implies

$$\begin{aligned} \frac{V^m(I_1)}{V^*(I_1)} &= \frac{f_F + (\alpha - 1)c_F + (\alpha - 1)(f_R + c_R)}{f_R + (\alpha - 1)c_R + (\alpha - 1)(f_F + c_F)} \\ &= 1 + \frac{(\alpha - 2)(f_R - f_F)}{f_R + (\alpha - 1)(f_F + c_R + c_F)}. \end{aligned}$$

Therefore,

$$\text{RATIO(Myopic)} \geq 1 + \frac{(\alpha - 2)(f_R - f_F)}{f_R + (\alpha - 1)(f_F + c_R + c_F)}.$$

The proof of Theorem 7 is completed. □

3.8.8 Proof of Theorem 8

In this section, we present the proof of Theorem 8. Similar to the proof of Theorem 7, we also need to construct several instances. Note all the instances have exactly one unit of inventory for all the items; that is, $I_{0,i} = 1$ for all $i = 1, \dots, n$. We specify the number of items and the order sequences under Assumptions (A)-(C) separately. Section 3.8.14 provides a summary

of the instances used in this proof.

Proof. Proof of Theorem 8. Let \mathcal{A} be the set of all deterministic online algorithms. From Yao's minimax principle, it suffices to construct two order sequences I_1 and I_2 with associated probabilities p_1 and $1 - p_1$, and then prove the expected competitive ratio is at least LB for any deterministic online algorithm; namely,

$$\inf_{A \in \mathcal{A}} \left[p_1 \frac{V^A(I_1)}{V^*(I_1)} + (1 - p_1) \frac{V^A(I_2)}{V^*(I_2)} \right] \geq LB, \quad (3.43)$$

where $V^A(I_i)$ and $V^*(I_i)$ represent the total cost under I_i incurred by A and OPT , respectively, and LB represents the desired lower bound in (3.9) under Assumptions (A), (B), and (C), respectively.

Under Assumption (A). We want to show

$$\inf_{A \in \mathcal{A}} \left[p_1 \frac{V^A(I_1)}{V^*(I_1)} + (1 - p_1) \frac{V^A(I_2)}{V^*(I_2)} \right] \geq 1 + \left(\frac{f_F + (\beta + 1)c_F + c_R}{f_R} + \frac{c_F}{c_R - c_F} \right)^{-1}. \quad (3.44)$$

We construct the following order sequences I_1 and I_2 . Fix a sufficiently large integer $m > 0$ and let $n = m(\beta + 1)$. For I_1 , only one m -item order is placed: $o_1 = \{1, \dots, m\}$. For I_2 , one m -item order in period one is placed: $o_1 = \{1, \dots, m\}$, followed by m $(\beta + 1)$ -item orders: $o_{i+1} = \{i, m + (i - 1)\beta + 1, \dots, m + i\beta\}$ for $i = 1, \dots, m$. Note under order sequence I_1 , OPT only uses the FDC to fulfill the first order; under order sequence I_2 , OPT uses the RDC to fulfill the first order and then uses the FDC to fulfill the remaining m orders. Hence, the corresponding costs under OPT are

$$V^*(I_1) = f_F + mc_F, \quad V^*(I_2) = f_R + mc_R + m(f_F + (\beta + 1)c_F). \quad (3.45)$$

For any deterministic policy $A \in \mathcal{A}$, let L^A denote the number of items fulfilled by the FDC

in period $t = 1$. Note $0 \leq L^A \leq m$. In addition,

$$\begin{aligned} V^A(I_1) &= \mathbb{I}(L^A > 0) f_F + \mathbb{I}(L^A < m) f_R + L^A c_F + (m - L^A) c_R \\ &\geq m c_R - L^A (c_R - c_F), \end{aligned} \quad (3.46)$$

and because using the FDC is cheaper whenever the requested item is available from period $t = 2$ to $t = m + 1$, we have

$$\begin{aligned} V^A(I_2) &\geq \mathbb{I}(L^A > 0) f_F + \mathbb{I}(L^A < m) f_R + L^A f_R + m (f_F + (\beta + 1) c_F + c_R) \\ &\geq L^A f_R + m (f_F + (\beta + 1) c_F + c_R). \end{aligned} \quad (3.47)$$

We next choose $p_1 \in [0, 1]$ to be

$$p_1 = \frac{f_R (f_F + m c_F)}{(c_R - c_F) [f_R + m (f_F + (\beta + 1) c_F + c_R)] + f_R (f_F + m c_F)}. \quad (3.48)$$

Under this random-order sequence, the first m -item order appears with probability one and the $m(\beta + 1)$ -item orders may or may not appear with appearing probability $1 - p_1$. From (3.45), (3.46), and (3.47), for any $A \in \mathcal{A}$,

$$\begin{aligned} &p_1 \frac{V^A(I_1)}{V^*(I_1)} + (1 - p_1) \frac{V^A(I_2)}{V^*(I_2)} \\ &\geq p_1 \frac{m c_R - L^A (c_R - c_F)}{f_F + m c_F} + (1 - p_1) \frac{L^A f_R + m (f_F + (\beta + 1) c_F + c_R)}{f_R + m (f_F + (\beta + 1) c_F + c_R)} \\ &= \left[-p_1 \frac{c_R - c_F}{f_F + m c_F} + (1 - p_1) \frac{f_R}{f_R + m (f_F + (\beta + 1) c_F + c_R)} \right] L^A \\ &\quad + p_1 \frac{m c_R}{f_F + m c_F} + (1 - p_1) \frac{m (f_F + (\beta + 1) c_F + c_R)}{f_R + m (f_F + (\beta + 1) c_F + c_R)}. \end{aligned} \quad (3.49)$$

From (3.48), we can verify the coefficient of L^A above is equal to zero, and then the right-

hand side of (3.49) converges to

$$\begin{aligned} & 1 + \frac{f_R(c_R - c_F)}{(c_R - c_F)(f_F + (\beta + 1)c_F + c_R) + f_R c_F} \\ & = 1 + \left(\frac{f_F + (\beta + 1)c_F + c_R}{f_R} + \frac{c_F}{c_R - c_F} \right)^{-1} \quad \text{as } m \rightarrow \infty, \end{aligned}$$

completing the proof of (3.44).

Under Assumption (B). We want to show that for each $n \geq 2$,

$$\begin{aligned} & \inf_{A \in \mathcal{A}} \left[p_1 \frac{V^A(I_1)}{V^*(I_1)} + (1 - p_1) \frac{V^A(I_2)}{V^*(I_2)} \right] \\ & \geq 1 + \left(\frac{f_F + n c_F}{f_R - f_F + n(c_R - c_F)} + \frac{f_R + n(f_F + c_R + c_F)}{(n - 1)(f_R - f_F)} \right)^{-1}. \end{aligned} \quad (3.50)$$

We construct the following two order sequences I_1 and I_2 . For I_1 , only one order is placed: $o_1 = \{1, 2, \dots, n\}$. For I_2 , $o_1 = \{1, 2, \dots, n\}$ is followed by n single-item orders: $o_{i+1} = \{i\}$. Note OPT only uses the FDC to fulfill the first order under order sequence I_1 , and first uses the RDC to fulfill the first order and then uses the FDC to fulfill the remaining n orders under order sequence I_2 . Hence, the corresponding costs under OPT are

$$V^*(I_1) = f_F + n c_F, \quad V^*(I_2) = f_R + n(f_F + c_R + c_F). \quad (3.51)$$

For any $A \in \mathcal{A}$, let L^A denote the number of items fulfilled by the FDC in period $t = 1$. Note $0 \leq L^A \leq n$. In addition,

$$\begin{aligned} V^A(I_1) &= \mathbb{I}(L^A > 0) f_F + \mathbb{I}(L^A < n) f_R + L^A c_F + (n - L^A) c_R, \\ V^A(I_2) &\geq \mathbb{I}(L^A > 0) f_F + \mathbb{I}(L^A < n) f_R + (n - L^A) f_F + L^A f_R + n(c_R + c_F), \end{aligned} \quad (3.52)$$

where the inequality results from the fact that using the FDC is always cheaper whenever

the requested item is available from period $t = 2$ to $t = n + 1$. From (3.52), we have

$$\begin{aligned}
V^A(I_1) &\geq \begin{cases} f_R + L^A c_F + (n - L^A) c_R & \text{if } 0 \leq L^A < n; \\ f_F + n c_F & \text{if } L^A = n, \end{cases} \\
V^A(I_2) &\geq \begin{cases} (L^A + 1) f_R + (n - L^A) f_F + n(c_R + c_F) & \text{if } 0 \leq L^A < n; \\ f_F + n f_R + n(c_R + c_F) & \text{if } L^A = n. \end{cases}
\end{aligned} \tag{3.53}$$

We next choose $p_1 \in [0, 1]$ to be

$$p_1 = \frac{(n - 1)(f_R - f_F)(f_F + n c_F)}{[f_R - f_F + n(c_R - c_F)][f_R + n(f_F + c_R + c_F)] + (n - 1)(f_R - f_F)(f_F + n c_F)}. \tag{3.54}$$

Under this random-order sequence, the first n -item order appears with probability one, and the n single-item orders may or may not appear with appearing probability $1 - p_1$. From (3.51), (3.52), and (3.53), for any $A \in \mathcal{A}$, if $0 \leq L^A < n$,

$$\begin{aligned}
&p_1 \frac{V^A(I_1)}{V^*(I_1)} + (1 - p_1) \frac{V^A(I_2)}{V^*(I_2)} \\
&\geq p_1 \frac{f_R + n c_R - L^A(c_R - c_F)}{f_F + n c_F} + (1 - p_1) \frac{L^A(f_R - f_F) + f_R + n(f_F + c_R + c_F)}{f_R + n(f_F + c_R + c_F)} \\
&= \left[-p_1 \frac{c_R - c_F}{f_F + n c_F} + (1 - p_1) \frac{f_R - f_F}{f_R + n(f_F + c_R + c_F)} \right] L^A \\
&\quad + p_1 \frac{f_R + n c_R}{f_F + n c_F} + (1 - p_1);
\end{aligned} \tag{3.55}$$

if $L^A = n$,

$$\begin{aligned}
p_1 \frac{V^A(I_1)}{V^*(I_1)} + (1 - p_1) \frac{V^A(I_2)}{V^*(I_2)} &\geq p_1 \frac{f_F + n c_F}{f_F + n c_F} + (1 - p_1) \frac{f_F + n f_R + n(c_R + c_F)}{f_R + n(f_F + c_R + c_F)} \\
&= p_1 + (1 - p_1) \frac{f_F + n(f_R + c_R + c_F)}{f_R + n(f_F + c_R + c_F)}.
\end{aligned} \tag{3.56}$$

It can be checked that the coefficient of L^A in (3.55) is non-negative such that

$$\begin{aligned}
& \min_{0 \leq L^A < n} \left[p_1 \frac{V^A(I_1)}{V^*(I_1)} + (1-p_1) \frac{V^A(I_2)}{V^*(I_2)} \right] \\
& \geq \left\{ \left[-p_1 \frac{c_R - c_F}{f_F + nc_F} + (1-p_1) \frac{f_R - f_F}{f_R + n(f_F + c_R + c_F)} \right] L^A \right. \\
& \quad \left. + p_1 \frac{f_R + nc_R}{f_F + nc_F} + (1-p_1) \right\} \Big|_{L^A=0} \\
& = p_1 \frac{f_R + nc_R}{f_F + nc_F} + (1-p_1),
\end{aligned}$$

which is equal to the lower bound in (3.56) due to (3.54). Hence, for each $n \geq 2$,

$$\begin{aligned}
& \inf_{A \in \mathcal{A}} \left[p_1 \frac{V^A(I_1)}{V^*(I_1)} + (1-p_1) \frac{V^A(I_2)}{V^*(I_2)} \right] \\
& \geq p_1 \frac{f_R + nc_R}{f_F + nc_F} + (1-p_1) \\
& = 1 + \left(\frac{f_F + nc_F}{f_R - f_F + n(c_R - c_F)} + \frac{f_R + n(f_F + c_R + c_F)}{(n-1)(f_R - f_F)} \right)^{-1},
\end{aligned}$$

completing the proof of (3.50).

Under Assumption (C). Without loss of generality, we assume $\alpha \geq 2$; otherwise, the result holds trivially. We want to show that for each $1 \leq \alpha' \leq \alpha - 1$,

$$\begin{aligned}
& \inf_{A \in \mathcal{A}} \left[p_1 \frac{V^A(I_1)}{V^*(I_1)} + (1-p_1) \frac{V^A(I_2)}{V^*(I_2)} \right] \\
& \geq 1 + \left(\frac{f_R + \alpha'(f_F + c_F + c_R)}{(\alpha' - 1)(f_R - f_F)} + \frac{f_F + \alpha'c_F}{f_R - f_F - \alpha'(c_F - c_R)} \right)^{-1}.
\end{aligned} \tag{3.57}$$

We construct the following order sequences I_1 and I_2 . Let $n = \alpha'$. For I_1 , one α' -item order is placed: $o_1 = \{1, 2, \dots, \alpha'\}$. For I_2 , it has the same order in period $t = 1$ as I_1 but with additional α' single-item orders: $o_{i+1} = \{i\}$ for $i = 1, \dots, \alpha'$. Note OPT only uses the FDC to fulfill all the orders under order sequence I_1 , and uses the RDC to fulfill the first order

and then uses the FDC to fulfill the remaining orders under order sequence I_2 . Hence, the corresponding costs under OPT are

$$\begin{aligned} V^*(I_1) &= f_F + \alpha' c_F, & V^*(I_2) &= f_R + \alpha' c_R + \alpha' (f_F + c_F) \\ & & &= f_R + \alpha' (f_F + c_F + c_R). \end{aligned} \quad (3.58)$$

For any $A \in \mathcal{A}$, similar to Observation 2, if A uses both DCs in a certain period, we can construct a better policy A' only using the RDC in that period without impacting future fulfillment. Hence, we can restrict ourselves to $\mathcal{A}' \subseteq \mathcal{A}$, where \mathcal{A}' includes all the deterministic online policies that use only one of the two DCs in each period.

Given $A \in \mathcal{A}'$, let w^A denote the total number of times that the FDC is used in period $t = 1$. Note $0 \leq w^A \leq 1$. In addition,

$$V^A(I_1) = w^A (f_F + \alpha' c_F) + (1 - w^A) (f_R + \alpha' c_R), \quad (3.59)$$

and because using the FDC to ship the remaining orders is cheaper than using the RDC,

$$\begin{aligned} &V^A(I_2) \\ &\geq w^A (f_F + \alpha' c_F) + (1 - w^A) (f_R + \alpha' c_R) + w^A \alpha' (f_R + c_R) + (1 - w^A) \alpha' (f_F + c_F) \\ &= w^A [f_F + \alpha' (f_R + c_F + c_R)] + (1 - w^A) [f_R + \alpha' (f_F + c_F + c_R)]. \end{aligned} \quad (3.60)$$

We next choose $p_1 \in [0, 1]$ to be

$$p_1 = \frac{(\alpha' - 1) (f_R - f_F) (f_F + \alpha' c_F)}{(\alpha' - 1) (f_R - f_F) (f_F + \alpha' c_F) + [f_R - f_F - \alpha' (c_F - c_R)] [f_R + \alpha' (f_F + c_F + c_R)]}. \quad (3.61)$$

Under this random-order sequence, the m α' -item orders appear with probability one and the last $m\alpha'$ -item order may or may not appear with appearing probability $1 - p_1$. From

(3.58), (3.59), and (3.60), for any $A \in \mathcal{A}$,

$$\begin{aligned}
& p_1 \frac{V^A(I_1)}{V^*(I_1)} + (1-p_1) \frac{V^A(I_2)}{V^*(I_2)} \\
& \geq p_1 \cdot \frac{w^A (f_F + \alpha' c_F) + (1-w^A) (f_R + \alpha' c_R)}{f_F + \alpha' c_F} \\
& \quad + (1-p_1) \cdot \frac{w^A [f_F + \alpha' (f_R + c_F + c_R)] + (1-w^A) [f_R + \alpha' (f_F + c_F + c_R)]}{f_R + \alpha' (f_F + c_F + c_R)} \quad (3.62) \\
& = \left[-p_1 \cdot \frac{f_R - f_F - \alpha' (c_F - c_R)}{f_F + \alpha' c_F} + (1-p_1) \cdot \frac{(\alpha' - 1) (f_R - f_F)}{f_R + \alpha' (f_F + c_F + c_R)} \right] w^A \\
& \quad + p_1 \cdot \frac{f_R + \alpha' c_R}{f_F + \alpha' c_F} + (1-p_1) \cdot \frac{f_R + \alpha' (f_F + c_F + c_R)}{f_R + \alpha' (f_F + c_F + c_R)}.
\end{aligned}$$

From (3.61), we can verify the coefficient of w^A above is equal to 0, and then the right-hand side of (3.62) is equal to

$$1 + \left(\frac{f_R + \alpha' (f_F + c_F + c_R)}{(\alpha' - 1) (f_R - f_F)} + \frac{f_F + \alpha' c_F}{f_R - f_F - \alpha' (c_F - c_R)} \right)^{-1}.$$

The proof of (3.57) is completed. In conclusion, we complete the proof of Theorem 8. \square

3.8.9 Proof of Theorem 9

Proof. Proof of Theorem 9. **Under Assumption (B).** We proceed by a case analysis.

When $\beta \geq 2$, it suffices to show that

$$1 + \frac{f_R}{f_F + c_F} \leq 2 \left(1 + \frac{f_R - f_F}{f_F + c_F + c_R} \right),$$

which is equivalent to

$$f_R (f_F + c_F - c_R) + (f_F + c_F) (c_F + c_R - f_F) \geq 0. \quad (3.63)$$

From the definition of β in (3.5), $\beta \geq 2$ is equivalent to

$$f_F + c_F \geq c_R. \quad (3.64)$$

Combining (3.64) with the fact that $f_R \geq f_F$, to prove (3.63), it suffices to prove that

$$f_F (f_F + c_F - c_R) + (f_F + c_F) (c_F + c_R - f_F) \geq 0. \quad (3.65)$$

It is not difficult to verify that (3.65) holds.

When $\beta = 1$ (equivalent to $f_F + c_F < c_R$), it suffices to show

$$1 + \frac{f_R + f_F}{f_F + c_F + c_R} \leq 2 \left(1 + \frac{f_R - f_F}{f_F + c_F + c_R} \right),$$

which clearly holds.

Under Assumption (C). It is not difficult to verify that the result holds when $\alpha \leq 2$ (implying $f_R + 2c_R < f_F + 2c_F$), because in this case, the upper bound on $RATIO(Myopic)$ is no greater than 2. Without loss of generality, we assume $\alpha \geq 3$, which implies

$$f_R + 2c_R \geq f_F + 2c_F. \quad (3.66)$$

We want to prove

$$2 \cdot \frac{(\alpha - 1)(f_R + c_R + c_F) + f_F}{f_R + (\alpha - 1)(f_F + c_R + c_F)} \geq \frac{f_R + c_R}{f_F + c_F}, \quad (3.67)$$

which is equivalent to

$$(\alpha - 1) \cdot [(f_F + c_F - c_R)(f_R + c_R) + 2c_F(f_F + c_F)] \geq (f_R + c_R)f_R - 2(f_F + c_F)f_F. \quad (3.68)$$

Note the coefficient of $(\alpha - 1)$ in (3.68) is non-negative. Hence, to prove (3.68), from the

definition of α in (3.2), it suffices to show that

$$\left(\frac{f_R - f_F}{c_F - c_R} - 1\right) \cdot [(f_F + c_F - c_R)(f_R + c_R) + 2c_F(f_F + c_F)] \geq (f_R + c_R)f_R - 2(f_F + c_F)f_F,$$

which is equivalent to

$$\begin{aligned} & (f_R + c_R - f_F - c_F) [(f_F + c_F - c_R)(f_R + c_R) + 2c_F(f_F + c_F)] \\ & \geq (f_R + c_R)f_R(c_F - c_R) - 2(c_F - c_R)(f_F + c_F)f_F. \end{aligned} \quad (3.69)$$

We further expand the left-hand side of (3.69):

$$\begin{aligned} & (f_R + c_R) [(f_R + c_R)(f_F + c_F - c_R) - (f_F + c_F)(f_F + c_F - c_R) + 2c_F(f_F + c_F)] \\ & \quad - 2(f_F + c_F)^2 c_F \\ & \geq (f_R + c_R)f_R(c_F - c_R) - 2(c_F - c_R)(f_F + c_F)f_F. \end{aligned} \quad (3.70)$$

After a straightforward calculation, (3.70) is equivalent to

$$(f_R + c_R) \left[(f_R - f_F + 2c_R)f_F + c_F^2 + 2c_F c_R - c_R^2 \right] \geq 2(f_F + c_F)(c_F^2 + f_F c_R). \quad (3.71)$$

It suffices to prove the following inequality after replacing $(f_R + c_R)$ and $(f_R - f_F + 2c_R)$ in (3.71) by their lower bounds $(f_F + 2c_F - c_R)$ and $2c_F$, respectively (due to (3.66)):

$$(f_F + 2c_F - c_R) \left(2c_F f_F + c_F^2 + 2c_F c_R - c_R^2 \right) \geq 2(f_F + c_F)(c_F^2 + f_F c_R), \quad (3.72)$$

because both of their coefficients are non-negative due to Assumption (C). After decomposing $(f_F + 2c_F - c_R)$ into $(f_F + c_F)$ and $(c_F - c_R)$, (3.72) can be simplified to

$$(f_F + c_F)(c_F - c_R)(2f_F - c_F + c_R) + (c_F - c_R) \left(2c_F f_F + c_F^2 + 2c_F c_R - c_R^2 \right) \geq 0. \quad (3.73)$$

Finally, the coefficient of $(c_F - c_R)$ in (3.73) is equal to

$$(f_F + c_F)(2f_F - c_F + c_R) + (2c_F f_F + c_F^2 + 2c_F c_R - c_R^2) = f_F(2f_F + 3c_F + c_R) + c_R(3c_F - c_R),$$

which is non-negative due to the assumption that $c_F \geq c_R$. The proof of (3.67) is completed.

Tightness of the bounds. Under Assumption (B), let $c_F = 0$, $c_R = f_F$, and $f_R > 2f_F$. Then, the upper and lower bounds are $1 + \frac{f_R}{f_F}$ and $1 + \frac{f_R - f_F}{2f_F}$, respectively. Their ratio is exactly 2. Under Assumption (C), let $c_F = 1$, $c_R = 0$, $f_F = \epsilon$, and $f_R = M + \epsilon$, where $\epsilon > 0$ and M is a positive integer. Then, $\alpha = M + 1$, and the upper and lower bounds are $1 + \frac{M-1}{1+\epsilon}$ and $1 + \frac{(M-1)M}{M(2+\epsilon)+\epsilon}$, respectively. Hence, as $M \rightarrow \infty$, their ratio converges to $\frac{2+\epsilon}{1+\epsilon}$, which converges to 2 as $\epsilon \rightarrow 0$.

The proof of Theorem 9 is completed. □

3.8.10 Proof of Theorem 10

Proof. Proof of Theorem 10. Assume $I_{0,i} = 1$ for all $i = 1, \dots, n$. To prove this result, we construct two order sequences I_1 and I_2 . For I_1 , one order is placed: $o_1 = \{1, 2, \dots, n\}$. For I_2 , $o_1 = \{1, 2, \dots, n\}$ is followed by n single-item orders: $o_{i+1} = \{i\}$. Note OPT only uses the FDC to fulfill the first order under order sequence I_1 , and first uses the RDC to fulfill the first order and then uses the FDC to fulfill the remaining n orders under order sequence I_2 . Hence, the corresponding costs under OPT are

$$V^*(I_1) = f_F, \quad V^*(I_2) = f_R + n f_F.$$

Any online deterministic algorithm has only two options: using the FDC to fulfill the first order (called algorithm A), or using the RDC to fulfill the first order (called algorithm B).

Then,

$$\begin{aligned} V^A(I_1) &= f_F, & V^A(I_2) &= f_F + nf_R, \\ V^B(I_1) &= f_R, & V^B(I_2) &= f_R + nf_F. \end{aligned}$$

It follows that

$$\min \left\{ \frac{V^A(I_2)}{V^*(I_2)}, \frac{V^B(I_1)}{V^*(I_1)} \right\} = \min \left\{ \frac{f_F + nf_R}{f_R + nf_F}, \frac{f_R}{f_F} \right\},$$

which converges to $\frac{f_R}{f_F}$ as $n \rightarrow \infty$, completing the first part of Theorem 10.

For the second part of the theorem, we apply Yao's minimax principle and assign probability $1/2$ to instances I_1 and I_2 , respectively. Under this random-order sequence, the first order $o_1 = \{1, 2, \dots, n\}$ always appears, and the last n single-item orders appear with probability $1/2$. From similar calculations,

$$\frac{1}{2} \left[\frac{V^B(I_1)}{V^*(I_1)} + \frac{V^B(I_2)}{V^*(I_2)} \right] = \frac{1}{2} \left[\frac{f_R}{f_F} + 1 \right],$$

and

$$\frac{1}{2} \left[\frac{V^A(I_1)}{V^*(I_1)} + \frac{V^A(I_2)}{V^*(I_2)} \right] = \frac{1}{2} \left[1 + \frac{f_F + nf_R}{f_R + nf_F} \right],$$

which converges to $\frac{1}{2} \left[1 + \frac{f_R}{f_F} \right]$ as $n \rightarrow \infty$, completing the second part of Theorem 10. \square

3.8.11 Proof of Theorem 11

Proof. Proof of Theorem 11. Because $\tilde{C}_{LP}^*(T)$ is a lower bound of $C^*(T)$, it suffices to prove

$$\lim_{T \rightarrow \infty} \frac{C^{LPR}(T)}{\tilde{C}_{LP}^*(T)} = 1. \quad (3.74)$$

Recall that $D^{\mathcal{J},t} = 1$ if an order type \mathcal{J} arrives in period t , and 0 otherwise. Let $\{X_{R,i}^{\mathcal{J},t}\}_t$ and $\{X_{F,i}^{\mathcal{J},t}\}_t$ be sequences of i.i.d. Bernoulli random variables with parameter $x_{R,i}^{\mathcal{J},*}$ and $x_{F,i}^{\mathcal{J},*}$, respectively.

First, from the definition of the LP Rounding policy, we can see that when the FDC inventory is sufficient and no FDC stock-out occurs, the LP Rounding policy incurs the same expected variable costs (both RDC and FDC) as the optimal LP solution. Similarly, when no FDC stock-out occurs, LP Rounding also incurs the same expected fixed costs as the optimal LP solution; namely, for each order type \mathcal{J} with size m ,

$$\mathbb{E} \left[\max_{i \in \mathcal{J}} \{X_{R,i}^{\mathcal{J},t}\} \right] = \max_{i \in \mathcal{J}} \{x_{R,i}^{\mathcal{J},*}\}, \quad \mathbb{E} \left[\max_{i \in \mathcal{J}} \{X_{F,i}^{\mathcal{J},t}\} \right] = \max_{i \in \mathcal{J}} \{x_{F,i}^{\mathcal{J},*}\}.$$

To see this, without loss of generality, assume $\{x_{R,i}^{\mathcal{J},*}\}_i$ are ordered in ascending order. Consequently, $\{x_{F,i}^{\mathcal{J},*}\}_i$ are in descending order. By our LP Rounding scheme, (3.23), and (3.24),

$$\begin{aligned} \mathbb{E} \left[\max_{i \in \mathcal{J}} \{X_{R,i}^{\mathcal{J},t}\} \right] &= x_{R,1}^{\mathcal{J},*} + \sum_{i=2}^m (x_{R,i}^{\mathcal{J},*} - x_{R,i-1}^{\mathcal{J},*}) = x_{R,m}^{\mathcal{J},*} = \max_{i \in \mathcal{J}} \{x_{R,i}^{\mathcal{J},*}\}, \\ \mathbb{E} \left[\max_{i \in \mathcal{J}} \{X_{F,i}^{\mathcal{J},t}\} \right] &= \sum_{i=2}^m (x_{R,i}^{\mathcal{J},*} - x_{R,i-1}^{\mathcal{J},*}) + 1 - x_{R,m}^{\mathcal{J},*} = 1 - x_{R,1}^{\mathcal{J},*} = x_{F,1}^{\mathcal{J},*} = \max_{i \in \mathcal{J}} \{x_{F,i}^{\mathcal{J},*}\}. \end{aligned}$$

Hence, we conclude that when no FDC stock-out occurs, LP Rounding incurs the same variable and fixed costs as the optimal LP solution.

By contrast, when item i is out of stock at the FDC, LP Rounding has to fulfill that item from the RDC and incurs an additional cost, which is upper bounded by $c_{R,i} + f_R$. Therefore, we achieve the following bound on the gap between $C^{LPR}(T)$ and $\tilde{C}_{LP}^*(T)$:

$$C^{LPR}(T) - \tilde{C}_{LP}^*(T) \leq \sum_i (c_{R,i} + f_R) \mathbb{E} \left[\left(\sum_{t=1}^T \sum_{\mathcal{J} \ni i} D^{\mathcal{J},t} X_{F,i}^{\mathcal{J},t} - \theta_i T \right)^+ \right].$$

The remaining is to show that for each i , the term $\mathbb{E} \left[\left(\sum_{t=1}^T \sum_{\mathcal{J} \ni i} D^{\mathcal{J},t} X_{F,i}^{\mathcal{J},t} - \theta_i T \right)^+ \right]$ dimin-

ishes as $T \rightarrow \infty$. To see this,

$$\begin{aligned}
& \mathbb{E} \left[\left(\sum_{t=1}^T \sum_{\mathcal{J} \ni i} D^{\mathcal{J},t} X_{F,i}^{\mathcal{J},t} - \theta_i T \right)^+ \right] \\
& \leq \mathbb{E} \left[\left(\sum_{t=1}^T \sum_{\mathcal{J} \ni i} \left(D^{\mathcal{J},t} X_{F,i}^{\mathcal{J},t} - \lambda^{\mathcal{J}} x_{F,i}^{\mathcal{J},*} \right) \right)^+ \right] + \mathbb{E} \left[\left(\sum_{t=1}^T \sum_{\mathcal{J} \ni i} \lambda^{\mathcal{J}} x_{F,i}^{\mathcal{J},*} - \theta_i T \right)^+ \right] \\
& = \mathbb{E} \left[\left(\sum_{t=1}^T \sum_{\mathcal{J} \ni i} \left(D^{\mathcal{J},t} X_{F,i}^{\mathcal{J},t} - \lambda^{\mathcal{J}} x_{F,i}^{\mathcal{J},*} \right) \right)^+ \right] \\
& \leq \sum_{t=1}^T \sum_{\mathcal{J} \ni i} \mathbb{E} \left[\left(D^{\mathcal{J},t} X_{F,i}^{\mathcal{J},t} - \lambda^{\mathcal{J}} x_{F,i}^{\mathcal{J},*} \right)^+ \right] \\
& \leq \sum_{t=1}^T \sum_{\mathcal{J} \ni i} \sqrt{\text{VAR}(D^{\mathcal{J},t} X_{F,i}^{\mathcal{J},t})} \leq M\sqrt{T},
\end{aligned}$$

where the second equality holds because $T \sum_{\mathcal{J} \ni i} \lambda^{\mathcal{J}} x_{F,i}^{\mathcal{J},*} \leq \theta_i T$ due to constraint (3.21), the second to the last inequality comes from the fact that $\mathbb{E} \left[D^{\mathcal{J},t} X_{F,i}^{\mathcal{J},t} \right] = \lambda^{\mathcal{J}} x_{F,i}^{\mathcal{J},*}$, and the last inequality follows because the variance of a bounded random variable $D^{\mathcal{J},t} X_{F,i}^{\mathcal{J},t}$ is bounded. Here, M is a constant independent of T .

In summary, we conclude that $C^{LPR}(T) - \tilde{C}_{LP}^*(T) \leq O(\sqrt{T})$, whereas $\tilde{C}_{LP}^*(T)$ grows linearly in T , completing the proof of (3.74). The proof of Theorem 11 is completed. \square

3.8.12 Other proofs

In this section, we provide the remaining proofs.

Proof. Proof of Corollary 1. From (3.3) and Assumption (A), we have

$$f_R + \alpha c_R > f_F + \alpha c_F, \quad f_R < f_F,$$

which implies

$$\left(\frac{f_R}{f_F + \alpha c_F} \vee \frac{f_F}{f_R + \alpha c_R} \right) \leq \frac{f_F}{f_F + \alpha c_F} \leq 1.$$

Therefore, we conclude that $\text{RATIO}(\text{Myopic}) \leq 2$, completing the proof. \square

Proof. Proof of Corollary 2. Fix a sufficiently small $\epsilon > 0$ and let $c_F = \epsilon^2$, $c_R = \epsilon$, $f_R = 1$, and $f_F = 1 + \epsilon$. From (3.5),

$$\beta = \left\lfloor \frac{f_F}{c_R - c_F} \right\rfloor + 1 \leq \frac{1 + \epsilon}{\epsilon(1 - \epsilon)} + 1.$$

Hence, from (3.9), the lower bound under Assumption (A) is at least

$$1 + \left(\frac{f_F + (\beta + 1)c_F + c_R}{f_R} + \frac{c_F}{c_R - c_F} \right)^{-1} \geq 1 + \left(1 + 2\epsilon + 2\epsilon^2 + \frac{2 + \epsilon}{1 - \epsilon}\epsilon \right)^{-1},$$

which converges to 2 as $\epsilon \rightarrow 0$. The proof is completed. \square

Proof. Proof of Corollary 3. **Under Assumption (B).** Let us explicitly compute

$$\inf_{n \geq 2, n \in \mathbb{Z}} \left[\frac{f_F + nc_F}{f_R - f_F + n(c_R - c_F)} + \frac{f_R + n(f_F + c_R + c_F)}{(n - 1)(f_R - f_F)} \right]. \quad (3.75)$$

Define

$$G(x) \triangleq \frac{f_F + xc_F}{f_R - f_F + x(c_R - c_F)} + \frac{f_R + x(f_F + c_R + c_F)}{(x - 1)(f_R - f_F)} \quad x \geq 2.$$

Then,

$$G'(x) = \frac{c_F f_R - c_R f_F}{[f_R - f_F + x(c_R - c_F)]^2} - \frac{f_R + f_F + c_R + c_F}{(x - 1)^2 (f_R - f_F)}.$$

We consider the following two cases.

Case I: $c_F f_R - c_R f_F \leq 0$. Then, $G(x)$ is non-increasing on $[2, \infty)$ such that the minimizer of (3.75) is $+\infty$.

Case II: $c_F f_R - c_R f_F > 0$. Setting $G'(\bar{x}_*) = 0$ implies

$$\bar{x}_* = \frac{\sqrt{(c_F f_R - c_R f_F) \cdot (f_R - f_F)} + (f_R - f_F) \cdot \sqrt{f_R + f_F + c_R + c_F}}{\sqrt{(c_F f_R - c_R f_F) \cdot (f_R - f_F)} - (c_R - c_F) \cdot \sqrt{f_R + f_F + c_R + c_F}}.$$

It can be checked that $G(x)$ is non-increasing on $[2, \bar{x}_*]$ and non-decreasing on $[\bar{x}_*, +\infty)$ such that the minimizer is \bar{x}_* . After rounding, we define x_* as either $\lfloor \bar{x}_* \rfloor$ or $\lceil \bar{x}_* \rceil$, whichever is better. It follows that the minimizer of (3.75) is equal to $\max\{x_*, 2\}$. Therefore, the minimizer of (3.75) is

$$n_* = \begin{cases} +\infty & \text{if } c_F f_R - c_R f_F \leq 0; \\ \max\{x_*, 2\} & \text{if } c_F f_R - c_R f_F > 0. \end{cases}$$

As an immediate consequence, n_* grows on the order of $\sqrt{f_R}$ as $f_R \rightarrow \infty$ such that

$$1 + \sup_{n \geq 2, n \in \mathbb{Z}} \left(\frac{f_F + n c_F}{f_R - f_F + n(c_R - c_F)} + \frac{f_R + n(f_F + c_R + c_F)}{(n-1)(f_R - f_F)} \right)^{-1}$$

goes to infinity on the order of $\sqrt{f_R}$.

Under Assumption (C). Let us explicitly compute

$$\min_{\alpha' \in \{1, \dots, \alpha-1\}} \left[\frac{f_R + \alpha' (f_F + c_F + c_R)}{(\alpha' - 1)(f_R - f_F)} + \frac{f_F + \alpha' c_F}{f_R - f_F - \alpha' (c_F - c_R)} \right]. \quad (3.76)$$

Define

$$F(x) \triangleq \frac{f_R + x(f_F + c_F + c_R)}{(x-1)(f_R - f_F)} + \frac{f_F + x c_F}{f_R - f_F - x(c_F - c_R)} \quad x \in (1, \alpha - 1).$$

Then,

$$F'(x) \triangleq -\frac{f_R + f_F + c_F + c_R}{(x-1)^2 (f_R - f_F)} + \frac{f_R c_F - f_F c_R}{[f_R - f_F - x(c_F - c_R)]^2},$$

which is non-decreasing in x on $(1, \alpha - 1)$. Hence, $F(x)$ is convex in x on $(1, \alpha - 1)$. Setting $F'(\bar{x}_*) = 0$ implies

$$\bar{x}_* = \frac{\sqrt{f_R + f_F + c_F + c_R} \cdot (f_R - f_F) + \sqrt{f_R c_F - f_F c_R} \cdot \sqrt{f_R - f_F}}{\sqrt{f_R + f_F + c_F + c_R} \cdot (c_F - c_R) + \sqrt{f_R c_F - f_F c_R} \cdot \sqrt{f_R - f_F}}.$$

After rounding, we define x_* as either $\lfloor \bar{x}_* \rfloor$ or $\lceil \bar{x}_* \rceil$, whichever is better. Therefore, the minimizer of (3.76) is equal to

$$\alpha'_* = \begin{cases} 1 & \text{if } x_* < 1; \\ x_* & \text{if } 1 \leq x_* \leq \alpha - 1; \\ \alpha - 1 & \text{if } x_* > \alpha - 1. \end{cases}$$

As an immediate consequence, α and α'_* grow on the order of f_R and $\sqrt{f_R}$, respectively, as $f_R \rightarrow \infty$, such that

$$1 + \max_{\alpha' \in \{1, \dots, \alpha - 1\}} \left(\frac{f_R + \alpha' (f_F + c_F + c_R)}{(\alpha' - 1)(f_R - f_F)} + \frac{f_F + \alpha' c_F}{f_R - f_F - \alpha' (c_F - c_R)} \right)^{-1}$$

goes to infinity on the order of $\sqrt{f_R}$. The proof is completed. □

3.8.13 Other numerical experiments

In this section, we present numerical experiments under Assumptions (A) and (B). The simulation procedure is the same as described in Section 3.6. Again, unless specified otherwise, the model parameters $n, \overline{|o|}, \underline{|o|}, \theta, f_F, f_R, c_F, c_R$ are fixed in this section to provide a clear graphic demonstration. The trends, however, hold generally and are consistent with the ones under Assumptions (C) and (D) as discussed in Section 3.6.

3.8.13.1 Comparing Myopic with LP Rounding

Figure 3.7 demonstrates how the observed performance ratios change as T varies, holding other parameters constant. The maximum and average performance ratios of both Myopic and LP Rounding are low, and their performance ratios decrease and stabilize when T increases. Additionally, LP Rounding is outperformed by Myopic when T is small and outperforms Myopic when T is sufficiently large, consistent with its asymptotic optimality. These trends match what we observe in Section 3.6.1. In the remainder of this subsection, we show our simulation results based on two different T 's, representing small and large time horizons, respectively.

Figures 3.8 and 3.9 compare Myopic with LP Rounding under Assumptions (A) and (B) by varying the RDC fixed cost f_R . The number of different items is small ($n = 10$) in Figure 3.8 and is large ($n = 50$) in Figure 3.9. The trends here again match those in Section 3.6.1. Specifically, both ratios grow with f_R , and both policies perform better for larger T . Finally, after fixing the order-type characteristics ($\overline{|o|}$, $\underline{|o|}$), the gap between LP Rounding and Myopic increases when the total number of items n increases. As the number of items n becomes relatively large, order split becomes less frequent under Myopic, and thus, it performs better. Hence, under Assumptions (A) and (B), Myopic is also valuable and competitive even when demand forecasting is available.

Next, the impact of initial FDC inventory in Figure 3.10 exhibits the same trend under Assumption (A). Both policies perform well when FDC has limited or a sufficiently large amount of inventory, and Myopic's average performance has a competitive advantage over LP Rounding when the time horizon T is small.

3.8.13.2 Delay in order fulfillment

In Figure 3.11, we compare Myopic with Delay (as defined in Section 3.6.2) under Assumption (A). Although Delay is expected to perform no worse than Myopic and performs better with

longer delay intervals, the marginal benefit of Delay decreases as T^d increases. Most of the gap between Myopic and OPT is closed by implementing a delay policy of less than 40 periods.

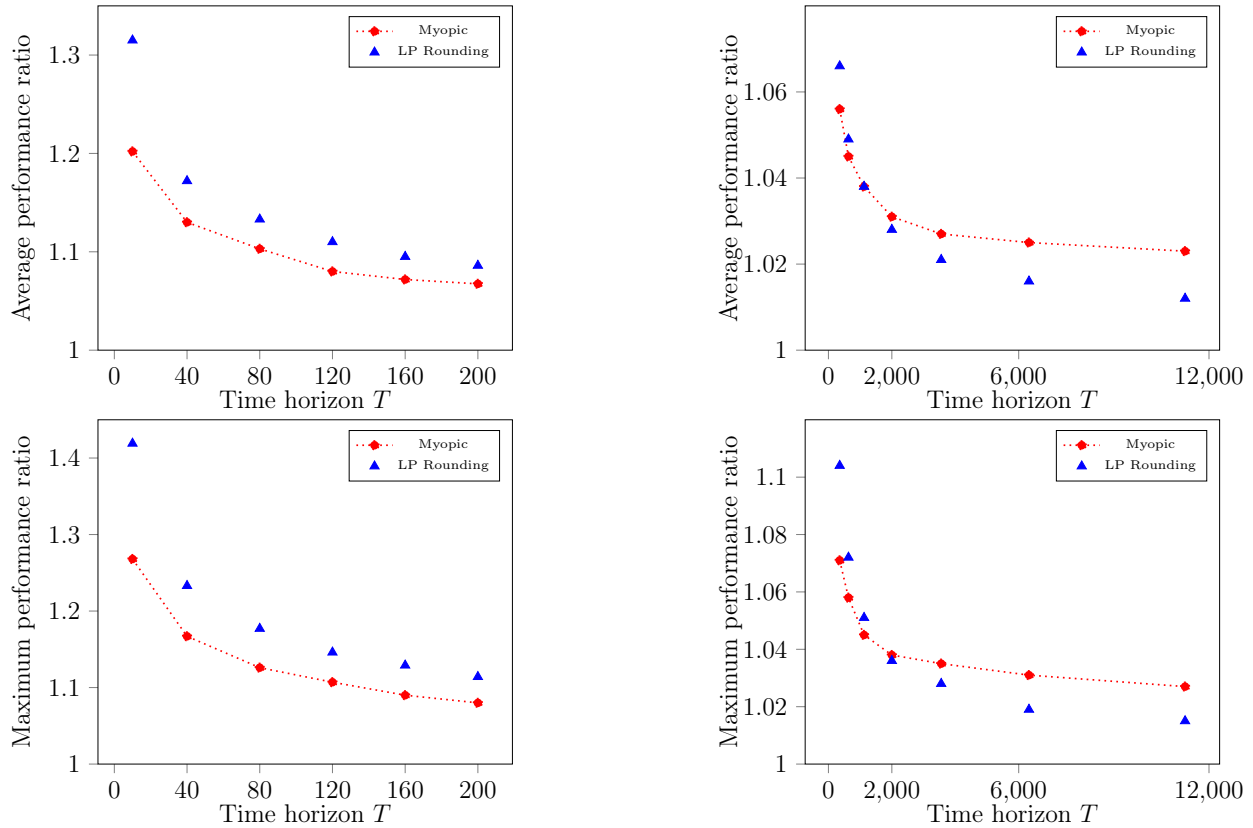


Figure 3.7: Performance w.r.t. time horizon T , holding $n = 10$, $\overline{|o|} = 5$, $\underline{|o|} = 5$, $\theta = 0.8$, $f_F = 100$, $f_R = 80$, $c_F = 0$, $c_R = 50$. Left panel: T from 10 to 200; right panel: T from 356 to 11,246.

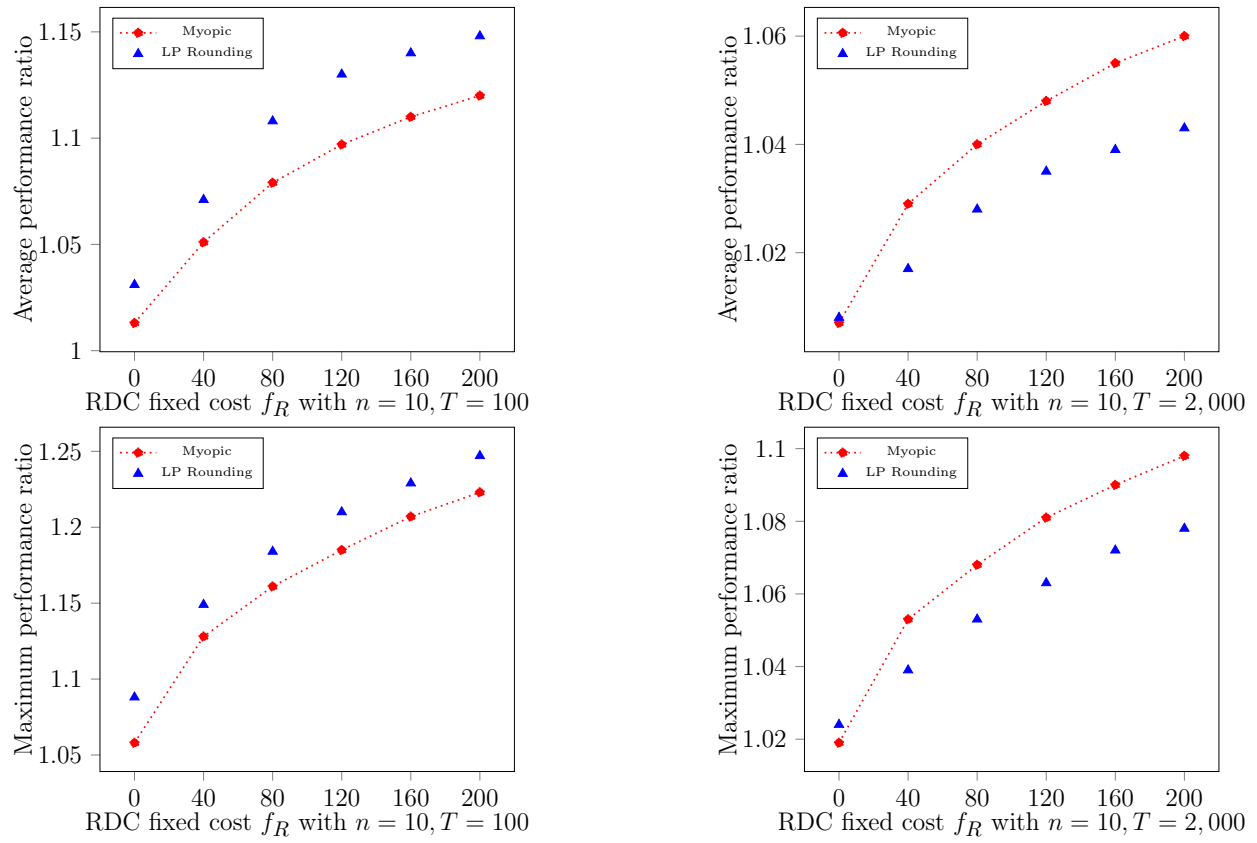


Figure 3.8: Performance w.r.t. RDC fixed cost f_R , holding $\overline{|o|} = 5, \underline{|o|} = 5, \theta = 0.8, f_F = 100, c_F = 0, c_R = 50$. Left panels: $T = 100$; right panels: $T = 2,000$.

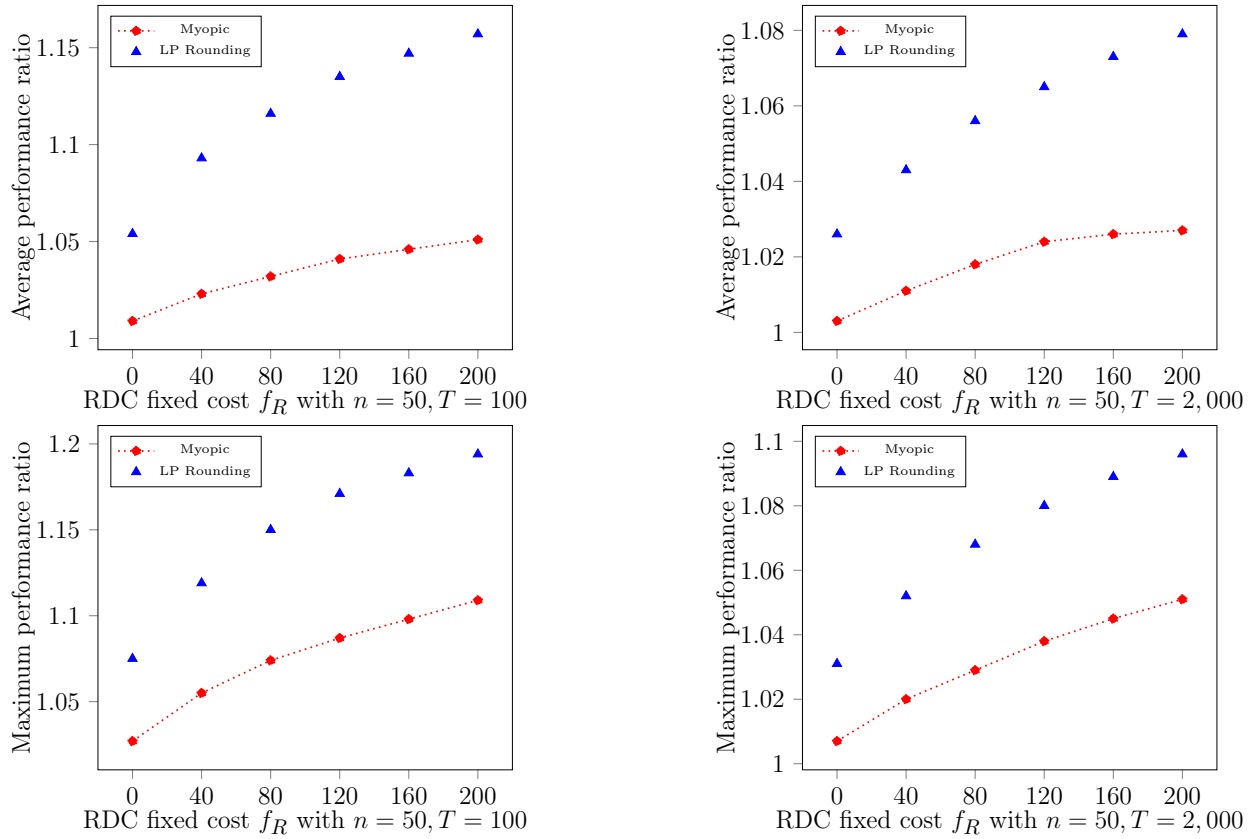


Figure 3.9: Performance w.r.t. RDC fixed cost f_R , holding $\overline{|o|} = 5$, $\underline{|o|} = 5$, $\theta = 0.8$, $f_F = 100$, $c_F = 0$, $c_R = 50$. Left panels: $T = 100$; right panels: $T = 2,000$.




Figure 3.10: Performance w.r.t. FDC inventory scaling parameter θ , holding $n = 10$, $\overline{|o|} = 5$, $\underline{|o|} = 5$, $f_F = 100$, $f_R = 80$, $c_F = 0$, $c_R = 50$. Left panel: $T = 100$; right panel: $T = 2,000$.

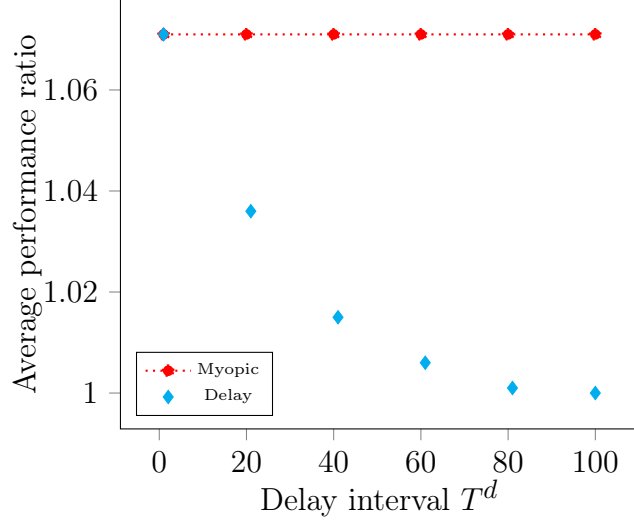


Figure 3.11: Performance w.r.t. order fulfillment delay T^d , holding $T = 100, n = 10, \overline{|o|} = 5, \underline{|o|} = 5, f_F = 100, f_R = 80, c_F = 0, c_R = 50$.

3.8.13.3 Family of threshold policies

Last, we analyze Assumption (A) and propose a family of threshold policies. Recall that Myopic can be viewed as a threshold policy, where the two order-size thresholds, α and β as defined in (3.2) and (3.5), respectively, represent whether the FDC is utilized conditional on its inventory level. Specifically, under Assumption (A), when all the requested items are available at the FDC, Myopic uses the FDC if and only if the order size is greater than or equal to α . When some of the requested items are unavailable at the FDC, the order is split between the two DCs if and only if the number of the remaining FDC in-stock items is greater than or equal to β . Here, Myopic exclusively depends on the cost parameters, whose estimation is possibly inaccurate (or even unavailable) in practice. This observation inspires us to consider the following family of threshold policies that is robust against the sensitivity of the cost parameters. Each threshold policy in this family is associated with two order-size thresholds (analogous to α and β associated with Myopic). We are interested in comparing Myopic with the best policy in this family, denoted as Threshold, which is associated with thresholds α^* and β^* .

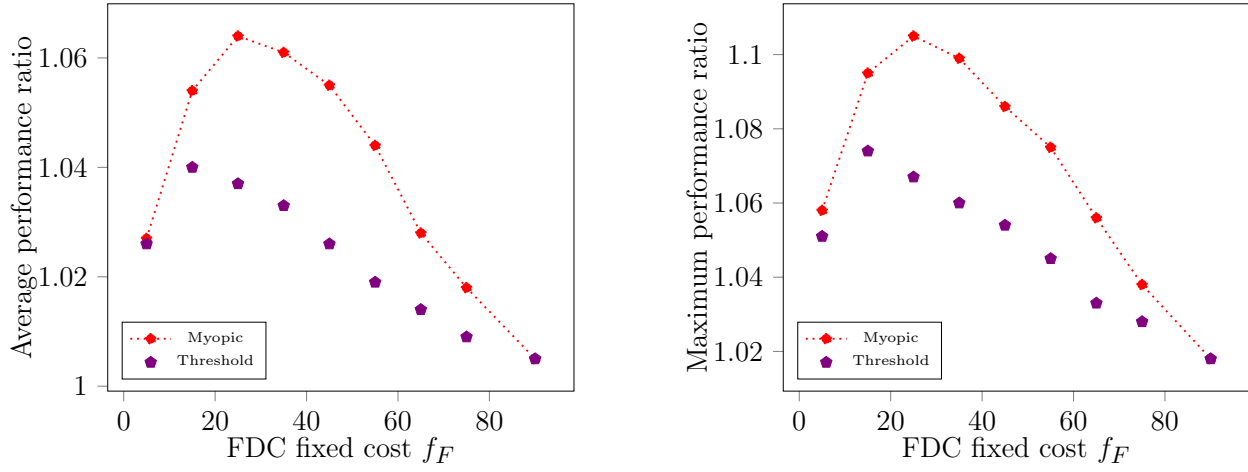


Figure 3.12: Performance w.r.t. FDC fixed cost f_F , holding $T = 100, n = 20, \overline{|o|} = 8, \underline{|o|} = 5, \theta = 0.5, f_R = 4, c_F = 1, c_R = 15$.

Table 3.10: A snapshot of the Myopic and Threshold policy parameters under Assumption (A)

FDC fixed cost f_F	Myopic		Threshold	
	α	β	α^*	β^*
5	1	1	1	2
25	2	2	4	4
45	3	4	5	5
65	5	5	6	6
90	7	7	7	7

We compare the performance ratios of Myopic with those of Threshold by varying FDC cost f_F and report the results in Figure 3.12 and Table 3.10. First, from Table 3.10, one advantage of Threshold is the ability to protect the FDC inventory and to reduce order splits by setting higher thresholds (imagine higher β^* saves on the FDC fixed cost from less order split during stockout events and the saved FDC inventory consequently can be used for the potentially larger orders showing up later). Second, the gap between Myopic and Threshold first increases and then decreases as f_F increases. In particular, when f_F is too large, using the FDC at all is no longer optimal and protecting its inventory is no longer sensible. Hence, both Myopic and Threshold perform similarly and close to OPT. The above phenomena are reflected in both Table 3.10 and Figure 3.12, where the greatest differences between Myopic and Threshold occur when f_F is in the middle range.

3.8.14 Summary of instances

In this section, we summarize all the instances used in the paper. Recall that all the instances have exactly one unit of inventory for all the items.

I_1	Items ($n = m\alpha, m \rightarrow \infty$)			
Orders	$1, \dots, \alpha$	$\alpha + 1, \dots, 2\alpha$...	$(m - 1)\alpha + 1, \dots, m\alpha$
o_1	$1, \dots, 1$	$0, \dots, 0$...	$0, \dots, 0$
o_2	$0, \dots, 0$	$1, \dots, 1$...	$0, \dots, 0$
...
o_m	$0, \dots, 0$	$0, \dots, 0$...	$1, \dots, 1$
o_{m+1}	$1, \dots, 1$	$1, \dots, 1$...	$1, \dots, 1$

Table 3.11: Proof of Theorem 7, Assumption (A).

I_2	Items ($n = m + m\beta, m \rightarrow \infty$)										
Orders	1	2	...	m	$m+1$...	$m+\beta$...	$m+(m-1)\beta+1$...	$m+m\beta$
o_1	1	1	...	1	0	...	0	...	0	...	0
o_2	1	0	...	0	1	...	1	...	0	...	0
...
o_{m+1}	0	0	...	1	0	...	0	...	1	...	1

Table 3.12: Proof of Theorem 7, Assumption (A).

I_1	Items ($n \rightarrow \infty$)				I_2	Items ($n = \beta + 1$)			
Orders	1	2	...	n	Orders	1	2	...	n
o_1	1	1	...	1	o_1	1	0	...	0
o_2	1	0	...	0	o_2	1	1	...	1
o_3	0	1	...	0					
...					
o_{n+1}	0	0	...	1					

Table 3.13: Proof of Theorem 7, Assumption (B).

I_1	Items ($n = \alpha - 1$)			
Orders	1	2	...	$\alpha - 1$
o_1	1	1	...	1
o_2	1	0	...	0
o_3	0	1	...	0
...
o_{n+1}	0	0	...	1

Table 3.14: Proof of Theorem 7, Assumption (C).

I_1	Items ($n = m + m\beta$, $m \rightarrow \infty$ or $m = \alpha$)						
Orders	1	2	...	m	$m + 1$...	$m + m\beta$
o_1	1	1	...	1	0	...	0

Table 3.15: Proof of Theorem 8, Assumption (A).

I_2	Items ($n = m + m\beta$, $m \rightarrow \infty$ or $m = \alpha$)										
Orders	1	2	...	m	$m + 1$...	$m + \beta$...	$m + (m - 1)\beta + 1$...	$m + m\beta$
o_1	1	1	...	1	0	...	0	...	0	...	0
o_2	1	0	...	0	1	...	1	...	0	...	0
...
o_{m+1}	0	0	...	1	0	...	0	...	1	...	1

Table 3.16: Proof of Theorem 8, Assumption (A).

I_1	Items ($n \geq 2$)				I_2	Items ($n \geq 2$)			
Orders	1	2	...	n	Orders	1	2	...	n
o_1	1	1	...	1	o_1	1	1	...	1
					o_2	1	0	...	0
					o_3	0	1	...	0
				
					o_{n+1}	0	0	...	1

Table 3.17: Proofs of Theorem 8, Assumption (B) and Theorem 10.

I_1	Items ($n = \alpha'$, $1 \leq \alpha' \leq \alpha - 1$)				I_2	Items ($n = \alpha'$, $1 \leq \alpha' \leq \alpha - 1$)			
Orders	1	2	...	α'	Orders	1	2	...	α'
o_1	1	1	...	1	o_1	1	1	...	1
					o_2	1	0	...	0
					o_3	0	1	...	0
				
					o_{n+1}	0	0	...	1

Table 3.18: Proof of Theorem 8, Assumption (C).

REFERENCES

- Acimovic, Jason, Vivek F. Farias. 2019. The fulfillment-optimization problem. *Operations Research & Management Science in the Age of Analytics*. INFORMS, 218–237.
- Acimovic, Jason, Stephen C. Graves. 2015. Making better fulfillment decisions on the fly in an online retail environment. *Manufacturing & Service Operations Management* **17**(1) 34–51.
- Acimovic, Jason, Stephen C Graves. 2017. Mitigating spillover in online retailing via replenishment. *Manufacturing & Service Operations Management* **19**(3) 419–436.
- Akturk, M. Selim, Elif Gorgulu. 1999. Match-up scheduling under a machine breakdown. *European Journal of Operational Research* **112**(1) 81–97.
- Amazon. 2020. URL <https://www.aboutamazon.com/amazon-fulfillment/our-fulfillment-centers>.
- Andrews, John M, Vivek F Farias, Aryan I Khojandi, Chad M Yan. 2019. Primal–dual algorithms for order fulfillment at urban outfitters, inc. *INFORMS Journal on Applied Analytics* **49**(5) 355–370.
- Asadpour, Arash, Xuan Wang, Jiawei Zhang. 2020. Online resource allocation with limited flexibility. *Management Science* **66**(2) 642–666.
- Ball, Michael O, Maurice Queyranne. 2009. Toward robust revenue management: Competitive analysis of online booking. *Operations Research* **57**(4) 950–963.
- Bartholdi III, John J., Kevin R. Gue. 2000. Reducing labor costs in an ltl crossdocking terminal. *Operations Research* **48**(6) 823–832.
- Bean, James C., John R. Birge, John Mittenthal, Charles E. Noon. 1991. Matchup scheduling with multiple resources, release dates and disruptions. *Operations Research* **39**(3) 470–483.
- Berg, Natalie, Miya Knights. 2021. *Amazon: How the world’s most relentless retailer will continue to revolutionize commerce*. Kogan Page Publishers.
- Borodin, Allan, Ran El-Yaniv. 2005. *Online computation and competitive analysis*. Cambridge university press.
- Buchbinder, Niv, Tracy Kimbrel, Retsef Levi, Konstantin Makarychev, Maxim Sviridenko. 2013. Online make-to-order joint replenishment model: Primal-dual competitive algorithms. *Operations Research* **61**(4) 1014–1029.
- Caggiano, Kathryn E., Peter L. Jackson, John A. Muckstadt, James A. Rappold. 2007. Optimizing service parts inventory in a multiechelon, multi-item supply chain with time-based customer service-level agreements. *Operations Research* **55**(2) 303–318.

- Cattani, Kyle D., Gilvan C. Souza, Shengqi Ye. 2014. Shelf loathing: cross docking at an online retailer. *Production and Operations Management* **23**(5) 893–906.
- Chen, Annie I., Stephen C. Graves. 2021. Item aggregation and column generation for online-retail inventory placement. *Manufacturing & Service Operations Management* **23**(5) 1062–1076.
- Chen, Fangruo, Jing-Sheng Song. 2001. Optimal policies for multiechelon inventory problems with markov-modulated demand. *Operations Research* **49**(2) 226–234.
- Chen, Xi, Will Ma, David Simchi-Levi, Linwei Xin. 2020. Assortment planning for recommendations at checkout under inventory constraints. *Available at SSRN 2853093* .
- Chopra, Sunil. 2017. Seven-Eleven Japan Co. *Kellogg School of Management Cases* .
- Clark, Andrew J., Herbert Scarf. 1960. Optimal policies for a multi-echelon inventory problem. *Management Science* **6**(4) 475–490.
- Colliers. 2020. URL <https://www.colliers.com/en/download-research?itemId=bbb452d6-8d01-4560-8392-d79653e64a44>.
- Dastin, Jeffrey. 2020. URL <https://www.reuters.com/article/us-amazon-com-delivery/amazon-adds-warehouse-network-closer-to-cities-to-speed-up-same-day-delivery-idUSKBN20Q0T3>.
- Deming, W. Edwards. 1975. On some statistical aids toward economic production. *Interfaces* **5**(4) 1–15.
- DeValve, Levi, Yehua Wei, Di Wu, Rong Yuan. 2021. Understanding the value of fulfillment flexibility in an online retailing environment. *Manufacturing & Service Operations Management* **to appear**.
- Dođru, Mustafa K., Ton De Kok, Geert-Jan van Houtum. 2009. A numerical study on the effect of the balance assumption in one-warehouse multi-retailer inventory systems. *Flexible services and manufacturing journal* **21**(3) 114–147.
- Dođru, Mustafa K., Martin I. Reiman, Qiong Wang. 2010. A stochastic programming based inventory policy for assemble-to-order systems with application to the W model. *Operations Research* **58**(4-part-1) 849–864.
- Dođru, Mustafa K., Martin I. Reiman, Qiong Wang. 2017. Assemble-to-order inventory management via stochastic programming: Chained BOMs and the M-system. *Production and Operations Management* **26**(3) 446–468.
- Drent, Melvin, Joachim Arts. 2021. Expediting in two-echelon spare parts inventory systems. *Manufacturing & Service Operations Management* **23**(6) 1431–1448.
- Durrett, Richard. 1999. *Essentials of Stochastic Processes*, vol. 1. Springer.

- Dutkay, Dorin Ervin, Constantin P. Niculescu, Florin Popovici. 2013. Stirling’s formula and its extension for the Gamma function. *The American Mathematical Monthly* **120**(8) 737–740.
- Elmachtoub, Adam N, Retsef Levi. 2016. Supply chain management with online customer selection. *Operations Research* **64**(2) 458–473.
- Eppen, Gary, Linus Schrage. 1981. Centralized ordering policies in a multi-warehouse system with lead times and random demand. Leroy B. Schwarz, ed., *TIMS Studies in the Management Sciences, Vol. 16: Multi-Level Production/Inventory Control Systems: Theory and Practice*. North-Holland, 51–67.
- Federgruen, Awi, Paul Zipkin. 1984a. Allocation policies and cost approximations for multilocation inventory systems. *Naval Research Logistics Quarterly* **31**(1) 97–129.
- Federgruen, Awi, Paul Zipkin. 1984b. Approximations of dynamic, multilocation production and inventory problems. *Management Science* **30**(1) 69–84.
- Federgruen, Awi, Paul Zipkin. 1984c. Computational issues in an infinite-horizon, multiechelon inventory model. *Operations Research* **32**(4) 818–836.
- FedEx. 2021. Fedex service guide .
- Ge, Deng, Yi Pan, Zuo-Jun Max Shen, Di Wu, Rong Yuan, Chao Zhang. 2019. Retail supply chain management: a review of theories and practices. *Journal of Data, Information and Management* **1**(1) 45–64.
- Golrezaei, Negin, Hamid Nazerzadeh, Paat Rusmevichientong. 2014. Real-time optimization of personalized assortments. *Management Science* **60**(6) 1532–1551.
- Govindarajan, Aravind, Amitabh Sinha, Joline Uichanco. 2021. Joint inventory and fulfillment decisions for omnichannel retail networks. *Naval Research Logistics (NRL)* **68**(6) 779–794.
- Harsha, Pavithra, Shivaram Subramanian, Joline Uichanco. 2019. Dynamic pricing of omnichannel inventories: honorable mention—2017 m&som practice-based research competition. *Manufacturing & Service Operations Management* **21**(1) 47–65.
- Jasin, Stefanus, Amitabh Sinha. 2015. An LP-based correlated rounding scheme for multi-item ecommerce order fulfillment. *Operations Research* **63**(6) 1336–1351.
- Lawson, David G., Evan L. Porteus. 2000. Multistage inventory management with expediting. *Operations Research* **48**(6) 878–893.
- Lei, Yanzhe, Stefanus Jasin, Amitabh Sinha. 2018. Joint dynamic pricing and order fulfillment for e-commerce retailers. *Manufacturing & Service Operations Management* **20**(2) 269–284.

- Li, Xiaobo, Yun Fong Lim, Fang Liu. 2019. Optimal policies and heuristics to match supply with demand for online retailing. *Available at SSRN 3280752* .
- Lim, Yun Fong, Song Jiu, Marcus Ang. 2021. Integrating anticipative replenishment allocation with reactive fulfillment for online retailing using robust optimization. *Manufacturing & Service Operations Management* **23**(6) 1616–1633.
- Ma, Will, David Simchi-Levi. 2020. Algorithms for online matching, assortment, and pricing with tight weight-dependent competitive ratios. *Operations Research* **68**(6) 1787–1803.
- Ma, Will, David Simchi-Levi, Chung-Piaw Teo. 2021. On policies for single-leg revenue management with limited demand information. *Operations Research* **69**(1) 207–226.
- Macy's. 2020. URL https://www.customerservice-macys.com/app/answers/detail/a_id/284/~why-is-my-order-arriving-in-multiple-shipments%3F#:~:text=Why%20is%20my%20order%20arriving,Availability%20of%20items%20ordered.
- Mak, Ho-Yin, Zuo-Jun Shen. 2012. Risk diversification and risk pooling in supply chain design. *IIE Transactions* **44**(8) 603–621.
- Mamani, Hamed, Kamran Moinzadeh. 2014. Lead time management through expediting in a continuous review inventory system. *Production and Operations Management* **23**(1) 95–109.
- Moinzadeh, Kamran, Prabhu K. Aggarwal. 1997. An information based multiechelon inventory system with emergency orders. *Operations Research* **45**(5) 694–701.
- Moinzadeh, Kamran, Charles P. Schmidt. 1991. An (S-1, S) inventory system with emergency orders. *Operations Research* **39**(2) 308–321.
- Muckstadt, John, Michael G. Klein, Peter L. Jackson, Robert M. Gougelet, Nathaniel Hupert. 2021. Application of supply chain principles to pandemic planning. *Available at SSRN 3794252* .
- Muckstadt, John A., L. Joseph Thomas. 1980. Are multi-echelon inventory methods worth implementing in systems with low-demand-rate items? *Management Science* **26**(5) 483–494.
- Muharremoglu, Alp, John N. Tsitsiklis. 2008. A single-unit decomposition approach to multiechelon inventory systems. *Operations Research* **56**(5) 1089–1103.
- Nocedal, Jorge, Stephen J. Wright. 1999. *Numerical Optimization*. Springer.
- Peng, Peng, Lawrence V. Snyder, Andrew Lim, Zuli Liu. 2011. Reliable logistics networks design with facility disruptions. *Transportation Research Part B: Methodological* **45**(8) 1190–1211.

- Perakis, Georgia, Divya Singhvi, Yiannis Spantidakis. 2020. Leveraging the newsvendor for inventory distribution at a large fashion e-retailer with depth and capacity constraints. *Available at SSRN 3632459* .
- Press, William H, Brian P Flannery, Saul A Teukolsky, William T Vetterling. 1992. Incomplete gamma function, error function, chi-square probability function, cumulative Poisson function. *Numerical Recipes in Fortran*. Cambridge University Press, New York City, 209–215.
- Qin, Hengle, Jun Xiao, Dongdong Ge, Linwei Xin, Jianjun Gao, Simai He, Haodong Hu, John Gunnar Carlsson. 2022. Jd.com: Operations research algorithms drive intelligent warehouse robots to work. *INFORMS Journal on Applied Analytics* **52**(1) 42–55.
- Reiman, Martin I., Qiong Wang. 2015. Asymptotically optimal inventory control for assemble-to-order systems with identical lead times. *Operations Research* **63**(3) 716–732.
- Roundy, Robin. 1985. 98%-effective integer-ratio lot-sizing for one-warehouse multi-retailer systems. *Management science* **31**(11) 1416–1430.
- Shang, Kevin H., Jing-Sheng Song. 2003. Newsvendor bounds and heuristic for optimal policies in serial supply chains. *Management Science* **49**(5) 618–638.
- Shannon, Claude E. 1948. A mathematical theory of communication. *The Bell System Technical Journal* **27**(3) 379–423. doi:10.1002/j.1538-7305.1948.tb01338.x.
- Shen, Max, Christopher S Tang, Di Wu, Rong Yuan, Wei Zhou. 2020. Jd.com: Transaction-level data for the 2020 msom data driven research challenge. *Manufacturing & Service Operations Management* .
- Shen, Xiaobei, Yimin Yu, Jing-Sheng Song. 2022. Optimal policies for a multi-echelon inventory problem with service time target and expediting. *Manufacturing & Service Operations Management* .
- Simchi-Levi, David, William Schmidt, Yehua Wei. 2014. From superstorms to factory fires: Managing unpredictable supply chain disruptions. *Harvard Business Review* **92**(1-2) 96–101.
- Simchi-Levi, David, William Schmidt, Yehua Wei, Peter Yun Zhang, Keith Combs, Yao Ge, Oleg Gusikhin, Michael Sanders, Don Zhang. 2015. Identifying risks and mitigating disruptions in the automotive supply chain. *Interfaces* **45**(5) 375–390.
- Snyder, Lawrence V., Maria P. Scaparra, Mark S. Daskin, Richard L. Church. 2006. Planning for disruptions in supply chain networks. *Models, Methods, and Applications for Innovative Decision Making*. INFORMS, 234–257.
- Song, Jing-Sheng, Susan H. Xu, Bin Liu. 1999. Order-fulfillment performance measures in an assemble-to-order system with stochastic leadtimes. *Operations Research* **47**(1) 131–149.

- Soper, Spencer. 2020. URL <https://finance.yahoo.com/news/amazon-plans-put-1-00-090053785.html>.
- Statista. 2020a. URL <https://www.statista.com/statistics/272391/us-retail-e-commerce-sales-forecast/>.
- Statista. 2020b. URL <https://www.statista.com/chart/17207/amazon-shipping-and-fulfillment-costs/>.
- Stein, Charles M. 1981. Estimation of the mean of a multivariate normal distribution. *The Annals of Statistics* 1135–1151.
- Sun, Libo, Guodong Lyu, Yugang Yu, Chung-Piaw Teo. 2020. Fulfillment by amazon versus fulfillment by seller: An interpretable risk-adjusted fulfillment model. *Naval Research Logistics (NRL)* **67**(8) 627–645.
- Talluri, Kalyan T., Garrett Van Ryzin. 2004. *The Theory and Practice of Revenue Management*, vol. 1. Springer.
- Treharne, James T., Charles R. Sox. 2002. Adaptive inventory control for nonstationary demand and partial information. *Management Science* **48**(5) 607–624.
- Van Belle, Jan, Paul Valckenaers, Dirk Cattrysse. 2012. Cross-docking: State of the art. *Omega* **40**(6) 827–846.
- Van der Vaart, H.R. 1961. Some extensions of the idea of bias. *The Annals of Mathematical Statistics* **32** 436–447.
- Waller, Matthew A., C. Richard Cassady, John Ozment. 2006. Impact of cross-docking on inventory in a decentralized retail supply chain. *Transportation Research Part E: Logistics and Transportation Review* **42**(5) 359–382.
- Wang, Xinshang, Van-Anh Truong. 2018. Multi-priority online scheduling with cancellations. *Operations Research* **66**(1) 104–122.
- Ward, Tom. 2020. URL <https://corporate.walmart.com/newsroom/2021/01/27/from-ground-breaking-to-breaking-ground-walmart-begins-to-scale-market-fulfillment-centers>.
- Wei, Lai, Roman Kapuscinski, Stefanus Jasin. 2021. Shipping consolidation across two warehouses with delivery deadline and expedited options for e-commerce and omni-channel retailers. *Manufacturing & Service Operations Management* **23**(6) 1634–1650.
- Wulfraat, Marc. 2020. URL https://mwpv1.com/html/amazon_com.html.
- Xu, Ping Josephine, Russell Allgor, Stephen C. Graves. 2009. Benefits of reevaluating real-time order fulfillment decisions. *Manufacturing & Service Operations Management* **11**(2) 340–355.

- Xu, Zhen, Hailun Zhang, Jiheng Zhang, Rachel Q Zhang. 2020. Online demand fulfillment under limited flexibility. *Management Science* **66**(10) 4667–4685.
- Yao, Andrew Chi-Chin. 1977. Probabilistic computations: Toward a unified measure of complexity. *18th Annual Symposium on Foundations of Computer Science (sfcs 1977)*. IEEE Computer Society, 222–227.
- Zipkin, Paul. 1984. On the imbalance of inventories in multi-echelon systems. *Mathematics of Operations Research* **9**(3) 402–423.