Practical Algorithms for Managing Uncertain Demand in Complex Systems

by

Levi DeValve

Business Administration
Duke University

Date: ____________________

Approved:

____
Aleksandar Pekeč, Supervisor

____
Yehua Wei

____
Alessandro Arlotto

____
Jing-Sheng (Jeannette) Song

Dissertation submitted in partial fulfillment of the
requirements for the degree of Doctor of Philosophy
in the Department of Business Administration
in the Graduate School of
Duke University

2019

ABSTRACT

Practical Algorithms for Managing Uncertain Demand in Complex Systems

by

Levi DeValve

Business Administration
Duke University

Date: ______________________

Approved:

____
Aleksandar Pekeč, Supervisor

____
Yehua Wei

____
Alessandro Arlotto

____
Jing-Sheng (Jeannette) Song

An abstract of a dissertation submitted in partial fulfillment of the
requirements for the degree of Doctor of Philosophy
in the Department of Business Administration
in the Graduate School of
Duke University

2019

Abstract

This dissertation considers complex operational settings with uncertain demand. We consider three classes of problems, and show that for each one, a well designed policy can give close to optimal performance for either cost or profit metrics. To prove our results we use a variety of techniques from the theory of optimization and applied probability.

First, we study assemble-to-order (ATO) problems, where a set of products are assembled from a set of common components. ATO problems with general structure and integrality constraints are well known to be difficult to solve, and we provide new insight into these issues by establishing worst-case approximation guarantees through various primal-dual analyses and LP rounding. First, we relax the one-period ATO problem using a natural newsvendor decomposition and use the dual solution for the relaxation to derive a lower bound on optimal cost, providing a tight approximation guarantee that grows with the maximum product size in the system. Then, we present LP rounding algorithms that achieve both asymptotic optimality as demand grows large, and a 1.8 approximation factor for any problem instance. Finally, we demonstrate that our one-period LP rounding results can be extended to analyze dynamic ATO problems. Specifically, we use our rounding scheme to develop an asymptotically optimal *integral* policy for dynamic ATO problems with backlogging and identical component lead-times.

In the second class of problems, we consider a two-sided market with heterogeneous and stochastic values of matches between different units of supply and demand (such as those embedded in a network). With limited resources, a service provider can only choose a subset of locations in the market where matches can be made, and wants to choose a subset maximizing the expected value of matches. This problem captures decisions in a wide range of application domains, including those faced by franchise service providers, transportation-service platforms, and online advertisers. We demonstrate that the problem is NP-hard to approximate with a factor better than 1.58, and that the objective function does not enjoy the classic submodularity property, which would allow the application of known approximation guarantees for submodular function maximization. To overcome these hurdles, we propose a generalization of submodularity, called γ-*cover modularity*, which allows us to prove approximation guarantees of 3 and 2.31 for an exchange algorithm and

greedy algorithm, respectively. In order to establish these results, we prove general approximation guarantees for γ-cover modular functions.

For the final problem, we apply the classical operations management strategy of flexibility design to the context of online retailing fulfillment. More specifically, our method helps online retailers to make the strategic selection of a flexible fulfillment network and to improve their tactical decisions on which distribution center to use to fulfill online orders. To derive our results and insights, we use LP-based asymptotic analysis, stochastic programming, concepts from revenue management, and numerical simulations. We collaborate with a large online retailer to propose a data-driven online resource allocation model which captures the key features of the retailer's fulfillment network. We identify settings where the naive greedy policy performs poorly, and propose a *spillover limit* policy that is asymptotically optimal under a general setting. Finally, we use the spillover limit policy and a stochastic programming allocation policy to evaluate the benefit of a proposed flexible fulfillment network using data from our industrial collaborator. We estimate that our proposed flexible fulfillment network decreases the lost sales plus fulfillment costs of our industrial partner by one percent, which equates to a profit improvement on the order of tens of millions of U.S. dollars.

Acknowledgements

I acknowledge a deep debt and express immense gratitude to a host of people who have supported me throughout the long path to reach this milestone. Without the aid of these many, I would not have achieved this goal, nor much else in my life. I have been tremendously blessed, and owe very much to the help of others.

I would first like to thank my advisors, Saša Pekeč and Yehua Wei. Their unique combination of insight, creativity, technical skill, experience, patience, and encouragement have provided both necessary and sufficient conditions for me to succeed academically. I will be forever grateful for the immense amount of time and energy they each have expended on my behalf, and count myself very fortunate to have had the privilege of such expert investment in my life.

As the PhD program coordinator, Saša originally recruited me to Duke by convincing me it would be a great place to do strong methodological research and raise my family. Duke delivered high above expectations on both dimensions, and I thank Saša for being an integral part of bringing me here. We began working together early on, and have enjoyed all the ups and downs of research together. I appreciate Saša's quick wit, deep mathematical expertise, knack for story telling, and compassion for people. These qualities have cultivated my research, but more importantly they have enriched my life, and I am incredibly grateful for having the chance to work with Saša.

I began working with Yehua during my second year at Duke, and immediately felt both very supported and incredibly challenged. These simultaneous conditions provided the structure and motivation for me to be stretched technically well beyond what I thought possible. We read and discussed countless papers, attended CS classes, and tried many failed (and a few successful) proofs together. These were productive endeavors, but also incredibly fun, due to Yehua's good humor, easy smile, genuine technical brilliance, and quick intuition. I aspire to Yehua's ability to both conceive a proof, and simply explain its practical impact, and I am undeservedly fortunate to have the opportunity to learn from him.

I would also like to express my deepest gratitude to my committee members, Alessandro Arlotto and Jeannette Song. Talking with Alessandro is a genuine treat; I can always count on deep insight, an update on current events, and a hearty laugh. His support and advice for my research and career has been invaluable. I am also deeply indebted to Jeannette for numerous very helpful conversations and her continued support of my research pursuits. I always enjoy our discussions, and I am grateful for and impressed by her deep expertise on a range of topics.

There were of course many other faculty members at Duke who greatly enriched my experience. I thank Peng Sun and David Brown for their optimization classes that proved foundational to much of my research, and Bob Winkler, Alessandro Arlotto, Bob Nau, Bora Keskin, Jeannette Song, and Robert Swinney for their classes and great conversations. I am also thankful for many very helpful conversations and good advice from Alex Belloni, Ali Makhdoumi, Ozan Candogan, Santiago Balseiro, Jim Smith, and many others. I am also extremely fortunate to have met so many great friends in the Duke Ph.D. program, including Huseyin Gurkan, Xinchang Xie, Dip Chakraborty, Mingliu Chen, Chen Chen, Andrew Frazelle, Yunke Mai, Asa Palley, Safak Yucel, Sajad Modaresi, and many others, with whom I spent many a Fuqua Friday, played basketball, attended classes, and had a lot of fun.

I also acknowledge the formative years I spent in consulting at ZS Associates, and the support I received from many great mentors there, including Daniel Brox, John Sadlow, Adam Schneider, and Megan Smith. My life has also forever been shaped by the mentoring and friendship I experienced with Dan, Michael, Ross, and others during my time with Cru. And all of this would not have been possible without the great education I received from UConn, and support from many faculty members, including Tom Roby, Alvaro Lozano-Robledo, Keith Conrad, Bill Abikoff, Alex Teplyaev, and Ben Steinhurst, among others.

To my wife, Shannon, I could not possibly express in words the extent of my gratitude for you. I admire your wisdom, caring heart, intuition, and understanding of people, and greatly appreciate your smile, companionship, and never-ending work in caring for our daughters. You are the best woman, wife, and mother, without qualifier, and I am thankful for the profound improvements your presence has made in every area of my life. Neither this dissertation, nor much else positive in my life, would have been possible without your unwavering support. To our daughters, your arrivals and lives have coincided with the entirety of my graduate studies, and your presence has been a source of unending joy and rejuvenation throughout this time. Your endless energy, relentless curiosity, and authentic friendships with one another are inspiring, and I am incredibly blessed to leave graduate school with my quiver full.

To my father and mother, your enduring support and encouragement have helped carry me through many milestones in life, and I will never be able to thank you enough for teaching me the responsibility, faith, independence, creativity, and work ethic required to be a man. To my brothers, Caleb, Jacob, and Seth, thank you for your shaping influence while growing up, your current support of me and my family, and most of all, your friendships. To Chris and Karen, thanks for being amazing in-laws and for all your support of our family during my studies. I am also thankful for the many friendships cultivated at Summit Church in Durham, including Stephen, Jordan, David, Andrew, Brandon, and many others. I am fortunate to have many incredible people in my life, including those who I have mentioned here, and many more I do not have space to mention by name: teachers, coaches, mentors, friends; and I thank all of them for their support.

Most importantly, I acknowledge Jesus Christ as the source of everything I am thankful for, who gave his life for mine to make me his own. I thank God, who "gives everyone life and breath and everything else," "for from him and through him and for him are all things. To him be the glory forever!"

Contents

List of Figures

List of Tables

Chapter 1

Introduction

This dissertation considers complex operational settings with uncertain demand. As novel strategies and technologies transform the way customers engage with businesses, new and complex operational decisions emerge as well, whose efficient, real-time solutions are essential for delivering value to customers. The complexities of the systems we consider arise from various sources, including multiple, inter-related products, components, warehouses, and locations. Demand uncertainty in these systems is a consequence of the fact that many operational decisions need to be made today in anticipation of an unknown customer that will arrive tomorrow. This dissertation attempts to answer the question: "How well can relatively straightforward and practical algorithms perform in such complex settings?"

We consider three classes of problems, and show that for each one, a well designed policy can give close to optimal performance for either cost or profit metrics. To prove our results we use a variety of techniques from the theory of optimization and applied probability. The goal of this research is to provide a theoretical basis for understanding the performance of intuitive methods of dealing with complex decisions, in order to provide managers with practical decision-making strategies that can be implemented with confidence.

1.1 Assemble-to-Order Systems

In Chapters 2 and 3, we study assemble-to-order (ATO) problems from the literature, which model the fulfillment of demand for products that are made from a set of common components. In the one-period problem considered in Chapter 2, component inventories are ordered before demand is realized (i.e., the first stage problem), then, after demand arrives, components are allocated to products for fulfillment (i.e., the second stage problem). A prominent example faced by large-scale e-retailers, like JD.com or Amazon, is fulfilling orders consisting of multiple items, while other examples include automobile manufacturing and laptop assembly. ATO problems with general structure and integrality constraints are well known to be difficult to solve, stemming from the fact that the second stage problem is not just NP-hard to solve, but also NP-hard to approximate with a constant factor.

We overcome these hurdles by developing a new rounding algorithm based on the linear programming (LP) relaxation of the problem, which we use to approximate the first stage cost and expected second stage cost together. In particular, our LP rounding algorithm for the ATO problem provides both asymptotic optimality in large demand regimes and a constant approximation factor of 1.8 for any problem instance; each of which are the first provable approximation guarantee of their kind. To prove these results, we leverage the complementary slackness conditions of the optimal LP solution to bound the cost of the rounded solution in terms of the optimal dual cost. We also show through simulations that the algorithm is effective across a wide range of problem instances, typically providing much better performance than the worst case guarantee. Finally, we show that the ATO problem is NP-hard through a reduction from the classic set cover problem, establishing the first formal complexity results for ATO.

We further investigate these issues in a dynamic setting in Chapter 3. In particular, we demonstrate that our LP rounding analysis for the two-stage problem can be extended to analyze dynamic ATO problems. First, we use our rounding scheme to develop the first asymptotically optimal *integral* policy for dynamic ATO problems with backlogging and identical component lead-times. Second, we show that our solution for the one-period problem can also be used to design a well-performing policy for a lost sales model where allocation decisions need to be made in an online fashion. These results demonstrate that our analysis of the one-period problem in Chapter 2 provides useful insight into managing a range of dynamic ATO systems.

The research in these chapters was conducted under the supervision of Saša Pekeč and Yehua Wei and the analysis is largely included in DeValve et al. (2019b).

1.2 Matching Supply and Demand in a Resource Constrained Service Network

Chapter 4 considers a market where both demand and supply are uncertain. This phenomenon occurs naturally in many two-sided market settings, such as those faced by franchise service providers, transportation-service platforms, and online advertisers. We capture the fundamental aspects of these settings with a model of uncertain supply and demand embedded in a network. With limited resources, a service provider can only choose a subset of locations in the market where matches between supply and demand can be made, and wants to choose a subset maximizing the expected value of matches. We demonstrate that the problem is NP-hard to approximate with a factor better than $1/(1-e^{-1}) \approx 1.58$, and that the objective function does not enjoy the classic submodularity property, which would allow the application of known approximation guarantees for submodular function maximization.

To overcome these obstacles, we propose a generalization of submodularity, called *γ-cover modularity*, which allows us to prove approximation factors of 3 and $2/(1-e^{-2}) \approx 2.31$ for an exchange algorithm and greedy algorithm, respectively. In order to establish these results, we prove general approximation guarantees for γ-cover modular functions. These results provide the insight that selecting locations in a locally optimal fashion on the network will perform well relative to the global optimum, which can be helpful to managers looking for an implementable solution. We also discuss a variety of applications domains where these results could be helpful.

The research in this chapter was conducted under the supervision of Saša Pekeč and Yehua Wei and the analysis is largely included in DeValve et al. (2019a).

1.3 Understanding the Value of Fulfillment Flexibility in an Online Retailing Environment

Chapter 5 considers a problem that is motivated by our work with an industrial partner, one of the largest global e-retailers, who seeks to identify an implementable strategy that can efficiently manage order fulfillment across several warehouses at a low cost. Our research applies the classical operations management strategy of flexibility design to the context of online retailing fulfillment. Together with our partner, we develop a general method that will help online retailers to make the strategic selection of a flexible fulfillment network and to improve their tactical decisions on which distribution center to use to fulfill online orders.

We propose a data-driven online resource allocation model which captures the key features of the retailer's fulfillment network. We identify settings where the naive greedy policy performs poorly, motivating our development of a *spillover limit* policy that places limits on the amount of spillover demand the policy can fulfill. This policy is not only intuitive and straightforward (giving it potential for wide adoption), but we also demonstrate that it is asymptotically optimal under a general setting. Finally, we use the spillover limit fulfillment policy and a stochastic programming allocation policy to evaluate the benefit of a proposed flexible fulfillment network using data from our industrial collaborator. In particular, we estimate that our proposed flexible fulfillment network will decrease the lost sales plus fulfillment costs by one percent. Since the current operations of our industrial partner incur very large costs in these areas, this translates to a profit improvement on the order of tens of millions of U.S. dollars.

The research in this chapter was conducted under the supervision of Yehua Wei and in collaboration with Di Wu and Rong Yuan, and the analysis is largely included in DeValve et al. (2019c).

Chapter 2

One-Period Assemble-to-Order Systems

2.1 Introduction

In order to gain a competitive edge, firms facing uncertain demand are regularly attempting to reduce their demand-supply mismatch. To accomplish this, firms may employ the assemble-to-order (ATO) strategy, which takes advantage of demand pooling for a set of products that use common components. Intuitively, the strategy stocks inventory for components, and waits until demand arrives to assemble products for delivery. Compared with a make-to-stock model, an ATO system can provide cost savings while maintaining service levels, as components can be flexibly allocated to products when demand arrives. Many major manufacturing companies, including Dell (Kapuscinski et al., 2004), Amazon (Xu et al., 2009), and BMW (Muller, 2010), have leveraged this strategy to gain an operational advantage.

Despite the prevalence of ATO, it is often difficult to achieve the strategy's full benefits in practice. As a result, researchers have sought to optimize inventory and allocation decisions in various models of ATO systems (Song and Zipkin, 2003). One such model that is often used as the basis for studying more complex settings is known as the one-period ATO model, which is formulated as a two-stage stochastic program. In the first stage of the one-period model, component inventories are ordered before demand is realized. Then, after demand arrives, the second stage problem is to allocate components to products for assembly. The one-period problem is often used to design heuristics for more realistic dynamic ATO settings, i.e., models where inventory and allocation decisions are made dynamically over time (Song and Zipkin, 2003). In particular, in a series of pioneering works, Doğru, Reiman, and Wang (2010) and Reiman and Wang (2015) proved that for general ATO structures, solutions for one-period models can be used to construct asymptotically optimal fractional policies for dynamic ATO models with identical component lead times.

However, a significant theoretical hurdle in ATO settings is the integrality constraint, that is, the inventory and allocation decisions are required to take integer values. Even for the one-period model, the second stage allocation decision for general ATO structures is a generalization of the classical set-cover problem, which is not just NP-hard to solve, but also NP-hard to approximate with a constant factor (Feige, 1998). As a result, typical approaches for solving ATO problems numerically often resort to various heuristics (e.g., Lu and Song, 2005; van Jaarsveld and Scheller-Wolf, 2015). Recently, researchers have started to identify special ATO structures for which the second stage problem can be solved efficiently using notions of discrete convexity (Zipkin, 2016; Doğru et al., 2017). Unfortunately, for general ATO structures these techniques break down, and finding provably good solutions for the one-period model has remained a challenging open problem.

In this chapter, we provide a fresh perspective for understanding ATO models by applying a primal-dual analysis to the integer linear programming formulation of the one-period ATO problem. While primal-dual analysis is widely used to derive approximation algorithms for various combinatorial optimization problems, it has not received much attention in the ATO literature. We leverage this approach to prove the first approximation guarantees of various heuristics for the one-period problem for general ATO structures. Specifically, we develop an LP rounding algorithm that achieves both asymptotic optimality as demand grows large, and constant factor approximation for any problem instance. Furthermore, in Chapter 3, we show how to extend our LP rounding results for the one-period problem to design an asymptotically optimal integral policy in a dynamic model for general ATO structures.

Next, we briefly describe the key findings in this chapter.

Newsvendor decomposition. We first consider a natural decomposition of the one-period ATO model wherein each component solves a separate newsvendor problem to determine its order quantity, an approach that has been used as a heuristic in the ATO literature (see e.g., Lu and Song, 2005; van Jaarsveld and Scheller-Wolf, 2015). We develop a linear programming (LP) representation of such newsvendor decomposition heuristics, which allows us to use the dual to characterize a lower bound on optimal cost in terms of problem primitives (which provides a key component of our subsequent analysis). Furthermore, we use our primal-dual analysis to characterize an approximation guarantee that grows at the same rate as the maximum number of components required by a product in the system. We also establish that this approximation guarantee is asymptotically tight, implying that a newsvendor decomposition approach cannot provide a worst-case approximation guarantee with a constant factor.

LP rounding. We propose and analyze a rounding algorithm based on the linear programming (LP) relaxation of the one-period ATO problem. Our LP rounding algorithm provides both asymptotic optimality in large demand regimes and constant factor approximation for any problem instance; each of which are the first provable approximation guarantee of their kind. To establish asymptotic optimality, we use the lower-bound obtained from the newsvendor decomposition to characterize an approximation factor in terms of primitives of the underlying demand distribution. To demonstrate a constant approximation factor of 1.8, we leverage the complementary slackness conditions of the optimal LP solution to bound the rounded second stage cost with both the first and second stage LP cost. Our rounding algorithm picks one of two rounding schemes depending on a comparison of the LP shortage and ordering costs, providing insight on how the relative values of costs in the LP solution may impact rounding decisions. Finally, we show that the LP relaxation can be solved by a stochastic subgradient method, which implies that our LP rounding algorithm can be implemented efficiently even if the number of random scenarios increases exponentially with the number of products.

We note that a constant factor approximation guarantee for ATO is significant, since the deterministic second stage problem with M components cannot be approximated better than a $\log(M)$ factor (Feige, 1998; Moshkovitz, 2015). Our algorithm overcomes this hurdle by approximating the first stage cost and expected second stage cost together. Many existing approximation approaches for two-stage stochastic integer programs adapt known algorithms for deterministic problems to their stochastic counterparts, often enlarging the approximation factor in the process (e.g., set cover, facility location and others, see Shmoys and Swamy, 2006; Ravi and Sinha, 2006). Our technique, in contrast, develops a new rounding scheme that balances first and second stage costs, and reduces the approximation factor to a constant.

Computational complexity. We prove that the one-period ATO problem is NP-hard to solve exactly. To the best of our knowledge, this is the first computational complexity result for this problem. The proof of NP-hardness constructs a reduction from the classic set-cover problem, and requires an intermediate reduction from a novel class of set-cover instances that may be of independent interest.

The rest of this chapter is organized as follows. The remainder of the introduction discusses relevant literature, while Section 2.2 introduces the model and our primal-dual approach, and Section 2.3 considers newsvendor heuristics. In Section 2.4, we develop LP rounding schemes for the one-period problem and prove approximation guarantees. Section 2.5 presents and discusses simulation results for heuristics considered in this chapter. In Section 2.6 we establish the computational complexity of ATO problems. We provide concluding remarks in Section 2.7.

2.1.1 Related Literature

Early work studying inventory pooling in systems with shared components includes Baker et al. (1986), and Gerchak et al. (1988) who consider the impact of component commonality on optimal stocking levels. There has since grown an extensive literature on this topic exploring a broad range of managerial issues. Glasserman and Wang (1998) characterize how inventory levels change with lead times, Plambeck and Ward (2006) study joint pricing and production capacity decisions in a queuing model, and Plambeck and Ward (2007) identify conditions allowing a multidimensional ATO control problem to separate by components. Song and Zhao (2009) demonstrate that simple policies may fail to capitalize on the advantages of ATO, Lu et al. (2010) consider no-holdback allocation rules, and Bernstein et al. (2011) explore delaying component allocation. Competitive issues have also been modeled, including Nagarajan and Bassok (2008) and Nagarajan and Sošić (2009), who consider bargaining and coalition formation among suppliers in an assembly system. An excellent overview of several streams of research is provided in Song and Zipkin (2003).

Recent work in inventory theory has established that straightforward policies can be asymptotically optimal in complex dynamic systems (e.g., Huh et al., 2009; Xin and Goldberg, 2016). This has been a fruitful approach for ATO systems, with several researchers demonstrating how to derive good dynamic policies from the solution of a one-period model. Solutions of one-period ATO problems are often used to set base stock levels for dynamic inventory replenishment (e.g. Lu and Song, 2005; van Jaarsveld and Scheller-Wolf, 2015). Further, Doğru et al. (2010); Reiman and Wang (2015); Doğru et al. (2017) use the solution of a single period stochastic program to develop a dynamic policy that is asymptotically optimal when component lead time becomes large. These works demonstrate that some of the principal difficulties of managing dynamic ATO systems (i.e., the combinatorial allocation problem and demand uncertainty) can be overcome by considering appropriate one-period models, which we build upon in Chapter 3 of this dissertation.

Another theme in the ATO literature has been to study specific sublcasses of ATO systems (e.g., Doğru et al., 2010; Nadar et al., 2014; Lu et al., 2015). For the single period problem, efficient algorithms have been identified for two ATO system subclasses. Zipkin (2016) shows that a discrete convexity property, called $L^\natural$-convexity, allows using known algorithms for a class of assembly structures called "tree families." This algorithm can also be carried out on a broader class of assembly structures, and has been observed to perform well numerically, but no general performance guarantees are established. Doğru et al. (2017), leverages $L^\natural$-convexity to provide an efficient algorithm for assembly structures they call "chained BOM." They observe that $L^\natural$-convexity fails to hold for more general assembly structures, and the algorithm will not work in these cases.

We build on a broad literature of primal-dual approaches for approximating a wide variety of problems that include, e.g., Steiner network, facility location, joint replenishment, and online linear programming (Williamson et al., 1995; Jain and Vazirani, 2001; Levi et al., 2006; Agrawal et al., 2014). There is a vast literature applying rounding techniques to operations problems, with examples including facility location (Shmoys et al., 1997), lot sizing (Teo and Bertsimas, 2001), survivable network design (Jain and Vazirani, 2001), scheduling unrelated machines (Schulz and Skutella, 2002), the one-warehouse multiretailer problem (Levi et al., 2008), and deterministic inventory control (Nagarajan and Shi, 2016). Approximating stochastic optimization problems has also received increased attention recently, with a sampling of work on specific problems including service provision (Dye et al., 2003), network design (Gupta et al., 2007), stochastic inventory control (Levi et al., 2007), and non-linear newsvendors (Halman et al., 2012).

Our rounding schemes require solving a stochastic program, a topic with a broad literature on techniques (see e.g., Kleywegt et al., 2002; Shapiro et al., 2009). One class of algorithms finds a solution by following a stochastic subgradient, mimicking classic gradient descent methods (e.g., Nemirovski et al., 2009; Shapiro et al., 2009). We demonstrate that a stochastic subgradient method from Shapiro et al. (2009) has a simple and intuitive implementation in the ATO setting, and runs in polynomial time independent of the number of demand scenarios.

2.2 One-Period Model

In this section we introduce general notation used throughout the chapter and formally define the one-period ATO model. We consider an ATO system using M components (indexed by i) to manufacture N products (indexed by j). A particular system is comprised of three sets of inputs: an assembly structure, a demand distribution, and cost parameters.

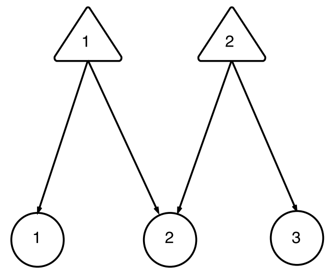

Figure 2.1: The M System

Assembly structure. The assembly structure is described by an integer matrix A, with $a_{ij} \in \mathbb{Z}_+$ (where $\mathbb{Z}_+$ denotes the set of non-negative integers) representing the number of units of component i needed to assemble product j. If $a_{ij} > 0$, we say that component i *serves* product j, and we use the notation $\mathcal{N}(j) = \{i | a_{ij} > 0\}$ to denote the set of components serving product j. Similarly, let $\mathcal{N}(i) = \{j | a_{ij} > 0\}$ denote the set of products that component i serves. Let $n_i = \sum_j a_{ij}$ denote the number of units of component i required to make one unit of all products it serves, and $m_j = \sum_i a_{ij}$ denote the number of components required to make one unit of product j. Let $m = \max_j \{m_j\}$ denote the maximum number of components required by any product in the system.

For the special case of a binary matrix A, the assembly structure can be represented by a bipartite graph where the sets of components and products form the bipartition, and $a_{ij} = 1$ denotes an edge between component i and product j. In that case the sets $\mathcal{N}(i)$ and $\mathcal{N}(j)$ represent the neighbors of a component node and product node in the graph, respectively. As an example, the M-system is depicted in Figure 2.1, consisting of two components (represented by triangles) and three products (represented by circles).

Demand distribution. We use Ω to denote the set consisting of all possible stochastic demand scenarios. In a demand scenario $\omega \in \Omega$, we use $d_{j\omega} \in \mathbb{Z}_+$ to denote demand for product j, and $\mathbf{d}_\omega = (d_{1\omega}, \ldots, d_{N\omega})$ to denote the N-vector of product demands. We assume Ω is a finite set. Moreover, we use $\mu_\omega \geq 0$ to denote the probability of demand scenario ω, so $\sum_{\omega \in \Omega} \mu_\omega = 1$. Let $\bar{d}_j = \sum_\omega \mu_\omega d_{j\omega}$ denote the expected demand for product j. At times, it will be convenient to use D_i to denote a random variable representing the aggregate demand for component i, i.e., a realization of D_i for scenario ω is $D_{i\omega} = \sum_j a_{ij} d_{j\omega}$. Let $\bar{D}_i = \sum_j a_{ij} \bar{d}_j$ denote the expected demand for component i.

For an event E, its indicator function is denoted $\mathbb{1}_{\{E\}}$ and its probability is denoted $\mathbb{P}[E]$. The notation $\mathbb{E}[X]$ denotes the expectation of a random variable X, and we use $(x)^+ = \max(x, 0)$ to denote the positive part of a number x.

Cost parameters. There are two sets of cost parameters: an inventory ordering cost for each component and a demand shortage cost for each product. Let $c_i > 0$ represent the cost to order one unit of component i; and $p_j > 0$ represent the cost of one unit of demand shortage for product j. Alternatively, we can interpret p_j as the profit loss due to not selling one unit of product j. We define the product markup, $\gamma_j = p_j / \sum_i a_{ij} c_i$, as the ratio of the product's shortage cost to its inventory costs. Thus, we can equivalently write the shortage cost as $p_j = \gamma_j \sum_i a_{ij} c_i$. Let the maximum and minimum markups in the system be denoted as $\bar{\gamma} = \max_j\{\gamma_j\}$ and $\underline{\gamma} = \min_j\{\gamma_j\}$ We say a system has identical markups if $\gamma_j = \underline{\gamma}$ for all products j.

2.2.1 Formulation for One-Period Model

The dynamics of the problem take place over two stages. In the first stage, prior to product demand realization, component inventories are ordered. Let $r_i \in \mathbb{Z}_+$ represent the inventory quantity ordered for component i in the first stage. Then in the second stage product demands are realized and products are assembled from available components to fill demand. Let $y_{j\omega} \in \mathbb{Z}_+$ denote the shortage of product j in demand scenario ω.

At the start of the first stage, consider the problem of finding the component order quantities, r_i, which minimize the ordering cost plus expected shortage cost. This problem can be cast as the following stochastic integer program with recourse:

$$\min_{\mathbf{r},\mathbf{y}} \sum_i r_i c_i + \sum_{j,\omega} \mu_\omega p_j y_{j\omega} \tag{ATO}$$

$$\text{s.t. } r_i + \sum_j a_{ij} y_{j\omega} \geq \sum_j a_{ij} d_{j\omega}, \ \forall\, i, \omega \tag{2.1}$$

$$y_{j\omega} \leq d_{j\omega}, \ \forall\, j, \omega \tag{2.2}$$

$$\mathbf{r}, \mathbf{y} \geq 0, \ \text{integer}. \tag{2.3}$$

The sum $\sum_i r_i c_i$ is the ordering cost, while $\sum_{j,\omega} \mu_\omega p_j y_{j\omega}$ is the expected shortage cost. The inventory constraint (2.1) ensures that, in each scenario, the demand filled for products using component i does not exceed its inventory. The demand constraint (2.2) ensures product j's shortage does not exceed demand in scenario ω, which is equivalent to requiring non-negative filled demand. We note that without constraint (2.2), in some scenario, the model may choose a shortage quantity for a low cost product higher than its actual demand, in order to reduce the shortage of other products with higher costs. The integer constraint is a reflection of the indivisibility of the units of production, and that customers demand only whole product units.

We briefly mention here the connections between (ATO) and related dynamic problems. Several studies of dynamic ATO systems (Lu and Song, 2005; Doğru et al., 2010; Reiman and Wang, 2015; van Jaarsveld and Scheller-Wolf, 2015; Doğru et al., 2017) solve a problem similar to (ATO) in order to set basestock levels. In particular, c_i and p_j represent the holding and back order cost of component i and product j respectively, and the optimal quantity r_i determines the base stock level for component i, while the variable y_j estimates the back orders for product j over some appropriately chosen period. Thus, the single-period problem (ATO) is critical to understanding dynamic problems, a connection that we explore further in Chapter 3.

A few issues present challenges to obtaining exact solutions for (ATO). First, consider the second stage allocation problem in a given demand scenario ω, after the inventory decisions r_i have been made:

$$\begin{aligned} \min \sum_j & p_j y_{j\omega} \\ \text{s.t.} \sum_j & a_{ij} y_{j\omega} \geq \sum_j a_{ij} d_{j\omega} - r_i, \ \forall\ i, \\ & y_{j\omega} \leq d_{j\omega}, \ \forall\ j, \\ & y_{j\omega} \geq 0, \ \text{integer}, \ \forall\ j. \end{aligned} \tag{2.4}$$

This problem is a covering integer program (CIP) with multiplicity constraints (Kolliopoulos and Young, 2001; Kolliopoulos, 2003), where the inventory constraints for each component i correspond to "covering" constraints, and the demand constraints for each product j correspond to "multiplicity" constraints. CIP's include the classic set-cover problem as a special case, which is well known to be NP-hard to solve exactly, and even NP-hard to approximate with better than an $O(\log M)$ factor (Feige, 1998; Moshkovitz, 2015), where M is the number of elements in the set cover instance (corresponding, in (2.4), to the number of components). Further, the multiplicity constraints in (2.4) (i.e., $y_{j\omega} \leq d_{j\omega}$) make the problem in some sense even more challenging, as Kolliopoulos and Young (2001) provide examples where even a logarithmic approximation factor cannot be achieved with the linear relaxation. This issue is further magnified by the fact that solving (ATO) requires determining the allocation in (2.4) for every demand scenario in Ω, and Ω may be large. Thus, simply calculating the optimal expected shortage cost for a given inventory vector is computationally intractable in general.

Interestingly, in Section 2.4, we demonstrate that the inapproximability of the second stage cost can be overcome by bounding errors in its approximation in terms of both the first and second stage costs. Our development of these approximation bounds relies heavily on the dual of the LP relaxation of (ATO), which we introduce next.

2.2.2 Dual Problem

To derive lower bounds on the optimal cost, we consider the dual of the LP relaxation of (ATO).

$$\max_{\boldsymbol{\theta},\boldsymbol{\pi}} \sum_{i,\omega} \theta_{i\omega} \sum_{j} a_{ij} d_{j\omega} - \sum_{j,\omega} d_{j\omega}\pi_{j\omega} \tag{DATO}$$

$$\text{s.t.} \sum_{i} a_{ij}\theta_{i\omega} - \pi_{j\omega} \le \mu_\omega p_j,\ \forall\, j,\omega \tag{2.5}$$

$$\sum_{\omega} \theta_{i\omega} \le c_i,\ \forall\, i \tag{2.6}$$

$$\boldsymbol{\theta},\boldsymbol{\pi} \ge 0. \tag{2.7}$$

The dual variable $\theta_{i\omega}$ prices the cost of the inventory constraint (2.1) for component i in scenario ω, while dual variable $\pi_{j\omega}$ prices the gain from demand constraint (2.2) allowing no more shortage than $d_{j\omega}$ for product j in scenario ω. The constraint (2.5) corresponds to the dual pricing of the shortage costs for each product in every demand scenario. Similarly, the constraint (2.6) corresponds to the accumulation of the dual pricing of the ordering cost for each component across all demand scenarios. We call (2.5) and (2.6) the shortage cost and ordering cost constraints, respectively.

Let OPT^{LP} represent the optimal cost of the LP relaxation of (ATO). As the optimal cost of a relaxation, we have $OPT^{LP} \le OPT$. Further, by weak duality, any feasible dual solution to (DATO) provides a lower bound on OPT^{LP}, and hence on OPT as well.

With the dual problem (DATO) introduced, we now give a brief, high-level overview of our primal-dual approach for proving approximation guarantees for (ATO). Consider a primal solution, $(r_i, y_{j\omega})$ and dual solution $(\theta_{i\omega}, \pi_{j\omega})$, to the LP relaxation of (ATO) and (DATO), respectively. If these solutions satisfy the well known complementary slackness conditions:

$$r_i > 0 \implies \sum_{\omega} \theta_{i\omega} = c_i, \ \forall \ i, \tag{2.8}$$

$$y_{j\omega} > 0 \implies \sum_{i} a_{ij}\theta_{i\omega} - \pi_{j\omega} = \mu_\omega p_j, \ \forall \ j, \omega, \tag{2.9}$$

$$\theta_{i\omega} > 0 \implies r_i + \sum_{j} a_{ij} y_{j\omega} = \sum_{j} a_{ij} d_{j\omega}, \ \forall \ i, \omega, \tag{2.10}$$

$$\pi_{j\omega} > 0 \implies y_{j\omega} = d_{j\omega}, \ \forall \ j, \omega, \tag{2.11}$$

then they are clearly optimal in their respective programs. Next we describe two directions in which we leverage conditions (2.8-2.11).

First, it is well known that the gap between a primal and dual solution is bounded if they satisfy conditions (2.8-2.11) approximately, i.e., the equalities are relaxed to inequalities with appropriate multiplicative factors adjusting the constants on the right-hand sides. Such relaxed conditions are known as approximate complementary slackness (ACS), and have been used widely to prove approximation guarantees for integer programs (Vazirani, 2013). In Section 2.3.1, we develop an extension of classical ACS that allows greater flexibility in meeting the conditions for all four sets of variable/constraint pairs simultaneously. We then characterize a primal-dual pair satisfying these conditions to establish an approximation guarantee for a commonly used newsvendor heuristic in Section 2.3.

We further use the complementary slackness conditions (2.8-2.11) in a different way when studying LP rounding algorithms in Section 2.4. In particular, we note that the conditions (2.8-2.11) provide a wealth of information about an optimal primal-dual pair. We leverage this information in order to bound the change in cost resulting from rounding the primal LP solution. For example, if an optimal LP solution has positive shortage for a product j in scenario ω, then by condition (2.9) the associated primal cost, $\mu_\omega p_j$, is a function of the dual variables, which permits a direct comparison with the dual cost. This technique allows us to characterize an intuitive cost trade-off for an LP rounding scheme introduced in Section 2.4.

Thus, we approach the problem of approximating (ATO) with the following general procedure. First, we find an efficient way to produce an integer primal solution to (ATO). Next, we construct a corresponding feasible dual solution for (DATO). Finally, we compare the costs of these two solutions using (approximate) complementary slackness conditions, providing a bound on the cost of the primal solution relative to OPT.

2.3 Newsvendor Decomposition

Before introducing our main LP-rounding algorithms in Section 2.4, we first discuss and analyze a class of decomposition heuristics. This natural approach attempts to overcome (ATO)'s complexities by considering the ordering decision for each component individually. For example, consider ordering r_i units of inventory for component i, which faces uncertain demand accumulated from all products it serves (denoted by the random variable D_i), and has ordering cost c_i. Assuming the availability of other components, it is clear that demand for component i can be filled when $D_i \leq r_i$, and otherwise there is a shortage of $D_i - r_i$. If we let q_i be an estimate of the per unit cost of this shortage, we obtain a heuristic order quantity for component i by solving the following univariate problem

$$\min_{r_i \geq 0} \left\{ c_i r_i + q_i \mathbb{E}[(D_i - r_i)^+] \right\} \tag{2.12}$$

This is the classic newsvendor problem, whose solution is well known as the "critical quantile" of the demand distribution, and can be computed efficiently. We call this general method a *newsvendor decomposition*, since it decomposes (ATO) into separate newsvendor problems. There are two issues associated with this heuristic: i) other components are not always available when component i is, and ii) the shortage cost estimate q_i may not accurately reflect the true product-specific shortage costs in a given scenario. Despite these issues, some researchers have demonstrated numerically that a newsvendor decomposition can work well in certain settings (e.g., Lu and Song, 2005; van Jaarsveld and Scheller-Wolf, 2015).

We approach newsvendor heuristics from a primal-dual perspective. Initially, we demonstrate that, for a particular choice of q_i, the optimal dual solution for (2.12) easily translates into a feasible dual solution for (DATO), which allows us to prove two key results. First, we translate the dual solution into a lower bound on the optimal cost of (ATO) in terms of primitives of the demand distribution, shedding light on system parameters that influence the achievable cost of the system.

Consider setting the shortage cost for component i as its ordering cost times the minimum markup in the system,

$$q_i = \underline{\gamma} c_i, \tag{2.13}$$

which we call the *constant markup* shortage costs. Intuitively, this can be thought of as a conservative estimate for "splitting" the true shortage cost of a product, p_j, among its products, since $\sum_i a_{ij} q_i =$

$\underline{\gamma} \sum_i a_{ij} c_i \leq \gamma_j \sum_i a_{ij} c_i = p_j$. From a technical perspective, this choice is helpful because it allows us to develop a useful solution for the original dual problem (DATO). To demonstrate this, consider the following LP representation of the newsvendor problem (2.12)

$$\min_{r_i, y_{i\omega} \geq 0} \left\{ c_i r_i + \underline{\gamma} c_i \sum_{\omega} \mu_\omega y_{i\omega} \,\middle|\, y_{i\omega} + r_i \geq D_{i\omega},\ \forall \omega \right\}, \tag{2.14}$$

for some $1 \leq i \leq M$, whose dual is

$$\max_{\theta_{i\omega} \geq 0} \left\{ \sum_{\omega} \theta_{i\omega} D_{i\omega} \,\middle|\, \theta_{i\omega} \leq \mu_\omega \underline{\gamma} c_i, \forall \omega,\ \sum_{\omega} \theta_{i\omega} \leq c_i \right\}. \tag{2.15}$$

The well known "critical quantile" solution for the newsvendor problem balances ordering cost with expected shortage cost to give

$$r_i^{\underline{\gamma}} = \min \left\{ r \in \mathbb{Z}_+ \,\middle|\, \mathbb{P}\left[D_i > r\right] < \frac{1}{\underline{\gamma}} \right\}, \tag{2.16}$$

as the optimal order quantity for (2.14). As suggested by the objective value in (2.12), the corresponding optimal shortages for (2.14) are $y_{i\omega}^{\underline{\gamma}} = (D_{i\omega} - r_i^{\underline{\gamma}})^+$. The simple structure of the dual problem (2.15) allows us to also characterize the associated optimal dual solution:

$$\theta_{i\omega}^{\underline{\gamma}} = \begin{cases} \mu_\omega \underline{\gamma} c_i, & r_i^{\underline{\gamma}} < D_{i\omega} \\ z_i \mu_\omega \underline{\gamma} c_i, & r_i^{\underline{\gamma}} = D_{i\omega} \\ 0, & r_i^{\underline{\gamma}} > D_{i\omega} \end{cases}, \tag{2.17}$$

where $z_i \in [0, 1]$ is defined by

$$z_i = \mathbb{1}_{\{r_i^{\underline{\gamma}} > 0\}} \left(\frac{1 - \underline{\gamma} \mathbb{P}\left[D_i > r_i^{\underline{\gamma}}\right]}{\underline{\gamma} \mathbb{P}\left[D_i = r_i^{\underline{\gamma}}\right]} \right). \tag{2.18}$$

Intuitively, the dual variable $\theta_{i\omega}^{\underline{\gamma}}$ captures scenario ω's contribution to the expected marginal shortage cost of inventory level $r_i^{\underline{\gamma}}$. It is easily verified by inspection that this primal and dual solution are feasible and satisfy complementary slackness, and hence are optimal in their respective programs. Our first result observes that the dual variables (2.17) also constitute a feasible solution for (DATO). Indeed, feasibility for constraint (2.6) is inherited from the second constraint in the LP (2.15), while the first constraint in the same program implies that for any product j,

$$\sum_i a_{ij} \theta_{i\omega}^{\underline{\gamma}} \leq \mu_\omega \underline{\gamma} \sum_i a_{ij} c_i,$$

$$\leq \mu_\omega p_j,$$

which satisfies constraint (2.5) for $\pi^{\gamma}_{j\omega} = 0$. Thus we have proven the following Lemma.

Lemma 2.1. *The dual variables* (2.17)*, together with* $\pi^{\gamma}_{j\omega} = 0, \forall j, \omega$*, are feasible for* (DATO).

This result, together with weak duality, demonstrates that summing the optimal costs of the LP (2.14) across all components i, provides a lower bound on the optimal cost of the original problem (ATO). Our next result writes this cost in terms of primitives of the underlying demand distribution. Let

$$\hat{\gamma} = \min(1, \underline{\gamma}), \quad \gamma^* = \frac{1}{2}\min(1, (\underline{\gamma} - 1)^+). \tag{2.19}$$

We note that in many real world ATO systems, a product's markup will be larger than one, i.e., $\gamma_j > 1$, since otherwise the product's ordering cost would outweigh its shortage cost, leaving no incentive to order inventory to fill its demand. Indeed, in the dynamic backlog model discussed in Section 3.2, the cost parameters imply $\gamma_j > 1$ for all j. In such a setting we have $\underline{\gamma} > 1$, and we see from (2.19) that $\hat{\gamma} = 1$ and $\gamma^* > 0$. Further, if $\underline{\gamma} \geq 2$ (i.e. each product's shortage cost is at least twice the cost of ordering its required components) then we have $\hat{\gamma} = 1$ and $\gamma^* = 1/2$, i.e., the values in (2.19) are constants. Intuitively, a product's markup has an impact on its demand fill rate, and so in ATO settings with requirements for high service levels, the values in (2.19) will be constants. We state the newsvendor cost lower bound in terms of the markup parameters defined in (2.19) in order to demonstrate its applicability outside of the large markup parameter region.

Proposition 2.1. *The constant markup newsvendor provides the following lower bound on the optimal cost of* (ATO)

$$\sum_i c_i \left(\hat{\gamma}\bar{D}_i + \gamma^* \mathbb{E}\left[|D_i - \bar{D}_i| \right] \right) \leq OPT.$$

Before presenting the proof of Proposition 2.1, we present the following lower bound on the absolute deviation of a random variable from a given quantity. We note that this bound is tight when X is a Bernoulli distribution with probability of success approaching zero and $s = 0$.

Lemma 2.2. *For a random variable X with mean $\mu < \infty$, and any real value s, we have*

$$\mathbb{E}[|X - s|] \geq \frac{1}{2}\mathbb{E}[|X - \mu|].$$

Proof. For any z, the following identities are easily verified by checking the positive and negative cases,

$$|z| = (z)^+ + (-z)^+, \quad (z)^+ = \frac{|z| + z}{2}.$$

Using these identities, we prove the claim by consider two cases. First, if $s \geq \mu$, then we have

$$\begin{aligned}\mathbb{E}[|X-s|] = \mathbb{E}[(X-s)^+ + (s-X)^+] &\geq \mathbb{E}[(\mu - X)^+] \\ &= \frac{1}{2}\mathbb{E}[|\mu - X| + \mu - X] \\ &= \frac{1}{2}\mathbb{E}[|X-\mu|].\end{aligned}$$

Similarly, if $s < \mu$ then

$$\begin{aligned}\mathbb{E}[|X-s|] = \mathbb{E}[(X-s)^+ + (s-X)^+] &\geq \mathbb{E}[(X-\mu)^+] \\ &= \frac{1}{2}\mathbb{E}[|X-\mu| + X - \mu] \\ &= \frac{1}{2}\mathbb{E}[|X-\mu|].\end{aligned}$$

□

Proof of Proposition 2.1. By Lemma 2.1 and weak duality, summing the optimal costs of (2.14) across all components i gives

$$\begin{aligned}OPT &\geq \sum_i c_i \left(r_i^\gamma + \underline{\gamma}\mathbb{E}[(D_i - r_i^\gamma)^+] \right), \\ &= \sum_i c_i \left(\hat{\gamma}\bar{D}_i + \mathbb{E}[r_i^\gamma - \hat{\gamma}D_i] + \underline{\gamma}\mathbb{E}[(D_i - r_i^\gamma)^+] \right),\end{aligned}$$

so that the claim follows if we show

$$\mathbb{E}[r_i^\gamma - \hat{\gamma}D_i] + \underline{\gamma}\mathbb{E}[(D_i - r_i^\gamma)^+] \geq \gamma^* \mathbb{E}\left[|D_i - \bar{D}_i|\right], \tag{2.20}$$

for all i. Consider two cases, the first being $\underline{\gamma} < 1$, so that $\hat{\gamma} = \underline{\gamma}$ and $\gamma^* = 0$, reducing (2.20) to

$$\mathbb{E}[r_i^\gamma - \underline{\gamma}D_i] + \underline{\gamma}\mathbb{E}[(D_i - r_i^\gamma)^+] \geq 0.$$

We note that $\underline{\gamma} < 1$ implies $1 < 1/\underline{\gamma}$, so that the newsvendor inventory quantity for all i is $r_i^\gamma = 0$, by the critical quantile definition in (2.16) (since $\mathbb{P}[D_i > 0] \leq 1$ always). Thus, it is clear that the left-hand side of the above expression is equal to zero in this case, and hence (2.20) holds.

Otherwise, if $\underline{\gamma} \geq 1$ then we have $\hat{\gamma} = 1$. Further, note that $r_i^\gamma - D_i = (r_i^\gamma - D_i)^+ - (D_i - r_i^\gamma)^+$, for all realizations of D_i, so that the left-hand side of (2.20) is

$$\begin{aligned}\mathbb{E}[r_i^\gamma - \hat{\gamma}D_i] + \underline{\gamma}\mathbb{E}[(D_i - r_i^\gamma)^+] &= \mathbb{E}[(r_i^\gamma - D_i)^+] + (\underline{\gamma} - 1)\mathbb{E}[(D_i - r_i^\gamma)^+], \\ &\geq \min(1, \underline{\gamma} - 1)\mathbb{E}[(r_i^\gamma - D_i)^+ + (D_i - r_i^\gamma)^+], \\ &= 2\gamma^*\mathbb{E}[|D_i - r_i^\gamma|],\end{aligned}$$

and (2.20) follows by Lemma 2.2. □

Intuitively, the quantity $\mathbb{E}\left[|D_i - \bar{D}_i|\right]$ (i.e., the expected absolute deviation from the mean) is a measure of the variability of demand for component i. Indeed, if D_i were normally distributed, this quantity would be proportional to the standard deviation. Thus, Proposition 2.1 shows that the optimal cost of (ATO) grows with average component demand, as well as the variability of component demand (as long as $\underline{\gamma} > 1$). This lower bound is easy to compute, and so provides a practical reference point for comparing the cost of a given heuristic to optimal. In Section 2.4, we formally use such a comparison to prove an asymptotically optimal approximation guarantee for an LP rounding algorithm.

Before turning to LP rounding, however, we wish to consider the performance of the newsvendor heuristic as an approximation algorithm. We take a brief aside in the next section to develop a tool to carry out this analysis.

2.3.1 Approximate Complementary Slackness

In this section we provide an extension to the standard approximate complementary slackness (ACS) conditions that have been used to bound the performance of a wide variety of LP-based algorithms (Vazirani, 2013). Our extension essentially provides more flexibility in applying the conditions to the four sets of variable/constraint pairs in the (ATO) and (DATO) problems. Consider constants κ and ν, which we refer to as *primal ACS factors* and constants δ and η, referred to as *dual ACS factors*. Then, given primal and dual solutions $(r_i, y_{j\omega})$ and $(\theta_{i\omega}, \pi_{j\omega})$ feasible to (ATO) and (DATO), respectively, consider the following extended ACS conditions

$$y_{j\omega} > 0 \implies \sum_{i\in\mathcal{N}(j)} \theta_{i\omega} - \pi_{j\omega} \geq \frac{\mu_\omega p_j}{\kappa}, \ \forall\, j, \omega, \tag{2.21}$$

$$r_i > 0 \implies \sum_{\omega} \theta_{i\omega} \geq \frac{c_i}{\nu}, \ \forall\, i, \tag{2.22}$$

$$\theta_{i\omega} > 0 \implies r_i + \frac{\kappa}{\nu} \sum_{j\in\mathcal{N}(i)} y_{j\omega} \leq \frac{\delta}{\nu} \sum_{j\in\mathcal{N}(i)} d_{j\omega}, \ \forall\, i, \omega, \tag{2.23}$$

$$\pi_{j\omega} > 0 \implies y_{j\omega} \geq \frac{\eta}{\kappa} d_{j\omega}, \ \forall\, j, \omega. \tag{2.24}$$

Intuitively, these extended conditions mimic standard ACS by specifying bounds on constraint slackness, but in a way that allows handling different sets of constraints flexibly. In particular, (2.21) is the ACS condition for the primal variables $y_{j\omega}$, and bounds the slackness of the associated dual shortage cost constraints (2.5) with the factor κ. Likewise, (2.22) is the ACS condition for the primal variables r_i, bounding slackness of the dual ordering cost cosntraints (2.6) with ν. The key innovation of our approach is separating these sets of variables to allow for different bounds on the slackness of their associated dual constraints. Similar to standard ACS, the bounds on constraint slackness in (2.21-2.22) easily translate into a bound on the cost of the primal solution,

$$\begin{aligned}\sum_i r_i c_i + \sum_{j,\omega} \mu_\omega p_j y_{j\omega} &\leq \nu \sum_i r_i \sum_\omega \theta_{i\omega} + \kappa \sum_{j,\omega} y_{j\omega} \left(\sum_{i \in \mathcal{N}(j)} \theta_{i\omega} - \pi_{j\omega} \right), \\ &= \sum_{i,\omega} \theta_{i\omega} \left(\nu r_i + \kappa \sum_{j \in \mathcal{N}(i)} y_{j\omega} \right) - \kappa \sum_{j,\omega} y_{j\omega} \pi_{j\omega}.\end{aligned} \tag{2.25}$$

Next, (2.23) is the ACS condition for dual variables $\theta_{i\omega}$, which bounds slackness in the primal inventory constraints (2.1) with factor δ, while accounting for the differing primal factors from (2.21-2.22). Finally, (2.24) is the ACS condition for dual variables $\pi_{j\omega}$, bounding slackness in primal demand constraints (2.2) with factor η, while accounting for the primal factor κ from (2.21). A direct application of (2.23-2.24) on the final right-hand side of expression (2.25) gives

$$\sum_{i,\omega} \theta_{i\omega} \left(\nu r_i + \kappa \sum_{j \in \mathcal{N}(i)} y_{j\omega} \right) - \kappa \sum_{j,\omega} y_{j\omega} \pi_{j\omega} \leq \delta \sum_{i,\omega} \theta_{i\omega} \sum_{j \in \mathcal{N}(i)} d_{j\omega} - \eta \sum_{j,\omega} d_{j\omega} \pi_{j\omega}. \tag{2.26}$$

Finally, combining (2.25) and (2.26) we obtain the following bound on the primal cost in terms of the dual solution.

Proposition 2.2. *Let* $(r_i, y_{j\omega})$ *and* $(\theta_{i\omega}, \pi_{j\omega})$ *be a primal-dual pair feasible to* (ATO) *and* (DATO)*, respectively. If* $(r_i, y_{j\omega})$ *and* $(\theta_{i\omega}, \pi_{j\omega})$ *satisfy the extended ACS conditions (2.21-2.24) then*

$$\sum_i r_i c_i + \sum_{j,\omega} \mu_\omega p_j y_{j\omega} \leq \delta \sum_{i,\omega} \theta_{i\omega} \sum_{j \in \mathcal{N}(i)} d_{j\omega} - \eta \sum_{j,\omega} d_{j\omega} \pi_{j\omega}. \tag{2.27}$$

We use these conditions to prove performance guarantees for several heuristics for (ATO). We note that while the extended ACS conditions do not always provide the simplest proofs for the performance guarantees we derive, the conditions provide a systematic method for studying the effectiveness of heuristics. Therefore, the advantage of the extended ACS conditions is that they offer a standard primal-dual framework for analyzing the effectiveness of any (existing or new) heuristics proposed for the single period ATO model.

We also highlight the connection with the standard ACS conditions, which bound dual slackness with a factor α and primal slackness with a factor β. In the extended ACS conditions, letting $\delta = \eta = \alpha\beta$ and $\kappa = \nu = \alpha$, gives the standard ACS conditions and implies an approximation factor of $\alpha\beta$. Thus, the extended conditions essentially allow the standard α and β factors to vary for each set of primal and dual constraints, thereby allowing greater flexibility in achieving the ACS conditions.

2.3.2 Constant Markup Newsvendor Guarantees

With the extended ACS conditions of Section 2.3.1, we are now ready to prove performance guarantees for the constant markup newsvendor heuristic. Specifically, we consider using the inventory quantities determined by the newsvendor heuristic (2.16) for each component, and ask how close these can come to achieving optimal cost.

Bound for General Systems

We first demonstrate that for general systems, the approximation factor of the newsvendor heuristic grows with m, the maximum number of components required by any product in the system ($m = \max_j\{\sum_i a_{ij}\}$), which we show with both an upper and lower bound. To state the upper bound, denote the lowest component ordering cost by $\underline{c} = \min_i\{c_i\}$ and the largest by $\bar{c} = \max_i\{c_i\}$.

Proposition 2.3. *The constant markup newsvendor inventory quantities r_i^γ, $1 \leq i \leq M$, provide the following approximation factor for* (ATO)

$$1 + \left(\frac{\bar{\gamma}\bar{c}}{\underline{\gamma}\underline{c}}\right) m. \tag{2.28}$$

Proof. Recall that the optimal shortage quantities for the newsvendor LP (2.14) are $y_{i\omega}^\gamma = (D_{i\omega} - r_i^\gamma)^+$. We use these quantities to define feasible shortages for the original problem (ATO) as follows

$$y_{j\omega}^\gamma = d_{j\omega} \mathbb{1}_{\{\sum_i a_{ij} y_{i\omega}^\gamma > 0\}}. \tag{2.29}$$

These shortages clearly satisfy the demand constraint (2.2). For the inventory constraint (2.1), we observe that when $r_i^\gamma \geq \sum_j a_{ij} d_{j\omega}$, the constraint holds trivially since shortage quantities are positive. Otherwise, if $r_i^\gamma < \sum_j a_{ij} d_{j\omega}$, then by definition we have $y_{i\omega}^\gamma > 0$ which implies $\sum_i a_{ij} y_{i\omega}^\gamma > 0$ for all $j \in \mathcal{N}(i)$, and so $y_{j\omega}^\gamma = d_{j\omega}$ for all $j \in \mathcal{N}(i)$. Thus we have $\sum_j a_{ij} y_{j\omega}^\gamma = \sum_j a_{ij} d_{j\omega}$, and constraint (2.1) holds since $r_i^\gamma \geq 0$.

We establish the bound by demonstrating the ACS conditions for factors $\kappa = \left(\frac{\bar{\gamma}\bar{c}}{\underline{\gamma}\underline{c}}\right) m$, $\nu = 1$, and $\delta = \eta = 1 + \left(\frac{\bar{\gamma}\bar{c}}{\underline{\gamma}\underline{c}}\right) m$. If $y^{\gamma}_{j\omega} > 0$, then by (2.29) there is some $i' \in \mathcal{N}(j)$ such that $y^{\gamma}_{i'\omega} > 0$, which implies $r^{\gamma}_{i'} < D_{i'\omega}$. Then, by definition in (2.17), we have $\theta^{\gamma}_{i'\omega} = \mu_\omega \underline{\gamma} c_{i'}$, meaning that at least component i' contributes toward the dual accounting in the shortage cost constraint (2.5). To obtain a lower bound the amount of this contribution, observe that $p_j = \gamma_j \sum_i a_{ij} c_i \leq \bar{\gamma}\bar{c}m$, so that, together with $a_{i'j} \geq 1$, we have

$$\begin{aligned}\sum_i a_{ij}\theta^{\gamma}_{i\omega} &\geq \frac{p_j}{\bar{\gamma}\bar{c}m}\mu_\omega \underline{\gamma} a_{i'j} c_{i'} \\ &\geq \frac{\mu_\omega p_j}{\left(\frac{\bar{\gamma}\bar{c}}{\underline{\gamma}\underline{c}}\right) m}.\end{aligned}$$

and it follows that ACS condition (2.21) is satisfied for $\kappa = \left(\frac{\bar{\gamma}\bar{c}}{\underline{\gamma}\underline{c}}\right) m$. By complementary slackness for the LP (2.14), we have $\sum_\omega \theta^{\gamma}_{i\omega} = c_i$ if $r^{\gamma}_i > 0$, so ACS condition (2.22) holds for $\nu = 1$. When $\theta^{\gamma}_{i\omega} > 0$, its definition in (2.17) implies $r^{\gamma}_i \leq D_{i\omega}$, and since $y^{\gamma}_{j\omega} \leq d_{j\omega}$, we have

$$r^{\gamma}_i + \kappa \sum_{j \in \mathcal{N}(i)} y^{\gamma}_{j\omega} \leq (1+\kappa) D_{i\omega},$$

which satisfies ACS condition (2.23) for $\delta = 1 + \kappa = 1 + \left(\frac{\bar{\gamma}\bar{c}}{\underline{\gamma}\underline{c}}\right) m$. Finally, since $\pi^{\gamma}_{j\omega} = 0$, ACS condition (2.24) is trivially satisfied for $\eta = 1 + \kappa$ as well. Thus, by Proposition 2.2 the primal cost of r^{γ}_i and $y^{\gamma}_{j\omega}$ is bound by $1 + \left(\frac{\bar{\gamma}\bar{c}}{\underline{\gamma}\underline{c}}\right) m$ times the dual cost. □

The proof of Proposition 2.3 essentially makes a comparison between the primal and dual costs. We demonstrate that the complementary slackness condition for the primal shortage variable, (2.9), holds approximately when the right-hand side is scaled by a factor proportional to m, the largest product size in the system. This leads to an approximation factor that scales linearly with m, which intuitively results from the decomposition: when fitting the decomposed solutions back into the original problem, there is a chance for mis-coordination on each component a product uses.

Next we construct a family of examples that realizes such mis-coordination for every product in every demand scenario. In this family of examples, the approximation factor of the newsvendor heuristic is at least m, thus demonstrating asymptotic tightness of the guarantee in Proposition 2.3.

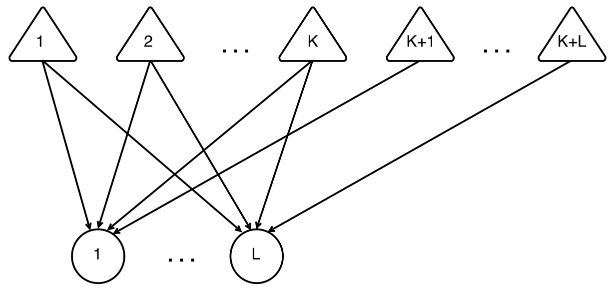

Figure 2.2: The assembly structure of the systems in Examples 2.1 and 2.3

Example 2.1 ((K, L, ϵ) system)**.** *Consider a system with $K + L$ components and $L \geq 2$ products. For each product j let $a_{ij} = 1$ if $1 \leq i \leq K$ or $i = K + j$ and $a_{ij} = 0$ otherwise. In words, each product requires the first K components, plus one dedicated component from the remaining L components, as illustrated in Figure 2.2. Note that the maximum number of components serving a product is $m = K + 1$.*

Let $d \geq 1$ be a positive integer, and let $\mathbf{e}_k$ represent an L-vector with a one in the k^{th} element and zero elsewhere. There are $L+1$ demand scenarios, indexed by $1 \leq k \leq L+1$. The first L scenarios have $\mathbf{d}_k = d\mathbf{e}_k$ (i.e., demand d for product $j = k$) and each occur with probability $0 < \mu \leq \frac{1}{L}$, while the last scenario has zero demand for all products and occurs with the remaining probability, $1 - L\mu$. Finally, costs are $c_i = 1$ for all i, and $p_j = \frac{1-\epsilon}{\mu}(K + 1)$ for all j, for some $0 < \epsilon \leq 1 - \frac{1}{L}$. With these costs the products all have markup $\gamma = \frac{1-\epsilon}{\mu}$.

We state the resulting lower bound formally in the next proposition.

Proposition 2.4. *For any integer $K \geq 1$ and $\bar{\epsilon} \in (0, K)$, there exists an ATO system with $m = K+1$ such that the constant markup newsvendor inventory quantities, r_i^{γ}, $1 \leq i \leq M$, have cost at least $(m - \bar{\epsilon})OPT$.*

Proof. What we show is that for any $K \geq 1$ and $\bar{\epsilon} \in (0, K)$, there exists ϵ_0 and L_0, such that for all $\epsilon \in (0, \epsilon_0)$ and all $L \geq L_0$, the constant markup newsvendor inventory quantities, r_i^{γ}, $1 \leq i \leq M$, for the (K, L, ϵ) system have cost at least $(K + 1 - \bar{\epsilon})OPT$. This implies the claim since $m = K + 1$ in the (K, L, ϵ) system.

The markup parameter is $\gamma = \frac{1-\epsilon}{\mu}$. Upon inspection of the demands it is clear that each component sees demand for at most d units in any demand scenario, and so the optimal newsvendor quantity must be an integer between zero and d. We note that $\epsilon \leq 1 - 1/L$ implies $L \geq 1/(1-\epsilon)$, so that for $1 \leq i \leq K$, ordering a quantity $r < d$ implies

$$\mathbb{P}[D_i > r] = L\mu \geq \frac{\mu}{1-\epsilon} = \frac{1}{\underline{\gamma}},$$

which violates the optimality condition for $r_i^{\underline{\gamma}}$ in equation (2.16). Since $\mathbb{P}[D_i > d] = 0$ satisfies the condition, we have $r_i^{\underline{\gamma}} = d$ for all $1 \leq i \leq K$. Meanwhile, for $K+1 \leq i \leq K+L$ we have

$$\mathbb{P}[D_i > 0] = \mu < \frac{\mu}{1-\epsilon} = \frac{1}{\underline{\gamma}},$$

and so $r_i^{\underline{\gamma}} = 0$ by (2.16). Thus, for each product j, we will always be missing the component $K+j$, and hence cannot fill any demand in any scenario. So, the optimal shortages given the inventory quantities $r_i^{\underline{\gamma}}$ must be $\mathbf{y}_k = \mathbf{d}_k$, in each demand scenario k. Therefore, the ordering cost plus the minimal shortage cost associated with the inventory quantities $r_i^{\underline{\gamma}}$ is

$$dK + \frac{1-\epsilon}{\mu}(K+1)d(L\mu) = dK + (1-\epsilon)dL(K+1).$$

Meanwhile, since demand never arrives for multiple products simultaneously, we can fill all demand (i.e. incur no shortage costs) by ordering d units of each component, with cost $d(K+L)$. Let $\epsilon_0 = \frac{\bar{\epsilon}}{K+1}$, and for $\epsilon \in (0, \epsilon_0)$ let $L_0 = \frac{K(K-\bar{\epsilon})}{\bar{\epsilon}-\epsilon(K+1)}$. Then for $\epsilon \in (0, \epsilon_0)$ and $L \geq L_0$, compare the cost under the constant markup newsvendor inventory quantities to the feasible cost of filling all demand, which is an upper bound on OPT:

$$\begin{aligned} dK + (1-\epsilon)dL(K+1) &\geq (K+1-\bar{\epsilon})d(K+L), \\ &\geq (K+1-\bar{\epsilon})OPT. \end{aligned}$$

□

Proposition 2.4 demonstrates that the decentralization of the newsvendor heuristic can cause substantial mis-coordination between the components. In Example 2.1, (i.e., the family of examples proving Proposition 2.4), the newsvendor heuristic orders large quantities for several components that serve many products, but fails to order anything for a set of product specific components. This allows no demand to be filled in any scenario, driving the poor performance of the example.

We note that the family of examples of Proposition 2.4 have $\bar{\gamma} = \underline{\gamma}$ and $\bar{c} = \underline{c}$ so that the bound of Proposition 2.3 is $1+m$, thus demonstrating asymptotic tightness as m grows large. In Section 2.3.3

we show that this limitation on performance persists for an alternate family of newsvendor heuristics. This suggests that the system parameter m may be a structural obstacle for the performance of newsvendor decomposition approaches in general, and motivates our consideration of LP rounding heuristics in Section 2.4.

We further note that in Example 2.1, the average demand in the system can be chosen to be arbitrarily large, so that the lower bound of Proposition 2.1 also grows large. One might hope for better performance in such a scaling regime, as large demand intuitively implies more room for error in satisfying the integrality constraints. However, this example shows that the newsvendor heuristic can perform poorly, even in settings with large demand. In contrast, we demonstrate in Section 2.4 that our LP rounding heuristic is asymptotically optimal in such a large demand setting.

Bounds for Special Systems

In this section we consider bounds for two special types of systems, in order to show that the newsvendor heuristic can perform well in some settings. Intuitively, the first type of system has a hierarchical assembly structure, while the second has highly positively correlated demand. We prove that the constant markup newsvendor heuristic can perform well on these types of systems. We note that a non-hierarchical structure and negatively correlated demand are precisely the features driving poor performance in the systems in Example 2.1 (and also Example 2.3, to be introduced later).

Product laminar systems. We first consider a special class of assembly structures, called *product laminar systems*, in which products form a hierarchy in terms of the components they use. This hiearchy is with respect to set inclusion and is characterized by the sets $\mathcal{N}(j)$ having no non-trivial partial intersections:

Definition 2.1. *A system is product laminar if all $a_{ij} \in \{0,1\}$ and for any two products j, j', $\mathcal{N}(j) \cap \mathcal{N}(j') \in \{\mathcal{N}(j), \mathcal{N}(j'), \emptyset\}$.*

This laminar structure gives rise to an interesting property. We say an ATO system possesses the *balanced inventory* property if, in every scenario, unfulfilled demand for a product j coincides with an inventory shortage for all its components (i.e., $r_i < D_{i\omega}$ for all $i \in \mathcal{N}(j)$). Informally, this property allows for an unambiguous accounting of unfulfilled product demand in terms of the newsvendor shortage costs for each component. In terms of (ATO) and (DATO) formulations, the balanced inventory property facilitates dual cost accounting using extended ACS, allowing us to prove an approximation guarantee for laminar systems. We illustrate a product laminar system and the balanced inventory property in the following example.

Example 2.2. *Consider the M-system consisting of two components and three products, as illustrated in Figure 2.1. The sets of components used to make products 1 and 3 do not intersect and are both contained in the set of components used to make product 2. Thus, the M-system is product laminar. To illustrate the balanced inventory property, consider an M-system with one unit of inventory for each component. First, if demand for product 2 (the product served by both components) is greater than one, then there is an inventory shortage for both components. Hence all products have a shortage for all their components, and the condition holds. Otherwise, if demand for product 2 is at most one, then any shortage for either component is induced by its dedicated product (i.e. product 1 for component 1 or product 3 for component 2). Thus, the shortage can be wholly allocated to the dedicated product (which trivially has a shortage on all its components), without needing to assign a shortage to product 2.*

Intuitively, the balanced inventory property is useful because it eliminates concern that a product's shortage might be induced by high demand for another product with a shared component. This means the impact of demand uncertainty can remain localized to a single product, allowing coordination across components in the newsvendor approximation, and facilitating shortage cost accounting in the dual solution. Our approximation guarantee relies on proving that the balanced inventory property holds for general product laminar systems that have the same inventory level for all components. For assembly structures with a product laminar structure, the following approximation guarantee holds for all cost parameters and demand distributions.

Proposition 2.5. *Let $\bar{r} = \max_i\{r_i^{\gamma}\}$ and $\underline{r} = \min_i\{r_i^{\gamma}\} > 0$. In a product laminar system, setting each component's inventory level to $\bar{r}$ provides the approximation factor for* (ATO)

$$\max\left(\frac{\bar{\gamma}}{\underline{\gamma}}, \frac{\bar{r}}{\underline{r}}\right). \tag{2.30}$$

Intuitively, this result holds because the product laminar structure naturally aligns the optimal shortage costs in each component's newsvendor LP (2.14), as long as both the markups in the system and the newsvendor inventory quantities do not vary too widely. Before presenting a proof, we demonstrate that these conditions are met in some reasonable settings, and so optimality follows from Proposition 2.5. Recall that a system has identical markups if $\gamma_j = \underline{\gamma}$ for all products j.

Corollary 2.1. *If a product laminar system has: (i) identical markups, and (ii) identical aggregate demand distribution for each component, then the constant markup newsvendor inventory quantities* $r_i^{\underline{\gamma}}$, $1 \leq i \leq M$, *are optimal.*

Proof. Since each component has identical aggregate demand distribution, the constant markup shortage costs induce the same optimal newsvendor quantity for each component, i.e., $r_i^{\underline{\gamma}} = r_{i'}^{\underline{\gamma}}$ for all i, i'. Thus, by identical markups, the approximation factor from Proposition 2.5 is one. $\square$

As a simple example, consider an identical markup M-system (see Figure 2.1) where products 1 and 3 have identical marginal demand distributions independent of product 2. Then the aggregate demand distribution of each component is identical, so Corollary 2.1 establishes optimality of the constant markup newsvendor. From this perspective, Proposition 2.5 demonstrates that the heuristic is robust to small deviations from the optimality conditions of Corollary 2.1.

We now present a proof of Proposition 2.5, and begin with the following two properties of product laminar systems.

Lemma 2.3. *Given a product laminar system and a component i, for products $j, j' \in \mathcal{N}(i)$ with $|\mathcal{N}(j)| \leq |\mathcal{N}(j')|$, we have $\mathcal{N}(j) \subseteq \mathcal{N}(j')$.*

Proof. By $j, j' \in \mathcal{N}(i)$ we have $i \in \mathcal{N}(j), \mathcal{N}(j')$, so that $\mathcal{N}(j) \cap \mathcal{N}(j') \neq \emptyset$. Now $|\mathcal{N}(j) \cap \mathcal{N}(j')| \leq \min(|\mathcal{N}(j)|, |\mathcal{N}(j')|) = |\mathcal{N}(j)|$. Therefore, by the definition of a product laminar system, we must have $\mathcal{N}(j) \cap \mathcal{N}(j') = \mathcal{N}(j)$, so that $\mathcal{N}(j) \subseteq \mathcal{N}(j')$. $\square$

Lemma 2.4. *Given a product laminar system, a component i and an arbitrary subset of products S, consider an element $j \in \mathcal{N}(i) \setminus S$ with the smallest $|\mathcal{N}(j)|$. Then for any $i' \in \mathcal{N}(j)$ we have $\mathcal{N}(i) \setminus S \subseteq \mathcal{N}(i')$.*

Proof. Consider any $j' \in \mathcal{N}(i) \setminus S$, then we have $j, j' \in \mathcal{N}(i)$. Further, $|\mathcal{N}(j)| \leq |\mathcal{N}(j')|$, so by Lemma 2.3, we have $\mathcal{N}(j) \subseteq \mathcal{N}(j')$. Since $i' \in \mathcal{N}(j)$ we have $i' \in \mathcal{N}(j')$ which implies $j' \in \mathcal{N}(i')$. $\square$

For an N-vector of product demands $\mathbf{d} \geq 0$ and an M-vector of component inventories $\mathbf{r} \geq 0$, define the following sets

$$I(\mathbf{d}, \mathbf{r}) = \{i | r_i < \sum_{j \in \mathcal{N}(i)} d_j\}, \tag{2.31}$$

$$J(\mathbf{d}, \mathbf{r}) = \{j | \mathcal{N}(j) \subseteq I(\mathbf{d}, \mathbf{r})\}. \tag{2.32}$$

The set $I(\mathbf{d}, \mathbf{r})$ consists of all components with a shortage under demand $\mathbf{d}$, while $J(\mathbf{d}, \mathbf{r})$ is the set of products whose components all have shortages. Now consider the following condition on an inventory vector $\mathbf{r}$

$$r_i \geq \sum_{j \in \mathcal{N}(i) \setminus J(\mathbf{d}, \mathbf{r})} d_j, \ \forall\ 1 \leq i \leq M, \ \forall\ \mathbf{d} \geq 0. \tag{2.33}$$

We show below that this condition holds for product laminar systems with a constant inventory quantity, hence formalizing the balanced inventory property discussed above. First, the following Lemma formalizes the intuition from the main body that we can cover the shortages of each component by leaving demand unfilled only for products that have a shortage for all their components.

Lemma 2.5. *For an inventory vector* $\mathbf{r}$ *that satisfies condition* (2.33)*, for any* $\mathbf{d} \geq 0$ *there exists* $\mathbf{y} \geq 0$ *that satisfies the following properties*

$$y_j \leq d_j, \forall\ 1 \leq j \leq N, \tag{2.34}$$

$$y_j = 0, \ \forall\ j \notin J(\mathbf{d}, \mathbf{r}) \tag{2.35}$$

$$r_i + \sum_{j \in \mathcal{N}(i)} y_j \geq \sum_{j \in \mathcal{N}(i)} d_j, \ \forall\ 1 \leq i \leq M, \tag{2.36}$$

$$r_i + \sum_{j \in \mathcal{N}(i)} y_j = \sum_{j \in \mathcal{N}(i)} d_j, \ \forall\ i \in I(\mathbf{d}, \mathbf{r}). \tag{2.37}$$

Proof. By condition (2.33), equations (2.34-2.36) can always be satisfied by some $\mathbf{y}$. We therefore can focus on cases where $I(\mathbf{d}, \mathbf{r})$ is non-empty, since otherwise (2.37) holds trivially. Note also that when $I(\mathbf{d}, \mathbf{r})$ is non-empty, $J(\mathbf{d}, \mathbf{r})$ must also be non-empty for condition (2.33) to hold.

Suppose S is the set of nonnegative demand vectors such that equations (2.34-2.37) cannot be satisfied simultaneously. Let $\mathbf{d}' \in \arg\min_{\tilde{\mathbf{d}} \in S} \sum_j \tilde{d}_j$. Consider a $j^* \in J(\mathbf{d}', \mathbf{r})$ with $d'_{j^*} > 0$. Such a j^* must exist in order for the condition (2.33) to hold for $i \in I(\mathbf{d}', \mathbf{r})$. Let $\mathbf{y}'$ be the vector such that

$$y'_j = 0, \forall j \neq j^*, \text{and } y'_{j^*} = \min\{d_{j^*}, \min_{i \in \mathcal{N}(j^*)} \sum_{\tilde{j} \in \mathcal{N}(i)} d'_{\tilde{j}} - r_i\}. \tag{2.38}$$

Note that $j^* \in J(\mathbf{d}', \mathbf{r})$ implies that $y'_{j^*} > 0$. Then, by our assumption, we must have $\mathbf{d}' - \mathbf{y}' \notin S$. This implies that we can find $\mathbf{y}^*$ such that

$$y_j^* \leq d'_j - y'_j, \forall\ 1 \leq j \leq N, \tag{2.39}$$

$$y_j^* = 0, \forall\ j \notin J(\mathbf{d}' - \mathbf{y}', \mathbf{r}) \tag{2.40}$$

$$r_i + \sum_{j \in \mathcal{N}(i)} y_j^* \geq \sum_{j \in \mathcal{N}(i)} (d'_j - y'_j), \forall\ 1 \leq i \leq M, \tag{2.41}$$

$$r_i + \sum_{j \in \mathcal{N}(i)} y_j^* = \sum_{j \in \mathcal{N}(i)} (d'_j - y'_j), \forall\ i \in I(\mathbf{d}' - \mathbf{y}', \mathbf{r}). \tag{2.42}$$

Note that because $\mathbf{d}' - \mathbf{y}' \leq \mathbf{d}'$, by the definition (2.31), we have $I(\mathbf{d}' - \mathbf{y}', \mathbf{r}) \subseteq I(\mathbf{d}', \mathbf{r})$. Further, since $j \in J(\mathbf{d}' - \mathbf{y}', \mathbf{r})$ implies $\mathcal{N}(j) \subseteq I(\mathbf{d}' - \mathbf{y}', \mathbf{r}) \subseteq I(\mathbf{d}', \mathbf{r})$, which implies $j \in J(\mathbf{d}', \mathbf{r})$, we also have $J(\mathbf{d}' - \mathbf{y}', \mathbf{r}) \subseteq J(\mathbf{d}', \mathbf{r})$. Then it is clear that $\mathbf{y}^* + \mathbf{y}'$ satisfies equations (2.34-2.36) for $\mathbf{d}'$. Next we show that $\mathbf{y}^* + \mathbf{y}'$ also satisfies equation (2.37), which gives us a contradiction and implies that S must be empty. If $I(\mathbf{d}' - \mathbf{y}', \mathbf{r}) = I(\mathbf{d}', \mathbf{r})$, we are done immediately by equation (2.42). If $I(\mathbf{d}' - \mathbf{y}', \mathbf{r}) \neq I(\mathbf{d}', \mathbf{r})$, then we must have $I(\mathbf{d}' - \mathbf{y}, \mathbf{r}) \subset I(\mathbf{d}', \mathbf{r})$. For any $i^* \in I(\mathbf{d}', \mathbf{r}) \setminus I(\mathbf{d}' - \mathbf{y}', \mathbf{r})$ we have both $\sum_{j \in \mathcal{N}(i^*)} d'_j > r_{i^*}$ and $\sum_{j \in \mathcal{N}(i^*)} (d'_j - y'_j) \leq r_{i^*}$. Since only $y'_{j^*} > 0$, it must be the case that $j^* \in \mathcal{N}(i^*)$ and $y'_{j^*} \geq \sum_{j \in \mathcal{N}(i^*)} d'_j - r_{i^*}$. Then by equation (2.38), we must have that

$$y'_{j^*} = \sum_{j \in \mathcal{N}(i^*)} d'_j - r_{i^*}. \tag{2.43}$$

Further, consider any $j \in \mathcal{N}(i^*)$. Because $i^* \notin I(\mathbf{d}' - \mathbf{y}', \mathbf{r})$, we cannot have $\mathcal{N}(j) \subseteq I(\mathbf{d}' - \mathbf{y}', \mathbf{r})$, so by equation (2.40), $y_j^* = 0$. Combining this with equations (2.38) and (2.43), we have that

$$r_{i^*} + \sum_{j \in \mathcal{N}(i^*)} (y_j^* + y'_j) = \sum_{j \in \mathcal{N}(i^*)} d'_j.$$

This implies that equation (2.37) is satisfied, and we are done. □

Next we show that for a product laminar system, a policy that sets the same inventory level for each component satisfies condition (2.33).

Lemma 2.6. *In a product laminar system, setting the same inventory level $\tilde{r}$ for all components satisfies condition* (2.33).

Proof. Let $\tilde{\mathbf{r}}$ denote an M-vector consisting of all $\tilde{r}$'s, and consider any $\mathbf{d} \geq 0$. For $i \notin I(\mathbf{d}, \tilde{\mathbf{r}})$ the inequality (2.33) holds trivially, so we focus on $i \in I(\mathbf{d}, \tilde{\mathbf{r}})$. We proceed by way of proof by contradiction. Assume that there exists $\mathbf{d}$ and $i \in I(\mathbf{d}, \tilde{\mathbf{r}})$ such that $\tilde{r} < \sum_{j \in \mathcal{N}(i) \setminus J(\mathbf{d}, \tilde{\mathbf{r}})} d_j$. Then choose $j \in \mathcal{N}(i) \setminus J(\mathbf{d}, \tilde{\mathbf{r}})$ with the smallest $|\mathcal{N}(j)|$. By Lemma 2.4, for any $i' \in \mathcal{N}(j)$ we have $\mathcal{N}(i) \setminus J(\mathbf{d}, \tilde{\mathbf{r}}) \subseteq \mathcal{N}(i')$. By assumption we have $\tilde{r} < \sum_{j \in \mathcal{N}(i) \setminus J(\mathbf{d}, \tilde{\mathbf{r}})} d_j$ and by $\mathcal{N}(i) \setminus J(\mathbf{d}, \tilde{\mathbf{r}}) \subseteq \mathcal{N}(i')$ we have $\sum_{j \in \mathcal{N}(i) \setminus J(\mathbf{d}, \tilde{\mathbf{r}})} d_j \leq \sum_{j \in \mathcal{N}(i')} d_j$. Therefore $\tilde{r} < \sum_{j \in \mathcal{N}(i')} d_j$ which implies $i' \in I(\mathbf{d}, \tilde{\mathbf{r}})$. Since this is true for any $i' \in \mathcal{N}(j)$ we have $\mathcal{N}(j) \subseteq I(\mathbf{d}, \tilde{\mathbf{r}})$ so that $j \in J(\mathbf{d}, \tilde{\mathbf{r}})$. But this contradicts the fact that $j \in \mathcal{N}(i) \setminus J(\mathbf{d}, \tilde{\mathbf{r}})$. □

Proof of Proposition 2.5. Let $\mathbf{r}_{\max}$ denote the M-vector consisting of $\bar{r}$ in each element. Since this is a product laminar system and $\mathbf{r}_{\max}$ sets the same inventory level for all components, Lemma 2.6 guarantees that condition (2.33) is satisfied. Therefore by Lemma 2.5, for every demand scenario $\mathbf{d}_\omega$ there is a vector of shortages $\mathbf{y}_\omega$ that satisfy equations (2.34-2.37). By (2.34) and (2.36) it is clear that $\mathbf{y}_\omega$ is feasible to (ATO). Choose such a $\mathbf{y}_\omega$ for each scenario ω and let it constitute the primal feasible shortage.

We establish the result using Proposition 2.2 with factors $\nu = 1$, $\kappa = \bar{\gamma}/\underline{\gamma}$, and $\delta = \eta = \max\left(\frac{\bar{\gamma}}{\underline{\gamma}}, \frac{\bar{r}}{\underline{r}}\right)$. First we consider primal complementary slackness. Since $r_i^{\gamma} \geq \underline{r} > 0$ for all i, by complementary slackness for the LP (2.14) we know that constraint (2.6) binds for all i, thus satisfying equation (2.22) with $\nu = 1$. By (2.35), if $y_{j\omega} > 0$ then $j \in J(\mathbf{d}_\omega, \mathbf{r}_{\max})$. Therefore, for any $i \in \mathcal{N}(j)$ we have $\sum_{j \in \mathcal{N}(i)} d_j > \bar{r} \geq r_i^{\gamma}$ which implies that $\theta_{i\omega}^{\gamma} = \mu_\omega \underline{\gamma} c_i$ by (2.17). Thus $\sum_{i \in \mathcal{N}(j)} \theta_{i\omega}^{\gamma} = \underline{\gamma} \mu_\omega \sum_{i \in \mathcal{N}(j)} c_i$. Thus, since $\bar{\gamma} \geq \gamma_j$ and $\pi_{j\omega}^{\gamma} = 0$ we have

$$\sum_{i \in \mathcal{N}(j)} \theta_{i\omega}^{\gamma} - \pi_{j\omega}^{\gamma} = \underline{\gamma} \mu_\omega \sum_{i \in \mathcal{N}(j)} c_i,$$

$$\geq \frac{\underline{\gamma}}{\bar{\gamma}} \mu_\omega p_j,$$

so equation (2.21) is satisfied with $\kappa = \bar{\gamma}/\underline{\gamma}$.

For dual complementary slackness, consider when $\theta_{i\omega}^{\gamma} > 0$. First, if $i \in I(\mathbf{d}_\omega, \mathbf{r}_{\max})$, then constraint (2.1) binds by (2.37). Thus, since $\nu = 1$ and $\kappa = \bar{\gamma}/\underline{\gamma}$, in this case we have

$$\nu \bar{r} + \kappa \sum_{j \in \mathcal{N}(i)} y_{j\omega} \leq \frac{\bar{\gamma}}{\underline{\gamma}} \left(\bar{r} + \sum_{j \in \mathcal{N}(i)} y_{j\omega} \right),$$

$$= \frac{\bar{\gamma}}{\underline{\gamma}} \sum_{j \in \mathcal{N}(i)} d_{j\omega}.$$

Next we consider the case when $i \notin I(\mathbf{d}_\omega, \mathbf{r}_{\max})$, so that for all $j \in \mathcal{N}(i)$ we have $j \notin J(\mathbf{d}_\omega, \mathbf{r}_{\max})$. Then by (2.35) we have $y_{j\omega} = 0$ for all $j \in \mathcal{N}(i)$. Further, by (2.17), when $\theta^{\gamma}_{i\omega} > 0$ we have $\sum_{j\in\mathcal{N}(i)} d_{j\omega} \geq r^{\gamma}_i \geq \underline{r}$, so that

$$\nu\bar{r} + \kappa \sum_{j\in\mathcal{N}(i)} y_{j\omega} = \bar{r},$$

$$\leq \frac{\bar{r}}{\underline{r}} \sum_{j\in\mathcal{N}(i)} d_{j\omega}.$$

Therefore (2.23) is satisfied with $\delta = \max\left(\frac{\bar{\gamma}}{\underline{\gamma}}, \frac{\bar{r}}{\underline{r}}\right)$. Finally, since $\pi^{\gamma}_{j\omega} = 0$, (2.24) is trivially satisfied for $\eta = \delta$. The result follows by Proposition 2.2. □

Comonotone demand. Next we consider arbitrary ATO structures with *comonotone* product demand. Intuitively, comonotonicity captures a situation where demand scenarios follow the same order for each product, i.e., a case of extreme positive correlation.

Definition 2.2. *Demand is comonotone if there exists a permutation of demand scenarios* $\omega_1, \omega_2, \ldots, \omega_k, \ldots$ *such that* $d_{j\omega_k} \leq d_{j\omega_{k+1}}$ *for all indices* k *and all products* j*.*

The advantage of this demand property is that it coordinates the newsvendor critical quantiles across components, giving the following result.

Proposition 2.6. *For comonotone demand, the constant markup newsvendor inventory quantities,* r^{γ}_i*,* $1 \leq i \leq M$*, provide an approximation factor of* $\bar{\gamma}/\underline{\gamma}$ *for* (ATO)*.*

Therefore, by Proposition 2.6, the constant markup newsvendor performs well when demand is comonotone and the markups do not vary too widely. Intuitively, this result implies that it is easier to coordinate the newsvendor solutions when demand is highly positively correlated.

To prove this result we first define the following marginal demand scenario and shortage quantities

$$k^* = \min\left\{k \,\middle|\, \sum_{l>k} \mu_{\omega_l} < \frac{1}{\underline{\gamma}}\right\}, \tag{2.44}$$

$$y_{j\omega_k} = (d_{j\omega_k} - d_{j\omega_{k^*}})^+, \tag{2.45}$$

and prove the following Lemma.

Lemma 2.7. *Under comonotone demand, we have* $\sum_j a_{ij} d_{j\omega_{k^*}} = r_i^\gamma$ *for all components* i.

Proof. For a demand scenario ω_k, recall the definition $D_{i\omega_k} = \sum_j a_{ij} d_{j\omega_k}$. For two indices k and k', we note that comonotone demand implies $k \leq k'$ if and only if $D_{i\omega_k} \leq D_{i\omega_{k'}}$.

Consider a component i and let $\tilde{r}_i = D_{i\omega_{k^*}}$. Consider a demand scenario k such that $D_{i\omega_k} > \tilde{r}_i = D_{i\omega_{k^*}}$. Then by comonotonicity of demand we would have a contradiction if $k \leq k^*$, so we must have $k > k^*$. Thus we have $\{\omega | D_{i\omega} > \tilde{r}_i\} \subseteq \{\omega_k | k > k^*\}$ so that $\mathbb{P}[D_{i\omega} > \tilde{r}_i] \leq \sum_{k>k^*} \mu_{\omega_k} < \frac{1}{\underline{\gamma}}$, where the last inequality follows from the definition of k^*, (2.44). Therefore, by the definition of r_i^γ (2.16) we must have $r_i^\gamma \leq \tilde{r}_i$.

Now for a contradiction, assume $r_i^\gamma < \tilde{r}_i$, so that $r_i^\gamma \leq \tilde{r}_i - 1$ and $\mathbb{P}[D_{i\omega} > \tilde{r}_i - 1] \leq \mathbb{P}[D_{i\omega} > r_i^\gamma] < \frac{1}{\underline{\gamma}}$. Then consider a demand scenario ω_k for which $k \geq k^*$, so that $k > k^* - 1$. By comonotone demand we must have $D_{i\omega_k} \geq D_{i\omega_{k^*}} > \tilde{r}_i - 1$, so that $\{\omega_k | k > k^* - 1\} \subseteq \{\omega | D_{i\omega} > \tilde{r}_i - 1\}$. Therefore $\sum_{k>k^*-1} \mu_{\omega_k} \leq \mathbb{P}[D_{i\omega} > \tilde{r}_i - 1] < \frac{1}{\underline{\gamma}}$ which contradicts the definition of k^*, (2.44). Thus, we must have $\tilde{r}_i = r_i^\gamma$. □

Proof of Proposition 2.6. To establish primal feasibility, it is clear the non-negativity and integrality constraints are satisfied. The primal constraint (2.1) is satisfied for $k \leq k^*$ since $D_{i\omega_k} \leq D_{i\omega_{k^*}} = r_i^\gamma$ by Lemma 2.7. Further, for $k > k^*$ we have $y_{j\omega} = d_{j\omega} - d_{j\omega_{k^*}}$ so that for any i

$$\begin{aligned} r_i^\gamma + \sum_j a_{ij} y_{j\omega_k} &= D_{i\omega_{k^*}} + \sum_j a_{ij} \left(d_{j\omega} - d_{j\omega_{k^*}} \right), \\ &= D_{i\omega_k}, \end{aligned}$$

and constraint (2.1) is satisfied with equality. Finally, constraint (2.2) is satisfied by the definition of $y_{j\omega}$ in (2.45).

Next, at the dual solution $\theta_{i\omega}^\gamma$, $\pi_{j\omega}^\gamma$, we claim the extended ACS conditions hold for $\kappa = \delta = \eta = \bar{\gamma}/\underline{\gamma}$ and $\nu = 1$. If $r_i^\gamma > 0$ then by complementary slackness for the LP (2.14), constraint (2.6) binds and equation (2.22) holds with $\nu = 1$. When $y_{j\omega_k} > 0$ we must have $d_{j\omega_k} > d_{j\omega_{k^*}}$, which by comonotone demand implies that $D_{i\omega_k} > D_{i\omega_{k^*}} = r_i^\gamma$ for all $i \in \mathcal{N}(j)$. Thus by (2.17) we have $\theta_{i\omega_k}^\gamma = \underline{\gamma} \mu_{\omega_k} c_i$ for all $i \in \mathcal{N}(j)$, and since $\pi_{j\omega}^\gamma = 0$, we have

$$\sum_i a_{ij} \theta_{i\omega_k}^\gamma - \pi_{j\omega_k}^\gamma = \underline{\gamma} \mu_{\omega_k} \sum_i a_{ij} c_i \geq \frac{\underline{\gamma}}{\bar{\gamma}} \mu_{\omega_k} p_j.$$

so that the extended ACS condition (2.21) holds for $\kappa = \bar{\gamma}/\underline{\gamma}$. If $\theta^{\gamma}_{i\omega_k} > 0$, then by (2.17), we have $D_{i\omega_k} \geq r_i^{\gamma} = D_{i\omega_{k^*}}$, so that $k \geq k^*$. Then by demand comonotonicity, we have $d_{j\omega_k} \geq d_{j\omega_{k^*}}$ for all j, so that $y_{j\omega_k} = d_{j\omega_k} - d_{j\omega_{k^*}}$ and constraint (2.1) binds, as discussed when establishing feasibility. Then, since $\nu = 1 \leq \bar{\gamma}/\underline{\gamma}$ we have

$$\nu r_i^{\gamma} + \kappa \sum_j a_{ij} y_{j\omega_k} \leq \frac{\bar{\gamma}}{\underline{\gamma}} \left(r_i^{\gamma} + \sum_j a_{ij} y_{j\omega_k} \right),$$

$$= \frac{\bar{\gamma}}{\underline{\gamma}} D_{i\omega_k},$$

so that (2.23) holds with $\delta = \bar{\gamma}/\underline{\gamma}$. Finally, since $\pi^{\gamma}_{j\omega} = 0$, (2.24) is trivially satisfied for $\eta = \delta$. The result follows by Proposition 2.2. □

2.3.3 Alternative Newsvendor Heuristic

In this section we consider an alternative newsvendor heuristic, i.e., a different choice of shortage cost assignment q_i. We first note that the optimality order quantity defined in equation (2.16) for the constant markup newsvendor can be extended to any q_i as follows

$$r_i^q = \min \left\{ r \in \mathbb{Z}_+ \;\middle|\; \mathbb{P}\left[D_i > r\right] < \frac{c_i}{q_i} \right\}. \tag{2.46}$$

We will introduce an alternative newsvendor heuristic to the constant markup heuristic considered in Section 2.3 and demonstrate that it suffers from the same limitation on performance. Intuitively, we can think of the constant markup shortage costs as a conservative estimate of the true cost of shortage for a component, since $q_i = \underline{\gamma} c_i \leq p_j$ for all $j \in \mathcal{N}(i)$, and indeed there may be a large gap between q_i and p_j. If we want a less conservative estimate, we could consider setting the newsvendor heuristic's shortage cost equal to some convex combination of the shortage costs in the original system, i.e., specify non-negative weights λ_{ij} such that $\sum_{j \in \mathcal{N}(i)} \lambda_{ij} = 1$ for all i and set

$$q_i = \sum_{j \in \mathcal{N}(i)} \lambda_{ij} p_j. \tag{2.47}$$

This cost assignment ensures that q_i is within the range of the relevant shortage costs in the original system, and the weights can be chosen to place q_i anywhere in this range. We call a newsvendor heuristic that uses costs of the form (2.47) a *weighted cost* newsvendor heuristic. (We note that a newsvendor heuristic suggested by Lu and Song (2005) takes this form, where the weights are determined by a product's contribution to the expected demand for component i). Next we introduce our family of examples demonstrating poor performance of the weighted cost newsvendor heuristic.

Example 2.3 (($K, L, \hat{\epsilon}$) system)**.** *Consider the same system and demand structure as in Example 2.1. Again assume $c_i = 1$ for all i. However, we modify shortage costs $p_j = \frac{1+\hat{\epsilon}}{L\mu}$ for some $0 < \hat{\epsilon} < L - 1$, for all j.*

Proposition 2.7. *For any $K \geq 1,\ L \geq 2$ and $\bar{\epsilon} \in (0, K)$, there exists ϵ_0 such that for all $\epsilon \in (0, \epsilon_0)$, the inventory quantities $r_i^q,\ 1 \leq i \leq M$, for any weighted cost newsvendor heuristic in the $(K, L, \hat{\epsilon})$ system have cost at least $(K + 1 - \bar{\epsilon})OPT$.*

Proof. Since the shortage cost of every product is the same, for any weights used in a weighted cost newsvendor we have all $q_i = \frac{1+\hat{\epsilon}}{L\mu}$, and since $c_i = 1$, we have $c_i/q_i = \frac{L\mu}{1+\hat{\epsilon}}$. Upon inspection of the demands it is clear that each component sees demand for at most d units in any demand scenario, and so the optimal newsvendor quantity must be an integer between zero and d. For $1 \leq i \leq K$, ordering a quantity $r < d$ implies

$$\mathbb{P}[D_i > r] = L\mu \geq \frac{L\mu}{1 + \hat{\epsilon}},$$

which violates the optimality condition for r_i^q in equation (2.46). Since $\mathbb{P}[D_i > d] = 0$ satisfies the condition, we have $r_i^q = d$ for all $1 \leq i \leq K$. Meanwhile, we note that $\hat{\epsilon} < L - 1$ implies $1 < L/(1 + \hat{\epsilon})$, so that for $K + 1 \leq i \leq K + L$ we have

$$\mathbb{P}[D_i > 0] = \mu < \frac{L\mu}{1 + \hat{\epsilon}},$$

and so $r_i^q = 0$ by (2.46). Thus, for each product j, we will always be missing the component $K + j$, and hence cannot fill any demand in any scenario. So, the optimal shortages given the inventory quantities r_i^q must be $\mathbf{y}_k = \mathbf{d}_k$, in each demand scenario k. Therefore, the ordering cost plus the minimal shortage cost associated with the inventory quantities r_i^q is

$$dK + \frac{1 + \hat{\epsilon}}{L\mu} d(L\mu) = d(K + 1 + \hat{\epsilon}).$$

Alternatively, we could order no inventory and incur the shortage penalty for all demand, at a cost of $\frac{1+\hat{\epsilon}}{L\mu} d(L\mu) = d(1 + \hat{\epsilon})$. Let $\epsilon_0 = \frac{\bar{\epsilon}}{K - \bar{\epsilon}}$. Then for $\hat{\epsilon} \in (0, \epsilon_0)$, compare the cost under the weighted cost newsvendor inventory quantities to the feasible cost of ordering nothing, which is an upper bound on OPT:

$$\begin{aligned} d(K+1+\hat{\epsilon}) &\geq (K+1-\bar{\epsilon})d(1+\hat{\epsilon}), \\ &\geq (K+1-\bar{\epsilon})OPT \end{aligned}$$

□

2.4 LP Rounding

We now consider another natural approximation approach for (ATO): rounding a solution of the LP relaxation to a feasible integer solution. In the following sections, we propose two such rounding schemes and establish a few approximation guarantees. First, in Section 2.4.1, we prove an approximation factor that scales with the inverse of both mean component demand and expected absolute deviation from mean component demand, demonstrating asymptotic optimality as demand gets large. Then, in Section 2.4.2, we prove an approximation factor of 1.8, independent of any parameters of the ATO system. Finally, in Section 2.4.3, we show that the LP relaxation of (ATO) can be solved efficiently using a stochastic subgradient method.

To the best of our knowledge, our results provide the first constant factor approximation algorithm and the first asymptotically optimal algorithm for (ATO) based on an LP rounding scheme. The existing ATO research contains a few examples of using LP rounding to determine base-stock levels for a dynamic problem under the assumption of first-come-first-served second stage allocation (Lu et al., 2005; van Jaarsveld and Scheller-Wolf, 2015), with rounding schemes implemented in a heuristic fashion (i.e. round to the nearest integer), and numerical studies validating the procedures. These heuristics can be adapted to a one-period problem like (ATO) by imagining that demand for each product arrives all at once (which we do in our numerical simulations of Section 2.5 in order to have a point of comparison for our rounding schemes). Other works use rounding to approximate values outside of the context of an LP solution (e.g., rounding inventory or basestock levels in a dynamic setting as in Glasserman and Wang, 1998; Song, 2002; DeCroix et al., 2009).

Before presenting our results, we briefly review the extensive literature using rounding to find approximate solutions in operations problems. Examples include filtering and rounding of Shmoys et al. (1997) (facility location problems), randomized rounding of Teo and Bertsimas (2001) (lot sizing problems) and Schulz and Skutella (2002) (scheduling unrelated machines), and iterative rounding techniques of Jain and Vazirani (2001) (survivable network design). Other examples of rounding in the operations management literature include the one-warehouse multi-retailer (OWMR) problem (Levi et al., 2008) and deterministic inventory problems (Nagarajan and Shi, 2016).

2.4.1 Asymptotic Optimality For Large Demand

To introduce our first rounding scheme, let $(r_i^{LP}, y_{j\omega}^{LP})$ denote an optimal solution to the LP relaxation of (ATO). Then consider rounding to an integer solution as follows

$$r_i^I = \lfloor r_i^{LP} \rfloor, \tag{2.48}$$

$$y_{j\omega}^I = \lceil y_{j\omega}^{LP} \rceil. \tag{2.49}$$

It is clear that rounding up the shortages in (2.49) maintains the demand constraint (2.2) (i.e., $y_{j\omega} \leq d_{j\omega}$), since the LP solution satisfies this constraint and demand is integral. Similarly, rounding down the inventory quantity cannot violate the inventory constraint (2.1) (i.e., $r_i + \sum_j a_{ij} y_{j\omega} \geq \sum_j a_{ij} d_{j\omega}$), since the LP solution satisfied this constraint, the shortage quantities are rounded up, and the right-hand side is integral.

We now demonstrate that the rounding (2.48-2.49) performs well as demand scales in any given system. Define the following

$$\bar{d} = \hat{\gamma} \left(\frac{\sum_i c_i \bar{D}_i}{\sum_i c_i n_i} \right),$$

$$\rho = \gamma^* \left(\frac{\sum_i c_i \mathbb{E}\left[|D_i - \bar{D}_i| \right]}{\sum_i c_i n_i} \right).$$

Recall that $\bar{D}_i = \sum_j a_{ij} \bar{d}_j$ and $n_i = \sum_j a_{ij}$. Thus, the quantity $\bar{d}$ is proportional to a weighted average of the expected demand in the system, while ρ is proportional to a weighted average of the expected absolute deviation from mean component demand. Together, these values represent the mean, and variability of demand in the system. Recall that for systems with large markups ($\underline{\gamma} \geq 2$), the expression (2.19) implies that $\hat{\gamma} = 1$ and $\gamma^* = 1/2$, i.e., these factors are just constants that scale the demand quantities of interest. Our next result uses these demand quantities to characterize an approximation factor for the rounding scheme (2.48-2.49).

Theorem 2.1. *The LP rounding (2.48-2.49) produces a feasible solution to* (ATO) *and provides the following approximation factor*

$$1 + \frac{1}{\bar{d} + \rho}$$

Proof. First we prove feasibility of (2.48-2.49). Since $d_{j\omega}$ is an integer, it is clear that $0 \leq y_{j\omega}^{LP} \leq d_{j\omega}$ implies $0 \leq y_{j\omega}^{I} \leq d_{j\omega}$. Further, we have

$$r_i^I + \sum_j a_{ij} y_{j\omega}^I > r_i^{LP} - 1 + \sum_j a_{ij} y_{j\omega}^{LP},$$

$$\geq -1 + \sum_j a_{ij} d_{j\omega},$$

and since the right- and left-hand sides are integers and the inequality is strict, we have $r_i^I + \sum_j a_{ij} y_{j\omega}^I \geq \sum_j a_{ij} d_{j\omega}$, thus satisfying the inventory constraint (2.1).

To prove the approximation factor, it is clear that $r_i^I \leq r_i^{LP}$. Further, since we only round up in (2.49) if $0 < y_{j\omega}^{LP} < d_{j\omega}$, we have

$$y_{j\omega}^I \leq y_{j\omega}^{LP} + \mathbb{1}_{\{0 < y_{j\omega}^{LP} < d_{j\omega}\}}.$$

Thus, we can write the cost of the primal solution as

$$\sum_i c_i r_i^I + \sum_{j,\omega} \mu_\omega p_j y_{j\omega}^I, \leq \sum_i c_i r_i^{LP} + \sum_{j,\omega} \mu_\omega p_j y_{j\omega}^{LP} + \sum_{j,\omega} \mu_\omega p_j \mathbb{1}_{\{0 < y_{j\omega}^{LP} < d_{j\omega}\}},$$

$$\leq OPT + \sum_{j,\omega} \mu_\omega p_j \mathbb{1}_{\{0 < y_{j\omega}^{LP} < d_{j\omega}\}}.$$

Thus, we focus on bounding $\sum_{j,\omega} \mu_\omega p_j \mathbb{1}_{\{0 < y_{j\omega}^{LP} < d_{j\omega}\}}$. Let $\theta_{i\omega}^{LP}$, $\pi_{j\omega}^{LP}$ be an optimal dual solution. By the complementary slackness conditions, if $y_{j\omega}^{LP} < d_{j\omega}$ then $\pi_{j\omega}^{LP} = 0$, and if $y_{j\omega}^{LP} > 0$ then $\sum_i a_{ij} \theta_{i\omega}^{LP} = \mu_\omega p_j$, and so we have

$$\sum_{j,\omega} \mu_\omega p_j \mathbb{1}_{\{0 < y_{j\omega}^{LP} < d_{j\omega}\}} = \sum_{j,\omega} \mathbb{1}_{\{0 < y_{j\omega}^{LP} < d_{j\omega}\}} \sum_i a_{ij} \theta_{i\omega}^{LP},$$

$$\leq \sum_i \sum_\omega \theta_{i\omega}^{LP} \sum_j a_{ij},$$

$$\leq \sum_i c_i n_i,$$

where the first equality follows from complementary slackness, the first inequality follows from interchanging the order of summation and $\mathbb{1}_{\{0 < y_{j\omega}^{LP} < d_{j\omega}\}} \leq 1$, and the final inequality from dual constraint (2.6).

Thus, the cost of the primal solution is bounded by

$$\begin{aligned} OPT + \sum_i c_i n_i &= OPT + \frac{\left(\sum_i c_i n_i\right) \sum_i c_i \left(\hat{\gamma}\bar{D}_i + \gamma^* \mathbb{E}\left[|D_i - \bar{D}_i|\right]\right)}{\sum_i c_i \left(\hat{\gamma}\bar{D}_i + \gamma^* \mathbb{E}\left[|D_i - \bar{D}_i|\right]\right)}, \\ &\leq \left(1 + \frac{\sum_i c_i n_i}{\hat{\gamma}\sum_i c_i \bar{D}_i + \gamma^* \sum_i c_i \mathbb{E}\left[|D_i - \bar{D}_i|\right]}\right) OPT, \\ &= \left(1 + \frac{1}{\bar{d} + \rho}\right) OPT, \end{aligned}$$

where the inequality follows from Proposition 2.1. □

Thus, in any asymptotic regime where $\bar{d} + \rho$ becomes large, the rounding scheme (2.48-2.49) becomes close to optimal. For example, consider a simple scaling regime where, for a given ATO system, the assembly structure and costs remain fixed, while product demands are scaled up by a positive integer β. Then it is clear that if $\bar{d}$ were the average demand in the original system, this quantity becomes $\beta\bar{d}$ in the scaled system. So, for large β, the approximation factor of Theorem 2.1 approaches one. In Section 3.2, we discuss another regime where demand scales as a result of summing a growing number of independent draws from identical demand distributions, which also results in asymptotic optimality.

2.4.2 Constant Factor Approximation

Next we enhance our LP rounding algorithm to establish an approximation guarantee that does not depend on any system parameters. We start our analysis by introducing a critical insight for the approximability of (ATO): approximation of the second stage cost can be achieved using both the first and second stage LP cost. In particular, we use the complimentary slackness condition (2.9) to demonstrate that the shortage cost resulting from rounding shortages up in (2.49), is less than the ordering cost plus shortage cost in the original LP solution. This gives the following result.

Proposition 2.8. *The LP rounding (2.48-2.49) produces a solution with cost less than*

$$2\sum_i c_i r_i^{LP} + \sum_{j,\omega} \mu_\omega p_j y_{j\omega}^{LP} \tag{2.50}$$

Proof. It is clear that $r_i^I \leq r_i^{LP}$. Further, since the demand quantities are integers and $y_{j\omega}^{LP} \leq d_{j\omega}$, rounding up shortages as in (2.49) satisfies

$$y_{j\omega}^I \leq d_{j\omega} \mathbb{1}_{\{y_{j\omega}^{LP} > 0\}}.$$

Let $\theta_{i\omega}^{LP}$, $\pi_{j\omega}^{LP}$ be an optimal dual solution. Next we claim that for all j, ω

$$\sum_i a_{ij}\theta_{i\omega}^{LP} - \pi_{j\omega}^{LP} \geq 0,$$

which we prove by contradiction. If $\sum_i a_{ij}\theta_{i\omega}^{LP} < \pi_{j\omega}^{LP}$, then $\pi_{j\omega}^{LP}$ can be strictly decreased by setting it equal to $\sum_i a_{ij}\theta_{i\omega}^{LP}$, which remains feasible since the left-hand side of constraint (2.5) would evaluate to zero, and this is the only constraint $\pi_{j\omega}$ appears in. This would strictly increase the objective, hence the original solution was not optimal, a contradiction.

Finaly, by complimentary slackness condition (2.9) we have

$$\mu_\omega p_j \mathbb{1}_{\{y_{j\omega}^{LP}>0\}} = \mathbb{1}_{\{y_{j\omega}^{LP}>0\}} \left(\sum_i a_{ij}\theta_{i\omega}^{LP} - \pi_{j\omega}^{LP} \right),$$

so that the we can bound the cost of the rounded solution by

$$\begin{aligned}
\sum_i c_i r_i^{LP} + \sum_{j,\omega} \mu_\omega p_j d_{j\omega} \mathbb{1}_{\{y_{j\omega}^{LP}>0\}} &= \sum_i c_i r_i^{LP} + \sum_{j,\omega} d_{j\omega} \mathbb{1}_{\{y_{j\omega}^{LP}>0\}} \left(\sum_i a_{ij}\theta_{i\omega}^{LP} - \pi_{j\omega}^{LP} \right), \\
&\leq \sum_i c_i r_i^{LP} + \sum_{j,\omega} d_{j\omega} \left(\sum_i a_{ij}\theta_{i\omega}^{LP} - \pi_{j\omega}^{LP} \right), \\
&= 2\sum_i c_i r_i^{LP} + \sum_{j,\omega} \mu_\omega p_j y_{j\omega}^{LP},
\end{aligned}$$

where the final equality follows from strong duality. □

This result readily implies an approximation factor of 2, but we note that it allows room for improvement. In particular, if the expected shortage cost of the LP solution ($\sum_{j,\omega} \mu_\omega p_j y_{j\omega}^{LP}$) is large relative to the ordering cost ($\sum_i c_i r_i^{LP}$), then the approximation factor implied by (2.50) will be less than 2. To take advantage of this observation, we consider an alternative rounding scheme that provides good performance in the opposite case, i.e., when ordering cost is large relative to shortage cost.

For any factor $\alpha > 1$ consider rounding as follows

$$r_i^I = \left\lfloor \alpha r_i^{LP} \right\rfloor, \tag{2.51}$$

$$y_{j\omega}^I = \min\left(\left\lfloor \frac{\alpha}{\alpha - 1} y_{j\omega}^{LP} \right\rfloor, d_{j\omega} \right). \tag{2.52}$$

We note that in many cases, scaling up before rounding will result in an integer larger than the original LP variable, and so (2.51-2.52) may effectively behave like a rounding up procedure.

Further, we highlight that the factor α characterizes a trade-off between approximating the first and second stage costs when rounding the LP solution. In particular, if α is close to one, then (2.51) approximates the LP ordering quantities well, but $\alpha/(\alpha-1)$ is large so (2.52) may not be a good approximation for the LP shortage quantities. The opposite effect occurs when α is large, demonstrating the trade-off inherent in the rounding. Intuitively, this scaling trade-off allows us to take advantage of situations in which the ordering cost is much larger than the shortage cost. This trade-off is summarized in the following bound on the cost of the rounded solution.

Proposition 2.9. *The LP rounding (2.51-2.52) produces a feasible solution to* (ATO) *with cost less than*

$$\alpha \sum_i c_i r_i^{LP} + \frac{\alpha}{\alpha - 1} \sum_{j,\omega} \mu_\omega p_j y_{j\omega}^{LP}.$$

Proof. Given feasibility, it is clear from (2.51-2.52) that $r_i^I \leq \alpha r_i^{LP}$ and $y_{j\omega}^I \leq \alpha/(\alpha-1) y_{j\omega}^{LP}$, which implies the result. To establish feasibility, since $d_{j\omega}$ is an integer, it is clear that all quantities are positive integers and that the demand constraint (2.2) is satisfied. Consider constraint (2.1) for some i and ω, and let $J_\omega = \{j | y_{j\omega}^{LP} \geq (1 - \frac{1}{\alpha}) d_{j\omega}\}$. For $j \in J_\omega$, we have $\alpha/(\alpha-1) y_{j\omega}^{LP} \geq d_{j\omega}$, so the rounding in (2.52) sets $y_{j\omega}^I = d_{j\omega}$. Further, we have

$$\begin{aligned}\sum_j a_{ij} d_{j\omega} &\leq \sum_j a_{ij} y_{j\omega}^{LP} + r_i^{LP}, \\ &< \left(1 - \frac{1}{\alpha}\right) \sum_{j \notin J_\omega} a_{ij} d_{j\omega} + \sum_{j \in J_\omega} a_{ij} d_{j\omega} + r_i^{LP},\end{aligned}$$

where the first inequality follows from feasibility of the LP solution, and the second from the definition of J_ω, and the fact that $y_{j\omega}^{LP} \leq d_{j\omega}$. Rearranging this inequality implies that

$$\sum_{j \notin J_\omega} a_{ij} d_{j\omega} < \alpha r_i^{LP}$$

Since $\sum_{j \notin J_\omega} a_{ij} d_{j\omega}$ is an integer, by the definition of the rounding in (2.51) we have

$$\sum_{j \notin J_\omega} a_{ij} d_{j\omega} \leq r_i^I.$$

Since $y^I_{j\omega} = d_{j\omega}$ for $j \in J_\omega$, and $y^I_{j\omega} \geq 0$ always, we have

$$\sum_j a_{ij} d_{j\omega} = \sum_{j \notin J_\omega} a_{ij} d_{j\omega} + \sum_{j \in J_\omega} a_{ij} y^I_{j\omega},$$
$$\leq r^I_i + \sum_j a_{ij} y^I_{j\omega},$$

which establishes feasibility for the inventory constraint (2.1). □

Note that Proposition 2.9 establishes that the LP rounding (2.51-2.52) with $\alpha = 2$ provides an alternative 2-approximation for (ATO). We next establish that the approximation factor can be improved by combining the rounding procedures and corresponding bounds of Propositions 2.8 and 2.9.

Theorem 2.2. *The minimum cost solution from the LP rounding schemes (2.48-2.49), and (2.51-2.52) with $\alpha = 3/2$, provides a 1.8 approximation.*

Proof. To ease notation, let $R = \sum_i c_i r^{LP}_i$ and $Y = \sum_{j,\omega} \mu_\omega p_j y^{LP}_{j\omega}$. We first prove a bound for any α, which we then minimize to obtain the α with the best approximation guarantee.

For a given α, comparing the minimum cost solution between rounding according to (2.48-2.49) and (2.51-2.52), with the cost of the LP solution, gives the following ratio

$$\frac{\min\left(2R + Y, \alpha R + \frac{\alpha}{\alpha-1} Y\right)}{R + Y}.$$

We may exclude from consideration any $\alpha \geq 2$, because in that case the minimizer is trivially $2R+Y$, giving no improvement over the bound of Proposition 2.8. So, for $\alpha < 2$, the approximation factor of the algorithm is the maximum of the above ratio over all possible R and Y, which is equivalent to

$$\max_{R+Y=1} \left\{ \min \left(2R + Y, \alpha R + \frac{\alpha}{\alpha - 1} Y \right) \right\}.$$

Note that the function inside the maximization is a piecewise linear concave function, whose maximum is achieved when $2R+Y = \alpha R + \frac{\alpha}{\alpha-1} Y$. Substituting into the objective, with $R+Y = 1$, this gives an approximation factor of

$$\frac{3\alpha - \alpha^2}{3\alpha - \alpha^2 - 1},$$

which, using elementary calculus, is minimized at $\alpha = 3/2$, with a value 1.8. □

We note an important feature of the rounding algorithm in Theorem 2.2 is that the minimum cost solution between the rounding schemes (2.48-2.49), and (2.51-2.52) (with any α) maintains the approximation guarantee of Theorem 2.1. In light of this, we can view Theorems 2.1 and 2.2 as approximation guarantees for a single LP rounding algorithm (i.e., the minimum cost solution between (2.48-2.49), and (2.51-2.52)). Thus, our LP rounding algorithm is provably robust, with performance guarantees in both large and small demand regimes for general assembly structures.

We note that the factor $\alpha = 3/2$ in Theorem 2.2 is chosen to minimize the worst case approximation guarantee of the algorithm. For a given problem instance, however, a different scaling factor may produce better performance. Thus, in our simulations in Section 2.5, we implement a heuristic search for a local minimum value of α. Our rounding scheme (2.51-2.52) provides an intuitive understanding of the trade-off between first and second stage costs, which can be used to balance these costs appropriately.

Further, we highlight that choosing the minimum cost solution between the two rounding schemes can be interpreted as a comparison between the ordering and shortage costs of the LP solution. In particular, for a given α, we choose the rounding (2.48-2.49) when $2R + Y \leq \alpha R + \frac{\alpha}{\alpha-1} Y$, which is equivalent to

$$(2 - \alpha)(\alpha - 1) \leq \frac{Y}{R}. \tag{2.53}$$

Thus, we can view (2.53) as a threshold policy, where we choose which rounding to use based on comparing the ratio Y/R to the threshold $(2 - \alpha)(\alpha - 1)$. If the LP ordering cost is relatively small compared to the shortage cost, we round inventory quantities down, otherwise we scale them up before rounding. The intuition of this approach is that, when ordering cost is small relative to shortage cost, decreasing the ordering cost and increasing the shortage cost will not increase overall cost too much. Conversely, when ordering cost is relatively large, increasing ordering cost slightly allows us to control any subsequent increases in shortage cost caused by rounding.

We close this subsection by noting that the rounding result of Theorem 2.2 appears reminiscent of the result in Levi et al. (2008) for the OWMR problem, as they also achieve a 1.8 approximation factor by taking the minimum cost solution of two rounding schemes. However, the rounding schemes used are quite different, as well as the problem itself, since Levi et al. (2008) use a random shift procedure to decide the timing of joint ordering decisions with fixed costs, while we use deterministic rounding to decide inventory and allocation levels.

2.4.3 Stochastic Subgradient Method

In Sections 2.4.1 and 2.4.2 we show how to round an optimal solution of the LP relaxation of (ATO) to obtain a feasible integral solution and associated bound on cost. This LP rounding approach assumes that the LP solution is readily available, which is the case if the number of demand scenarios in Ω is not too large. However, describing demand uncertainty could result in exponential growth of the LP's size (e.g. if demand is independent across products), and off-the-shelf LP solvers may not scale well.

Nonetheless, the stochastic programming literature offers various *stochastic approximations* (i.e., gradient methods) to overcome this difficulty (see e.g., Shapiro et al., 2009). While much of the theory applies to strongly convex functions (e.g., Nemirovski et al., 2009), the LP relaxation of (ATO) is neither strongly convex nor differentiable, motivating our use of the subgradient method in Shapiro et al. (2009). We show that with high probability the algorithm provides a near optimal solution to the LP relaxation of (ATO) in polynomial time independent of the number of demand scenarios.

The algorithm treats the objective value as a function of the first stage inventory decisions and uses the second stage dual LP to calculate subgradients. Typically, this approach would give little advantage over solving the original LP directly since the second stage LP still requires variables and constraints for each demand scenario. The key insight is that solving the second stage LP for a random sequence of demand scenarios maintains convergence in expectation (and therefore in probability). This offers potentially large computational savings as the LP solved in each iteration has only one demand scenario instead of (possibly exponentially) many. In this way, the algorithm provides a simple solution procedure: draw a sample, solve a small LP, update inventory, repeat.

For purposes of the algorithm, we think of the objective of (ATO) as a function of the first stage decision variables. To represent this function in a compact form consider the second stage LP for a given demand scenario, ω, and first stage vector $\mathbf{r}$

$$g_\omega(\mathbf{r}) = \left\{ \begin{array}{l} \min_{\mathbf{y}} \ \sum_j p_j y_{j\omega} \\ \\ \text{s.t. } r_i + \sum_j a_{ij} y_{j\omega} \geq \sum_j a_{ij} d_{j\omega}, \ \forall\, i \\ \\ \qquad 0 \leq y_{j\omega} \leq d_{j\omega}, \ \forall\, j \end{array} \right\} \tag{2.54}$$

Then define the function $f_\omega(\mathbf{r}) = \sum_i r_i c_i + g_\omega(\mathbf{r})$, and consider the problem

$$\min_{\mathbf{r}\geq 0} \ \mathbb{E}[f_\omega(\mathbf{r})], \tag{2.55}$$

which is equivalent to solving the LP relaxation of (ATO). A stochastic approximation method for this problem requires calculating a subgradient of $f_\omega(\mathbf{r})$ in successive iterations. To obtain a subgradient, consider the dual of the second stage problem defining $g_\omega(\mathbf{r})$ in (2.54)

$$\max_{\boldsymbol{\theta},\boldsymbol{\pi}} \sum_i \theta_{i\omega} \left(\sum_j a_{ij} d_{j\omega} - r_i \right) - \sum_j \pi_{j\omega} d_{j\omega} \tag{SDATO}$$

$$\text{s.t.} \sum_i a_{ij}\theta_{i\omega} - \pi_{j\omega} \leq p_j, \ \forall \ j \tag{2.56}$$

$$\boldsymbol{\theta}, \boldsymbol{\pi} \geq 0. \tag{2.57}$$

By strong duality $g_\omega(\mathbf{r})$ equals the optimal cost of (SDATO), allowing use of the optimal dual solution as a subgradient. In particular, for given $\mathbf{r}$ and demand scenario ω, let $\theta_{i\omega}(\mathbf{r})$ and $\pi_{j\omega}(\mathbf{r})$ denote an optimal solution to (SDATO). Further, letting $\boldsymbol{\theta}_\omega(\mathbf{r})$ represent the vector in $\mathbb{R}^M$ with $\theta_{i\omega}(\mathbf{r})$ in element i, a standard duality argument establishes that $\mathbf{c} - \boldsymbol{\theta}_\omega(\mathbf{r})$ is a subgradient of $f_\omega(\mathbf{r})$, which we prove formally in the following Lemma.

Lemma 2.8. *For any first stage vector $\mathbf{r}'$, the vector $\mathbf{c} - \boldsymbol{\theta}_\omega(\mathbf{r}')$ is a subgradient of f_ω at $\mathbf{r}'$.*

Proof. Consider another first stage vector $\mathbf{r}$. Since the constraints in the dual problem do not depend on $\mathbf{r}'$, $\theta_{i\omega}(\mathbf{r}')$, $\pi_{j\omega}(\mathbf{r}')$ is feasible to the dual problem for $\mathbf{r}$ as well, and hence weak duality gives

$$f_\omega(\mathbf{r}) \geq \sum_i c_i r_i + \sum_i \theta_{i\omega}(\mathbf{r}') \left(\sum_j a_{ij} d_{j\omega} - r_i \right) - \sum_j \pi_{j\omega}(\mathbf{r}') d_{j\omega}.$$

Therefore subtracting $f_\omega(\mathbf{r}')$ gives $f_\omega(\mathbf{r}) - f_\omega(\mathbf{r}') \geq \sum_i \left(c_i - \theta_{i\omega}(\mathbf{r}')\right)(r_i - r_i')$, satisfying the definition of a subgradient. □

To run the algorithm, we choose a number of iterations, T, and constant step size, h. In each iteration, t, we draw a random demand scenario $\psi_t \in \Omega$ (according to the probabilities μ_ω), solve the second stage dual problem, (SDATO), and update the first stage inventory vector in the direction of the negative subgradient. To obtain a bound on the initial distance from the optimal solution, we define the set

$$\mathcal{P}_{ATO} = \{\mathbf{r} \in \mathbb{R}^M \big|\ 0 \le r_i \le \sum_j p_j \bar{d}_j / c_i, 1 \le i \le M\}.$$

In the following Lemma we show that an LP optimal inventory vector, $\mathbf{r}^*$, lies in $\mathcal{P}_{ATO}$.

Lemma 2.9. *There exists an optimal solution to the LP relaxation of* (ATO) *in the convex set* $\mathcal{P}_{ATO}$*. Further, we have*

$$\|\mathbf{r} - \mathbf{r}^0\|^2 \le \frac{1}{4}\left(\sum_j p_j \bar{d}_j\right)^2 \sum_i \frac{1}{c_i^2},\ \forall\ \mathbf{r} \in \mathcal{P}_{ATO}.$$

Proof. Let r_i^* and $y_{j\omega}^*$ denote an optimal LP solution. Consider the feasible solution that sets $r_i = 0, \forall i$ and $y_{j\omega} = d_{j\omega}, \forall j, \omega$, which has cost $\sum_j p_j \bar{d}_j$. The optimal solution must have cost less than this feasible cost, so, for any i', we have

$$\begin{aligned} c_{i'} r_{i'}^* &\le \sum_i c_i r_i^* + \sum_{j,\omega} \mu_\omega p_j y_{j\omega}^*, \\ &\le \sum_j p_j \bar{d}_j, \end{aligned}$$

which establishes that $\mathbf{r}^* \in \mathcal{P}_{ATO}$. For the second claim, since $\mathbf{r} \in \mathcal{P}_{ATO}$, we have $0 \le r_i \le \sum_j p_j \bar{d}_j / c_i$, so that $(r_i - r_i^0)^2 \le \frac{1}{4}\left(\sum_j p_j \bar{d}_j / c_i\right)^2$. Summing over the components gives the bound. □

Thus, without loss of generality we add the constraint $\mathbf{r} \in \mathcal{P}_{ATO}$ to the LP relaxation of (ATO). Let $P(\mathbf{r})$ denote the Euclidean projection of a vector $\mathbf{r} \in \mathbb{R}^M$ onto $\mathcal{P}_{ATO}$, noting that the box structure reduces this projection to a simple element-wise max and min procedure. With this notation we present Algorithm 1.

To analyze convergence, we bound the initial distance from optimal, as well as the norm of the subgradients. For the former, consider the following initial inventory vector

$$\mathbf{r}^0 = (r_1^0, \ldots, r_M^0),\ \text{s.t.}\ r_i^0 = \frac{1}{2}\sum_j p_j \bar{d}_j / c_i,\ 1 \le i \le M,$$

Algorithm 1 Stochastic Subgradient

Initialize $\mathbf{r}^0$, number of iterations T, step size h
for $t = 0, \dots, T-1$ **do**
  Draw random scenario ψ_t
  Solve (SDATO) for $\mathbf{r}^t$ and ψ_t to obtain $\boldsymbol{\theta}_{\psi_t}(\mathbf{r}^t)$
  Update $\mathbf{r}^{t+1} \leftarrow P\left(\mathbf{r}^t - h\left[\mathbf{c} - \boldsymbol{\theta}_{\psi_t}(\mathbf{r}^t)\right]\right)$
end for
$\bar{\mathbf{r}} \leftarrow \frac{1}{T}\sum_{t=0}^{T-1} \mathbf{r}^t$.

and define the following quantity

$$R = \frac{1}{4}\left(\sum_j p_j \bar{d}_j\right)^2 \sum_i \frac{1}{c_i^2}.$$

Since $\mathbf{r}^* \in \mathcal{P}_{ATO}$, it is straightforward to show that $\|\mathbf{r}^* - \mathbf{r}^0\|^2 \leq R$. Finally, we define the quantity

$$G = \sum_i \hat{c}_i^2,$$

where $\hat{c}_i = \max\left\{c_i, \max_{j\in\mathcal{N}(i)}\{p_j\} - c_i\right\}$. In the following Lemma, we show that the norm of the subgradient is bound by G.

Lemma 2.10. *For any demand scenario ω and any $\mathbf{r} \in \mathcal{P}_{ATO}$, we have*

$$\|\mathbf{c} - \boldsymbol{\theta}_\omega(\mathbf{r})\|^2 \leq \sum_i \hat{c}_i^2$$

Proof. Let $p_{i\max} = \max_{j\in\mathcal{N}(i)}\{p_j\}$, so $\hat{c}_i = \max\{c_i, p_{i\max} - c_i\}$. First we claim that $\theta_{i\omega}(\mathbf{r}) \leq p_{i\max}$. Assume to the contrary that there is an i such that $\theta_{i\omega}(\mathbf{r}) > p_{i\max}$. Then for all $j \in \mathcal{N}(i)$, we have $\sum_{i'} a_{i'j}\theta_{i'\omega}(\mathbf{r}) - p_j > 0$, so that the optimal π variable must be $\pi_{j\omega}(\mathbf{r}) = \sum_{i'} a_{i'j}\theta_{i'\omega}(\mathbf{r}) - p_j$. Therefore, the variables $\theta_{i\omega}(\mathbf{r})$ and $\pi_{j\omega}(\mathbf{r})$ for $j \in \mathcal{N}(i)$ contribute the following amount to the dual objective value

$$\theta_{i\omega}(\mathbf{r})\sum_j a_{ij}d_{j\omega} - \sum_{j\in\mathcal{N}(i)}\left(a_{ij}\theta_{i\omega}(\mathbf{r}) + \sum_{i'\neq i} a_{i'j}\theta_{i'\omega}(\mathbf{r}) - p_j\right)d_{j\omega} - \theta_{i\omega}(\mathbf{r})r_i,$$

$$= \sum_{j\in\mathcal{N}(i)}\left(p_j - \sum_{i'\neq i} a_{i'j}\theta_{i'\omega}(\mathbf{r})\right)d_{j\omega} - \theta_{i\omega}(\mathbf{r})r_i.$$

Consider instead setting $\theta_{i\omega} = p_{i\max}$. Then we have $\theta_{i\omega} + \sum_{i'\neq i} a_{i'j}\theta_{i'\omega}(\mathbf{r}) \geq p_j$ for all $j \in \mathcal{N}(i)$ so it is feasible to set $\pi_{j\omega} = a_{ij}\theta_{i\omega} + \sum_{i'\neq i} a_{i'j}\theta_{i'\omega}(\mathbf{r}) - p_j$ for $j \in \mathcal{N}(i)$. Now the variables $\theta_{i\omega}$ and $\pi_{j\omega}$

for $j \in \mathcal{N}(i)$ will contribute the following amount to the objective value

$$\theta_{i\omega} \sum_{j} a_{ij} d_{j\omega} - \sum_{j \in \mathcal{N}(i)} \left(a_{ij}\theta_{i\omega} + \sum_{i' \neq i} a_{i'j}\theta_{i'\omega}(\mathbf{r}) - p_j \right) d_{j\omega} - \theta_{i\omega} r_i,$$

$$= \sum_{j \in \mathcal{N}(i)} \left(p_j - \sum_{i' \neq i} a_{i'j}\theta_{i'\omega}(\mathbf{r}) \right) d_{j\omega} - p_{i\,\max} r_i.$$

Since we assumed $\theta_{i\omega}(\mathbf{r}) > p_{i\,\max}$, this is an increase in the objective value, and so $\theta_{i\omega}(\mathbf{r})$ was not optimal. Therefore, we have $0 \leq \theta_{i\omega}(\mathbf{r}) \leq p_{i\,\max}$. Thus, if $\theta_{i\omega}(\mathbf{r}) \leq c_i$, we have $|c_i - \theta_{i\omega}(\mathbf{r})| \leq c_i$, and if $\theta_{i\omega}(\mathbf{r}) > c_i$ we have $|c_i - \theta_{i\omega}(\mathbf{r})| \leq p_{i\,\max} - c_i$, so that $|c_i - \theta_{i\omega}(\mathbf{r})| \leq \max\{c_i, p_{i\,\max} - c_i\} = \hat{c}_i$. Thus $(c_i - \theta_{i\omega}(\mathbf{r}))^2 \leq \hat{c}_i^2$ and summing over i gives the claim. □

With these bounds we can state the convergence result for Algorithm 1. Note that the algorithm's sampling procedure generates a random solution $\bar{\mathbf{r}}$, and let $STO(\bar{\mathbf{r}})$ represent its cost, i.e., $STO(\bar{\mathbf{r}}) = \mathbb{E}[f_\omega(\bar{\mathbf{r}})]$. Recall that OPT^{LP} denotes the optimal cost of the LP relaxation of (ATO). Then we have the following result

Proposition 2.10. *For any $\delta, \epsilon > 0$, let the number of iterations and step size be*

$$T = \left\lceil \frac{RG}{\delta^2 \epsilon^2} \right\rceil, \quad h = \sqrt{\frac{R}{GT}}.$$

Then Algorithm 1 produces $\bar{\mathbf{r}}$ satisfying

$$\mathbb{P}\left[STO(\bar{\mathbf{r}}) - OPT^{LP} \geq \delta \right] \leq \epsilon$$

Proof. Given the bound on the diameter of $\mathcal{P}_{ATO}$ from Lemma 2.9 and the bound on the norm of the subgradients from Lemma 2.10, Proposition 2.10 follows immediately from the analysis of Shapiro et al. (2009) (see pages 235-236). □

Therefore we can obtain a close to optimal fractional solution with high probability in a polynomial number of iterations that is independent of the number of demand scenarios. Each iteration requires solving an LP that has $M + N$ variables and N constraints, and so the entire procedure can be carried out in polynomial time that is independent of the number of demand scenarios.

We note that while there exist other algorithms for solving the stochastic program, the focus of this chapter is not on determining the best such algorithm. We present Algorithm 1 primarily to demonstrate that efficient solutions are available for the LP relaxation of (ATO) that can be used as input for the rounding schemes in Sections 2.4.1 and 2.4.2. Furthermore, we find the algorithm attractive for a simple managerial interpretation it provides, which we describe next.

Consider a manager who is contemplating an order for a particular inventory vector $\mathbf{r}^0$. She wants to know how $\mathbf{r}^0$ might perform and so randomly draws a demand scenario from Ω. She solves the second stage allocation under this demand scenario and makes note of which products have a shortage and which components have a surplus. Then she adjusts upward the inventory of components serving the short products and adjusts downward the inventory of components with a surplus. Then she draws another sample, solves the allocation again, and makes another adjustment. After repeating this process several times, she averages the inventory vectors, figuring that should account for the variety of scenarios she has seen.

The procedure described in the previous paragraph roughly corresponds to Algorithm 1 via an interpretation of the dual variables. From the update in Step 5 of Algorithm 1, we see that, roughly speaking, r_i increases if $\theta_{i\omega}$ is large and decreases if $\theta_{i\omega}$ is small. Note that $\theta_{i\omega}$ is the dual variable for the primal inventory constraint $r_i + \sum_j a_{ij} y_{j\omega} \geq \sum_j a_{ij} d_{j\omega}$. Thus, if this constraint is slack, i.e. there is a surplus of component i, then by complementary slackness $\theta_{i\omega} = 0$ and the algorithm will decrease the inventory quantity. Further, when component i is causing large shortages for the products it serves, the dual solution will require a large value for the $\theta_{i\omega}$ variable. Then the algorithm will increase the inventory quantity r_i. Therefore, the intuitive procedure described above can lead to effective inventory quantities when implemented according to the details of Algorithm 1 and Proposition 2.10.

2.5 Simulation Results

In this section we compare algorithms for (ATO) using numerical simulations.

2.5.1 Setup

We consider a total of 200 problem instances, i.e., all combinations of 5 assembly structures, 10 demand distributions, and 4 sets of cost parameters, which we describe next.

Assembly Structures. We consider five assembly structures:

- M-system of Figure 2.1 (2 components, 3 products)

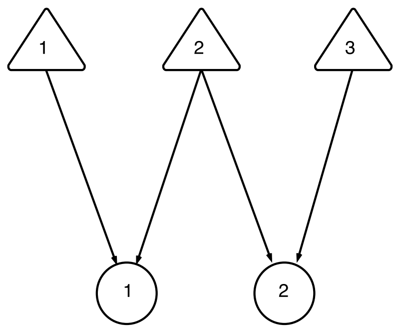

$$\begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 \\ 1 & 1 & 1 & 1 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 \end{bmatrix}$$

Figure 2.3: The W System and LS System

$$\begin{bmatrix} 1 & 0 & 1 & 0 & 1 & 0 & 1 \\ 0 & 1 & 1 & 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 1 & 0 & 1 & 1 & 0 & 1 & 0 \\ 0 & 1 & 1 & 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 1 & 0 & 0 & 1 \end{bmatrix}$$

Figure 2.4: The SC System

- W system on the left side of Figure 2.3 (3 components, 2 products)
- "KL" structure of Figure 2.2 with K=3 and L=5, (8 components, 5 products)
- "LS" system, whose bipartite graph incidence matrix (components as rows, products as columns) is on the right side of Figure 2.3 (example from Lu and Song, 2005) (14 components, 6 products)
- "SC" system in Figure 2.4, which corresponds to an instance of set cover (i.e., components are elements, products are sets) from a family of examples known to have a large integrality gap (Vazirani, 2013) (7 components, 7 products)

We choose these examples to give our algorithms a range of inputs, including structures that are simple (M, W), somewhat realistic (LS), and theoretically challenging (KL, SC).

Demand. For each problem instance, we construct the set of demand scenarios Ω by drawing 10,000 random samples from a specified distribution, then assign each scenario in Ω equal probability. For ease of implementation, we specify several continuous distributions, which are rounded to integers. Individual product demand is capped at a maximum of 20, and floored at 0. We consider the following distributions:

- Multivariate normal distributions; with various combinations of means either 10 or 13, variances either 10 or 20, and zero, positive, and negative correlation (6 distributions).
- Independent exponential distributions, with means 10 or 13 (2 distributions).
- I.I.D. discrete uniform on the integers $\{0, \ldots, 20\}$ (1 distribution).
- I.I.D. Bernoulli random variables with success probability 1/2 (1 distribution).

Costs. We consider 4 sets of cost parameters. For each set, we have $c_i = 1$ for all i. To maintain consistency across systems, we model the shortage costs through the markup parameters γ_j as follows:

- Identical markups of 1.1.
- Identical markups of 3.
- Markups 2 and 1.5 for different products.
- Markups 3 and 1.1 for different products.

2.5.2 Heuristics and Computation

For each problem instance we run 5 heuristics. The first three are newsvendor heuristics: i) constant markup newsvendor (CM); ii) weighted cost newsvendor (WC), with weights determined by a product's contribution to average demand for a component; and iii) full cost newsvendor (FC), which uses the full shortage cost, p_j, of product j for each newsvendor subproblem (as in van Jaarsveld and Scheller-Wolf, 2015). The next two heuristics are LP rounding schemes that operate on the solution produced by the stochastic subgradient method of Section 2.4.3. These are iv) myopic rounding (MY), which rounds inventory quantities to the nearest integer and fills demand in each scenario on a first-come-first-served (FCFS) basis according to a random permutation of the products (i.e., an imitation of the dynamic policies of Lu et al., 2005 and van Jaarsveld and Scheller-Wolf, 2015 in the one-period problem, included to have a point of comparison with existing techniques); and

v) our rounding scheme of Theorem 2.2 (RD), with a heuristic search to find a good value for the scaling parameter α.

We also directly solve the LP relaxation of (ATO) (including all 10,000 demand scenarios) to obtain a lower bound on optimal cost, and a benchmark running time for comparison with the stochastic subgradient method. All LPs were solved with Gurobi version 7.0.1 using concurrent simplex and barrier methods. All simulations were implemented in Python 3.5 on an Intel i7-6700 CPU at 3.4 GHz and 16 GB of RAM.

Algorithm 2 Heuristic Stopping Criterion

On first run, initialize $\bar{\mathbf{r}}^{-1} = \mathbf{r}^0$, epoch length s, tolerance level ϵ_0
Receive input: period t and $\bar{\mathbf{r}}^t = \frac{1}{t}\sum_{t'=0}^{t-1} \mathbf{r}^{t'}$
if $t \bmod s = 0$ **then**
    if $\frac{T-t}{s}\|\bar{\mathbf{r}}^t - \bar{\mathbf{r}}^{-1}\|_\infty \le \epsilon_0$ **then**
        Set $T = t$ and exit outer "for" loop of Algorithm 1
    end if
    Set $\bar{\mathbf{r}}^{-1} = \bar{\mathbf{r}}^t$
end if

The order quantities for the CM and WC newsvendor heuristics are computed by iterating through a sorted list of the realizations of D_i until the unique value satisfying the definition in equation (2.16) (or (2.46)) is found. We estimate the second stage shortage cost for the resulting inventory vectors by solving the second stage LP (2.4) for each demand scenario, and round up the solution. The FC newsvendor heuristic has an optimal quantile definition similar to (2.46) that incorporates different costs for each product, and is efficiently computed using binary search.

The MY and RD heuristics use the stochastic subgradient procedure to solve the LP, following Algorithm 1 of Section 2.4.3. While the running time bound of Proposition 2.10 is theoretically important, it is well known that such subgradient algorithms can exhibit slow convergence in practice. Therefore, we implement a heuristic stopping criteria, Algorithm 2, at the end of each iteration of the "for" loop in Algorithm 1. The basic idea is to monitor the change in the average solution across time epochs. Specifically, let $\bar{\mathbf{r}}^t = \frac{1}{t}\sum_{t'=0}^{t-1} \mathbf{r}^{t'}$ be the average solution up to period t and choose an epoch length $s < T$. At iterations $t_k = ks$ for $k = 1, 2, \ldots$, check the change in the average solution, projected over the remaining iterations, i.e., consider $\frac{T-t_k}{s}(\bar{\mathbf{r}}^{t_k} - \bar{\mathbf{r}}^{t_{k-1}})$. If this projected change is below some predefined threshold ϵ_0 in a chosen norm, terminate the algorithm. This approach performs well numerically, since the rounding scheme erases small deviations in the output $\bar{\mathbf{r}}$. We chose a tolerance level of $\epsilon_0 = 1$, and an epoch length of $s = 100$.

For the tolerance level of the stochastic subgradient convergence result in Proposition 2.10 we chose $\epsilon = .5$. To choose a reasonable level for δ, we leverage the lower bound on cost calculated by the constant markup newsvendor heuristic (per Lemma 2.1). We set δ equal to 20% of this lower bound, meaning the tolerance level is calibrated to at most 20% of the optimal cost. We initialize the algorithm at the CM heuristic solution. Since we use the CM heuristic to implement the RD heuristic, we include the CM run time in the RD run time.

Once the stochastic subgradient algorithm produces a first stage inventory vector, we solve the resulting second stage LP (2.4) for each demand scenario to obtain the LP shortage quantities (this computation time is also included in the RD run time). We round the resulting LP solution using both (2.48-2.49) and (2.51-2.52), and use the minimum cost solution per Theorem 2.2. For the rounding scheme (2.51-2.52) we use a heuristic golden section search (Press et al., 2002) to find a local minimizer for the parameter α between 1 and 2. The computation time for this search is also included in the RD run time.

For the MY heuristic, we round the first stage inventory vector produced by the stochastic subgradient algorithm to the nearest integer. Then, for each scenario we consider the product demands arriving in a random permutation, and greedily fill demand in this order until components run out (to implement a FCFS allocation rule).

2.5.3 Results

The overall optimality gaps and run times for the simulations are summarized in Figure 2.5, with details for each system given in Tables 2.1 and 2.2. The optimality gaps are given in terms of percentage above optimal (i.e., $COST/OPT - 1$). Note that we report upper bounds on the optimality gap as we compare heuristic performance to the solution of the LP relaxation of (ATO).

The RD heuristic performs very well overall, with an average optimality gap of 1.2% (compared to 10-13% for the newsvendor heuristics, and 5.9% for the other rounding heurisitic, MY), and worst-case gap of 7.4% across all 200 instances (compared to 62-153% for newsvendor and 40% for MY). We note that across systems, the worst rounding performance occurs with the Bernoulli demand distribution, whose small mean (and variability) make the guarantee of Theorem 2.1 less favorable. The worst case gap for RD among the other 9 demand distributions is around 4%.

In terms of run time, the newsvendor heuristics are very fast, with averages around 1 second.

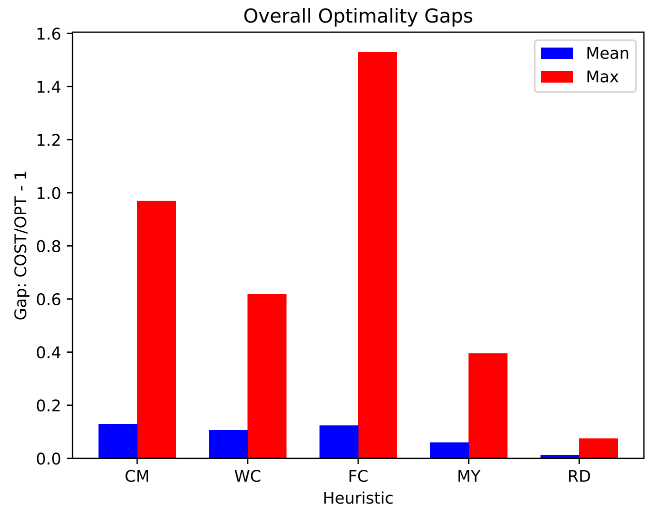

Figure 2.5: Relative optimality gap by heuristic

Table 2.1: Average Gaps and Run Times by Heuristic

Avg. Gap & Time			Gap %					Time (Seconds)				
Syst.	Comp.	Prod.	CM	WC	FC	MY	RD	CM	WC	FC	MY	RD
M	2	3	3.6	5.5	1.2	11.8	0.5	0.7	0.7	0.8	2.5	1.4
W	3	2	11.7	9.1	10.7	0.9	0.8	0.7	0.7	0.8	2.4	1.5
LS	14	6	17.7	16.2	20.8	4.9	1.7	1.2	1.1	2.1	8.0	6.2
KL	8	5	27.2	11.8	19.8	2.0	1.2	0.9	0.9	1.4	4.0	2.6
SC	7	7	4.2	10.2	9.1	10.0	1.6	1.1	0.9	1.5	4.5	2.8
		Average	12.9	10.6	12.3	5.9	1.2	0.9	0.8	1.3	4.3	2.9

The RD algorithm takes longer since it needs to solve the LP, although its run time using the stochastic subgradient method is about half that of solving the full LP on average. We also note that the stochastic subgradient algorithm can require significantly less memory than solving the full LP, since it maintains only a few first stage vectors of length M and a single second stage vector of length N.

In summary, these simulations suggest that while the newsvendor heuristics can provide a reasonable level of accuracy in a very short run time, our rounding scheme provides a significant improvement in accuracy with a modest increase in run time. Further, in addition to providing the first approximation guarantees, our rounding scheme improves upon the numerical performance of heuristic rounding approaches considered in the literature. Finally, these simulations serve to further illustrate a main point of the chapter, that integrality constraints in (ATO) need not hinder our ability to find good solutions. In particular, we show that the rounding scheme we develop

Table 2.2: Maximum Gaps and Run Times by Heuristic

Max. Gap & Time			Gap %					Time (Seconds)				
Syst.	Comp.	Prod.	CM	WC	FC	MY	RD	CM	WC	FC	MY	RD
M	2	3	59.3	16.3	16.3	39.6	1.8	0.7	0.8	1.0	2.7	1.6
W	3	2	44.9	61.9	74.2	3.8	2.7	0.7	0.7	1.1	2.7	1.6
LS	14	6	66.8	50.5	104	14.8	7.4	1.4	1.2	4.2	16	14
KL	8	5	97.0	36.9	153	9.3	4.4	1.1	1.0	2.3	5.6	4.2
SC	7	7	31.4	48.6	48.6	25.8	6.3	1.2	1.0	2.8	5.9	4.2
		Maximum	97.0	61.9	153	39.6	7.4	1.4	1.2	4.2	16	14

performs very well.

2.6 Computational Complexity

In the model introduction of Section 2.2, we observe a connection between the NP-hard set cover problem and the second stage allocation problem of (ATO). In this section, we build upon this observation to formally demonstrate that solving (ATO) is NP-hard.

2.6.1 A Class of Set Cover Problems

In order to establish the computational complexity of (ATO), we first consider a class of set cover instances (which we show are NP-hard) that will aid the reductions for ATO problems. First, we recall the classic definition of the set cover problem (Vazirani, 2013). For a ground set U, with typical element $i \in U$, let $\mathcal{S}$ denote a family of subsets of U, i.e., $S_j \subseteq U$ for all $S_j \in \mathcal{S}$. The set cover problem is, given U and $\mathcal{S}$, find the smallest number of subsets whose union covers U. Let J denote the set of indices of the sets $S_j \in \mathcal{S}$, so that $|J| = |\mathcal{S}|$.

We want to consider families of subsets with a partially overlapping structure, and follow Lubell (1966) in making the following definition of Sperner families.

Definition 2.3 (Sperner Family)**.** *A family of subsets $\mathcal{S}$ of a ground set U is called Sperner if no set $S_j \in \mathcal{S}$ contains another $S_{j'} \in \mathcal{S}$, $j \neq j'$.*

We make the following definition to describe the set of sets that an element i of U is contained in.

Definition 2.4 (Inverse Family)**.** *For a family of subsets $\mathcal{S}$ of a ground set U, its inverse family is the family $\mathcal{U}$ of subsets of J representing the sets in $\mathcal{S}$ that contain each element of U, i.e., for each*

$i \in U$, $T_i \in \mathcal{U}$ is defined as

$$T_i = \{j \in J | i \in S_j\}.$$

The families of subsets we are interested in are those whose inverse family is Sperner, which we define as follows.

Definition 2.5 (Inverse Sperner Family)**.** *A family of subsets $\mathcal{S}$ of a ground set U is called inverse Sperner if its inverse family is Sperner.*

This definition will be essential to our proof of NP-hardness for (ATO). In (ATO), the components correspond to the elements of the set cover instance, and an inverse Sperner family gives us the property that, for any two components $i \neq i'$, we have

$$|\mathcal{N}(i) \cap \mathcal{N}(i')| \leq |\mathcal{N}(i)| - 1.$$

This property will allow us to construct a set of demand scenarios where each component i has demand $|\mathcal{N}(i)| - 1$ in at least one scenario, and never has more demand than this in any scenario. This in turn provides a means to encode a set cover instance with the second stage allocation problem of (ATO). First we need to demonstrate that any set cover instance can be reduced to set cover for an inverse Sperner family. To that end, the *inverse Sperner set cover* problem is to solve set cover for a ground set U and a family of subsets $\mathcal{S}$ which is inverse Sperner.

Lemma 2.11. *Given any set cover instance U and $\mathcal{S}$, we can construct in polynomial time an instance of inverse Sperner set cover U' and $\mathcal{S}'$, with the following properties: i) $|U'| \leq |U|$, and ii) a collection of sets from $\mathcal{S}'$ is a cover for U' if and only if the corresponding collection of sets from $\mathcal{S}$ is a cover for U. Specifically, let $K \subseteq J$ denote a subset of set indices and we claim that*

$$\bigcup_{j \in K} S'_j = U', \text{ if and only if } \bigcup_{j \in K} S_j = U.$$

Proof. Given a set cover instance, U and $\mathcal{S}$, we construct the inverse Sperner set cover instance U' and $\mathcal{S}'$ as follows. Begin with $U' = U$. For each $i \in U'$, if there exists another $i' \in U'$ such that $T_{i'} \subseteq T_i$, then remove i from U' and call i' the *removing witness* of i. If i was already a removing witness for another element i'', then update i' to be the removing witness of i'' as well. After constructing the set U', update $\mathcal{S}'$ to be the collection of sets $S'_j = U' \cap S_j$. Then, by construction, $\mathcal{S}'$ is an inverse Sperner family of subsets of U'. Constructing the sets T_i takes $O(|U||\mathcal{S}|)$ time, while the construction of U' and $\mathcal{S}'$ takes $O(|U|^2|\mathcal{S}|)$ time. To construct U', we only removed elements from U, and therefore $|U'| \leq |U|$.

Let $K \subseteq J$ denote the indices a cover in $\mathcal{S}'$ for U', i.e., $\cup_{j \in K} S'_j = U'$. Note that since $S'_j \subseteq S_j$ for all j, we have $U' = \cup_{j \in K} S'_j \subseteq \cup_{j \in K} S_j$, so that the elements $i \in U'$ are clearly covered by $\cup_{j \in K} S_j$. Then, consider some $i \in U \setminus U'$ that was removed from U when constructing U' and let i' be its removing witness, so that $i' \in U'$ and $T_{i'} \subseteq T_i$. Letting $j' \in K$ be the index of a set that covers i' in the optimal solution K, we have $j' \in T_{i'} \subseteq T_i$. Therefore $i \in S_{j'}$, so that i is covered by $\cup_{j \in K} S_j$.

Next, let $L \subseteq J$ denote a cover for U, i.e., $\cup_{j \in L} S_j = U$. Then we have

$$\begin{aligned} \bigcup_{j \in L} S'_j &= \bigcup_{j \in L} S_j \cap U', \\ &= U', \end{aligned}$$

where the first line follows from the definition of S'_j and the last from the fact that $\cup_{j \in L} S_j$ covers U and $U' \subseteq U$. □

Lemma 2.12. *Any set cover instance can be reduced in polynomial time to an instance of inverse Sperner set cover.*

Proof. With the correspondence of Lemma 2.11, it is clear that an optimal subset of indices K^* that covers U' must also be an optimal subset of indices that covers U. Thus, having solved set cover for U' and $\mathcal{S}'$, the only computation required to produce an optimal solution for U and $\mathcal{S}$ is mapping the indices of K^* back to the sets in $\mathcal{S}$, which takes time $O(|\mathcal{S}|)$. □

2.6.2 Complexity of (ATO)

With Lemma 2.12 in hand, we now prove our main result showing that (ATO) can be used to solve any instance of inverse Sperner set cover.

Proposition 2.11. *Any inverse Sperner set cover instance can be reduced in polynomial time to an instance of* (ATO).

Proof. Consider a given instance of inverse Sperner set cover, U and $\mathcal{S}$. Then create an instance $\mathcal{I}$ of (ATO) as follows. Let there be $M = |U|$ components and $N = |\mathcal{S}|$ products, and for each $i \in U$, let $a_{ij} = 1$ for each j such that $i \in S_j$, and let $a_{ij} = 0$ otherwise. Choose some $0 < \epsilon < 1$ and let

$$\delta = \frac{\epsilon}{2M^2 + \epsilon},$$

and let the cost of each product be $p_j = 2M/(1-\delta)$, and the cost of each component be $c_i = 1$. Let there be $M+1$ demand scenarios, indexed by $k = 1, \dots, M+1$. The first M scenarios each occur with probability $\mu_k = (1-\delta)/M$, and each have unit demand localized to all but one of the products served by a single component. In particular, for scenario $k \in \{1, \dots, M\}$, consider component $i = k$ and choose a proper subset of the products it serves, $\mathcal{D}(k) \subset \mathcal{N}(i)$, of size $|\mathcal{D}(k)| = |\mathcal{N}(i)| - 1$, and let these products have unit demand, i.e., $d_{jk} = 1$ for $j \in \mathcal{D}(k)$, while all other products have zero demand in this scenario. The final scenario, $k = M+1$, occurs with probability $\mu_k = \delta$ and has unit demand for all products, i.e., $d_{jk} = 1$ for all j. We call scenario $k = M+1$ the *pivotal* scenario. Constructing this instance $\mathcal{I}$ of (ATO) takes $O(|U||\mathcal{S}|)$ time.

Optimal inventory claim. We establish the reduction by demonstrating that the second stage allocation in the pivotal scenario of $\mathcal{I}$ is equivalent to the set cover problem for U and $\mathcal{S}$. As the key step in establishing this equivalence, we claim that the optimal first stage solution for the instance $\mathcal{I}$ of (ATO) is to set, for each component i

$$r_i^* = |N(i)| - 1. \tag{2.58}$$

Encoding set cover. Given optimal inventory quantities (2.58), since $d_{jk} = 1$ for all products in the pivotal demand scenario, the shortage of component i in this scenario is

$$\begin{aligned}\sum_j a_{ij} d_{jk} - r_i^* &= |N(i)| - (|N(i)| - 1), \\ &= 1.\end{aligned}$$

Further, $0 \le y_{jk} \le d_{jk} = 1$ together with the integrality requirement is equivalent to $y_{jk} \in \{0, 1\}$. Finally, in the pivotal scenario $k = M+1$, the objective coefficient for the shortage variable y_{jk} for any j is $\mu_k p_j = \delta 2M/(1-\delta) = \epsilon/M$. Thus, given the optimal inventory quantities in (2.58), finding the optimal shortage cost in the pivotal scenario requires solving the following problem

$$\begin{aligned}\min\ & \frac{\epsilon}{M} \sum_j y_{jk} \\ \text{s.t. } & \sum_j a_{ij} y_{jk} \ge 1, \ \forall\, i, \\ & y_{j\omega} \in \{0, 1\}, \ \forall\, j.\end{aligned} \tag{2.59}$$

Given the definition of a_{ij}, and the uniform scaling of the objective function, this problem is equivalent to the standard integer programming formulation for the set cover instance U and $\mathcal{S}$ (see e.g., Vazirani, 2013). With this equivalence, given an optimal solution for instance $\mathcal{I}$ of (ATO), we can construct an optimal set cover for U and $\mathcal{S}$ by choosing all sets such that $y_{jk} = 1$ in the pivotal demand scenario $k = M + 1$. This construction of the optimal set cover takes time $O(|\mathcal{S}|)$.

Unique optimal solution for an alternative problem. Now, to prove the claim (2.58) for instance $\mathcal{I}$ of (ATO), consider an alternate integer program, denoted (with abuse of notation) by $\mathcal{I}'$, which is equivalent to instance $\mathcal{I}$ except the variables and constraints corresponding to the pivotal demand scenario are removed. Thus, in problem $\mathcal{I}'$ there are M scenarios indexed by $k = 1, \ldots, M$ as described by the construction of $\mathcal{I}$, whose shortage variables have the same objective coefficients as in $\mathcal{I}$. We first show that the inventory quantities in (2.58) constitute the unique optimal solution for problem $\mathcal{I}'$, which we then use to show that they are also optimal for instance $\mathcal{I}$ of (ATO).

To begin, note that the coefficient of the variable y_{jk} is $\mu_k p_j$, which for $k = 1, \ldots, M$ is equal to

$$\frac{1-\delta}{M}\left(\frac{2M}{1-\delta}\right) = 2.$$

Thus, we can write the problem $\mathcal{I}'$ as

$$\min_{\mathbf{r},\mathbf{y}} \sum_i r_i + 2\sum_j \sum_{k \le M} y_{jk} \tag{$\mathcal{I}'$}$$

$$\text{s.t. } r_i + \sum_j a_{ij} y_{jk} \ge \sum_j a_{ij} d_{jk}, \ \forall\, i, k \le M \tag{2.60}$$

$$y_{jk} \le d_{jk}, \ \forall\, j, k \le M \tag{2.61}$$

$$\mathbf{r}, \mathbf{y} \ge 0, \ \text{integer}. \tag{2.62}$$

Now consider a solution to $\mathcal{I}'$ that sets inventory quantities as in (2.58) for all components and sets all shortage quantities to zero, i.e., $y^*_{jk} = 0$ for all products j and scenarios k. We first show that this solution is feasible and optimal for $\mathcal{I}'$, and then demonstrate that it is the only optimal solution. The solution is clearly non-negative, integral, and satisfies the demand constraint (2.61), and so we focus on the inventory constraint (2.60). Consider a component i and scenario k. Let $i' = k$ be the component for which we selected the set $\mathcal{D}(k) \subset \mathcal{N}(i')$ of products to have unit demand in scenario k. Demand for component i in scenario k is

$$\begin{aligned}\sum_j a_{ij} d_{jk} &= |\mathcal{N}(i) \cap \mathcal{D}(k)|, \\ &\leq |\mathcal{N}(i) \cap \mathcal{N}(i')|, \\ &\leq |\mathcal{N}(i)| - 1, \\ &= r_i^*,\end{aligned}$$

where the first line follows from the definition of a_{ij} and d_{jk}, the second from $\mathcal{D}(k) \subset \mathcal{N}(i')$, the third from the inverse Sperner property, and the fourth from setting inventory quantities as in (2.58). Thus, setting shortages to zero is feasible for constraint (2.60). To see that this solution is optimal, consider the dual of the LP relaxation of $(\mathcal{I}')$,

$$\begin{aligned}\max_{\boldsymbol{\theta},\boldsymbol{\pi}} & \sum_{i,k} \theta_{ik} \sum_j a_{ij} d_{jk} - \sum_{j,k} d_{jk} \pi_{jk} \\ \text{s.t.} & \sum_i a_{ij} \theta_{ik} - \pi_{jk} \leq 2, \ \forall \ j, k \leq M \\ & \sum_\omega \theta_{ik} \leq 1, \ \forall \ i \\ & \boldsymbol{\theta}, \boldsymbol{\pi} \geq 0.\end{aligned}$$

Consider setting all $\pi_{jk}^* = 0$ and

$$\theta_{ik}^* = \begin{cases} 1, & \text{if } i = k \\ 0, & \text{otherwise} \end{cases}$$

This solution is feasible to the dual, and satisfies complementary slackness, since for each i we have $\sum_\omega \theta_{ik}^* = 1$, while for $\theta_{ik}^* > 0$ we have

$$\begin{aligned}\sum_j a_{ij} d_{jk} &= |\mathcal{N}(i)| - 1, \\ &= r_i^*,\end{aligned}$$

establishing optimality of the primal solution r_i^* for the LP relaxation and thus also for the original problem $\mathcal{I}'$. To demonstrate uniqueness of this optimal solution for the LP relaxation (and thus also for the original integer program), we use a characterization of Mangasarian (1979), which adds small perturbations to the objective coefficients of the LP. In particular, consider $q_i \in \mathbb{R}$ and $q_{jk} \in \mathbb{R}$

and for $\alpha > 0$ consider the following perturbed version of the LP relaxation of $\mathcal{I}'$

$$\begin{aligned}
\min_{\mathbf{r},\mathbf{y}} \; & \sum_i (1+\alpha q_i) r_i + \sum_j \sum_{k \le M} (2+\alpha q_{jk}) y_{jk} \\
\text{s.t. } & r_i + \sum_j a_{ij} y_{jk} \ge \sum_j a_{ij} d_{jk}, \;\; \forall\, i, k \le M \\
& y_{jk} \le d_{jk}, \;\; \forall\, j, k \le M \\
& \mathbf{r}, \mathbf{y} \ge 0.
\end{aligned} \tag{2.63}$$

Then, Theorem 1 in Mangasarian (1979) says that an optimal solution, r_i^*, y_{jk}^*, for the LP relaxation of $\mathcal{I}'$ is unique if, for any perturbations $q_i \in \mathbb{R}$ and $q_{jk} \in \mathbb{R}$, there exists $\alpha > 0$ such that it remains an optimal solution to (2.63). To see that this is the case, consider the dual of the LP relaxation of (2.63),

$$\begin{aligned}
\max_{\boldsymbol{\theta},\boldsymbol{\pi}} & \sum_{i,k} \theta_{ik} \sum_j a_{ij} d_{jk} - \sum_{j,k} d_{jk} \pi_{jk} \\
\text{s.t. } & \sum_i a_{ij} \theta_{ik} - \pi_{jk} \le 2 + \alpha q_{jk}, \;\; \forall\, j, k \le M \\
& \sum_\omega \theta_{ik} \le 1 + \alpha q_i, \;\; \forall\, i \\
& \boldsymbol{\theta}, \boldsymbol{\pi} \ge 0,
\end{aligned}$$

and set all $\pi'_{jk} = 0$ and

$$\theta'_{ik} = \begin{cases} 1 + \alpha q_i, & \text{if } i = k \\ 0, & \text{otherwise} \end{cases}$$

and let

$$\begin{aligned}
\beta &= \max\left(\max_i\{|q_i|\}, \max_{j,k}\{|q_{jk}|\}, \max_{i,j,k}\{|q_i - q_{jk}|\} \right), \\
\alpha &= \frac{1}{\beta},
\end{aligned}$$

from which it is straightforward to observe that $1 + \alpha q_i \ge 0$ for all i, $2 + \alpha q_{jk} \ge 0$ for all j, k, and $1 + \alpha q_i \le 2 + \alpha q_{jk}$ for all i, j, k. These properties directly imply feasibility of the dual solution θ'_{ik}, π'_{jk}. Finally the primal and dual solutions satisfy complementary slackness, since for each i we have $\sum_\omega \theta^*_{ik} = 1 + \alpha q_i$, while for $\theta^*_{ik} > 0$ we have $\sum_j a_{ij} d_{jk} = r_i^*$. Thus, r_i^*, y_{jk}^* is the unique optimal solution for $\mathcal{I}'$ and has objective value $\sum_i r_i^*$.

Proof of optimal inventory claim. Now we are ready to show that the optimal solution to instance $\mathcal{I}$ of (ATO) must have first stage optimal solution r_i^* for all i. To do this, we first construct an upper bound on the cost of the optimal solution with inventory quantities r_i^*. As in the optimal solution of problem $\mathcal{I}'$, for scenarios $k = 1, \ldots, M$, we set $y_{jk} = 0$ for all products j, incurring no shortage cost. In the pivotal scenario $k = M + 1$, since there are M constraints in the second stage allocation problem (2.59), each with right hand side equal to one, we need to set at most M shortage variables to one in order to cover these constraints. Thus, since $\mu_{M+1} = \delta$ and $p_j = 2M/(1-\delta)$ for all j, the cost of the optimal shortage variables in the pivotal scenario is at most

$$\delta M \left(\frac{2M}{1-\delta} \right) = \left(\frac{\epsilon M}{2M^2 + \epsilon} \right) \left(\frac{2M^2 + \epsilon}{M} \right) = \epsilon.$$

Thus, the optimal cost for instance $\mathcal{I}$ of (ATO) is at most

$$\sum_i r_i^* + \epsilon. \tag{2.64}$$

Now assume that there exists an optimal solution to instance $\mathcal{I}$ of (ATO), denoted by r_i' for $i = 1, \ldots, M$ and y_{jk}' for $j = 1, \ldots, N$ and $k = 1, \ldots, M+1$, such that $r_i' \neq r_i^*$ for at least one component i. Then, note that since r_i' for $i = 1, \ldots, M$ and y_{jk}' for $j = 1, \ldots, N$ and $k = 1, \ldots, M+1$ are feasible for instance $\mathcal{I}$ of (ATO), then r_i' for $i = 1, \ldots, M$ and y_{jk}' for $j = 1, \ldots, N$ and $k = 1, \ldots, M$ are feasible for problem $\mathcal{I}'$. Further, by the upper bound (2.64), we must have

$$\begin{aligned} \sum_i r_i^* + \epsilon &\geq \sum_i r_i' + 2 \sum_j \sum_{k \leq M} y_{jk}' + \sum_j \mu_{M+1} p_j y_{jM+1}', \\ &\geq \sum_i r_i' + 2 \sum_j \sum_{k \leq M} y_{jk}', \end{aligned}$$

where the last line follows since $y_{jM+1}' \geq 0$. Now, since r_i' and y_{jk}' are integers, the final expression is also an integer. Further, since $\sum_i r_i^*$ is an integer and $\epsilon < 1$, we have

$$\sum_i r_i' + 2 \sum_j \sum_{k \leq M} y_{jk}' \leq \sum_i r_i^*, \tag{2.65}$$

which contradicts r_i^* and y_{jk}^* being the unique optimal solution to $(\mathcal{I}')$. □

2.7 Conclusion

The present work leverages a primal-dual approach to provide a better understanding for managing ATO systems. It is well known that optimizing inventory and allocation policies for ATO systems is computationally challenging, which leads us to focus on the performance of heuristics. We derive approximation guarantees for newsvendor decomposition and LP rounding heuristics for the one-period ATO problem, and extend our results to design integral policies for a dynamic problem. For the one-period problem we demonstrate that, although the newsvendor decomposition provides an intuitive heuristic solution procedure, its worst-case performance can deteriorate as the system becomes more complex. However, we show that the newsvendor decomposition provides a useful lower bound on optimal cost in terms of primitives of the underlying demand distribution. These findings motivate us to consider rounding algorithms based on the optimal solution of the LP relaxation of the ATO problem. We show that the LP relaxation can be solved efficiently using a stochastic subgradient method, even with exponentially many scenarios. Using the LP solution, we carefully design a rounding algorithm in order to achieve superior performance guarantees compared to the newsvendor decomposition. In particular, we use our primal-dual approach to analyze the rounding scheme, thereby establishing a constant approximation factor of 1.8 for any ATO problem instance. Further, we leverage the aforementioned newsvendor lower bound to demonstrate that our rounding is asymptotically optimal as demand grows large. In closing, we believe our analysis using a primal-dual approach offers a fresh perspective for understanding various ATO models.

Chapter 3

Dynamic Assemble-to-Order Systems

3.1 Introduction

We extend the analysis of Chapter 2 to a dynamic setting by leveraging our results for the one-period problem. In particular, we use our rounding algorithm and the fractional policy of Reiman and Wang (2015) to propose an efficiently implementable integral dynamic policy for general ATO systems with identical component lead times. We establish the asymptotic optimality of our integral policy as lead time grows large by using our analysis of both the newsvendor decomposition and LP rounding for the one-period model to extend the framework of Reiman and Wang (2015). Additionally, in Section 3.3, we demonstrate that our solution for the one-period problem can be used to design a good policy for an alternative online allocation problem with lost sales. Thus, our analysis demonstrates that our approach for the one-period problem provides an important building block for analyzing dynamic settings.

The remainder of this chapter is organized as follows. In Section 3.2, we develop an integral policy for a continuous review system with identical component lead times. Section 3.3 analyzes a dynamic online allocation problem, and Section 3.4 offers concluding remarks. We exclude a second literature review in this chapter, as relevant work has already been discussed in the literature review of Chapter 2.

3.2 Continuous Review Backlog Model

In this section we leverage our results for the one-period problem to analyze a dynamic ATO model. We focus on a model of continuous review ATO systems with product backlogs, inventory holding costs, and a common ordering lead time of L for all components, as considered by Doğru et al. (2010); Reiman and Wang (2015); Doğru et al. (2017). We use our LP rounding algorithm, as well as the fractional policy of Reiman and Wang (2015), to design an integral policy for the dynamic system. Then, we use our newsvendor and LP rounding analysis for the one-period problem to show that our policy is asymptotically optimal as lead time grows large, thereby generalizing the result established by Reiman and Wang (2015) for fractional policies.

3.2.1 Stochastic Programming Formulation

In this section we introduce the stochastic program Reiman and Wang (2015) use to develop their dynamic policy, and demonstrate its connections with (ATO). Specifically, we show that our approximation guarantee of Theorem 2.1 for (ATO) yields a similar approximation guarantee for the stochastic program of Reiman and Wang (2015), despite differences in the structure of the objective function. This insight then allows us to establish the main result of this subsection in Corollary 3.1: our LP rounding is asymptotically optimal for the stochastic program of Reiman and Wang (2015) as the lead time grows large.

In the dynamic model, a backlog cost b_j is charged for unfulfilled demand for product j, and a holding cost c_i is charged for unused inventory of component i. Using the notation of the (ATO) problem, $y_{j\omega}$ represents the unfulfilled demand for product j in scenario ω, so $d_{j\omega} - y_{j\omega}$ represents the filled demand, and $r_i - \sum_j a_{ij}(d_{j\omega} - y_{j\omega})$ is the unused inventory of component i. Thus, for a solution that is feasible to the constraints of (ATO), its expected holding and backlog cost is

$$\sum_{\omega} \mu_\omega \left[\sum_i c_i \left(r_i - \sum_j a_{ij}(d_{j\omega} - y_{j\omega}) \right) + \sum_j b_j y_{j\omega} \right]$$

$$= \sum_i c_i r_i - \sum_i c_i \bar{D}_i + \sum_j \mu_\omega \left(b_j + \sum_i a_{ij} c_i \right) y_{j\omega}.$$

Therefore, letting $p_j = b_j + \sum_i a_{ij} c_i$, we cast the following holding cost version of (ATO).

$$\begin{aligned} \min_{\mathbf{r},\mathbf{y}} \ & \sum_i r_i c_i + \sum_{j,\omega} \mu_\omega p_j y_{j\omega} - \sum_i c_i \bar{D}_i \\ \text{s.t. } & r_i + \sum_j a_{ij} y_{j\omega} \geq \sum_j a_{ij} d_{j\omega}, \ \forall\ i,\omega \\ & y_{j\omega} \leq d_{j\omega}, \ \forall\ j,\omega \\ & \mathbf{r},\mathbf{y} \geq 0, \ \text{integer.} \end{aligned} \tag{ATO$_H$}$$

The continuous relaxation of this stochastic integer program is fundamental to the policy development in Doğru et al. (2010); Reiman and Wang (2015); Doğru et al. (2017), and similar problems are considered by Lu and Song (2005); van Jaarsveld and Scheller-Wolf (2015); Zipkin (2016). Indeed, a key step in the analysis of Reiman and Wang (2015) solves the LP relaxation of (ATO$_H$)

to determine base stock levels, and Doğru et al. (2010, 2017) identify special cases of assembly structures for which (ATO$_H$) can be solved exactly.

We further observe that (ATO$_H$) represents an affine transformation of (ATO). In particular, for a given instance, the feasible solutions for (ATO$_H$) and (ATO) are identical, while the objective value of the former is equal to the latter less the quantity $\sum_i c_i \bar{D}_i$ (which is a constant with respect to the decision variables). Letting OPT_H denote the optimal cost of (ATO$_H$), we therefore have the following relationship between the optimal cost of (ATO$_H$) and (ATO)

$$OPT_H = OPT - \sum_i c_i \bar{D}_i. \tag{3.1}$$

Thus, from an exact optimization perspective, (ATO$_H$) and (ATO) are equivalent. We note, though, that since OPT_H is less than OPT, approximation factors for (ATO) do not directly hold for (ATO$_H$). Developing a constant factor approximation guarantee for (ATO$_H$) by directly applying the methods we develop in Chapter 2 is challenging, since most existing primal-dual techniques are well suited for problems with *linear* objective functions (e.g., (ATO)), as opposed to *affine* objective functions (e.g., (ATO$_H$)). However, we next show that the guarantee of Theorem 2.1 does translate well to (ATO$_H$).

Proposition 3.1. *The LP rounding (2.48-2.49) provides the following approximation factor for* (ATO$_H$),

$$1 + \frac{1}{\rho}.$$

Proof. We first note that for the cost structure of problem (ATO$_H$), for any product j the markup is

$$\gamma_j = \frac{b_j + \sum_i a_{ij} c_i}{\sum_i a_{ij} c_i} = \frac{b_j}{\sum_i a_{ij} c_i} + 1,$$

which is strictly larger than one. Thus, the smallest markup is $\underline{\gamma} > 1$, and we have $\hat{\gamma} = 1$ and $\gamma^* > 0$.

Since $\hat{\gamma} = 1$, by Proposition 2.1 and (3.1), we have

$$\begin{aligned} \gamma^* \sum_i c_i \mathbb{E}\left[|D_i - \bar{D}_i|\right] &\leq OPT - \sum_i c_i \bar{D}_i, \\ &= OPT_H. \end{aligned}$$

By this lower bound on OPT_H, together with (3.1) and $\hat{\gamma} = 1$, we can upper bound OPT as follows

$$
\begin{aligned}
OPT &= \left(\frac{OPT_H + \sum_i c_i \bar{D}_i}{OPT_H} \right) OPT_H, \\
&\leq \left(1 + \frac{\sum_i c_i \bar{D}_i}{\gamma^* \sum_i c_i \mathbb{E}[|D_i - \bar{D}_i|]} \right) OPT_H, \\
&= \left(1 + \frac{\bar{d}}{\rho} \right) OPT_H
\end{aligned}
$$

Let RND denote the objective value of the LP rounded solution in (ATO), and let RND_H denote objective value of this same solution in the program (ATO$_H$). Then, we have

$$
\begin{aligned}
RND_H &= RND - \sum_i c_i \bar{D}_i, \\
&\leq \left(1 + \frac{1}{\bar{d} + \rho} \right) OPT - \sum_i c_i \bar{D}_i, \\
&= \frac{1}{\bar{d} + \rho} OPT + OPT_H, \\
&\leq \left(1 + \frac{1}{\rho} \right) OPT_H,
\end{aligned}
$$

where the first inequality follows from Theorem 2.1. □

The intuition for Proposition 3.1 is that subtracting $\sum_i c_i \bar{D}_i$ from the optimal value of (ATO) removes this quantity from the lower bound of Proposition 2.1, providing a lower bound on OPT_H in terms of $\sum_i c_i \mathbb{E}\left[|D_i - \bar{D}_i|\right]$, which naturally leads to a bound in terms of ρ. Thus, for any regime where ρ grows large, the rounding (2.48-2.49) is asymptotically optimal for (ATO$_H$).

Next we show that ρ grows large for the particular demand structure modeled in Reiman and Wang (2015). Demand for all products arrives in continuous time, indexed by $t \geq 0$, according to a Poisson process with rate λ. Each time there is an arrival of the Poisson process, product demand is realized according to the multivariate random variable $\mathbf{S} = (S_1, \ldots, S_N)$. The realizations of the vector $\mathbf{S}$ are independent from each other for different arrivals of the Poisson process, as well as independent from the Poisson process itself, however correlation between demands for different products is allowed within a single realization of $\mathbf{S}$. We refer to this demand process as *compound Poisson demand*.

For two points in time, $t_2 > t_1 \geq 0$, let $\mathbf{D}(t_1, t_2)$ denote the product demand arriving in this time interval. In the model of Reiman and Wang (2015), all components have a common, deterministic lead time of length L. To design a dynamic policy, Reiman and Wang (2015) study the problem where demand in (ATO$_H$) is distributed as $\mathbf{D}(0, L)$, i.e., the demand arriving during a lead time L. We follow their asymptotic regime, letting $L \to \infty$ and show that the approximation factor for (ATO$_H$) approaches one.

Corollary 3.1. *When demand is distributed as* $\mathbf{D}(0, L)$*, the rounding (2.48-2.49) provides the following approximation factor for* (ATO$_H$)*,*

$$1 + O\left(L^{-\frac{1}{2}}\right).$$

Thus, the rounding scheme (2.48-2.49) is asymptotically optimal for (ATO$_H$) as L grows large in the demand model considered by Reiman and Wang (2015). In the following section, we leverage this result to extend the entire dynamic policy of Reiman and Wang (2015) to the setting with integer constraints.

Before presenting the proof of Corollary 3.1, we introduce the following notation. Under compound Poisson demand with rate λ, for time $t \geq 0$, let $N(t)$ denote the number of arrivals of the Poisson process up to time t, and let $\mathbf{S}(k)$ denote the kth realization of the product demand vector $\mathbf{S}$. Then, the cumulative demand for product j up to time t is

$$D_j(0, t) = \sum_{k=1}^{N(t)} S_j(k),$$

and for brevity, we will denote the aggregate demand for component i over a lead time (i.e., the demand D_i used to derive the approximation factor) by

$$D_i(L) = \sum_j a_{ij} D_j(0, L).$$

By standard properties of compound Poisson processes, if we consider the demand per unit time and denote the means as

$$\mu_j = \lambda \mathbb{E}[S_j], \ \forall j,$$

and covariances as

$$\sigma_{jj'}^2 = \lambda \mathbb{E}[S_j S_{j'}] = \lambda \left(\mathrm{Cov}(S_j, S_{j'}) + \mathbb{E}[S_j]\mathbb{E}[S_{j'}]\right), \ \forall j, j',$$

then $\mu_i = \sum_j a_{ij}\mu_j$ and $\sigma_i^2 = \sum_{j,j'} a_{ij}a_{ij'}\sigma_{jj'}^2$ are the mean and variance of demand for component i per unit time. Further, the mean and variance of demand for component i during the lead time L are

$$\mathbb{E}[D_i(L)] = L\mu_i,$$
$$\text{Var}[D_i(L)] = L\sigma_i^2.$$

Proof of Corollary 3.1. Under compound Poisson demand, the approximation factor of Proposition 3.1 is

$$1 + \frac{\sum_i c_i n_i}{\gamma^* \sum_i c_i \mathbb{E}[|D_i(L) - L\mu_i|]},$$

and we show that for any $0 < \epsilon < \sqrt{\frac{2}{\pi}}$, there exists L_0 such that for all $L \geq L_0$, this factor is less than

$$1 + \frac{\sum_i c_i n_i}{\gamma^* \left(\sqrt{\frac{2}{\pi}} - \epsilon\right) \sqrt{L} \sum_i c_i \sigma_i},$$

for which it will suffice to show that for any component i, there exists L_0^i such that for $L \geq L_0^i$ we have

$$\left(\sqrt{\frac{2}{\pi}} - \epsilon\right) \sqrt{L}\sigma_i \leq \mathbb{E}[|D_i(L) - L\mu_i|],$$

and letting $L_0 = \max_i\{L_0^i\}$. To prove this, we focus on a single component i, and consider centering $D_i(L)$ by its mean and scaling by its standard deviation:

$$X(L) = \frac{D_i(L) - L\mu_i}{\sqrt{L}\sigma_i},$$

so that the claim will follow by showing there exists L_0^i such that for $L \geq L_0^i$ we have

$$\sqrt{\frac{2}{\pi}} - \epsilon \leq \mathbb{E}[|X(L)|].$$

By the central limit theorem for compound Poisson processes (Robbins, 1948), as $L \to \infty$, the random variable $X(L)$ converges in distribution to a standard normal random variable (zero mean, unit variance), which we denote by Z. We then claim a small extension to Fatou's Lemma (Durrett, 2013, Exercise 3.2.4), namely

$$\mathbb{E}[|Z|] \leq \liminf_{L\to\infty} \mathbb{E}[|X(L)|], \tag{3.2}$$

which we will prove by contradiction. First note that Durrett (2013) Exercise 3.2.4 shows that if a countable sequence of random variables X_n converges to a random variable X in distribution, then

$$\mathbb{E}[|X|] \leq \liminf_{n \to \infty} \mathbb{E}[|X_n|], \tag{3.3}$$

since the absolute value is a continuous, non-negative function. Note that (3.2) is the same as (3.3) except that the limit is taken along an uncountable sequence of random variables.

Now, assume that (3.2) does not hold. Then, for some $\epsilon > 0$, there exists an unbounded set $\mathcal{L} = \{L | \mathbb{E}[|X(L)|] < \mathbb{E}[|Z|] - \epsilon\}$. Choose some countable, unbounded subset of $\mathcal{L}$, denoted $\mathcal{L}'$, and define $X_n = X(L_n)$, where L_n is the nth smallest element of $\mathcal{L}'$. Then, X_n is a countable subsequence of $X(L)$, and thus also converges in distribution to Z. However, the definition of $\mathcal{L}$ implies that $\mathbb{E}[|X_n|] < \mathbb{E}[|Z|] - \epsilon$ for all n, which contradicts (3.3). Therefore (3.2) holds along the uncountable sequence of random variables $X(L)$.

Thus, for $\epsilon > 0$, by (3.2) there exists L_0^i, such that for all $L \geq L_0^i$, we have

$$\begin{aligned} \mathbb{E}[|X(L)|] &\geq \mathbb{E}[|Z|] - \epsilon, \\ &= \sqrt{\frac{2}{\pi}} - \epsilon, \end{aligned}$$

where the final equality follows from the fact that Z is a standard normal. □

3.2.2 Asymptotically Optimal Policy

We now consider component replenishment and allocation policies for the dynamic model. To do so, we first introduce necessary notation for the various processes in the system. At any time $t \geq 0$, the following events (if they occur at time t) take place in this sequence: new demand arrives, component inventories ordered at time $t - L$ arrive, on hand inventory is allocated to clear backlogged demand, and new component inventory is ordered. Any unfilled demand is backlogged and unused components remain in inventory. Let $\mathbf{I}(t)$ and $\mathbf{B}(t)$ denote the on-hand inventory and backlog levels, respectively, at time t. Let $\mathbf{I}^-(t)$ and $\mathbf{B}^-(t)$ denote the on-hand inventory and backlog levels at time t after any components are received and any demand has arrived, but before any replenishment or allocation decision has been made. The decisions at time t are a component replenishment order, denoted by and $\mathbf{r}(t)$ (to be received at time $t + L$), and the demand to fill instantaneously at time t, denoted by $\mathbf{z}(t)$. The inventory and backlog levels before and after demand fulfillment are related as follows (recalling that A denotes the matrix of component requirements for each product),

$$\mathbf{I}(t) = \mathbf{I}^{-}(t) - A\mathbf{z}(t), \tag{3.4}$$

$$\mathbf{B}(t) = \mathbf{B}^{-}(t) - \mathbf{z}(t). \tag{3.5}$$

To help define allocation policies, we also consider component shortage levels, defined as

$$\mathbf{Q}(t) = A\mathbf{B}^{-}(t) - \mathbf{I}^{-}(t) \tag{3.6}$$

Now consider a replenishment and allocation policy π, which decides $\mathbf{r}(t)$ and $\mathbf{z}(t)$ at each time t, as described above. A feasible policy must be non-anticipatory (i.e., only use information at time t that is available at time t) and satisfy the following constraints

$$A\mathbf{z}(t) \leq \mathbf{I}^{-}(t), \tag{3.7}$$

$$\mathbf{z}(t) \leq \mathbf{B}^{-}(t), \tag{3.8}$$

$$\mathbf{r}(t), \mathbf{z}(t) \geq 0, \text{ integer}. \tag{3.9}$$

The first constraint guarantees that the policy only uses the inventory available to fill demand, the second that only existing backlogs can be cleared, and the third that only non-negative integer quantities can be chosen. Let $\mathbf{c}$, and $\mathbf{b}$ denote the vector of holding and backlog costs, respectively. For a given policy π, letting $\mathbf{B}^{\pi}(t)$ and $\mathbf{I}^{\pi}(t)$ denote the backlog and on-hand inventory levels, respectively, at time t, then long run average cost of this policy is

$$C^{\pi} = \limsup_{T\to\infty} \frac{1}{T} \int_{L}^{T+L} \mathbb{E}[\mathbf{b} \cdot \mathbf{B}^{\pi}(t) + \mathbf{c} \cdot \mathbf{I}^{\pi}(t)]dt.$$

An optimal policy minimizes C^{π} among all feasible policies π, and we denote the optimal cost by C^{*}. In this setting, Reiman and Wang (2015) develop a policy that is asymptotically optimal as the lead time, L, becomes large. Their policy implicitly assumes that inventory can be replenished and allocated in fractional units. While this is a highly non-trivial step in demonstrating the usefulness of their policy, a practical limitation is that real world applications typically require integrality constraints on both replenishment and allocation decisions. In many practical settings (e.g., automotive, computer, e-retail), suppliers will only ship whole units of components, and customers only demand whole units of products. It is well known that integrality constraints make many problems intractable, and it is not immediately clear how to extend the analysis of Reiman and Wang (2015) to this setting. In the following discussion we use our results for the single-period problem to adapt the dynamic policy of Reiman and Wang (2015) to handle integer constraints, while maintaining asymptotic optimality.

Our policy makes replenishment decisions according to a basestock policy, which keeps the net inventory position at a constant level, i.e., on-hand inventory, plus ordered but not yet arrived inventory, minus backlog remains at the base stock level for each component. Let $\mathbf{r}^{LP}$ denote an optimal first stage solution to (ATO$_H$). Then, we set base stock levels as

$$\mathbf{r}^I = \lfloor \mathbf{r}^{LP} \rfloor, \tag{3.10}$$

where the rounding down is performed on the vector component-wise. Next, we determine the allocation decisions by setting backlog targets at time $t \geq 0$ according to the following program

$$\mathbf{B}^*(t) = \operatorname{argmin} \left\{ \|\mathbf{B}\| \,\middle|\, \mathbf{B} \in \operatorname*{argmin}_{\mathbf{B} \geq 0} \{\mathbf{p} \cdot \mathbf{B} | A\mathbf{B} \geq \mathbf{Q}(t)\} \right\}, \tag{3.11}$$

where $\|\mathbf{B}\|$ denotes the Euclidean norm of the vector $\mathbf{B}$. The minimum norm solution is chosen in order to guarantee that $\mathbf{B}^*(t)$ obeys a uniform Lipschitz condition with respect to $\mathbf{Q}(t)$, which facilitates technical proofs bounding the evolution of the backlog targets over time. The intuition of the backlog targets is that solving the LP, $\min\{\mathbf{p} \cdot \mathbf{B} | A\mathbf{B} \geq \mathbf{Q}(t), \mathbf{B} \geq 0\}$, attempts to update the backlog levels in order to minimize the backlog cost, while taking into account some of the feasibility constraints.

In particular, it follows directly from (3.4-3.6) that $A\mathbf{B}(t) \geq \mathbf{Q}(t)$ and $\mathbf{B}(t) \geq 0$ are equivalent to $\mathbf{I}^-(t) \geq A\mathbf{z}(t)$ and $\mathbf{B}^-(t) \geq \mathbf{z}(t)$, respectively, thus satisfying feasibility constraints (3.7) and (3.8). However, the non-negativity and integrality requirements on $\mathbf{z}(t)$ in constraint (3.9) are not guaranteed by (3.11). Thus, in order to implement these backlog targets in a feasible allocation policy, we first round them according to

$$\mathbf{B}^I(t) = \lceil \mathbf{B}^*(t) \rceil. \tag{3.12}$$

We note that the rounding (3.10) and (3.12) mirrors the LP-rounding scheme (2.48-2.49), since the backlog targets $\mathbf{B}^*(t)$ are analogous to the shortage variables $y_{j\omega}$ in (ATO). Next, we accommodate the non-negativity constraint on $\mathbf{z}(t)$ by implementing a greedy allocation procedure that treats the backlog target $\mathbf{B}^I(t)$ as a lower bound on the true backlog implemented in the system. Algorithm 3 serves products with backlogs that are currently above their target level, corresponding to when $(B_j^-(t) - B_j^I(t))^+ > 0$ in the first argument of the minimization in (3.13). The second argument in the minimization of (3.13) is included to maintain feasibility for the inventory constraint

(3.7). We note that Algorithm 3 is an implementation of the Allocation Principle described in Reiman and Wang (2015), using integer backlog targets.

Algorithm 3 Allocation Procedure

1: Input $\mathbf{I}^-(t), \mathbf{B}^-(t)$, and $\mathbf{B}^I(t)$
2: Initialize $\mathbf{z} \leftarrow \mathbf{0}$
3: **for** $\forall j$ **do**

$$z_j = \min\left(\left(B_j^-(t) - B_j^I(t)\right)^+, \min_{i \in \mathcal{N}(j)} \left\{ \left\lfloor \frac{\left(I_i^-(t) - \sum_{j' \neq j} a_{ij'} z_{j'}\right)^+}{a_{ij}} \right\rfloor \right\} \right) \tag{3.13}$$

4: **end for**
5: $\mathbf{z}(t) \leftarrow \mathbf{z}$

Every time there is a demand arrival or inventory delivery (i.e. any time $\mathbf{I}^-(t)$ or $\mathbf{B}^-(t)$ changes) we resolve (3.11) to set new backlog targets and run Algorithm 3 to determine the backlogged demand to fill, $\mathbf{z}(t)$. It is straightforward to establish that $\mathbf{z}(t)$ output by Algorithm 3 is feasible, as z_j can only take non-negative integer values, and for each product the minimization in (3.13) guarantees that the solution remains feasible for constraints (3.7) and (3.8). Note that in Step 3, the products can be considered in any order; Reiman and Wang (2015) make the reasonable suggestion of prioritizing products with higher cost p_j.

Our next result establishes that this policy is asymptotically optimal. Let C_L^* denote the optimal long run average cost of the system with lead time L, and let C_L^I denote the long run average cost of the policy with base stock levels (3.10) and which allocates components according to Algorithm 3. Finally, we maintain the assumption of Reiman and Wang (2015) that for some $\delta > 0$, the random vector $\mathbf{S}$ (i.e., the product demand for each arrival) has a finite moment of order $2 + \delta$ for each j,

$$\eta_j = \mathbb{E}[S_j^{2+\delta}] < \infty, \ \forall j.$$

Theorem 3.1. *When demand is distributed according to compound Poisson demand with finite η_j for all j, a policy with base stock levels* (3.10)*, and allocations made according to Algorithm 3, is asymptotically optimal as $L \to \infty$, i.e.,*

$$\lim_{L \to \infty} \frac{C_L^I}{C_L^*} = 1.$$

This result demonstrates the applicability of our rounding results to the dynamic setting. The rounding scheme (2.48-2.49) provides the right intuition for rounding base stock levels and backlog targets to maintain asymptotic optimality for an integral policy in the model of Reiman and Wang (2015). In Section 3.2.3 we extend the result of Theorem 3.1 to a general family of base stock policies that also incorporate reservation levels in the allocation algorithm, which fully generalize the family of policies considered in Reiman and Wang (2015) to the integral setting. The proof of Theorem 3.1 then follows as a special case of Proposition 3.3 in Section 3.2.3.

3.2.3 Policy Analysis

In this section, we prove Theorem 3.1. In order to do this, we extend a broader family of policies developed by Reiman and Wang (2015) to the integral setting, and demonstrate asymptotic optimality for all policies in this family. In particular, Reiman and Wang (2015) consider two related stochastic programs, and allow setting base stock levels using any convex combination of the programs' solutions. We adapt this approach to the integral setting with an appropriate rounding procedure. Further, Reiman and Wang (2015) also incorporate component reservation quantities in their allocation procedure, which we also allow. Thus, we fully extend the analysis of Reiman and Wang (2015) to the integral setting. We note that the proof of Theorem 3.1 then follows as a special case of the proof of Proposition 3.3 at the end of this section.

Stochastic programs. To begin, we discuss the two stochastic programs considered by Reiman and Wang (2015). Let $\mathbf{D}$ denote a random vector representing demand. Recall that $\mathbf{c}$ and $\mathbf{b}$ are the vector of holding and backlog costs, respectively, and that $\mathbf{p} = \mathbf{b} + \mathbf{c} \cdot A$. Then the first stochastic program, i.e., equation (12) of Reiman and Wang (2015), is

$$\begin{aligned} &\min_{\mathbf{r} \geq 0} \mathbf{c} \cdot \mathbf{r} + \mathbf{b} \cdot \mathbb{E}[\mathbf{D}] - \mathbb{E}[\varphi_+(\mathbf{r}, \mathbf{D})], \\ &\text{s.t. } \varphi_+(\mathbf{r}, \mathbf{D}) = \max_{\mathbf{z} \geq 0} \{\mathbf{p} \cdot \mathbf{z} | \mathbf{z} \leq \mathbf{D}, A\mathbf{z} \leq \mathbf{r}\}. \end{aligned} \tag{3.14}$$

We claim that this program is equivalent with the LP relaxation of (ATO_H), considered in Section 3.2.1, which will be clear after a transformation of the decision variables. For a given realization of demand $\mathbf{D}$, define the shortage vector $\mathbf{y} = \mathbf{D} - \mathbf{z}$. Then it is clear that $0 \leq \mathbf{z} \leq \mathbf{D}$ if and only if $0 \leq \mathbf{y} \leq \mathbf{D}$. Further, with the constraint that $0 \leq \mathbf{y} \leq \mathbf{D}$, it is clear that $A\mathbf{z} \leq \mathbf{r}$ is equivalent to $A\mathbf{y} \geq A\mathbf{D} - \mathbf{r}$. Thus, with this change of variables, we have

$$\begin{aligned} \varphi_+(\mathbf{r}, \mathbf{D}) &= \max_{\mathbf{y} \geq 0} \{\mathbf{p} \cdot (\mathbf{D} - \mathbf{y}) | \mathbf{y} \leq \mathbf{D}, A\mathbf{y} \geq A\mathbf{D} - \mathbf{r}\}, \\ &= \mathbf{p} \cdot \mathbf{D} - \min_{\mathbf{y} \geq 0} \{\mathbf{p} \cdot \mathbf{y} | \mathbf{y} \leq \mathbf{D}, A\mathbf{y} \geq A\mathbf{D} - \mathbf{r}\}. \end{aligned}$$

We therefore restate the program (3.14) as the following equivalent program

$$\begin{aligned} &\min_{\mathbf{r}\geq 0} \mathbf{c}\cdot\mathbf{r} - \mathbf{c}\cdot A\mathbb{E}[\mathbf{D}] + \mathbb{E}[\psi_+(\mathbf{r},\mathbf{D})], \\ &\text{s.t. } \psi_+(\mathbf{r},\mathbf{D}) = \min_{\mathbf{y}\geq 0}\{\mathbf{p}\cdot\mathbf{y}|\mathbf{y}\leq\mathbf{D}, A\mathbf{y}\geq A\mathbf{D}-\mathbf{r}\}. \end{aligned} \tag{3.15}$$

This problem is equivalent to the LP relaxation of (ATO$_H$) since $\mathbf{c}\cdot A\mathbb{E}[\mathbf{D}] = \sum_i c_i\bar{D}_i$ and the constraints of the LP ψ_+ correspond to those of (ATO$_H$). Thus, the approximation guarantee of Corollary 3.1 also holds for (3.15), which we restate here for clarity.

Corollary 3.2. *When* $\mathbf{D} = \mathbf{D}(0, L)$*, the rounding (2.48-2.49) provides the following approximation factor for* (3.15)*,*

$$1 + O\left(L^{-\frac{1}{2}}\right).$$

The second stochastic program, i.e., equation (17) in Reiman and Wang (2015), is

$$\begin{aligned} &\min_{\mathbf{r}} \mathbf{c}\cdot\mathbf{r} + \mathbf{b}\cdot\mathbb{E}[\mathbf{D}] - \mathbb{E}[\varphi(\mathbf{r},\mathbf{D})], \\ &\varphi(\mathbf{r},\mathbf{D}) = \max_{\mathbf{z}}\{\mathbf{p}\cdot\mathbf{z}|\mathbf{z}\leq\mathbf{D}, A\mathbf{z}\leq\mathbf{r}\}. \end{aligned} \tag{3.16}$$

The difference between (3.14) and (3.16) is that decision variables $\mathbf{r}$ and $\mathbf{z}$ are allowed to be negative in the latter problem. Reiman and Wang (2015) show that the optimal cost of this program provides a lower bound on the optimal cost in their dynamic setting, which is a key step to showing asymptotic optimality. We now perform a similar transformation to the decision variables by letting $\mathbf{y} = \mathbf{D} - \mathbf{z}$, and by the same argument showing equivalence between problems (3.14) and (3.15), it is clear that (3.16) is equivalent to

$$\begin{aligned} &\min_{\mathbf{r}} \mathbf{c}\cdot\mathbf{r} - \mathbf{c}\cdot A\mathbb{E}[\mathbf{D}] + \mathbb{E}[\psi(\mathbf{r},\mathbf{D})], \\ &\text{s.t. } \psi(\mathbf{r},\mathbf{D}) = \min_{\mathbf{y}\geq 0}\{\mathbf{p}\cdot\mathbf{y}|A\mathbf{y}\geq A\mathbf{D}-\mathbf{r}\}. \end{aligned} \tag{3.17}$$

We note that (3.17) is a relaxation of (3.15), since the constraints $\mathbf{r}\geq 0$ and $\mathbf{y}\leq\mathbf{D}$ are removed. Our first result demonstrates that we can approximate an optimal integral solution to (3.17) with the same factor as (ATO$_H$) from Proposition 3.1.

Proposition 3.2. *The LP rounding (2.48-2.49) provides the following approximation factor for* (3.17)*,*

$$1 + \frac{1}{\rho}.$$

Proof. Writing the problem (3.17) in extensive form as the following LP

$$
\begin{aligned}
\min_{\mathbf{r},\mathbf{y}} \ & \sum_i r_i c_i + \sum_{j,\omega} \mu_\omega p_j y_{j\omega} - \sum_i c_i \bar{D}_i \\
\text{s.t. } & r_i + \sum_j a_{ij} y_{j\omega} \geq \sum_j a_{ij} d_{j\omega}, \ \forall\, i, \omega \\
& \mathbf{y} \geq 0,
\end{aligned}
$$

its dual problem is

$$
\begin{aligned}
\max_{\boldsymbol{\theta},\boldsymbol{\pi}} & \sum_{i,\omega} \theta_{i\omega} \sum_j a_{ij} d_{j\omega} - \sum_i c_i \bar{D}_i \\
\text{s.t.} & \sum_i a_{ij} \theta_{i\omega} \leq \mu_\omega p_j, \ \forall\, j, \omega \\
& \sum_\omega \theta_{i\omega} \leq c_i, \ \forall\, i \\
& \boldsymbol{\theta}, \boldsymbol{\pi} \geq 0.
\end{aligned}
\tag{3.18}
$$

The constant markup newsvendor solution $\theta^\gamma_{i\omega}$ defined in (2.17) is feasible to this problem, since by Lemma 2.1 it is feasible to (DATO) for $\pi^\gamma_{j\omega} = 0$, which clearly doesn't violate the constraints of (3.18). Further, the objective value of the constant markup newsvendor solution in (3.18) equals the objective value of the same solution in (DATO), less the quantity $\sum_i c_i \bar{D}_i$. The result then follows from arguments similar to those in the proof of Theorem 2.1 and Proposition 3.1. □

We also have the following result similar to Corollaries 3.1 and 3.2, whose proof follows from Proposition 3.2 and an identical argument to the proof of Corollary 3.1 in Section 3.2.1.

Corollary 3.3. *When* $\mathbf{D} = \mathbf{D}(0, L)$*, the rounding (2.48-2.49) provides the following approximation factor for* (3.17),

$$1 + O\left(L^{-\frac{1}{2}}\right).$$

Thus, from Corollaries 3.2 and 3.3, the LP rounding solutions for (3.15) and (3.17) approach optimal as L becomes large. This convergence of the cost of the integral solution to the optimal LP solution will facilitate our subsequent analysis of the dynamic integral policy.

Preliminaries. We now introduce a bit more notation, and reproduce the notation from Section 3.2.2 to enhance readability of the current section. For $t \geq 0$, let $\mathbf{d}(t)$ denote the demand arriving at time t. For two points in time, $t_2 > t_1 \geq 0$, $\mathbf{D}(t_1, t_2)$ denotes the demand arriving in this time interval. Demand arriving during the lead time preceding $t \geq L$ is denoted

$$\mathbf{D}(t) = \mathbf{D}(t - L, t).$$

We note that the memoryless property of the Poisson arrival process implies that $\mathbf{D}(t)$ has the same distribution for all times $t \geq L$ (i.e., it has the same distribution as $\mathbf{D}(0, L)$ from Corollaries 3.1, 3.2, and 3.3). Let $\mathbf{D}^{(L)}$ denote a random vector with this distribution in the system with lead time L.

At any time $t \geq -L$, a policy can order new component inventory, which will be received at time $t + L$. For $t \geq -L$, let $\mathbf{r}(t)$ denote the inventory replenishment order placed at time point t. Similar to the demand process, for two points in time, $t_2 > t_1 \geq -L$, let $\mathbf{R}(t_1, t_2)$ denote the amount of new orders placed in this time interval. New orders placed during the lead time preceding $t \geq 0$ are denoted

$$\mathbf{R}(t) = \mathbf{R}(t - L, t).$$

For $t \geq 0$, let $\mathbf{z}(t)$ denote the demand filled at time point t. For $t_2 > t_1 \geq 0$, let $\mathbf{Z}(t_1, t_2)$ denote the amount demand filled in this time interval. Demand filled during the lead time preceding $t \geq L$ is denoted

$$\mathbf{Z}(t) = \mathbf{Z}(t - L, t).$$

As in Section 3.2.2, $\mathbf{I}(t)$ and $\mathbf{B}(t)$ denote the inventory and backlog levels at time $t \geq 0$, and we have the following relationships

$$\mathbf{I}(t) = \mathbf{I}(t - L) + \mathbf{R}(t - L) - A\mathbf{Z}(t),$$
$$\mathbf{B}(t) = \mathbf{B}(t - L) + \mathbf{D}(t) - \mathbf{Z}(t).$$

At any time $t \geq 0$, the following events (if they occur at time t) take place in a specific sequence: new demand arrives, component inventories ordered at time $t - L$ arrive, on hand inventory is allocated to clear backlogged demand, and new component inventory is ordered. For time $t \geq 0$, let $\mathbf{B}(t^-)$ denote the backlog levels immediately before any of these events occur at time t, so that

$$\mathbf{B}(t) = \mathbf{B}(t^-) + \mathbf{d}(t) - \mathbf{z}(t). \tag{3.19}$$

Further, recall from Section 3.2.2 that $\mathbf{I}^-(t)$ and $\mathbf{B}^-(t)$ denote the inventory and backlog levels at time t after any components are received and any demand has arrived, but before any replenishment or allocation decision has been made, and obey the following

$$\mathbf{I}(t) = \mathbf{I}^{-}(t) - A\mathbf{z}(t), \tag{3.4 revisited}$$

$$\mathbf{B}(t) = \mathbf{B}^{-}(t) - \mathbf{z}(t). \tag{3.5 revisited}$$

Recall also the component shortage levels

$$\mathbf{Q}(t) = A\mathbf{B}^{-}(t) - \mathbf{I}^{-}(t) \tag{3.6 revisited}$$

Recall that a policy π decides at each point in time t, a replenishment order $\mathbf{r}(t)$ and demand fulfillment $\mathbf{z}(t)$, and a feasible policy must be non-anticipatory (i.e., only use information at time t that is available at time t) and satisfy the following constraints:

$$A\mathbf{z}(t) \leq \mathbf{I}^{-}(t), \tag{3.7 revisited}$$

$$\mathbf{z}(t) \leq \mathbf{B}^{-}(t), \tag{3.8 revisited}$$

$$\mathbf{z}(t) \geq 0, \ \mathbf{r}(t), \mathbf{z}(t) \text{ integer.} \tag{3.9 revisited}$$

The first constraint guarantees that the policy only uses the inventory available to fill demand, the second that only existing backlogs can be cleared, and the third that only non-negative integer quantities can be chosen. Recall that $\mathbf{c}$, and $\mathbf{b}$ denote the vector of holding and backorder costs, respectively, and letting $\mathbf{B}^{\pi}(t)$ and $\mathbf{I}^{\pi}(t)$ denote the backlog and inventory levels of policy π, the long run average cost of this policy is

$$C^{\pi} = \limsup_{T\to\infty} \frac{1}{T} \int_{L}^{T+L} \mathbb{E}[\mathbf{b} \cdot \mathbf{B}^{\pi}(t) + \mathbf{c} \cdot \mathbf{I}^{\pi}(t)]dt.$$

The next theorem provides a lower bound on the cost of any feasible policy, and follows from a combination of Theorems 1 and 2 in Reiman and Wang (2015), of which the former follows from Theorem 2.1 in Doğru et al. (2010). We note that the since the fractional model Doğru et al. (2010) and Reiman and Wang (2015) is a relaxation of our model with integer constraints, their lower bound result applies to our setting directly.

Theorem 3.2 (Reiman and Wang, 2015)**.** *Let* $\underset{\sim}{C}$ *be the optimal value of* (3.17) *when demand is distributed according to* $\mathbf{D}^{(L)}$*. Then for any feasible policy* π*, we have*

$$\underset{\sim}{C} \leq C^{\pi}.$$

Policy description. Here we describe our family of dynamic policies, adapted from Reiman and Wang (2015) to handle integrality constraints. We first describe the form of the replenishment policy and a few resulting relationships. The replenishment decisions are made according to a basestock policy, which keeps the net inventory position at a constant level, i.e., on-hand inventory, plus ordered but not yet arrived inventory, minus backlog remains at a constant level for a given component. If we let $\mathbf{r}$ be the base stock levels, then we have

$$\mathbf{r} = \mathbf{I}(t) + \mathbf{R}(t) - A\mathbf{B}(t). \tag{3.20}$$

Since replenishment decisions are allowed starting at time $t = -L$, we follow Reiman and Wang (2015) in assuming, without loss of generality, that at time $t = 0$ inventory is at the basestock level, $\mathbf{r}$. Then, for $t \geq L$, to maintain base stock levels $\mathbf{r}$, it is clear from (3.20) that new orders over the previous lead time must compensate for new demand over the same period, i.e., $\mathbf{R}(t) = A\mathbf{D}(t)$. Thus, substituting into (3.20) and rearranging, we can write the inventory levels for $t \geq L$ as

$$\mathbf{I}(t) = \mathbf{r} + A\mathbf{B}(t) - A\mathbf{D}(t). \tag{3.21}$$

Combining this with the definition of $\mathbf{Q}(t)$, $\mathbf{I}^{-}(t)$, and $\mathbf{B}^{-}(t)$, we see that under a base stock policy, the component shortages satisfy

$$\mathbf{Q}(t) = A\mathbf{D}(t) - \mathbf{r}. \tag{3.22}$$

Now we will define the base stock levels used by the policy. Let $\mathbf{r}^*$ denote an optimal first stage solution to (3.17) and $\mathbf{r}^o$ an optimal first stage solution to (3.15). Then for any $\alpha \in [0, 1]$, let the base stock levels be

$$\mathbf{r}^{\alpha} = \lfloor \alpha \mathbf{r}^* + (1 - \alpha)\mathbf{r}^o \rfloor, \tag{3.23}$$

where the rounding down is performed on the vector component-wise. Next, we adapt the allocation policy of Reiman and Wang (2015), who set backlog targets at time $t \geq 0$ by solving the following program

$$\mathbf{B}^*(t) = \operatorname{argmin} \left\{ \|\mathbf{B}\| \,\middle|\, \mathbf{B} \in \underset{\mathbf{B} \geq 0}{\operatorname{argmin}} \{\mathbf{p} \cdot \mathbf{B} | A\mathbf{B} \geq \mathbf{Q}(t)\} \right\}, \tag{3.24}$$

where $\|\mathbf{B}\|$ denotes the Euclidean norm of the vector $\mathbf{B}$. The minimum norm solution is chosen in order to guarantee that $\mathbf{B}^*(t)$ obeys a uniform Lipschitz condition with respect to $\mathbf{Q}(t)$, which facilitates technical proofs bounding the evolution of the backlog targets over time. The intuition of the backlog targets is that solving the LP, $\min\{\mathbf{p} \cdot \mathbf{B} | A\mathbf{B} \geq \mathbf{Q}(t), \mathbf{B} \geq 0\}$, attempts to update the backlog levels in order to minimize the backlog cost. Reiman and Wang (2015) demonstrate that this relaxation allows the cost of their dynamic policy to asymptotically approach the lower bound of (3.17) as the lead time grows.

However, the backlog targets $\mathbf{B}^*(t)$ are not necessarily feasible for constraint (3.9). In order to find a feasible allocation, we first round the backlog targets up, (as in (3.12) of Section 3.2.2)

$$\mathbf{B}^I(t) = \lceil \mathbf{B}^*(t) \rceil. \tag{3.25}$$

For time $t \geq 0$, let $\mathbf{B}^*(t^-)$ denote the fractional backlog targets defined in (3.24) immediately before any event occurs at time t (similar to the definition of $\mathbf{B}(t^-)$ above). Likewise, let $\mathbf{B}^I(t^-) = \lceil \mathbf{B}^*(t^-) \rceil$ denote the rounded backlog target immediately before any event occurs at time t.

Given the rounded backlog targets in (3.25), we follow Reiman and Wang (2015) in defining a feasible allocation procedure. Let $w_{ij} \geq 0$ be some integer quantity of component i inventory that will be held back from allocating to product j. Intuitively, the quantity w_{ij} represents an amount of component i we do not want to use on product j, in the hopes of reserving it for another, higher value product. Algorithm 4 serves products with backlogs that are currently above their target level, i.e., $(B_j^-(t) - B_j^u(t))^+ > 0$ in the first argument of the minimization in (3.26), while the second argument in the minimization of (3.26) maintains feasibility for the inventory constraint (3.7) (while accounting for any reservation quantities, w_{ij}).

Algorithm 4 Allocation Procedure

1: Input $\mathbf{I}^-(t), \mathbf{B}^-(t), \mathbf{B}^I(t)$, and $w_{ij}, \forall i, j$
2: Initialize $\mathbf{z} \leftarrow \mathbf{0}$
3: **for** $\forall j$ **do**

$$z_j = \min\left((B_j^-(t) - B_j^u(t))^+, \min_{i \in \mathcal{N}(j)} \left\{ \left\lfloor \frac{\left(I_i^-(t) - \sum_{j' \neq j} a_{ij'} z_{j'} - w_{ij}\right)^+}{a_{ij}} \right\rfloor \right\} \right) \tag{3.26}$$

4: **end for**
5: $\mathbf{z}(t) \leftarrow \mathbf{z}$

Every time there is a demand arrival or inventory delivery (i.e. any time $\mathbf{I}^-(t)$ or $\mathbf{B}^-(t)$ changes) we re-solve (3.24) to set new backlog targets and run Algorithm 4 to determine the backlogged demand to fill, $\mathbf{z}(t)$. It is clear that $\mathbf{z}(t)$ output by Algorithm 4 is feasible, as z_j can only take non-negative integer values, and (3.26) guarantees feasibility for constraints (3.7) and (3.8). The next Lemma characterizes a property of the allocation given by Algorithm 4 that will be useful in our later proofs. We note that this property is similar to the Allocation Principle defined in Reiman and Wang (2015).

Lemma 3.1. *For any time t, Algorithm 4 produces an allocation such that, for each j, at least one of the following conditions hold:*

$$z_j(t) \geq B_j^-(t) - B_j^u(t), \text{ or}$$
$$\exists i \in \mathcal{N}(j) \text{ s.t. } \sum_{j'} a_{ij'} z_{j'}(t) \geq I_i^-(t) + 1 - a_{ij} - w_{ij}$$

Proof. If $z_j(t) \geq B_j^-(t) - B_j^u(t)$ then the claim holds, so assume $z_j(t) < B_j^-(t) - B_j^u(t)$. Then by (3.26), there exists $i \in \mathcal{N}(j)$ such that

$$z_j(t) = \left\lfloor \frac{\left(I_i^-(t) - \sum_{j' \neq j} a_{ij'} z_{j'} - w_{ij}\right)^+}{a_{ij}} \right\rfloor \geq \frac{\left(I_i^-(t) - \sum_{j' \neq j} a_{ij'} z_{j'} - w_{ij}\right)^+ - a_{ij} + 1}{a_{ij}},$$

where the inequality follows from the fact that, for any two integers, $m \geq 0$ and $n > 0$, we have $\lfloor m/n \rfloor \geq (m - n + 1)/n$. This inequality implies that $a_{ij} z_j(t) + \sum_{j' \neq j} a_{ij'} z_{j'} \geq I_i^-(t) + 1 - a_{ij} - w_{ij}$. Since the allocations only increase during the algorithm, i.e., $z_{j'}(t) \geq z_{j'}$, this proves the claim. □

Asymptotic analysis. To demonstrate asymptotic optimality, we first restate Theorem 3 of Reiman and Wang (2015), which shows that the scaled optimal cost of the continuous versions of (3.15) and (3.17) converge to the same positive constant. Since we will be discussing a system as the lead time L grows large, we introduce notation for different aspects of the system with a given lead time. For the demand and backlog processes, we use a superscript of (L) to denote the process in the system with lead time L; for example, as defined above, we let $\mathbf{D}^{(L)}$ denote the distribution of demand arriving over a lead time of length L. For costs, we use a subscript of L, e.g., we let C_L^+ and $\underset{\sim}{C}_L$ denote the optimal values of (3.15) and (3.17), respectively, when demand is distributed according to $\mathbf{D}^{(L)}$.

Theorem 3.3 (Reiman and Wang, 2015)**.** *There exists $0 < C < \infty$ such that*

$$\lim_{L \to \infty} \frac{C_L^+}{\sqrt{L}} = \lim_{L \to \infty} \frac{\underset{\sim}{C}_L}{\sqrt{L}} = C.$$

The proof of Theorem 3.3 follows directly from the proof of Theorem 3 in Reiman and Wang (2015), since the statement regards only the programs (3.15) and (3.17), which are equivalent to the stochastic programs considered in the aforementioned paper.

Next we adapt the asymptotic analysis of Reiman and Wang (2015) to our setting with integer constraints. For the system with lead time L and demand distributed according to $\mathbf{D}^{(L)}$, let C_L^α denote the long run average cost of the policy with base stock levels (3.23) and which allocates components according to Algorithm 4. Further, let R_L^+ and $\underaccent{\tilde}{R}_L$ be the objective values corresponding to the rounded solutions of (3.15) and (3.17), respectively, both of which are rounded according to (2.48) and (2.49). Further, let $R_L^\alpha = \alpha \underaccent{\tilde}{R}_L + (1-\alpha) R_L^+$ denote the convex combination of these costs, weighted by the same α used to determine the basestock levels in (3.23). The following lemma provides a basis for establishing asymptotic optimality by comparing the cost of our policy directly to R_L^α.

Lemma 3.2. *If* $\lim_{L\to\infty} C_L^\alpha / R_L^\alpha$ *exists, then we have*

$$\lim_{L\to\infty} \frac{C_L^\alpha}{\underaccent{\tilde}{C}_L} = \lim_{L\to\infty} \frac{C_L^\alpha}{R_L^\alpha},$$

$$\lim_{L\to\infty} \frac{R_L^\alpha}{\sqrt{L}} = \lim_{L\to\infty} \frac{\underaccent{\tilde}{C}_L}{\sqrt{L}} = C.$$

Proof. First, we observe that Corollaries 3.2 and 3.3 imply

$$\lim_{L\to\infty} R_L^+/C_L^+ = 1, \quad \lim_{L\to\infty} \underaccent{\tilde}{R}_L/\underaccent{\tilde}{C}_L = 1$$

Then we claim that $\lim_{L\to\infty} R_L^\alpha/\underaccent{\tilde}{C}_L = 1$, from which the first limit in the lemma follows. To show this, note that

$$\lim_{L\to\infty} \frac{R_L^\alpha}{\underaccent{\tilde}{C}_L} = \alpha \lim_{L\to\infty} \frac{\underaccent{\tilde}{R}_L}{\underaccent{\tilde}{C}_L} + (1-\alpha) \lim_{L\to\infty} \frac{R_L^+}{\underaccent{\tilde}{C}_L},$$
$$= \alpha + (1-\alpha) \lim_{L\to\infty} \frac{R_L^+}{\underaccent{\tilde}{C}_L},$$

so that to prove the claim we must show $\lim_{L\to\infty} R_L^+/\underaccent{\tilde}{C}_L = 1$, which follows since

$$\lim_{L\to\infty} \frac{R_L^+}{\underaccent{\tilde}{C}_L} = \left(\lim_{L\to\infty} \frac{R_L^+}{C_L^+}\right)\left(\lim_{L\to\infty} \frac{C_L^+}{\sqrt{L}}\right)\left(\lim_{L\to\infty} \frac{\sqrt{L}}{\underaccent{\tilde}{C}_L}\right),$$
$$= 1,$$

where the final equality follows from Theorem 3.3. For the second limit in the lemma, note that

$$\begin{aligned}\lim_{L\to\infty}\frac{R_L^\alpha}{\sqrt{L}} &= \alpha\left(\lim_{L\to\infty}\frac{\underline{R}_L}{\underline{C}_L}\right)\left(\lim_{L\to\infty}\frac{\underline{C}_L}{\sqrt{L}}\right)+(1-\alpha)\left(\lim_{L\to\infty}\frac{R_L^+}{C_L^+}\right)\left(\lim_{L\to\infty}\frac{C_L^+}{\sqrt{L}}\right),\\ &= C,\\ &= \lim_{L\to\infty}\frac{\underline{C}_L}{\sqrt{L}},\end{aligned}$$

where the second equality follows from Corollaries 3.2 and 3.3, and Theorem 3.3, and the final equality from Theorem 3.3. □

With this result, we will show asymptotic optimality by demonstrating that C_L^α approaches R_L^α. Before pursuing this goal, we follow Reiman and Wang (2015) by first introducing several scaled quantities in order to facilitate discussion of the limits. For the demand and backlog processes, we let a superscript of (L) denote a quantity in the system with lead time L, a tilde (i.e., $\tilde{x}$) denotes a variable scaled by $\sqrt{L}$, and a hat (i.e., $\hat{x}$) denotes a scaled (by $\sqrt{L}$) and centered variable (by the appropriate mean). Accordingly, let

$$\begin{aligned}\tilde{\mathbf{d}}^{(L)}(t) &= \frac{\mathbf{d}^{(L)}(Lt)}{\sqrt{L}},\\ \tilde{\mathbf{D}}^{(L)}(t) &= \frac{\mathbf{D}^{(L)}(Lt)}{\sqrt{L}},\\ \hat{\mathbf{D}}^{(L)}(t) &= \frac{\mathbf{D}^{(L)}(Lt)-L\boldsymbol{\mu}}{\sqrt{L}},\\ \hat{\mathbf{D}}^{(L)}(t_1,t_2) &= \frac{\mathbf{D}^{(L)}(Lt_1,Lt_2)-L(t_2-t_1)\boldsymbol{\mu}}{\sqrt{L}},\ t_2>t_1\geq 0,\\ \tilde{\mathbf{B}}^{(L)}(t) &= \frac{\mathbf{B}^{(L)}(Lt)}{\sqrt{L}},\\ \tilde{\mathbf{B}}^{(L)*}(t) &= \frac{\mathbf{B}^{(L)*}(Lt)}{\sqrt{L}},\\ \tilde{\mathbf{B}}^{(L)I}(t) &= \frac{\mathbf{B}^{(L)I}(Lt)}{\sqrt{L}}\end{aligned}$$

Following the approach of Reiman and Wang (2015), our first step is to establish a sufficient condition in the following verification lemma. We note that this result is analogous to Lemma 6 in Reiman and Wang (2015), however the proof requires different techniques to handle the integrality constraints.

Lemma 3.3. *If a policy uses basestock levels defined in* (3.23) *and satisfies, for each product* j,

$$\lim_{L\to\infty} \sup_{t\geq 1} \left\{ |\mathbb{E}[\tilde{B}_j^{(L)}(t)] - \mathbb{E}[\tilde{B}_j^{(L)I}(t)]| \right\} = 0, \tag{3.27}$$

then, letting C_L^π *denote the policy's long run average cost, we have*

$$\lim_{L\to\infty} \frac{C_L^\pi - R_L^\alpha}{\sqrt{L}} = 0.$$

Proof. First, we claim that

$$\liminf_{L\to\infty} \frac{C_L^\pi - R_L^\alpha}{\sqrt{L}} \geq 0,$$

since otherwise there would exist $\epsilon > 0$ such that $(C_L^\pi - R_L^\alpha)/\sqrt{L} < -\epsilon$ for infinitely many L as $L \to \infty$. But this contradicts Lemma 3.2, since $\lim_{L\to\infty} R_L^\alpha/\sqrt{L} = \lim_{L\to\infty} \underline{C}_L/\sqrt{L}$ and by Theorem 3.2, $C_L^\pi \geq \underline{C}_L$. The remainder of the proof focuses on the limit superior. In order to bound R_L^α, define the following quantity

$$S_L^\alpha = \mathbf{c}\cdot\mathbf{r}^{(L)\alpha} + \mathbb{E}[\mathbf{p}\cdot\lceil\hat{\mathbf{y}}(\mathbf{r}^{(L)\alpha}, \mathbf{D}^{(L)})\rceil] - \mathbf{c}\cdot A\mathbb{E}[\mathbf{D}^{(L)}], \tag{3.28}$$

$$\text{where } \hat{\mathbf{y}}(\mathbf{r}^{(L)\alpha}, \mathbf{D}^{(L)}) = \text{argmin}\left\{ \|\mathbf{y}\| \,\middle|\, \mathbf{y} \in \underset{\mathbf{y}\geq 0}{\text{argmin}}\{\mathbf{p}\cdot\mathbf{y} | A\mathbf{y} \geq A\mathbf{D}^{(L)} - \mathbf{r}^{(L)\alpha}\} \right\}. \tag{3.29}$$

Let $\mathbf{1}^c$ and $\mathbf{1}^p$ denote vectors of ones of lengths M and N, respectively, i.e., the number of components and products. Then we claim that

$$S_L^\alpha \leq R_L^\alpha + \mathbf{c}\cdot\mathbf{1}^c + 2\mathbf{p}\cdot\mathbf{1}^p.$$

To demonstrate this, first note that

$$\begin{aligned}
\mathbf{c}\cdot\mathbf{r}^{(L)\alpha} &\leq \alpha\mathbf{c}\cdot\mathbf{r}^{(L)*} + (1-\alpha)\mathbf{c}\cdot\mathbf{r}^{(L)o},\\
&< \alpha\mathbf{c}\cdot(\lfloor\mathbf{r}^{(L)*}\rfloor + \mathbf{1}^c) + (1-\alpha)\mathbf{c}\cdot(\lfloor\mathbf{r}^{(L)o}\rfloor + \mathbf{1}^c),\\
&= \alpha\mathbf{c}\cdot\lfloor\mathbf{r}^{(L)*}\rfloor + (1-\alpha)\mathbf{c}\cdot\lfloor\mathbf{r}^{(L)o}\rfloor + \mathbf{c}\cdot\mathbf{1}^c.
\end{aligned}$$

Next, for a given realization of $\mathbf{D}^{(L)}$, let $\mathbf{y}^0(\mathbf{D}^{(L)})$ and $\mathbf{y}^*(\mathbf{D}^{(L)})$ be the optimal shortages for that demand scenario in (3.15) and (3.17), respectively. From their feasibility for the first constraint in those programs, we have

$$
\begin{aligned}
&A(\alpha \mathbf{y}^*(\mathbf{D}^{(L)}) + (1-\alpha)\mathbf{y}^0(\mathbf{D}^{(L)}) + \mathbf{1}^p) \\
&= \alpha A\mathbf{y}^*(\mathbf{D}^{(L)}) + (1-\alpha)A\mathbf{y}^0(\mathbf{D}^{(L)}) + A\mathbf{1}^p, \\
&\geq \alpha(A\mathbf{D}^{(L)} - \mathbf{r}^{(L)*}) + (1-\alpha)(A\mathbf{D}^{(L)} - \mathbf{r}^{(L)o}) + \mathbf{1}^c, \\
&= A\mathbf{D}^{(L)} - (\alpha\mathbf{r}^{(L)*} + (1-\alpha)\mathbf{r}^{(L)o} - \mathbf{1}^c), \\
&\geq A\mathbf{D}^{(L)} - \mathbf{r}^{(L)\alpha},
\end{aligned}
$$

so that $\alpha\mathbf{y}^*(\mathbf{D}^{(L)}) + (1-\alpha)\mathbf{y}^0(\mathbf{D}^{(L)}) + \mathbf{1}^p$ is feasible for the LP $\min_{\mathbf{y}\geq 0}\{\mathbf{p}\cdot\mathbf{y}|A\mathbf{y} \geq A\mathbf{D}^{(L)} - \mathbf{r}^{(L)\alpha}\}$ from the definition of $\hat{\mathbf{y}}(\mathbf{r}^{(L)\alpha}, \mathbf{D}^{(L)})$, which implies

$$
\begin{aligned}
\mathbf{p}\cdot\lceil\hat{\mathbf{y}}(\mathbf{r}^{(L)\alpha}, \mathbf{D}^{(L)})\rceil &\leq \mathbf{p}\cdot\hat{\mathbf{y}}(\mathbf{r}^{(L)\alpha}, \mathbf{D}^{(L)}) + \mathbf{p}\cdot\mathbf{1}^p, \\
&\leq \mathbf{p}\cdot(\alpha\mathbf{y}^*(\mathbf{D}^{(L)}) + (1-\alpha)\mathbf{y}^0(\mathbf{D}^{(L)}) + \mathbf{1}^p) + \mathbf{p}\cdot\mathbf{1}^p, \\
&\leq \alpha\mathbf{p}\cdot\lceil\mathbf{y}^*(\mathbf{D}^{(L)})\rceil + (1-\alpha)\mathbf{p}\cdot\lceil\mathbf{y}^0(\mathbf{D}^{(L)})\rceil + 2\mathbf{p}\cdot\mathbf{1}^p.
\end{aligned}
$$

Therefore, we have

$$
\begin{aligned}
S_L^\alpha &\leq \alpha\left(\mathbf{c}\cdot\lfloor\mathbf{r}^{(L)*}\rfloor + \mathbb{E}[\mathbf{p}\cdot\lceil\mathbf{y}^*(\mathbf{D}^{(L)})\rceil] - \mathbf{c}\cdot A\mathbb{E}[\mathbf{D}^{(L)}]\right) \\
&\quad + (1-\alpha)\left(\mathbf{c}\cdot\lfloor\mathbf{r}^{(L)o}\rfloor + \mathbb{E}[\mathbf{p}\cdot\lceil\mathbf{y}^0(\mathbf{D}^{(L)})\rceil] - \mathbf{c}\cdot A\mathbb{E}[\mathbf{D}^{(L)}]\right) \\
&\quad + \mathbf{c}\cdot\mathbf{1}^c + 2\mathbf{p}\cdot\mathbf{1}^p, \\
&= \alpha\underline{R}_L + (1-\alpha)R_L^+ + \mathbf{c}\cdot\mathbf{1}^c + 2\mathbf{p}\cdot\mathbf{1}^p, \\
&= R_L^\alpha + \mathbf{c}\cdot\mathbf{1}^c + 2\mathbf{p}\cdot\mathbf{1}^p.
\end{aligned}
$$

Next we consider the long run average cost of the policy, C_L^π. By (3.21) we have

$$
\begin{aligned}
\mathbb{E}[\mathbf{b}\cdot\mathbf{B}^{(L)}(Lt) + \mathbf{c}\cdot\mathbf{I}^{(L)}(Lt)] &= \mathbb{E}[\mathbf{b}\cdot\mathbf{B}^{(L)}(Lt) + \mathbf{c}\cdot(\mathbf{r}^{(L)\alpha} + A\mathbf{B}^{(L)}(Lt) - A\mathbf{D}^{(L)}(Lt))], \\
&= \mathbb{E}[\mathbf{p}\cdot\mathbf{B}^{(L)}(Lt)] + \mathbf{c}\cdot\mathbf{r}^{(L)\alpha} - \mathbf{c}\cdot A\mathbb{E}[\mathbf{D}^{(L)}].
\end{aligned}
$$

Therefore,

$$
C_L^\pi = \mathbf{c}\cdot\mathbf{r}^{(L)\alpha} - \mathbf{c}\cdot A\mathbb{E}[\mathbf{D}^{(L)}] + \limsup_{T\to\infty}\frac{1}{T}\int_1^{T+1}\mathbb{E}[\mathbf{p}\cdot\mathbf{B}^{(L)}(Lt)]dt,
$$

so that

$$
\begin{aligned}
& C_L^\pi - R_L^\alpha \\
& \leq \mathbf{c} \cdot \mathbf{r}^{(L)\alpha} - \mathbf{c} \cdot A\mathbb{E}[\mathbf{D}^{(L)}] - S_L^\alpha + \mathbf{c} \cdot \mathbf{1}^c + 2\mathbf{p} \cdot \mathbf{1}^p + \limsup_{T\to\infty} \frac{1}{T} \int_1^{T+1} \mathbb{E}[\mathbf{p} \cdot \mathbf{B}^{(L)}(Lt)]dt, \\
& = \mathbf{c} \cdot \mathbf{1}^c + 2\mathbf{p} \cdot \mathbf{1}^p - \mathbb{E}[\mathbf{p} \cdot \lceil \hat{\mathbf{y}}(\mathbf{r}^{(L)\alpha}, \mathbf{D}^{(L)}) \rceil] + \limsup_{T\to\infty} \frac{1}{T} \int_1^{T+1} \mathbb{E}[\mathbf{p} \cdot \mathbf{B}^{(L)}(Lt)]dt, \\
& = \mathbf{c} \cdot \mathbf{1}^c + 2\mathbf{p} \cdot \mathbf{1}^p + \limsup_{T\to\infty} \frac{1}{T} \int_1^{T+1} \left(\mathbb{E}[\mathbf{p} \cdot \mathbf{B}^{(L)}(Lt)] - \mathbb{E}[\mathbf{p} \cdot \lceil \hat{\mathbf{y}}(\mathbf{r}^{(L)\alpha}, \mathbf{D}^{(L)}) \rceil] \right) dt,
\end{aligned}
$$

Now, since π uses a basestock replenishment policy, by (3.22), the component shortage process obeys, for $t \geq 1$

$$\mathbf{Q}^{(L)}(Lt) = A\mathbf{D}^{(L)}(Lt) - \mathbf{r}^{(L)\alpha},$$

which has the same distribution as $A\mathbf{D}^{(L)} - \mathbf{r}^{(L)\alpha}$ due to the memoryless property of the Poisson process. Thus, by the definition of $\mathbf{B}^I(t)$ from (3.24-3.25), and the definition of $\hat{\mathbf{y}}(\mathbf{r}^{(L)\alpha}, \mathbf{D}^{(L)})$ in (3.29), this equivalence in distributions implies that, for all $t \geq 1$

$$\mathbb{E}[\mathbf{p} \cdot \mathbf{B}^{(L)I}(Lt)] = \mathbb{E}[\mathbf{p} \cdot \lceil \hat{\mathbf{y}}(\mathbf{r}^{(L)\alpha}, \mathbf{D}^{(L)}) \rceil].$$

Thus, we have shown that

$$\frac{C_L^\pi - R_L^\alpha}{\sqrt{L}} \leq \frac{\mathbf{c} \cdot \mathbf{1}^c + 2\mathbf{p} \cdot \mathbf{1}^p}{\sqrt{L}} + \limsup_{T\to\infty} \frac{1}{T} \int_1^{T+1} \mathbf{p} \cdot \left(\frac{\mathbb{E}[\mathbf{B}^{(L)}(Lt)] - \mathbb{E}[\mathbf{B}^{(L)I}(Lt)]}{\sqrt{L}} \right) dt,$$

the right-hand side of which converges to zero as $L \to \infty$, by condition (3.27) assumed in the lemma. □

Now we will show that our policy indeed achieves the sufficient condition of Lemma 3.3. In order to do so, we restate Lemmas 2 and 3 from Reiman and Wang (2015), giving bounds on extreme values of the demand process. The proofs of these lemmas follow directly from Reiman and Wang (2015), as they depend only on the details of the demand process, and not on the integrality constraints of a feasible policy.

Lemma 3.4 (Reiman and Wang, 2015)**.** *For each product j,*

$$\mathbb{E}\left[\sup_{t-1 \leq \tau \leq t} \tilde{d}_j^{(L)}(\tau) \right] \leq 3\lambda^{1/(2+\delta)}(1+\eta_j)L^{-\delta/(2(2+\delta))}.$$

Lemma 3.5 (Reiman and Wang, 2015)**.** *For each product j, and any positive constant κ,*

$$\mathbb{E}\left[\sup_{0\le\tau\le L^{-1/4}}|\hat{D}_j^{(L)}(0,\tau)|\right]\le(1+\sigma_{jj}^2)L^{-1/8},$$

$$\mathbb{E}\left[\sup_{L^{-1/4}\le\tau\le 1}\left(|\hat{D}_j^{(L)}(0,\tau)|-\sqrt{L}\tau\kappa\right)^+\right]\le\frac{\sigma_{jj}^2}{\kappa}L^{-1/4},$$

Next we restate Lemma 4 from Reiman and Wang (2015) which bounds the absolute value of the scaled and centered basestock levels for any lead time length. We make a slight adjustment to the version in Reiman and Wang (2015), by requiring $L\ge 1$ instead of $L>0$ which does not impact the asymptotic analysis as $L\to\infty$. For completeness we include a short proof, which follows directly from Lemma 4 in Reiman and Wang (2015). We note that the original Lemma 4 in Reiman and Wang (2015) holds in our setting because it deals only with the optimal solutions of (3.15) and (3.17).

Lemma 3.6. *There exists a finite constant K such that for all lead times $L\ge 1$, and all components i,*

$$\frac{|r_i^{(L)\alpha}-L\sum_j a_{ij}\mu_j|}{\sqrt{L}}\le K.$$

Proof. By Lemma 4 from Reiman and Wang (2015) there exists a finite constant K' such that for all lead times $L>0$ and all i,

$$\frac{|\alpha r_i^{(L)*}+(1-\alpha)r_i^{(L)0}-L\sum_j a_{ij}\mu_j|}{\sqrt{L}}\le K',$$

so that for $L\ge 1$ we have

$$\begin{aligned}\frac{|r_i^{(L)\alpha}-L\sum_j a_{ij}\mu_j|}{\sqrt{L}}&\le\frac{|\alpha r_i^{(L)*}+(1-\alpha)r_i^{(L)0}-L\sum_j a_{ij}\mu_j|+1}{\sqrt{L}},\\&\le K'+1,\end{aligned}$$

and let $K=K'+1$. □

Next we also restate Lemma 5 from Reiman and Wang (2015), which bounds the backlog target process in terms of the demand process. Again, the proof of this result directly follows from Reiman and Wang (2015), since our definition of the backlog target (3.24) (i.e., before rounding) is identical to the definition given in equations (19-20) of the aforementioned paper. Further, the proof relies only on the Lipschitz continuity of the minimum norm selection in (3.24), and the fact that a basestock policy is used. We will use this result to bound the rounded backlog targets, (3.25), that we use in our integer allocation policy.

Lemma 3.7 (Reiman and Wang, 2015)**.** *There exists a constant g that depends only on A and* $\mathbf{p}$ *such that for any $t_2 > t_1 \geq 1$, any $t \geq 1$, and each product j*

$$|\tilde{B}_j^{(L)*}(t_2) - \tilde{B}_j^{(L)*}(t_1)| \leq g \sum_l |\hat{D}_l^{(L)}(t_1, t_2) - \hat{D}_l^{(L)}(t_1 - 1, t_2 - 1)|,$$

$$|\tilde{B}_j^{(L)*}(t) - \tilde{B}_j^{(L)*}(t^-)| \leq g \sum_l |\tilde{d}_l^{(L)}(t) - \tilde{d}_l^{(L)}(t-1)|.$$

With this result we are ready to state and prove the asymptotic optimality result. Our proof follows a similar strategy to the proof of the analogous result in Reiman and Wang (2015), who bound backlog levels in the system in terms of demand, but we require care to account for the integrality constraint in several places. We require that the reservation levels w_{ij} from Algorithm 4 vanish asymptotically,

$$\lim_{L \to \infty} \frac{w_{ij}^{(L)}}{\sqrt{L}} = 0, \ \forall i, j. \tag{3.30}$$

Proposition 3.3. *A policy with base stock levels* (3.23)*, allocations made according to Algorithm 4, and reservation levels satisfying* (3.30)*, is asymptotically optimal as $L \to \infty$, i.e.,*

$$\lim_{L \to \infty} \frac{C_L^\alpha}{\underline{C}_L} = 1.$$

Proof. We prove the claim by demonstrating that condition (3.27) in Lemma 3.3 holds for the policy described. This is sufficient, since by Lemma 3.3 we then have $\lim_{L\to\infty}(C_L^\alpha - R_L^\alpha)/\sqrt{L} = 0$, which together with Lemma 3.2 gives

$$\begin{aligned} \lim_{L \to \infty} \frac{C_L^\alpha}{\underline{C}_L} &= \lim_{L \to \infty} \frac{C_L^\alpha}{R_L^\alpha}, \\ &= \lim_{L \to \infty} \left(\frac{\sqrt{L}}{R_L^\alpha} \right) \frac{C_L^\alpha - R_L^\alpha}{\sqrt{L}} + 1, \\ &= 1. \end{aligned}$$

In order to verify condition (3.27), for a given L and $t \geq 1$ define

$$\mathscr{S}_L^+(t) = \left\{ j \,\middle|\, \tilde{B}_j^{(L)}(t) > \tilde{B}_j^{(L)I}(t) \right\},$$

$$\mathscr{S}_L^-(t) = \left\{ j \,\middle|\, \tilde{B}_j^{(L)}(t) < \tilde{B}_j^{(L)I}(t) \right\},$$

as the sets of products, j, which have excess and deficit, respectively, of their current backlog level, $\tilde{B}_j^{(L)}(t)$, relative to their target, $\tilde{B}_j^{(L)I}(t)$.

Step 1: Bounding Backlog Excess by Deficit. We first demonstrate that for any product $j \in \mathscr{S}_L^+(t)$, its deviation from the backlog target can be bound in terms of the sum of deviations from the backlog target for all products in $j \in \mathscr{S}_L^-(t)$, i.e., a product can only have much backlog excess if other products have much backlog deficit. Then the result follows by a bound on the backlog deficit.

Consider some product $j' \in \mathscr{S}_L^+(t)$, and we claim that there exists some component $i' \in \mathcal{N}(j')$ such that

$$\begin{aligned} &\sum_{j \in \mathscr{S}_L^+(t)} a_{i'j} \left(B_j^{(L)}(Lt) - B_j^{(L)I}(Lt) \right) \\ &\leq \sum_{j \in \mathscr{S}_L^-(t)} a_{i'j} \left(B_j^{(L)I}(Lt) - B_j^{(L)}(Lt) \right) - (1 - a_{i'j'} - w_{i'j'}^{(L)}). \end{aligned} \tag{3.31}$$

To show this, we assume it is not true and demonstrate a contradiction with the allocation for product j' in Algorithm 4. Accordingly, assume that for all components $i' \in \mathcal{N}(j')$, the opposite inequality holds for (3.31), which, by the backlog relation (3.5), is equivalent to

$$\sum_j a_{i'j} z_j^{(L)}(Lt) < \sum_j a_{i'j} \left(B_j^{(L)-}(Lt) - B_j^{(L)I}(Lt) \right) + 1 - a_{i'j'} - w_{i'j'}^{(L)}. \tag{3.32}$$

Then, note that the definition of the backlog targets in (3.24) and (3.25) implies that

$$\begin{aligned} A\mathbf{B}^{(L)I}(Lt) &\geq A\mathbf{B}^{(L)*}(Lt), \\ &\geq \mathbf{Q}^{(L)}(Lt), \\ &= A\mathbf{B}^{(L)-}(Lt) - \mathbf{I}^{(L)-}(Lt), \end{aligned}$$

where the final equality follows from the definition (3.6). This implies that for component i',

$$\sum_j a_{i'j} \left(B_j^{(L)-}(Lt) - B_j^{(L)I}(Lt) \right) \leq I_{i'}^{(L)-}(Lt),$$

which combined with (3.32) gives, for all $i' \in \mathcal{N}(j')$,

$$\sum_j a_{i'j} z_j^{(L)}(Lt) < I_{i'}^{(L)-}(Lt) + 1 - a_{i'j'} - w_{i'j'}^{(L)}. \tag{3.33}$$

Now, by the by the backlog relation (3.5), and the definition of $\mathscr{S}_L^+(t)$, we have

$$\begin{aligned} z_{j'}^{(L)}(Lt) &= B_j^{(L)-}(Lt) - B_j^{(L)}(Lt), \\ &< B_j^{(L)-}(Lt) - B_j^{(L)I}(Lt), \end{aligned}$$

which cannot hold at the same time as (3.33) by Lemma 3.1. Thus, (3.31) holds and we note that by definition of $\mathscr{S}_L^+(t)$, the left-hand side is strictly positive. Thus, scaling by $\sqrt{L}$, we see that

$$\begin{aligned} &\frac{1 - a_{i'j'} - w_{i'j'}^{(L)}}{\sqrt{L}} < a_{i'j'} \left(\tilde{B}_{j'}^{(L)}(t) - \tilde{B}_{j'}^{(L)I}(t) \right) + \frac{1 - a_{i'j'} - w_{i'j'}^{(L)}}{\sqrt{L}} \\ &\leq \sum_{j \in \mathscr{S}_L^-(t)} a_{i'j} \left(\tilde{B}_j^{(L)I}(t) - \tilde{B}_j^{(L)}(t) \right). \end{aligned}$$

By (3.30) we have $\lim_{L\to\infty}(1 - a_{i'j'} - w_{i'j'}^{(L)})/\sqrt{L} = 0$, and it is clear that (3.27) will hold if we demonstrate that for any $\epsilon > 0$, we can choose L large enough so that for all $t \geq 1$ and $j \in \mathscr{S}_L^-(t)$ we have

$$\mathbb{E}[\tilde{B}_j^{(L)I}(t) - \tilde{B}_j^{(L)}(t)] < \epsilon. \tag{3.34}$$

To show this, for $t \geq 1$ and $j \in \mathscr{S}_L^-(t)$, define

$$t_j^{(L)} = \sup \left\{ \tau \,\middle|\, 0 \leq \tau \leq t \text{ and } \tilde{B}_j^{(L)I}(\tau) < \tilde{B}_j^{(L)-}(\tau) \right\}.$$

Then, we can write the expectation of interest as

$$\begin{aligned} &\mathbb{E}[\tilde{B}_j^{(L)I}(t) - \tilde{B}_j^{(L)}(t)] \\ &= \mathbb{E}\left[\left(\tilde{B}_j^{(L)I}(t) - \tilde{B}_j^{(L)}(t) \right) \mathbb{1}_{\{t_j^{(L)} < t-1\}} \right] + \mathbb{E}\left[\left(\tilde{B}_j^{(L)I}(t) - \tilde{B}_j^{(L)}(t) \right) \mathbb{1}_{\{t_j^{(L)} \geq t-1\}} \right], \end{aligned}$$

which allows us to establish the bound in two separate cases.

Step 2: Bounding Backlog Deficit Lasting Longer than a Lead Time. Consider the case when $t_j^{(L)} < t - 1$, which implies that $\tilde{B}_j^{(L)-}(\tau) - \tilde{B}_j^{(L)I}(\tau) \leq 0$ for all $\tau \in [L(t-1), Lt]$, so by the first argument of the minimization in (3.26) in Algorithm 4, no backlog for product j will be served during this lead time. Thus, the total backlog at time Lt must be at least as large as the demand arriving during this lead time,

$$\tilde{B}_j^{(L)}(t) \geq \tilde{D}_j^{(L)}(t) = \hat{D}_j^{(L)}(t) + \sqrt{L}\mu_j. \tag{3.35}$$

Next we observe that there must exist a component i' with $a_{i'j} > 0$ such that

$$a_{i'j} B_j^{(L)*}(Lt) \leq \left(Q_{i'}^{(L)}(Lt)\right)^+,$$

since otherwise we could strictly decrease $B_j^{(L)*}(Lt)$ without violating any constraint in the LP $\min_{\mathbf{B}\geq 0}\{\mathbf{p}\cdot\mathbf{B}|A\mathbf{B} \geq \mathbf{Q}^{(L)}(Lt)\}$, which violates the definition of $\mathbf{B}^{(L)*}(Lt)$ in (3.24). Further, by the component shortage relation (3.22), we have

$$Q_{i'}^{(L)}(Lt) = \sum_{l=1}^{N} a_{i'l} D_l^{(L)}(Lt) - r_{i'}^{(L)\alpha},$$

so that

$$\begin{aligned} a_{i'j} B_j^{(L)I}(Lt) &< a_{i'j}\left(B_j^{(L)*}(Lt) + 1\right), \\ &\leq \left(\sum_{l=1}^{N} a_{i'l} D_l^{(L)}(Lt) - r_{i'}^{(L)\alpha}\right)^+ + a_{i'j}. \end{aligned}$$

Thus, scaling all terms, and centering the demand and base stock level, we have

$$\begin{aligned} a_{i'j}\tilde{B}_j^{(L)I}(t) &\leq \left(\sum_{l=1}^{N} a_{i'l}\hat{D}_l^{(L)}(t) - \frac{r_{i'}^{(L)\alpha} - L\sum_{l=1}^{N} a_{i'l}\mu_j}{\sqrt{L}}\right)^+ + \frac{a_{i'j}}{\sqrt{L}}, \\ &\leq \bar{a}\sum_{l=1}^{N} |\hat{D}_l^{(L)}(t)| + K + \frac{a_{i'j}}{\sqrt{L}}, \end{aligned}$$

where the last inequality follows from Lemma 3.6. Let $\bar{\beta} = \bar{a}/\underline{a}$ and $\bar{\kappa} = K/\underline{a}$, so that

$$\tilde{B}_j^{(L)I}(t) \leq \bar{\beta}\sum_{l=1}^{N} |\hat{D}_l^{(L)}(t)| + \bar{\kappa} + L^{-1/2}.$$

Together with (3.35) this implies that

$$
\begin{aligned}
&\mathbb{E}\left[\left(\tilde{B}_j^{(L)I}(t) - \tilde{B}_j^{(L)}(t)\right) \mathbb{1}_{\{t_j^{(L)} < t-1\}}\right] \\
&\leq \mathbb{E}\left[\left(\bar{\beta}\sum_{l=1}^{N} |\hat{D}_l^{(L)}(t)| + \bar{\kappa} + |\hat{D}_j^{(L)}(t)| - \sqrt{L}\mu_j\right)^+\right] + L^{-1/2}, \\
&\leq \mathbb{E}\left[\left((\bar{\beta}+1)\sum_{l=1}^{N} |\hat{D}_l^{(L)}(t)| - (\sqrt{L}\mu_j - \bar{\kappa})\right)^+\right] + L^{-1/2}, \\
&\leq (\bar{\beta}+1)\sum_{l=1}^{N} \mathbb{E}\left[\left(|\hat{D}_l^{(L)}(t)| - \frac{\sqrt{L}\mu_j - \bar{\kappa}}{N(\bar{\beta}+1)}\right)^+\right] + L^{-1/2}, \\
&\leq \frac{N(\bar{\beta}+1)^2}{\sqrt{L}\mu_j - \bar{\kappa}} \sum_{l=1}^{N} \sigma_{ll}^2 + L^{-1/2},
\end{aligned}
\tag{3.36}
$$

where the last inequality follows from bounding the expectation using Chebyshev's inequality.

Step 3: Bounding Backlog Deficit Lasting Less than a Lead Time. Now consider the case when $t_j^{(L)} \geq t-1$. First we claim that

$$
\begin{aligned}
\tilde{B}_j^{(L)I}(t) - \tilde{B}_j^{(L)}(t) \leq \left(\tilde{B}_j^{(L)I}(t) - \tilde{B}_j^{(L)I}(t_j^{(L)})\right) - \left(\tilde{B}_j^{(L)}(t) - \tilde{B}_j^{(L)}(t_j^{(L)})\right), \\
+ |\tilde{B}_j^{(L)I}(t_j^{(L)}) - \tilde{B}_j^{(L)I}(t_j^{(L)-})|.
\end{aligned}
\tag{3.37}
$$

To show this, we will consider two cases. First, however, note that at any time t', if $B_j^{(L)I}(t') < B_j^{(L)-}(t')$, then the allocation policy in Algorithm 4 implies that $B_j^{(L)I}(t') \leq B_j^{(L)}(t')$. This is because $(B_j^{(L)-}(t') - B_j^{(L)I}(t'))^+ = B_j^{(L)-}(t') - B_j^{(L)I}(t')$ and so by the allocation in (3.26) of Algorithm 4 we have $z_j(t') \leq B_j^{(L)-}(t') - B_j^{(L)I}(t')$, which is equivalent to $B_j^{(L)I}(t') \leq B_j^{(L)-}(t') - z_j(t') = B_j^{(L)}(t')$, where the last equality follows from the backlog relation (3.5).

Now, consider the case where backlog for product j is served at time $Lt_j^{(L)}$, i.e., $z_j(Lt_j^{(L)}) > 0$. Then by the allocation in (3.26) of Algorithm 4 we must have had $0 < B_j^{(L)-}(Lt_j^{(L)}) - B_j^{(L)I}(Lt_j^{(L)})$ (since otherwise we could never have set $z_j > 0$). Thus, we have $B_j^{(L)I}(Lt_j^{(L)}) < B_j^{(L)-}(Lt_j^{(L)})$, which we showed above implies $B_j^{(L)I}(Lt_j^{(L)}) \leq B_j^{(L)}(Lt_j^{(L)})$. When scaled, this inequality implies (3.37).

Otherwise, if no backlog for product j is served at time $Lt_j^{(L)}$, then by (3.19) we have $\tilde{B}_j^{(L)}(t_j^{(L)}) =$

$\tilde{B}_j^{(L)}(t_j^{(L)-}) + d_j^{(L)}(t_j^{(L)}) \geq \tilde{B}_j^{(L)}(t_j^{(L)-})$, i.e., the backlog level increases. Now consider the largest $\tau < t_j^{(L)}$ for which demand or inventory arrives at time $L\tau$. Then, since our policy only recalculates backlog targets and clears backlog at the time of such an arrival, we have $\tilde{B}_j^{(L)I}(t_j^{(L)-}) = \tilde{B}_j^{(L)I}(\tau)$ and $\tilde{B}_j^{(L)}(t_j^{(L)-}) = \tilde{B}_j^{(L)}(\tau)$. Further, by the definition of $t_j^{(L)}$, we must have $\tilde{B}_j^{(L)I}(\tau) < \tilde{B}_j^{(L)-}(\tau)$ (since the last arrival before $Lt_j^{(L)}$ occurs at $L\tau$ and these quantities do not change when there are no arrivals). Thus, as shown above, we have $\tilde{B}_j^{(L)I}(\tau) \leq \tilde{B}_j^{(L)}(\tau)$. Combining these relations, we have

$$\tilde{B}_j^{(L)I}(t_j^{(L)-}) \;=\; \tilde{B}_j^{(L)I}(\tau) \;\leq\; \tilde{B}_j^{(L)}(\tau) \;=\; \tilde{B}_j^{(L)}(t_j^{(L)-}) \;\leq\; \tilde{B}_j^{(L)}(t_j^{(L)}),$$

which implies that

$$\begin{aligned} 0 &\leq \tilde{B}_j^{(L)}(t_j^{(L)}) - \tilde{B}_j^{(L)I}(t_j^{(L)-}), \\ &= \tilde{B}_j^{(L)}(t_j^{(L)}) - \tilde{B}_j^{(L)I}(t_j^{(L)}) + \tilde{B}_j^{(L)I}(t_j^{(L)}) - \tilde{B}_j^{(L)I}(t_j^{(L)-}), \\ &\leq \tilde{B}_j^{(L)}(t_j^{(L)}) - \tilde{B}_j^{(L)I}(t_j^{(L)}) + |\tilde{B}_j^{(L)I}(t_j^{(L)}) - \tilde{B}_j^{(L)I}(t_j^{(L)-})|, \end{aligned}$$

which implies (3.37).

Now we will bound each term on the right-hand side of (3.37) separately. To do so, first note that by the rounding of the backlog targets in (3.25), for any two times t_1 and t_2,

$$|B_j^{(L)I}(Lt_2) - B_j^{(L)I}(Lt_1)| \leq |B_j^{(L)*}(Lt_2) - B_j^{(L)*}(Lt_1)| + 1,$$

so that for the scaled version we have

$$|\tilde{B}_j^{(L)I}(t_2) - \tilde{B}_j^{(L)I}(t_1)| \leq |\tilde{B}_j^{(L)*}(t_2) - \tilde{B}_j^{(L)*}(t_1)| + L^{-1/2}.$$

Thus, for the last term in (3.37), by Lemma 3.7, we have,

$$\begin{aligned} &|\tilde{B}_j^{(L)I}(t_j^{(L)}) - \tilde{B}_j^{(L)I}(t_j^{(L)-})| \\ &\leq g \sum_l |\tilde{d}_l^{(L)}(t_j^{(L)}) - \tilde{d}_l^{(L)}(t_j^{(L)} - 1)| + L^{-1/2}, \\ &\leq \sup_{t-1 \leq \tau \leq t} \left\{ g \sum_l |\tilde{d}_l^{(L)}(\tau) - \tilde{d}_l^{(L)}(\tau - 1)| \right\} + L^{-1/2}, \\ &\leq g \sum_{l=1}^{N} \left(\sup_{t-1 \leq \tau_l \leq t} \left\{ \tilde{d}_l^{(L)}(\tau_l) \right\} + \sup_{t-2 \leq \tau_l' \leq t-1} \left\{ \tilde{d}_l^{(L)}(\tau_l') \right\} \right) + L^{-1/2}, \end{aligned}$$

where the second inequality follows from $t-1 \leq t_j^{(L)} \leq t$, and the third inequality from the fact that demand is non-negative. Thus, we can bound the expected value using Lemma 3.4,

$$\begin{aligned}
&\mathbb{E}\left[|\tilde{B}_j^{(L)I}(t_j^{(L)}) - \tilde{B}_j^{(L)I}(t_j^{(L)-})|\right] \\
&\leq g \sum_{l=1}^{N} \mathbb{E}\left[\sup_{t-1\leq\tau_l\leq t}\left\{\tilde{d}_l^{(L)}(\tau_l)\right\} + \sup_{t-2\leq\tau_l'\leq t-1}\left\{\tilde{d}_l^{(L)}(\tau_l')\right\}\right] + L^{-1/2}, \\
&\leq 6g\lambda^{1/(2+\delta)}L^{-\delta/(2(2+\delta))}\sum_{l=1}^{N}(1+\eta_j) + L^{-1/2}.
\end{aligned} \tag{3.38}$$

For the first term of (3.37), by Lemma 3.7, we have

$$\begin{aligned}
\tilde{B}_j^{(L)I}(t) - \tilde{B}_j^{(L)I}(t_j^{(L)}) &\leq |\tilde{B}_j^{(L)*}(t) - \tilde{B}_j^{(L)*}(t_j^{(L)})| + L^{-1/2}, \\
&\leq g\sum_{l}|\hat{D}_l^{(L)}(t_j^{(L)}, t) - \hat{D}_l^{(L)}(t_j^{(L)}-1, t-1)| + L^{-1/2}.
\end{aligned}$$

For the second term of (3.37), note that no backlog of product j is served between $t_j^{(L)}$ and t, so we have

$$\begin{aligned}
\tilde{B}_j^{(L)}(t) - \tilde{B}_j^{(L)}(t_j^{(L)}) &\geq \tilde{D}_j^{(L)}(t_j^{(L)}, t), \\
&= \hat{D}_j^{(L)}(t_j^{(L)}, t) + \sqrt{L}\mu_j(t - t_j^{(L)}).
\end{aligned}$$

Now let $\tau_t = t - t_j^{(L)}$, and since $t-1 \leq t_j^{(L)} \leq t$, we have $0 \leq \tau_t \leq 1$. Further due to the memoryless property of the Poisson process, for all products l, we have the following equivalences in distribution

$$\begin{aligned}
\hat{D}_l^{(L)}(t_j^{(L)}, t) &\overset{d}{=} \hat{D}_l^{(L)}(1, 1+\tau_t), \\
\hat{D}_l^{(L)}(t_j^{(L)}-1, t-1) &\overset{d}{=} \hat{D}_l^{(L)}(0, \tau_t).
\end{aligned}$$

Thus we can bound the first two terms of (3.37) by

$$\begin{aligned}
&\mathbb{E}\left[\left(\tilde{B}_j^{(L)I}(t) - \tilde{B}_j^{(L)I}(t_j^{(L)})\right) - \left(\tilde{B}_j^{(L)}(t) - \tilde{B}_j^{(L)}(t_j^{(L)})\right)\right], \\
&\leq \mathbb{E}\left[\left(g\sum_{l=1}^{N}|\hat{D}_l^{(L)}(1, 1+\tau_t) - \hat{D}_l^{(L)}(0,\tau_t)| + |\hat{D}_j^{(L)}(1, 1+\tau_t)| - \sqrt{L}\mu_j\tau_t\right)^+\right] + L^{-1/2}, \\
&\leq \mathbb{E}\left[\left(g\sum_{l=1}^{N}|\hat{D}_l^{(L)}(1, 1+\tau_t)| + g\sum_{l=1}^{N}|\hat{D}_l^{(L)}(0,\tau_t)| + |\hat{D}_j^{(L)}(1, 1+\tau_t)| - \sqrt{L}\mu_j\tau_t\right)^+\right] \\
&\quad + L^{-1/2}.
\end{aligned}$$

Let $\kappa = \mu_j/(2Ng+1)$ so that,

$$\left(g\sum_{l=1}^{N} |\hat{D}_l^{(L)}(1,1+\tau_t)| + g\sum_{l=1}^{N} |\hat{D}_l^{(L)}(0,\tau_t)| + |\hat{D}_j^{(L)}(1,1+\tau_t)| - \sqrt{L}\mu_j\tau_t \right)^+ ,$$

$$= \left(g\sum_{l=1}^{N} \left(|\hat{D}_l^{(L)}(1,1+\tau_t)| - \sqrt{L}\tau_t\kappa \right) \right.$$

$$\left. + g\sum_{l=1}^{N} \left(|\hat{D}_l^{(L)}(0,\tau_t)| - \sqrt{L}\tau_t\kappa \right) + |\hat{D}_j^{(L)}(1,1+\tau_t)| - \sqrt{L}\tau_t\kappa \right)^+ ,$$

$$\leq (g+1)\sum_{l=1}^{N} \left(|\hat{D}_l^{(L)}(1,1+\tau_t)| - \sqrt{L}\tau_t\kappa \right)^+ + g\sum_{l=1}^{N} \left(|\hat{D}_l^{(L)}(0,\tau_t)| - \sqrt{L}\tau_t\kappa \right)^+ ,$$

$$\leq \sup_{0\leq\tau\leq 1} \left\{ (g+1)\sum_{l=1}^{N} \left(|\hat{D}_l^{(L)}(1,1+\tau)| - \sqrt{L}\tau\kappa \right)^+ + g\sum_{l=1}^{N} \left(|\hat{D}_l^{(L)}(0,\tau)| - \sqrt{L}\tau\kappa \right)^+ \right\} ,$$

$$\leq (g+1)\sum_{l=1}^{N} \sup_{0\leq\tau_l\leq 1} \left\{ \left(|\hat{D}_l^{(L)}(1,1+\tau_l)| - \sqrt{L}\tau_l\kappa \right)^+ \right\}$$

$$+ g\sum_{l=1}^{N} \sup_{0\leq\tau_l'\leq 1} \left\{ \left(|\hat{D}_l^{(L)}(0,\tau_l')| - \sqrt{L}\tau_l'\kappa \right)^+ \right\} ,$$

and we have

$$\mathbb{E}\left[\left(\tilde{B}_j^{(L)I}(t) - \tilde{B}_j^{(L)I}(t_j^{(L)}) \right) - \left(\tilde{B}_j^{(L)}(t) - \tilde{B}_j^{(L)}(t_j^{(L)}) \right) \right] ,$$

$$\leq (g+1)\sum_{l=1}^{N} \mathbb{E}\left[\sup_{0\leq\tau_l\leq 1} \left\{ \left(|\hat{D}_l^{(L)}(1,1+\tau_l)| - \sqrt{L}\tau_l\kappa \right)^+ \right\} \right] ,$$

$$+ g\sum_{l=1}^{N} \mathbb{E}\left[\sup_{0\leq\tau_l'\leq 1} \left\{ \left(|\hat{D}_l^{(L)}(0,\tau_l')| - \sqrt{L}\tau_l'\kappa \right)^+ \right\} \right] + L^{-1/2}.$$

Now for a given product l, we have

$$\mathbb{E}\left[\sup_{0\le\tau\le 1}\left\{\left(|\hat{D}_l^{(L)}(0,\tau)| - \sqrt{L}\tau\kappa\right)^+\right\}\right],$$

$$\le \mathbb{E}\left[\sup_{0\le\tau\le L^{-1/4}}\left\{\left(|\hat{D}_l^{(L)}(0,\tau)| - \sqrt{L}\tau\kappa\right)^+\right\} + \sup_{L^{-1/4}\le\tau\le 1}\left\{\left(|\hat{D}_l^{(L)}(0,\tau)| - \sqrt{L}\tau\kappa\right)^+\right\}\right],$$

$$\le \mathbb{E}\left[\sup_{0\le\tau\le L^{-1/4}}\left\{|\hat{D}_l^{(L)}(0,\tau)|\right\}\right] + \mathbb{E}\left[\sup_{L^{-1/4}\le\tau\le 1}\left\{\left(|\hat{D}_l^{(L)}(0,\tau)| - \sqrt{L}\tau\kappa\right)^+\right\}\right],$$

$$\le (1+\sigma_{ll}^2)L^{-1/8} + \frac{\sigma_{ll}^2}{\kappa}L^{-1/4}.$$

Since $\hat{D}_l^{(L)}(1, 1+\tau_l)$ has the same distribution as $\hat{D}_l^{(L)}(0,\tau)$, the same bound applies, and we have,

$$\mathbb{E}\left[\left(\tilde{B}_j^{(L)I}(t) - \tilde{B}_j^{(L)I}(t_j^{(L)})\right) - \left(\tilde{B}_j^{(L)}(t) - \tilde{B}_j^{(L)}(t_j^{(L)})\right)\right],$$

$$\le (2g+1)\sum_{l=1}^{N}\left((1+\sigma_{ll}^2)L^{-1/8} + \frac{\sigma_{ll}^2}{\kappa}L^{-1/4}\right) + L^{-1/2}.$$

This, together with (3.37) and (3.38), shows that

$$\mathbb{E}\left[\left(\tilde{B}_j^{(L)I}(t) - \tilde{B}_j^{(L)}(t)\right)\mathbb{1}_{\{t_j^{(L)}\ge t-1\}}\right],$$

$$\le (2g+1)\sum_{l=1}^{N}\left((1+\sigma_{ll}^2)L^{-1/8} + \frac{\sigma_{ll}^2}{\kappa}L^{-1/4}\right) + 6g\lambda^{1/(2+\delta)}L^{-\delta/(2(2+\delta))}\sum_{l=1}^{N}(1+\eta_j)$$

$$+ 2L^{-1/2},$$

which, together with (3.36), makes it clear that L can be chosen large enough to satisfy (3.34), which completes the proof. □

From Proposition 3.3 it is clear that the policy described is asymptotically optimal, since by Theorem 3.2 we have $\underset{\sim}{C}_L \le C_L^* \le C_L^\alpha$, which implies $1 \le C_L^\alpha/C_L^* \le C_L^\alpha/\underset{\sim}{C}_L$, and thus

$$\lim_{L\to\infty}\frac{C_L^\alpha}{C_L^*} = \lim_{L\to\infty}\frac{C_L^\alpha}{\underset{\sim}{C}_L} = 1.$$

Finally, we note that the policy described in Theorem 3.1 of Section 3.2.2 is a special case of the policy described in Proposition 3.3, with $\alpha = 0$ for setting the base stock levels in (3.23) (which is equivalent to (3.10)), and reservation levels $w_{ij} = 0$ for all i, j (which clearly satisfies condition (3.30)). Thus, Theorem 3.1 follows as a special case of Proposition 3.3.

3.3 Online Allocation for a Lost Sales Model

In this section we consider an alternative dynamic ATO model from that of Section 3.2 and show how our approach for the one-period problem yields an online algorithm with an $O(\log m/\log\log m)$ competitive ratio. Specifically, we consider a finite horizon ATO model with lost sales, online allocation, and no replenishment. The model of online arrivals captures a setting with high uncertainty regarding the order in which products will be demanded. Considering such a setting guides the design of a robust allocation policy that hedges against these uncertainties. The policy makes critical use of the solution of the single period (ATO) problem, both to set policy parameters and serve as a lower bound on optimal cost.

We next describe the model in detail. Recall that the assembly structure is described by a matrix A, with a_{ij} denoting the number of units of component i required to assemble product j. Before the beginning of the first time period, a manager makes a single ordering decision, r_i, for each component i, at cost c_i per unit, and there is no replenishment after the start of the first period. There are T discrete time periods, indexed by t, with product demand arriving sequentially, i.e., one unit of product demand arrives in any single period t. The online model we adopt is comprised of two layers: an initial random draw from a base distribution over demand sample paths, followed by an adversarial permutation of this sample path. (Note that in our setting, purely adversarial demand is trivial and uninteresting, since the adversary could send zero demand after any positive ordering decision is revealed).

We describe the two layers of the demand model next.

Base demand distribution. In time period t, at most one unit of demand arrives for a single product. With probability q_{jt} demand arrives for one unit of product j, and no demand for any product arrives with probability $1 - \sum_j q_{jt}$. The arrivals are independent across time periods. Let ω denote a sample path of the base demand distribution, and let $d_{tj\omega} \in \{0, 1\}$ represent the demand for product j in period t on this sample path. Let μ_ω denote the probability of sample path ω, which can be calculated from the probabilities q_{jt}. This demand process is similar to that described in Adelman (2007) for a network revenue management model.

Adversarial permutation. Initially, nature draws a full sample path of demand realizations from the base demand distribution. An adversary then observes the full sample path and may permute the arrivals in any order, doing so to create the worst possible performance for the allocation policy in place. Initially, the allocation policy does not observe the full sample path or permutation, but demand is sequentially revealed, one unit at a time, in the order chosen by the adversary.

Each time demand for product j arrives, the policy must choose whether to fill this demand, subject to the availability of all components $i \in \mathcal{N}(j)$. If a unit of demand for product j is not filled for any reason, the demand is lost and the policy is charged a shortage cost of p_j. Thus, a policy is comprised of two sets of decisions. Initially, an order quantity r_i for each component i must be chosen before the start of the first period. Then, in each subsequent period, it must decide how to allocate remaining components to arriving product demand. The online policy must be non-anticipative, i.e., in time period t it must only use information available up to that time when making allocation decisions.

Let $d_{j\omega} = \sum_t d_{tj\omega}$ denote the total demand for product j on sample path ω. Then, given initial inventory quantities r_i, we can think of the allocation problem along a given sample path ω as an online version of the offline second stage problem (2.4), which we reproduce here for reference

$$\begin{aligned} \min \sum_j & p_j y_{j\omega} \\ \text{s.t.} \sum_j & a_{ij} y_{j\omega} \geq \sum_j a_{ij} d_{j\omega} - r_i, \ \forall\, i, \\ & y_{j\omega} \leq d_{j\omega}, \ \forall\, j, \\ & y_{j\omega} \geq 0, \ \text{integer}, \ \forall\, j. \end{aligned} \tag{2.4 revisited}$$

In the online version of (2.4), the demands $d_{j\omega}$ are revealed unit by unit. In particular, let $\hat{d}_{j\omega}$ denote the total demand for product j revealed by a given time period. Then initially all $\hat{d}_{j\omega} = 0$, and in each time period at most one $\hat{d}_{j\omega}$ is incremented by one, until all demands are at their final quantities in period T, i.e., at time period T we have $\hat{d}_{j\omega} = d_{j\omega}$. Each time a $\hat{d}_{j\omega}$ is incremented, the allocation policy must decide whether to also increment the variable $y_{j\omega}$ (and only this variable), and is required to do so in a way that maintains feasibility for all constraints in (2.4) for the currently revealed demand $\hat{d}_{j\omega}$. Choosing to increment $y_{j\omega}$ corresponds to not filling the arriving demand for product j. In this setting, the adversarial permutation corresponds to an adversary choosing the order in which the $\hat{d}_{j\omega}$ are incremented. We emphasize that the online policy does not know either the permutation, or the sample path ω, but it does know the distribution of the sample paths.

In order to assess the performance of a given policy, we adopt the notion of *competitive ratio*, which compares the performance of the policy to that of an optimal offline algorithm that knows the full sample path before making any allocation decisions. Specifically, like the online policy, the offline algorithm makes an ordering decision before any demand is realized, but the offline algorithm is allowed to wait until the end of period T before deciding what demand to fill, i.e., it solves the offline allocation problem (2.4). Thus, the adversary's permutation of demand doesn't impact the offline algorithm's performance, as it processes the entire sample path at once. This can be viewed as relaxing the online algorithm's non-anticipativity constraint after the initial ordering decision. Further, it is clear that the offline algorithm's combined ordering and allocation problem is equivalent to (ATO), where demand is distributed according to $d_{j\omega} = \sum_t d_{tj\omega}$ as described above. Thus, we maintain the use of OPT to denote the optimal offline cost for this demand distribution.

To formally define the competitive ratio in our setting, let σ_ω denote a permutation of a sample path ω. For a given online policy π, let $\mathbf{r}^\pi$ denote the ordering decision and $P^\pi(\sigma_\omega, \mathbf{r}^\pi)$ denote the lost sales cost when the demand sample path is ω and σ_ω is the permutation. Then the policy π has competitive ratio β if

$$\mathbf{c} \cdot \mathbf{r}^\pi + \mathbb{E}_\omega \left[\max_{\sigma_\omega} \{ P^\pi(\sigma_\omega, \mathbf{r}^\pi) \} \right] \leq \beta OPT.$$

We briefly discuss this definition in relation to similar concepts discussed in the literature. Some models of online resource allocation assume a given quantity of initial resources (i.e., components in our setting) and consider arbitrary demand arrival sequences with no distributional assumptions (e.g. Ball and Queyranne, 2009; Ma and Simchi-Levi, 2017). We consider the broader problem of deciding both initial inventory investment, as well as online allocation. Thus, some distributional information on demand is required in our setting, since any algorithm performs arbitrarily poorly if it orders inventory and then no demand arrives.

We also contrast our setting to the random permutation model of Agrawal et al. (2014), who assume that the columns of an adversarialy chosen LP constraint matrix arrive in a random order drawn from the uniform distribution over all permutations. Our model in some sense flips this order, since we consider an adversarial permutation of a random arrival sample path, as opposed to a random permutation of adversarial arrivals. As mentioned above, we note that adversarial arrivals in our setting are not interesting, since an adversary can create arbitrarily poor performance (e.g., after the ordering decision is made, the adversary can send zero arrivals if positive inventory is ordered for any component, creating an infinite competitive ratio). We also note that our results would hold in a model with a distribution over permutations of the random sample paths, since we consider the worst permutation for every sample path.

We further note that our base demand distribution is similar to demand in the network revenue management model of Adelman (2007), who considers a maximization version of the allocation problem. In the ATO setting, Bernstein et al. (2011) consider a similar model of an initial ordering decision, sequential allocation, and lost sales. To our knowledge, we are the first to consider the general online version of this ATO setting.

We now present our joint inventory investment and allocation policy, Algorithm 5. The intuition of the policy is to approximate the optimal offline solution to (ATO) for demand $d_{j\omega} = \sum_t d_{tj\omega}$, using the LP rounding approach developed in Section 2.4, then use this solution to find a good inventory vector and online allocation policy. In particular, the policy requires an input factor α, by which it scales up the inventory quantities from the LP rounding. The allocation procedure then determines the average fraction of lost demand for each product, and randomly declines to fill arriving demand with these probabilities. Let $\bar{d}_j = \sum_\omega \mu_\omega d_{j\omega} = \sum_t q_{jt}$ denote the expected cumulative demand for product j over the entire time horizon.

Algorithm 5 Inventory and Online Allocation

1: Approximate (ATO) using Theorem 2.2, let r_i^I, $y_{j\omega}^I$ denote the output solution
2: Set inventory levels $r_i' \leftarrow \lceil \alpha r_i^I \rceil,\ \forall i$
3: Set $\bar{y}_j \leftarrow \sum_\omega \mu_\omega y_{j\omega}^I / \bar{d}_j,\ \forall j$
4: Upon arrival of demand for product j do:

$$\begin{cases} \text{A: Leave demand unfulfilled,} & \text{w.p. } \bar{y}_j \\ \text{B: Fill demand if components available,} & \text{w.p. } 1 - \bar{y}_j \end{cases}$$

Proposition 3.4. *If $m \geq 3$, $a_{ij} \in \{0,1\}$, and markups and component costs are bounded such that $\gamma_j c_i / c_{i'} \leq C < \infty$ for all j and i, i', then the inventory and allocation policy of Algorithm 5 has competitive ratio*

$$O\left(\frac{\log m}{\log(\log m)}\right),$$

when $\alpha = O(\log m / \log(\log m))$.

To prove this result we first state two supporting lemmas. The first leverages a standard Chernoff bound to approximate the expected value of a sum of Bernoulli random variables above a given quantity.

Lemma 3.8. *Let $X = \sum_i X_i$, where X_i are independent Bernoulli random variables, $\mu = \mathbb{E}[X]$, and consider an integer $s > e\mu$. Then*

$$\mathbb{E}[(X-s)^+] \leq \left(\frac{e\mu}{s}\right)^s \left(\frac{s}{s-e\mu}\right)$$

Proof. Using a standard Chernoff bound (Hagerup and Rüb, 1990), for a non-negative integer x we have

$$\begin{aligned} \mathbb{P}[X > s+x] &= \mathbb{P}[X > \frac{s+x}{\mu}\mu], \\ &\leq \left(\frac{e^{\frac{s+x-\mu}{\mu}}}{\left(\frac{s+x}{\mu}\right)^{\frac{s+x}{\mu}}}\right)^{\mu}, \\ &\leq \left(\frac{e\mu}{s+x}\right)^{s+x}, \\ &\leq \left(\frac{e\mu}{s}\right)^s \left(\frac{e\mu}{s}\right)^x, \end{aligned}$$

where the first inequality follows from the Chernoff bound for independent Bernoulli random variables (eq. (5) in Hagerup and Rüb, 1990), the second from the fact that the exponential function is increasing, and the last from $x \geq 0$. Now, writing the expectation of interest as the sum of tail probabilities, we have

$$\begin{aligned} \mathbb{E}[(X-s)^+] &= \sum_{x=0}^{\infty} \mathbb{P}[X > s+x], \\ &\leq \left(\frac{e\mu}{s}\right)^s \sum_{x=0}^{\infty} \left(\frac{e\mu}{s}\right)^x, \\ &= \left(\frac{e\mu}{s}\right)^s \left(\frac{s}{s-e\mu}\right), \end{aligned}$$

where the final equality follows from the geometric sum for $s > e\mu$. □

We will use the next lemma to bound the decay rate of the exponential term arising in the Chernoff bound. For $x > e$, let

$$\alpha(x) = \frac{2\log x}{\log(\log x)}.$$

Lemma 3.9. *For $x > e$ we have $\alpha(x) \geq 2e$ and*

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

Proof. First, note that $\alpha(x) \geq 2e$ is equivalent to

$$f(x) = \log x - e\log(\log x) \geq 0.$$

We will consider two cases; the first is $e < x \leq e^e$. The second derivative of f is

$$f''(x) = \frac{e + \log x(e - \log x)}{x^2 \log^2 x},$$

which is non-negative over the domain because $e < x \leq e^e$ implies $1 < \log x \leq e$. So the function f is convex on the domain $e < x \leq e^e$, so checking the first derivative of f,

$$f'(x) = \frac{\log x - e}{x \log x},$$

we see that f is minimized when $\log x = e$, or $x = e^e$. Plugging back into the function we get $f(e^e) = 0$, hence we have shown $f(x) \geq 0$ when $e < x \leq e^e$. For the second case when $x > e^e$, note that the function value is zero at $x = e^e$ and the derivative is positive for $x > e^e$. Thus, the function only increases away from zero when $x > e^e$ which implies $f(x) \geq 0$.

For the second claim, taking logarithms and rearranging, $(e/\alpha(x))^{\alpha(x)} \leq 1/x$ is equivalent to

$$\frac{\log x}{e} \leq (\log \alpha(x) - 1)e^{\log \alpha(x) - 1}.$$

Now, consider the Lambert W function for real values $z \geq 0$, implicitly defined as $z = W(z)e^{W(z)}$, thus satisfying $z = W(ze^z)$. We will use three properties of this function; i) $W(z)$ is strictly increasing for $z \geq 0$, ii) $W(z)$ is positive for $z > 0$, and iii) $W(e) = 1$ (see e.g., Corless et al., 1996, for a detailed treatment). Since $W(z)$ is strictly increasing, our inequality is equivalent to

$$W\left(\frac{\log x}{e}\right) \leq \log \alpha(x) - 1.$$

Rearranging we get the equivalent inequality

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

Now, by definition of the W function, $e^{W(z)} = z/W(z)$, so this inequality is equivalent to

$$\frac{\log x}{W\left(\frac{\log x}{e}\right)} \leq \alpha(x).$$

Substituting in the expression for $\alpha(x)$, this condition is equivalent to

$$g(x) = 2W\left(\frac{\log x}{e}\right) - \log(\log x) \geq 0.$$

We will prove this claim for two cases; the first is $e < x < e^{e^2}$. Note that $e < x$ implies that $0 < (\log x)/e$, so the argument of the W function is strictly positive and thus the function value itself is positive. Further, note that at $x = e^{e^2}$ we have $W((\log e^{e^2})/e) = W(e) = 1$, and since $W(z)$ is strictly increasing, we have $W((\log x)/e) \leq 1$ for $x \leq e^{e^2}$. Using implicit differentiation, the second derivative of g is

$$g''(x) = \frac{(\log x + 1)\left(W\left(\frac{\log x}{e}\right) + 1 - W\left(\frac{\log x}{e}\right)^2 - W\left(\frac{\log x}{e}\right)^3\right) + 2W\left(\frac{\log x}{e}\right)}{x^2 \log^2 x \left(W\left(\frac{\log x}{e}\right) + 1\right)^3}.$$

Since $W((\log x)/e) \leq 1$, we have $W((\log x)/e)^2 \leq W((\log x)/e)$ and $W((\log x)/e)^3 \leq 1$, and since $W((\log x)/e) \geq 0$, the second derivative is non-negative over the domain $e < x < e^{e^2}$. Thus the function $g(x)$ is convex over this domain, and checking the first derivative,

$$g'(x) = \frac{W\left(\frac{\log x}{e}\right) - 1}{x \log x \left(W\left(\frac{\log x}{e}\right) + 1\right)},$$

we see that $g(x)$ is minimized at $W\left(\frac{\log x}{e}\right) = 1$, which we saw above occurs at $x = e^{e^2}$. Plugging this back in to g, we see that $g(e^{e^2}) = 0$, and hence $g(x) \geq 0$ for all $e < x < e^{e^2}$. For the second case when $x > e^{e^2}$, note that the derivative of g is positive, and since g is zero at $x = e^{e^2}$, the function only increases away from zero over the range $x > e^{e^2}$. Thus we have shown $g(x) \geq 0$ which establishes the claim. □

Proof of Proposition 3.4. From the rounding of the inventory quantities in Algorithm 5, if $r_i^I \geq 1$ we have

$$r_i' = \lceil \alpha r_i^I \rceil \leq \alpha r_i^I + 1 \leq (\alpha + 1) r_i^I,$$

while if $r_i^I = 0$ we trivially have $r_i' \leq (\alpha + 1) r_i^I$. Thus, the inventory ordering cost can be bound as

$$\sum_i c_i r_i' \leq (\alpha + 1) \sum_i c_i r_i^I. \tag{3.39}$$

Let $\hat{C} = \max(C, 1)$ and set α in Algorithm 5 as follows

$$\alpha = \frac{2\log(\hat{C}m)}{\log(\log(\hat{C}m))},$$

and note that since $m \geq 3$, we have $\hat{C}m > e$, so that by Lemma 3.9 we have $\alpha > 2e$. With this value of α, we will demonstrate that for any collection of arrival sequence permutations across all demand scenarios, the expected cost of Algorithm 5 is bounded by a linear function of α times OPT, the optimal value of the offline problem (ATO). Since we show this inequality for any arrival sequence permutations, it holds when the worst permutation is chosen for each demand scenario, and thus satisfies the competitive ratio definition. Then, since α is $O(\log m/\log(\log m))$, the result follows.

In the following, we use the subscript t to index random variables according to the time period that nature originally drew the demand in (as opposed to the adversarial order in which demand actually arrived). We define two random processes to account for the shortage incurred by the online algorithm. Let Y_{jt} be a Bernoulli random variable that takes value 1 if both of the following occur: i) there is demand for product j in period t, and ii) when this demand arrived the policy chose action A in Step 4 of Algorithm 5, i.e., it decided to leave the demand unfulfilled. Similarly, let Z_{jt} be a Bernoulli random variable that takes value 1 if both of the following occur: i) there is demand for product j in period t, and ii) when this demand arrived the policy chose action B in Step 4, i.e., it tried to fill the demand. Then we say the online algorithm incurs *type 1* shortage when $Y_{jt} = 1$, and *type 2* shortage when $Z_{jt} = 1$ and there is some component $i \in \mathcal{N}(j)$ with insufficient remaining inventory to fill the demand. For any arrival sequence permutations, the expected amount of type 1 shortage is

$$\begin{aligned} \mathbb{E}[\sum_t Y_{jt}] &= \sum_t q_{jt}\bar{y}_j, \\ &= \sum_\omega \mu_\omega y^I_{j\omega}, \end{aligned} \tag{3.40}$$

which is the expected shortage in the integral solution.

Now we focus on type 2 shortage, i.e., when the policy tentatively accepts demand for product j but does not have the inventory to fill it. For a given component i, let $Z^i_t = \sum_j a_{ij}Z_{jt}$ denote the demand for component i in period t. Note that Z^i_t is a Bernoulli random variable with success probability equal to $\sum_j a_{ij}q_{jt}(1 - \bar{y}_j)$. Further, for a given i, the Z^i_t are independent across time periods t. Thus, the total demand for component i across all time periods, $Z^i = \sum_t Z^i_t$ is a sum of independent Bernoulli random variables. Letting z_i denote the expected value of Z^i, we have

$$
\begin{aligned}
z_i &= \sum_t \sum_j a_{ij} q_{jt} (1 - \bar{y}_j), \\
&= \sum_\omega \mu_\omega \sum_j a_{ij} (d_{j\omega} - y^I_{j\omega}), \\
&\leq r^I_i,
\end{aligned}
$$

where the second equality follows from the definition of $\bar{y}_j$ and the fact that $\sum_t q_{jt} = \sum_\omega \mu_\omega d_{j\omega}$, and the final inequality follows from the fact that $\sum_j a_{ij}(d_{j\omega} - y^I_{j\omega}) \leq r^I_i$ for every demand scenario by the inventory constraint (2.1) of (ATO).

Next, consider a component i for which we ordered positive inventory, i.e., $r^I_i \geq 1$. For any arrival sequence permutations, the expected quantity of type 2 shortage that is caused by this component, i.e., the expected amount of Z^i that is above r'_i, is

$$
\begin{aligned}
\mathbb{E}[(Z^i - r'_i)^+] &\leq \mathbb{E}[(Z^i - \alpha r^I_i)^+], \\
&\leq \left(\frac{e z_i}{\alpha r^I_i}\right)^{\alpha r^I_i} \left(\frac{\alpha r^I_i}{\alpha r^I_i - e z_i}\right), \\
&\leq 2 \left(\frac{e}{\alpha}\right)^{\alpha}, \\
&\leq \frac{2}{Cm},
\end{aligned}
$$

where the second inequality follows from Lemma 3.8, the third from the facts that $\alpha > 2e$, $r^I_i \geq z_i$, and $r_i \geq 1$, and the last inequality from Lemma 3.9. Now, in the worst case, all the type 2 shortage caused by component i could occur on the most expensive product it serves, i.e., at cost $\hat{p}_i = \max_{j \in \mathcal{N}(i)}\{p_j\}$, so the expected shortage cost incurred by component i is at most $\hat{p}_i \mathbb{E}[(Z^i - r'_i)^+]$. Recall that $\bar{\gamma} = \max_j \gamma_j$ and $\bar{c} = \max_i c_i$ denote the maximum markup and component cost, respectively. Then, since $\bar{\gamma}\bar{c}/c_i \leq C$, for any i we have

$$
\begin{aligned}
\frac{\hat{p}_i}{Cm} &\leq \frac{\bar{\gamma}\bar{c}m}{Cm}, \\
&\leq c_i.
\end{aligned}
$$

Further, note that if $r^I_i = 0$ then the random variable Z^i is identically zero and there can be no type 2 shortage cost caused by this component. Therefore, for any arrival sequence permutations, the total type 2 shortage cost incurred by the algorithm is bound by

$$\begin{aligned}\sum_{i|r_i^I \geq 1} \hat{p}_i \mathbb{E}[(Z^i - \alpha r_i)^+] &\leq 2 \sum_{i|r_i^I \geq 1} \frac{\hat{p}_i}{Cm}, \\ &\leq 2 \sum_{i|r_i^I \geq 1} c_i, \\ &\leq 2 \sum_i c_i r_i^I.\end{aligned}$$

Combining this with the bounds on ordering cost and type 1 shortage cost in (3.39) and (3.40), respectively, we see that for any arrival sequence permutations, the expected cost of the algorithm is bound by

$$\begin{aligned}(\alpha+1)\sum_i c_i r_i^I + \sum_{j,\omega} \mu_\omega p_j y_{j\omega}^I + 2\sum_i c_i r_i^I, &\leq (\alpha+3)\left[\sum_i c_i r_i^I + \sum_\omega \mu_\omega y_{j\omega}^I\right], \\ &\leq 1.8(\alpha+3)OPT,\end{aligned}$$

where the final inequality follows from the approximation guarantee of Theorem 2.2. □

3.4 Conclusion

In this chapter, we have established that our approximation results for the one-period ATO problem extend to a dynamic setting. Specifically, we develop a dynamic integral policy that is asymptotically optimal for a backlogging model as lead time grows large, by applying our rounding analysis for the one-period problem to make replenishment and allocation decisions. Further, for a lost sales model, we use our one-period rounding solution to design an online allocation policy and characterize its competitive ratio.

Chapter 4

Matching Supply and Demand in a Resource Constrained Service Network

4.1 Introduction

Many markets feature central platforms at which supply and demand are matched. Examples include modern e-retailing platforms matching sellers and buyers (e.g., Amazon, e-Bay), ride-hailing platforms matching drivers and riders (e.g., Uber, Lyft), and online advertising networks matching impressions with advertisements (e.g., Google), as well as more traditional models of markets with intermediaries facilitating trade between buyers and sellers (e.g., financial markets). In this chapter, we consider the problem of how to locate a limited number of platforms in a two-sided market in order to maximize the expected value of matches made between supply and demand. We give a brief description of the main issues at play in this problem, before providing several examples of applications, and finally presenting a formal model.

In two-sided markets, each unit of supply is often heterogeneously suited to the needs of each unit of demand, resulting in variability in the value of matches made between the two sides of the market. For example, consumers of different demographics may provide significantly different value to various online advertisers depending on targeting strategies, while in a ride-hailing setting location may be the primary means of differentiation. Thus, in order to be effective at matching supply and demand, the platform must have access to various subsets of the population on each side of the market.

Further, in these settings often both the quantity and value of supply and demand can be highly uncertain. For example, online platforms do not know how much or which type of traffic their website will receive ahead of time, nor often times the advertisers who will bid in their real-time auction environments. Likewise the value on both the buying and selling side of financial markets is notoriously uncertain. For this reason, we consider maximizing the expected value of matches in the market.

Next we discuss several example of these types of matching decisions across a wide range of industries.

4.1.1 Application Examples

In this section we motivate the model we consider with several practical examples from various application domains.

Franchise service providers. Consider a franchise business that provides a service to customers which must be administered by licensed professionals (e.g., nursing home service or alternative chronic pain therapies, such as massage, chiropractors, or acupuncture). The demand for such services will clearly be geographically differentiated based the local population's age, income and other characteristics. Further, the supply of licensed professionals may also be limited in certain areas due to regulatory constraints, or a lack of local training programs. In such a setting the availability of different supply and demand resources may be modeled with a network structure, with edges denoting which resources could be matched with which potential site locations. When deciding on a set of new locations to pursue opening a franchise, the business must take into account the availability of both supply and demand in the network, which our model captures.

Markets with Intermediaries. There is a large body of research considering markets with intermediaries who facilitate transactions between buyers and sellers. For example, Blume et al. (2009) show that the optimal value of trades in a network of buyers and sellers who have access to each other via central traders can be computed through solving an appropriate linear program, while Nguyen (2015); Nguyen and Vohra (2017) consider coalitions and bargaining in trade networks facilitated by middlemen. Our model captures the problem of selecting a subset of traders or middlemen in such markets in order to maximize the value of trades that can occur in the network.

Transportation-service platforms. Another very active area of research concerns platforms who connect supply and demand of transportation services. For example, Shu et al. (2013); O'Mahony and Shmoys (2015); Freund et al. (2017) consider various aspects of locating and redistributing capacity for bike-sharing systems in major cities, while de Almeida Correia and Antunes (2012); He et al. (2017); Alijani et al. (2017) consider how to locate regions or markets in order to optimize a vehicle-sharing service. In such markets, our model considers choosing a subset of match-making locations in order to maximize the value of matches made by the platform.

Header-bidding advertising auctions. The recent rise in real-time bidding for online advertising impressions through the use of header-bidding strategies provides another application for our model. Header-bidding is a strategy employed by an increasing number of online publishers that implements real-time bidding for advertising impressions in the header of a web page, before the impressions pass to pre-negotiated direct orders (Vidakovic, 2017; Qin et al., 2017; Sayedi, 2018; Jauvion et al., 2018). In this market, publishers typically work with multiple "demand partners," who represent a network of advertisers who may value different customer impressions on the publisher's website. With the advent of header-bidding strategies, many publishers have noted the need to limit the number of demand partners they work with, due to speed issues associated with running the auction in the header of the web page before the page loads Vidakovic (2017); Prebid.org (2018). In this setting, our model considers choosing a limited number of demand partners in order to maximize the value of advertising matches made with customer impressions arriving to the publisher's page.

Process flexibility design. Our model also relates to research on process flexibility design. In this setting, there is a set of plants that can be used to produce a set of products, and the problem is to design a "flexibility structure" that connects plants with subsets of products in order to maximize the amount of plant capacity that is used to meet product demand. This line of research is often compares how close sparse flexibility structures, i.e., those with only a few plant-product connections, can come to the performance of full flexibility, i.e., a structure where every plant is connected to every product (Jordan and Graves, 1995b; Simchi-Levi and Wei, 2012; Chen et al., 2015). Our model can be used to capture a different but related question: what is the best flexibility structure with less than K plant-product connections (where K is some integer)?

4.1.2 Model

In this section we introduce the formal model we consider. The input data of the problem consists of a universe of location nodes U (with individual elements denoted by $i \in U$), an integer K, and a set of scenarios Ω, together with a probability measure μ_ω on elements $\omega \in \Omega$. Scenario $\omega \in \Omega$ is a realization of a set of supply nodes S_ω, demand nodes D_ω, and profits $p_{ijk\omega} \geq 0$ for all $j \in D_\omega$, $k \in S_\omega$, and $i \in U$. (To increase readability, we suppress the dependence of the supply and demand node indices on the scenario ω, and just use $j \in D_\omega$ and $k \in S_\omega$.) Let $n = |U|$ denote the number of location nodes.

The decision sequence of the problem is as follows. Before the scenario is realized, a set of location nodes $A \subseteq U$ with $|A| \le K$, are *opened*. Then a scenario is randomly drawn from Ω according to the probabilities μ_ω. After scenario ω is realized, a demand node $j \in D_\omega$ can be *matched* with supply node $k \in S_\omega$ at an open location node $i \in A$ to give profit $p_{ijk\omega} \ge 0$. Each supply and demand node can be used in at most one match at one open location node. The objective is to maximize the expected profit made by supply-demand matches at open facilities. We call this the *K-match* problem.

We highlight that our formulation with matching profits $p_{ijk\omega}$ is quite flexible to handle data in various settings. For example, if feasible matches are determined by a network and all matches have equal value, then the network structure can be encoded in the matching profits by setting $p_{ijk\omega} = 1$ if nodes j and k are connected to node i in the network in scenario ω, while $p_{ijk\omega} = 0$ otherwise. Our model also allows the much more general case where matches have differing values based on the combination of nodes involved.

To succinctly formulate the optimization problem, we define the profit objective in scenario $\omega \in \Omega$ as a function of the set $A \in U$ of open nodes as follows:

$$f_\omega(A) = \left\{ \begin{aligned} &\max \sum_{i \in A} \sum_{j \in D_\omega, k \in S_\omega} p_{ijk\omega} x_{ijk} \\ &\text{s.t.} \sum_{i \in A, k \in S_\omega} x_{ijk} \le 1,\ j \in D_\omega, \\ &\qquad \sum_{i \in A, j \in D_\omega} x_{ijk} \le 1,\ k \in S_\omega, \\ &\qquad x_{ijk} \in \{0,1\},\ i \in A,\ j \in D_\omega,\ k \in S_\omega. \end{aligned} \right\} \tag{4.1}$$

In the formulation of $f_\omega(A)$, the variable x_{ijk} is one if supply node k is matched with demand node k at location node i, and is zero otherwise. The first (second) constraint ensures that each demand (supply) node is matched with at most one supply (demand) node on at most one location node. With this definition of the matching value in scenario ω, we can state the K-match problem as follows:

$$\begin{aligned} &\max \mathbb{E}_\omega[f_\omega(A)] \\ &\text{s.t. } |A| \le K, \end{aligned} \tag{K-match}$$

where the expectation in the objective is over the random scenarios $\omega \in \Omega$, i.e., $\mathbb{E}_\omega[f_\omega(A)] = \sum_{\omega \in \Omega} \mu_\omega f_\omega(A)$.

4.1.3 Results

In this section we briefly describe our main results for the problem (K-match). In Section 4.2 we first demonstrate that it is NP-hard to approximate (K-match) with a factor better than $1/(1-e^{-1}) \approx 1.58$. We then develop two algorithms that provide constant factor approximations for (K-match). The first algorithm is an *exchange* algorithm, which starts with a feasible solution and exchanges elements in the current solutions with those not in the current solution until no such exchange increases the objective value (we define the algorithm formally in Section 4.3.1). We provide the following approximation guarantee for this algorithm.

Theorem 4.1. *For any $\epsilon > 0$, an exchange algorithm provides a $(3 - 2/K + \epsilon)$-approximation for* (K-match)*.*

The next algorithm we consider is a *greedy* algorithm, which starts with an empty solution and adds in each iteration the single element which provides the largest gain in the objective value (we define this algorithm formally in Section 4.3.2). We provide the following approximation guarantee for this algorithm.

Theorem 4.2. *A greedy algorithm provides a $2/(1 - e^{-2})$-approximation for* (K-match)*.*

We note that $2/(1 - e^{-2}) \approx 2.31$. In order to obtain these results, we develop a generalization of the classic submodularity property for set functions (Nemhauser et al., 1978; Fisher et al., 1978; Feige et al., 2011). In Section 4.3, we show that our generalization, called *γ-cover modularity*, implies general approximation factors for exchange and greedy algorithms applied to any maximization problem whose objective function satisfies this property (along with other standard regularity conditions). Then, in Section 4.4 we show that the objective function of (K-match) is γ-cover modularity for $\gamma = 2$, and use this to complete the proofs of Theorems 4.1 and 4.2.

4.1.4 Related Literature

In addition to the literature referenced in the application examples in Section 4.1.1, in this section we briefly review other literature related to the K-match problem. There are clear connections with the vast literature on facility location, which includes results on approximation algorithms, such as Cornuejols et al. (1977); Shmoys et al. (1997); Jain and Vazirani (2001), among many others. The main difference between our model and these papers is that our problem is two sided (i.e., matching supply and demand), whereas most traditional facility location problems are one sided (i.e., covering demand).

There are a few facility location models that incorporate non one sided demand fulfillment decisions, including Gourdin et al. (2000) who study branch and cut algorithms for a problem with matching between different demand locations to optimize routing costs, Zhu et al. (2010) who study a vehicle routing problem between facilities, customers, and suppliers, Caro et al. (2012) who select food processing activities and assign products to processes in order to transport them to different markets, and Cho et al. (2014) who study locating trauma centers and heliports in order to maximize coverage of demand for these emergency services. Our model differs from these papers in terms of the type of matching constraints we consider, (i.e., the typical bipartite matching polytope). Further, our focus is on worst case approximation guarantees, whereas these papers take the approach of validating heuristics through numerical simulation. Alijani et al. (2017) also consider a two sided facility location problem and provide approximation guarantees, but our setting is significantly different from theirs, which opens "virtual markets" that require a lower bound on flow to be viable, and incorporates a constraint on matching distance in an embedded metric space.

This chapter is also related to the literature on stochastic facility location. Louveaux and Peeters (1992) develop a dual method for a one sided facility location problem where demand, profit, and costs are stochastic, which they test numerically. Ravi and Sinha (2006) consider a stochastic version of the minimum cost, one sided facility location problem and adapt the LP rounding algorithm of Shmoys et al. (1997) to give a constant approximation guarantee. Ansari et al. (2017) study a queuing model for ambulance location with stochastic availability and propose an iterative optimize, then approximate approach, which they test with computational experiments. Our work builds on this stochastic facility location literature by considering two sided matching in a stochastic setting.

4.2 Computational Challenges

In this section we consider the computational challenges associated with solving (K-match). We show in Section 4.2.1 that (K-match) is NP-hard to approximate with a factor better than $1/(1-e^{-1}) \approx 1.58$. Then in Section 4.2.2, we further demonstrate that (K-match) does not enjoy the benefit of classical approximation results for maximizing set functions (e.g., Nemhauser et al., 1978), since its objective function is not submodular. First, however, we prove a positive result demonstrating that, for a given subset $A \subseteq U$ and scenario $\omega \in \Omega$, the value of the optimal matching can be efficiently computed via linear programming.

Lemma 4.1. *The constraint matrix of the LP defining $f_\omega(A)$ in (4.1) is totally unimodular for any $A \subseteq U$.*

Proof. We use a characterization of Tamir (1976) that says a matrix with elements a_{bc} is totally unimodular if every subset of rows J can be partitioned into two subsets $J_1 \cup J_2 = J$ such that $\sum_{b \in J_1} a_{bc} - \sum_{b \in J_2} a_{bc} \in \{-1, 0, 1\}$ for all columns c. The constraint matrix of (4.1) has a row for each $j \in D_\omega$ and each $k \in S_\omega$ and a column for every variable x_{ijk}. For a given $j \in D_\omega$ and $i \in A, j' \in D_\omega, k' \in S_\omega$ we can represent the constraint matrix by

$$a_{(j);(ij'k')} = \begin{cases} 1, & j = j' \\ 0, & j \neq j' \end{cases}.$$

Similarly, for a given $k \in S_\omega$ and $i \in A, j' \in D_\omega, k' \in S_\omega$, we can write

$$a_{(k);(ij'k')} = \begin{cases} 1, & k = k' \\ 0, & k \neq k' \end{cases}.$$

Any subset of the rows of the constraint matrix corresponds to the union of some subsets of D_ω and S_ω, i.e. $D' \cup S'$ where $D' \subseteq D_\omega$ and $S' \subseteq S_\omega$. Let the desired partition then be $J_1 = D'$ and $J_2 = S'$. Then for a given variable $x_{ij'k'}$,

$$\sum_{j \in D'} a_{(j);(ij'k')} - \sum_{k \in S'} a_{(k);(ij'k')} = \begin{cases} -1, & j' \notin D' \text{ and } k' \in S', \\ 0, & (j' \in D' \text{ and } k' \in S') \text{ or } (j' \notin D' \text{ and } k' \notin S'), \\ 1, & j' \in D' \text{ and } k' \notin S', \end{cases}$$

□

Thus, Lemma 4.1 guarantees that there exists an integer optimal solution to the linear relaxation of (4.1) (Theorem 13.2 in Papadimitriou and Steiglitz, 1998), which can be found efficiently via linear programming. Further, various sampling techniques (such as sample average approximation and stochastic gradient descent, Shapiro et al., 2009; Swamy and Shmoys, 2012) provide efficient approximation to arbitrary precision for the expected value of such linear programming objective functions. Thus, for a given subset A, the objective value of (K-match) can be computed efficiently. However, we demonstrate in the next section that choosing the optimal such subset is NP-hard, even to approximate.

4.2.1 Complexity

In this section we consider the computational complexity of the K-match problem. Using a reduction from the max k-cover problem, our main result is to provide a lower bound of $1/(1-e^{-1}) \approx 1.58$ on the approximation factor achievable for (K-match) by a polynomial time algorithm, unless $P = NP$, even when the supply and demand nodes are deterministic (i.e. $|\Omega| = 1$).

Proposition 4.1. *For any $\epsilon > 0$, no polynomial time algorithm for* (K-match) *can achieve approximation factor $1/(1-e^{-1}) - \epsilon$, unless $P = NP$, even when $|\Omega| = 1$.*

Proof. To prove the claim, consider the following problem.

Problem 4.1 (Max K-Cover)**.** *Given a set $S = \{s_1, \ldots, s_n\}$ of cardinality n, a collection of m subsets of S, $\mathcal{S} = \{S_1, \ldots, S_m\}$, and an integer $K \leq m$, choose K subsets in $\mathcal{S}$ to maximize the cardinality of their union.*

Given any instance of Problem 4.1, we will construct in polynomial time an instance of (K-match) that is an equivalent optimization formulation. The inapproximability result then follows by Theorem 5.3 of Feige (1998).

Consider an instance of Problem 4.1 with cardinality constraint K, and construct an instance of (K-match) as follows. Let there be only one demand scenario, so we drop the subscript ω. For each subset S_i, let there be a corresponding location node in U. For each element s_j of S let there be a corresponding demand node in D and supply node in S, which provide a profit of one if matched with each other on a location node $i \in U$ for which $s_j \in S_i$, and otherwise provide a profit of zero if involved in any other match.

First we claim that an optimal solution to the original Problem 4.1 can be used to construct, in polynomial time, a feasible solution to the constructed instance of (K-match) that has the same objective value. Given an optimal solution to Problem 4.1, select the K location nodes in the constructed instance of (K-match) that correspond to the K optimal subsets in Problem 4.1. Then go through these K optimal subsets in an arbitrary order, and for each $s_j \in S_i$, match the corresponding demand and supply node on location node i in the constructed instance of (K-match) if they have not already been matched on a previously considered subset. This provides an objective value for the constructed instance of (K-match) that is equal to the cardinality of the union of the K optimal subsets in Problem 4.1, i.e., the optimal value of Problem 4.1.

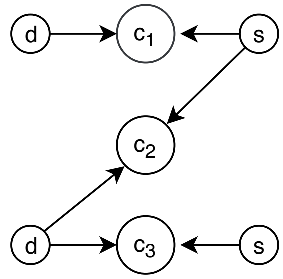

Figure 4.1: System demonstrating lack of submodularity.

Next we claim that an optimal solution of the constructed instance of (K-match) can be used to construct, in polynomial time, a feasible solution to Problem 4.1 that has objective value at least as large as the optimal value of (K-match). Given an optimal solution to the constructed instance of (K-match), choose the K subsets in Problem 4.1 that correspond to the K optimal location nodes selected in (K-match). Then it is clear that no demand node corresponding to an element outside of the union of these K subsets can be matched in (K-match), while every demand node corresponding to an element inside the union can be matched at most once. Thus the cardinality of the union of these K subsets is at least as large as the optimal number of matches made in the constructed instance of (K-match).

Therefore, we have shown that the optimal value of the constructed instance of (K-match) is equal to the optimal value of Problem 4.1, implying that a polynomial time approximation algorithm for (K-match) would provide a polynomial time approximation algorithm for Problem 4.1 with the same approximation factor. $\square$

4.2.2 Lack of Submodularity

In this section we consider the well known submodularity property for set functions. There are many equivalent definitions of submodularity; the most convenient one for our purposes describes the idea of decreasing marginal gains. In particular, given a universe of elements U, we say a set function $f : 2^U \to \mathbb{R}$ is submodular if $f(T \cup \{i\}) - f(T) \leq f(S \cup \{i\}) - f(S)$ for all $S \subset T \subset U$ and $i \notin T$. The intuitive explanation of submodularity is that adding element i to a small set increases the function more than adding it to a larger set. This property has been used to demonstrate approximation guarantees for several maximization problems (Nemhauser et al., 1978; Feige et al., 2011).

We show that the function f_ω defined in (4.1) fails to be submodular in general. To show this, consider the simple system depicted in Figure 4.1, which is described in detail in the following example.

Example 4.1. *Consider a system with three location nodes and two each of supply and demand nodes, as depicted in Figure 4.1. The profit of matching supply node k with demand node j on location node i is $p_{ijk} = 1$ if j and k are each connected to i as depicted in the graph in Figure 4.1 (otherwise the match has value $p_{ijk} = 0$). Formally, for triples denoted by (i, j, k), the following triples have $p_{ijk} = 1 : \{(1, 1, 1), (2, 2, 1), (3, 2, 2)\}$, while all remaining triples have $p_{ijk} = 0$.*

Consider the sets of location nodes $S = \{2\}$, $T = \{1, 2\}$, and the single location node $i = 3$. First consider the effect of adding element i to set S. It is clear that $f(S) = 1$, since only one supply node and one demand node are connected to it. Adding $i = 3$ to S does not increase the number of matches, since it only gives access to one more supply node, but no new demand nodes. Thus, $f(S \cup \{i\}) - f(S) = 0$. Next, consider the effect of adding i to T. Since T is only connected to one supply node, we have $f(T) = 1$. However, once we add $i = 3$ to T, all supply and demand nodes are available, and thus two matches can be made (one at $i = 1$ and one at $i = 3$). Therefore, we have $f(T \cup \{i\}) - f(T) = 2 > 1 = f(S \cup \{i\}) - f(S)$ and the function is not submodular.

4.3 Cover Modularity

In this section we will consider a generalization of the classic submodularity property, which preserves some of the classic approximation results. For a set function f over a universe of elements U, i.e., $f : 2^U \to \mathbb{R}$, and consider the following generalization of (K-match), where the objective function is replaced by f

$$\begin{aligned} &\max\ f(A) \\ &\text{s.t. } |A| \leq K, \end{aligned} \tag{4.2}$$

To analyze this problem we make the following definitions. Recall that $n = |U|$. The function f is called *monotone* if $f(S) \leq f(T)$ for all $S \subset T$. Define the marginal increase in the function f when adding the element $i \in U$ to the set S as $\rho_i(S) = f(S \cup i) - f(S)$. Nemhauser et al. (1978) derive the following key property of a submodular set function, which they use to prove approximation guarantees for the algorithms they consider.

$$f(S \cup T) - f(S) \leq \sum_{i \in T \setminus S} \rho_i(S),\ \forall S, T \in U.$$

This property is an intuitive consequence of the decreasing marginal gains interpretation of submodularity. In words it says that adding the elements in T to those in S increases the function no more than the aggregate increase each of those elements has when added to S individually. This follows naturally from submodularity by induction. We consider generalizing this condition as follows.

Definition 4.1. *For $\gamma \geq 1$, a set function f is called* γ-cover modular *if for all sets $S, T \subseteq U$,*

$$f(S \cup T) - \gamma f(S) \leq \sum_{i \in T \setminus S} \rho_i(S).$$

In words, γ-cover modularity says that we need to cover the value of set S with γ copies of itself in order to guarantee that the increase from adding the elements in T is less than the aggregate increase of each of those elements when added to S individually. This allows for controlled deviations from the decreasing marginal gains property, such as those observed in Example 4.1 of Section 4.2.2.

Thus, from Nemhauser et al. (1978) we know that a submodular function is 1-cover modular. We will show in Section 4.4 that the function f_ω defined in (4.1) is 2-cover modular. In the remainder of this section we will use the general γ-cover modular property to demonstrate approximation guarantees for various algorithms.

4.3.1 Exchange

The first algorithm we consider is an exchange algorithm, which starts with a solution of size K and swaps open and closed locations until no such swap produces a relative gain larger than ϵ/K. We state this algorithm formally in Algorithm 6.

Algorithm 6 Exchange

1: Choose some $\epsilon > 0$ and initialize $A = \emptyset$
2: Choose $i^* \in \text{argmax}_i\{f(\{i\})\}$ and add i^* to A
3: Choose any $K - 1$ elements of $U \setminus A$ to add to A
4: **while** $\max_{\{i' \in U \setminus A, i \in A\}}\{f(A \setminus \{i\} \cup \{i'\})\} \geq (1 + \frac{\epsilon}{K})f(A)$ **do**
5: Choose $i, i' \in \text{argmax}_{\{i' \in U \setminus A, i \in A\}}\{f(A \setminus \{i\} \cup \{i'\})\}$, add i' to A and remove i
6: **end while**

We require initialization with the best single element, as well as a small multiplicative increase in the objective value for each exchange in order to guarantee the algorithm has a polynomial run time, which we demonstrate formally below. The results we derive in this section for Algorithm 6 are predicated on the following conditions on the set function f.

Condition 4.1. *The set function f satisfies the following properties:*

i) Monotone

ii) γ-cover modular

iii) $\forall S \subseteq U$, $\exists i \in S$ s.t. $f(S)/|S| \leq f(S \setminus \{i\})/|S \setminus \{i\}|$

Condition 4.1 item i) is standard and a natural assumption in many settings, while item ii) is the key generalization of submodularity we are considering. Item iii) further requires that the average functional value per element in a set S must be non-decreasing upon the removal of at least one individual element from S. This property is easily derived from submodularity, but is also straightforward to show for functions defined as the linear combination of values attributable to individual elements (such as f_ω defined in (4.1)), by removing the element with the lowest value (see the proof of Lemma 4.4 in Section 4.4). We also note that Condition 4.1 item iii) is not needed to prove our approximation guarantee for the greedy algorithm in the Section 4.3.2.

Our first result establishes that Algorithm 6 requires only a polynomial number of evaluations of the function f, which implies a polynomial run time if f can be computed efficiently as well.

Lemma 4.2. *Under Condition 4.1, Algorithm 6 requires $O(\frac{nK^2 \log(K)}{\epsilon})$ evaluations of f.*

Proof. After Step 2, by monotonicity we have $f(A) \geq f(\{i^*\})$. Further, let A^* denote the optimal solution to (K-match), which by monotonicity is of size $|A^*| = K$. By item iii) in Condition 4.1 and induction, there must exist an element, $i' \in A^*$, such that $f(A^*)/K \leq f(\{i'\})$. Then, since $f(\{i^*\})$ is the largest value of any single element, we must have $f(A^*) \leq Kf(\{i'\}) \leq Kf(\{i^*\}) \leq Kf(A)$, so that $\frac{f(A^*)}{f(A)} \leq K$. Thus, the while loop in Steps 4 - 6 goes through at most $O(\frac{K \log(K)}{\epsilon})$ iterations (since each iteration increases the objective value by a multiplicative factor of ϵ/K), each requiring $O(nK)$ evaluations of f. Since Step 2 requires n evaluations of f, it is dominated by the other terms. $\square$

Next we prove an approximation guarantee for Algorithm 6 that is dependent upon the parameter γ in our definition of γ-cover modularity. We recall that the parameter $\epsilon > 0$ is required simply to ensure polynomial run time, and may be chosen to be a suitably small value.

Proposition 4.2. *Under Condition 4.1, Algorithm 6 provides the following approximation factor to* (4.2)

$$1 + \gamma - \frac{\gamma}{K} + \epsilon.$$

Proof. Let A^* denote the optimal solution to (4.2) and A the solution output by Algorithm 6. Consider $i \in A$ such that $f(A)/|A| \leq f(A \setminus \{i'\})/|A \setminus \{i'\}|$ (which exists by part iii) of Condition 4.1). Since $|A| = K$, this condition implies that $f(A) - f(A \setminus i') \leq f(A)/K$. Then, by monotonicity and γ-cover modularity, we have

$$\begin{aligned} f(A^*) &\leq f((A \setminus \{i'\}) \cup A^*), \\ &\leq \gamma f(A \setminus \{i'\}) + \sum_{i \in A^* \setminus (A \setminus \{i'\})} \rho_i(A \setminus \{i'\}), \\ &\leq \gamma f(A \setminus \{i'\}) + \sum_{i \in A^*} \left[f(A \setminus \{i'\} \cup \{i\}) - f(A \setminus \{i'\}) \right], \\ &\leq \gamma f(A \setminus \{i'\}) + \sum_{i \in A^*} \left[(1 + \frac{\epsilon}{K}) f(A) - f(A \setminus \{i'\}) \right], \\ &\leq \gamma f(A \setminus \{i'\}) + K \left[(1 + \frac{\epsilon}{K}) f(A) - f(A \setminus \{i'\}) \right], \end{aligned}$$

where the third inequality follows since $\rho_i(A \setminus \{i'\}) \geq 0$ by monotonicity, the fourth by the condition in the while loop of Step 4 of the algorithm, and the final inequality since $|A^*| \leq K$. Therefore, rearranging, we have

$$\begin{aligned} f(A^*) &\leq (\gamma + \epsilon) f(A) + (K - \gamma) \left[f(A) - f(A \setminus \{i'\}) \right], \\ &\leq \left(\gamma + \epsilon + \frac{K - \gamma}{K} \right) f(A), \end{aligned}$$

which gives the result. □

Thus, Algorithm 6 provides an approximation guarantee of about $1 + \gamma$ (for large values of K and small ϵ). For a submodular set function f, the parameter $\gamma = 1$, and so Proposition 4.2 gives an approximation guarantee of $(2K - 1)/K + \epsilon$ which is almost identical to the result in Nemhauser et al. (1978) for a similar exchange-type algorithm.

4.3.2 Greedy

Next we consider the greedy algorithm, which opens at each step the the location node which provides the largest marginal gain in the objective value. We state the procedure formally as

Algorithm 7.

Algorithm 7 Greedy

```
1: Initialize A_0 = ∅
2: for t = 1, ..., K do
3:     Choose i^t ∈ argmax_{i∈U\A_{t-1}} {f(A_{t-1} ∪ {i})}
4:     Set A_t = A_{t-1} ∪ {i^t}
5: end for
```

The greedy algorithm is notably simpler than the exchange algorithm. In particular, the greedy algorithm only needs to consider adding individual nodes to A, rather than considering pairs of nodes to exchange as in Algorithm 6. Further, the greedy algorithm only makes K additions by design, whereas the exchange algorithm will continue searching for pairs of elements to exchange until no such pair provides an improvement in value. These observations result in a significantly better worst case bound on the number of evaluations of f required by the algorithm.

Lemma 4.3. *Algorithm 7 requires $O(nK)$ evaluations of f.*

Proof. The For loop runs through K iterations, each evaluating f in Step 3 $O(n)$ times. □

To derive an approximation guarantee for the greedy algorithm, we require the following conditions on the set function f.

Condition 4.2. *The set function f satisfies the following properties:*

i) Monotone

ii) γ-cover modular

We do not require the average function value property of Condition 4.1 in order to prove the approximation guarantee for the greedy algorithm; we only require monotonicity and γ-cover modularity. Under these conditions, we have the following result.

Proposition 4.3. *Under Condition 4.2, Algorithm 7 provides the following approximation factor for* (4.2)

$$\frac{\gamma}{1-(1-\frac{\gamma}{K})^K} \leq \frac{\gamma}{1-e^{-\gamma}}.$$

Proof. Let A^* denote the optimal solution to (4.2) and A_t, i^t denote the greedy sets and elements constructed by Algorithm 7. Define $\rho_t = \rho_{i^t}(A_{t-1})$ as the marginal increase in the value of the solution in iteration t of the algorithm. A simple telescoping sum shows that $f(A_t) = \sum_{j=1}^{t} \rho_j$ (where for $t = 0$ the empty sum denotes zero). Then, for each $t = 1, \ldots, K$, we have

$$
\begin{aligned}
f(A^*) &\leq f(A_{t-1} \cup A^*), \\
&\leq \gamma f(A_{t-1}) + \sum_{i \in A^* \setminus A_{t-1}} \rho_i(A_{t-1}), \\
&\leq \gamma f(A_{t-1}) + \sum_{i \in A^* \setminus A_{t-1}} \rho_t, \\
&\leq \gamma \sum_{j=1}^{t-1} \rho_j + K\rho_t,
\end{aligned}
$$

where the first line follows by monotonicity, the second by γ-cover modularity, the third since the greedy algorithm guarantees that $\rho_i(A_{t-1}) \leq \rho_{i^t}(A_{t-1}) = \rho_t$ for all $i \in A^* \setminus A_{t-1}$, and the fourth since $|A^*| \leq K$. Now, given ρ_t for $t = 1, \ldots, K$, let $z = f(A^*)$ denote the true optimal value, and consider maximizing this value subject to the inequalities just derived. This problem can be written as the following linear program:

$$
\begin{aligned}
&\max\ z \\
&\text{s.t. } z \leq \gamma \sum_{j=1}^{t-1} \rho_j + K\rho_t,\ \forall t = 1, \ldots, K,
\end{aligned}
$$

whose dual is

$$
\begin{aligned}
\min\ &\sum_t \pi_t \left(\gamma \sum_{j=1}^{t-1} \rho_j + K\rho_t \right) \\
\text{s.t. } &\sum_t \pi_t = 1.
\end{aligned}
$$

We construct a feasible dual solution as follows. Let $\delta = \gamma/(1 - (1 - \frac{\gamma}{K})^K)$, and let

$$
\pi_t = \frac{\delta}{K} \left(1 - \frac{\gamma}{K}\right)^{K-t}.
$$

This solution is feasible to the dual since

$$\begin{aligned}
\sum_t \pi_t &= \sum_{t=1}^{K} \frac{\delta}{K} \left(1 - \frac{\gamma}{K}\right)^{K-t}, \\
&= \frac{\delta}{K} \sum_{j=0}^{K-1} \left(1 - \frac{\gamma}{K}\right)^{j}, \\
&= \frac{\delta}{K} \left(\frac{1 - \left(1 - \frac{\gamma}{K}\right)^{K}}{1 - \left(1 - \frac{\gamma}{K}\right)} \right), \\
&= 1,
\end{aligned}$$

by the definition of δ. Next, we consider rearranging the objective as follows

$$\sum_{t=1}^{K} \pi_t \left(\gamma \sum_{j=1}^{t-1} \rho_j + K\rho_t \right) = \sum_{j=1}^{K} \rho_j \left(\gamma \sum_{t=j+1}^{K} \pi_t + K\pi_j \right),$$

and for the feasible solution defined above, for each $j = 1, \ldots, K$ we have

$$\begin{aligned}
\gamma \sum_{t=j+1}^{K} \pi_t + K\pi_j &= \gamma \frac{\delta}{K} \sum_{l=0}^{K-j-1} \left(1 - \frac{\gamma}{K}\right)^{l} + K \frac{\delta}{K} \left(1 - \frac{\gamma}{K}\right)^{K-j}, \\
&= \gamma \frac{\delta}{K} \left(\frac{1 - \left(1 - \frac{\gamma}{K}\right)^{K-j}}{1 - \left(1 - \frac{\gamma}{K}\right)} \right) + \delta \left(1 - \frac{\gamma}{K}\right)^{K-j}, \\
&= \delta.
\end{aligned}$$

Thus, the dual objective value is $\delta \sum_j \rho_j = \delta f(A_K)$ where A_k is the final output of the greedy algorithm. Therefore, by weak duality we have $f(A^*) = z \leq \delta f(A_K)$, which establishes the first approximation factor in the statement of the result. The upper bound in terms of e follows from the elementary inequality $1+x \leq e^x$ for all $x \in \mathbb{R}$, since letting $x = -\gamma/K$ implies that $(1-\gamma/K)^K \leq e^{-\gamma}$ for all $K \geq \gamma$. □

We note that the approximation guarantee of Proposition 4.3 for the greedy algorithm is better than that of Proposition 4.2 for the exchange algorithm (for large enough K) for any $\gamma \geq 1$. This follows from the elementary inequality $1 + x \leq e^x$ for all $x \in \mathbb{R}$ (which itself can be seen by the Taylor expansion for e^x, or by convexity of e^x), since it implies that $\gamma/(1-e^{\gamma}) \leq 1+\gamma$ for all $\gamma \geq 1$. Thus, in terms of both worst case run time and approximation guarantee, the greedy algorithm outperforms the exchange algorithm.

4.4 Algorithms for K-match

In this section we apply the results of Section 4.3 to the problem (K-match). We do this by showing that the objective function of (K-match) satisfies the properties required by Conditions 4.1 and 4.2 in Section 4.4.1. Then in Section 4.4.2 we provide an example demonstrating that the analysis of the greedy algorithm is tight for (K-match).

4.4.1 K-match Cover Modularity

In this section we formally demonstrate the properties required by Conditions 4.1 and 4.2, thus completing the proofs of Theorems 4.1 and 4.2. First we show that the objective function of (K-match) satisfies both monotonicity and the non-decreasing average value property of Condition 4.1 item iii).

Lemma 4.4. *The objective function of* (K-match)*, $\sum_\omega \mu_\omega f_\omega(A)$ is monotone and satisfies $\forall S \subseteq U$, $\exists i \in S$ s.t. $f(S)/|S| \leq f(S \setminus \{i\})/|S \setminus \{i\}|$.*

Proof. For any scenario ω, we will show that $f_\omega(A)$, defined in (4.1), satisfies each property, from which the result follows by taking expectations. For monotonicity, note that the optimal solution to (4.1) for a set S is feasible for any set $T \supset S$, and thus the inequality $f_\omega(S) \leq f_\omega(T)$.

For the second property, let x^S_{ijk} denote the optimal solution to (4.1) for a set S. Then consider the element $i^* \in S$ which minimizes $\sum_{j \in D_\omega, k \in S_\omega} p_{ijk\omega} x^S_{ijk}$. We must have $\sum_{j \in D_\omega, k \in S_\omega} p_{ijk\omega} x^S_{i^*jk} \leq f_\omega(S)/|S|$, since otherwise we would have $\sum_{j \in D_\omega, k \in S_\omega} p_{ijk\omega} x^S_{ijk} > f_\omega(S)/|S|$ for all $i \in S$, which implies $f(S) = \sum_{i \in S, j \in D_\omega, k \in S_\omega} p_{ijk\omega} x^S_{ijk} > f_\omega(S)$, a contradiction. Then, since x^S_{ijk} for $i \in S \backslash \{i^*\}$ is feasible to (4.1) for $A = S \backslash \{i^*\}$, we have $f_\omega(S) - \sum_{j \in D_\omega, k \in S_\omega} p_{ijk\omega} x^S_{i^*jk} = \sum_{i \in S \backslash \{i^*\}, j \in D_\omega, k \in S_\omega} p_{ijk\omega} x^S_{ijk} \leq f_\omega(S \setminus \{i^*\})$. Therefore, $f_\omega(S) \leq f_\omega(S \setminus \{i^*\}) + \sum_{j \in D_\omega, k \in S_\omega} p_{ijk\omega} x^S_{i^*jk} \leq f_\omega(S \setminus \{i^*\}) + f_\omega(S)/|S|$. This is equivalent to $f_\omega(S)/|S| \leq f_\omega(S \setminus \{i^*\})/(|S| - 1)$. □

Next, we demonstrate the main result of this subsection, that the objective function of (K-match) is γ-cover modular with $\gamma = 2$.

Proposition 4.4. *The objective function of* (K-match)*, $\mathbb{E}_\omega[f_\omega(A)]$, is 2-cover modular.*

Proof. For any scenario ω, we will show that $f_\omega(A)$, defined in (4.1), is 2-cover modular, and the result then follows by taking expected values. With slight abuse of notation, in this proof let $\rho_i(S)$

denote the marginal value function for f_ω, i.e., $\rho_i(S) = f_\omega(S \cup \{i\}) - f_\omega(S)$. For a given set $A \subseteq U$, consider the dual of the LP relaxation of (4.1),

$$\begin{aligned} \min \ & \sum_{j \in D_\omega} \alpha_j + \sum_{k \in S_w} \alpha_k \\ \text{s.t. } & \alpha_j + \alpha_j \geq p_{ijk\omega}, \ \forall i \in A, \ j \in D_\omega, \ k \in S_\omega, \\ & \alpha_j, \alpha_k \geq 0, \ j \in D_\omega, \ k \in S_\omega. \end{aligned} \tag{4.3}$$

Then, for given $S, T \subseteq U$, we will demonstrate 2-cover modularity by constructing a feasible dual solution for the case when $A = S \cup T$. To do so, let x^S_{ijk} denote an optimal solution to (4.1) for the set S, so that $f_\omega(S) = \sum_{i \in S, j \in D_\omega, k \in S_\omega} p_{ijk\omega} x^S_{ijk}$. Then, define $\alpha^S_j = \sum_{i \in S, k \in S_\omega} p_{ijk\omega} x^S_{ijk}$ as the value of any match involving j, and $\alpha^S_k = \sum_{i \in S, j \in D_\omega} p_{ijk\omega} x^S_{ijk}$ as the value of any match involving k. We note that $\sum_{j \in D_\omega} \alpha^S_j + \sum_{k \in S_w} \alpha^S_k = 2 f_\omega(S)$. Further, we claim that α^S_j and α^S_k are feasible for (4.3) for $A = S$. To see this, assume to the contrary that there is an $i \in S$ and j, k such that $p_{ijk\omega} > \alpha^S_j + \alpha^S_k = \sum_{i \in S, k' \in S_\omega} p_{ijk'\omega} x^S_{ijk'} + \sum_{i \in S, j' \in D_\omega} p_{ij'k\omega} x^S_{ij'k}$. In other words, a match between j, k, and i would provide strictly higher value than the combined value of matches involving j and k in the x^S_{ijk} solution, contradicting its optimality.

Now for each $i \in T \setminus S$, define variables β^*_{ij} and β^*_{ik} as the optimal solution to the following LP

$$\begin{aligned} \min \ & \sum_{j \in D_\omega} \beta_{ij} + \sum_{k \in S_w} \beta_{ik} \\ \text{s.t. } & \beta_{ij} + \beta_{ik} \geq p_{ijk\omega} - (\alpha^S_j + \alpha^S_k), \ j \in D_\omega, \ k \in S_\omega, \\ & \beta_{ij}, \beta_{ik} \geq 0, \ j \in D_\omega, \ k \in S_\omega. \end{aligned} \tag{4.4}$$

Then define the dual solution to (4.3) for $A = S \cup T$ as $\alpha^{S \cup T}_j = \alpha^S_j + \sum_{i \in T \setminus S} \beta^*_{ij}$, and $\alpha^{S \cup T}_k = \alpha^S_k + \sum_{i \in T \setminus S} \beta^*_{ik}$. This solution is feasible to (4.3) for $A = S \cup T$ since for $i \in S$ and any j, k, we have

$$\alpha^{S \cup T}_j + \alpha^{S \cup T}_k \geq \alpha^S_j + \alpha^S_k \geq p_{ijk\omega},$$

since α^S_j and α^S_k are feasible for (4.3) for $A = S$. Further, for any $i \in T \setminus S$ and any j, k, we have

$$\alpha^{S \cup T}_j + \alpha^{S \cup T}_k \geq \alpha^S_j + \alpha^S_k + \beta^*_{ij} + \beta^*_{ik} \geq p_{ijk\omega},$$

by feasibility of the dual solution for the LP (4.4).

Next we claim that $\sum_{j\in D_\omega}\beta^*_{ij}+\sum_{k\in S_w}\beta^*_{ik}\le\rho_i(S)$ for $i\in T\setminus S$, which is equivalent to $\sum_{j\in D_\omega}\beta^*_{ij}+\sum_{k\in S_w}\beta^*_{ik}+f_\omega(S)\le f_\omega(S\cup\{i\})$. To see this, consider the dual of (4.4) for $i\in T\setminus S$:

$$\begin{aligned}\max\ &\sum_{j\in D_\omega,\ k\in S_\omega}(p_{ijk\omega}-(\alpha^S_j+\alpha^S_k))x_{ijk}\\ \text{s.t.}\ &\sum_{k\in S_\omega}x_{ijk}\le 1,\ j\in D_\omega,\\ &\sum_{j\in D_\omega}x_{ijk}\le 1,\ k\in S_\omega,\\ &x_{ijk}\ge 0,\ j\in D_\omega,\ k\in S_\omega.\end{aligned}\tag{4.5}$$

Following the same proof as Lemma 4.1, the constraint matrix of (4.5) is totally unimodular, thus there is an integer optimal solution, which is equal to $\sum_{j\in D_\omega}\beta^*_{ij}+\sum_{k\in S_w}\beta^*_{ik}$ by strong duality. Let x^*_{ijk} denote this optimal integer solution, and we observe that it must be binary since the program's constraints dictate that the variables be between zero and one. Thus, we have

$$\begin{aligned}&\sum_{j\in D_\omega}\beta^*_{ij}+\sum_{k\in S_w}\beta^*_{ik}+f_\omega(S)\\ =&\sum_{j\in D_\omega,\ k\in S_\omega}(p_{ijk\omega}-(\alpha^S_j+\alpha^S_k))x^*_{ijk}+\sum_{i'\in S,j\in D_\omega,\ k\in S_\omega}p_{i'jk\omega}x^S_{i'jk},\\ =&\sum_{j\in D_\omega,\ k\in S_\omega}p_{ijk\omega}x^*_{ijk}+\sum_{i'\in S,j\in D_\omega,\ k\in S_\omega}p_{i'jk\omega}x^S_{i'jk}(1-\sum_{k'\in S_\omega}x^*_{ijk'}-\sum_{j'\in S_\omega}x^*_{ij'k}),\\ \le&\sum_{j\in D_\omega,\ k\in S_\omega}p_{ijk\omega}x^*_{ijk}+\sum_{i'\in S,j\in D_\omega,\ k\in S_\omega}p_{i'jk\omega}x^S_{i'jk}(1-\sum_{k'\in S_\omega}x^*_{ijk'}-\sum_{j'\in S_\omega}x^*_{ij'k})^+.\end{aligned}$$

The last line is equal to the objective value in (4.1) for $A=S\cup\{i\}$ for the solution that sets $\hat{x}_{ijk}=x^*_{ijk}$ for i and $\hat{x}_{i'jk}=x^S_{i'jk}(1-\sum_{k'\in S_\omega}x^*_{ijk'}-\sum_{j'\in S_\omega}x^*_{ij'k})^+$ for $i'\in S$ and all j,k. Note that $\hat{x}_{i'jk}\le x^S_{i'jk}$ for $i'\in S$ and all j,k. We claim this solution is feasible for (4.1) for $A=S\cup\{i\}$. To see this, take the first constraint of (4.1) for some j and consider two cases, $\sum_{k'\in S_\omega}x^*_{ijk'}\in\{0,1\}$, which are exhaustive since the quantities $x^*_{ijk'}$ are binary. If $\sum_{k'\in S_\omega}x^*_{ijk'}=1$, then we have $\hat{x}_{i'jk}=0$ for all $i'\in S$ and $k\in S_\omega$, so that

$$\sum_{i'\in S\cup\{i\},k\in S_\omega}\hat{x}_{i'jk}=\sum_{k\in S_\omega}x^*_{ijk}=1,$$

and the constraint is satisfied. Otherwise, if $\sum_{k'\in S_\omega}x^*_{ijk'}=0$, then we have

$$\sum_{i'\in S\cup\{i\},k\in S_\omega}\hat{x}_{i'jk}=\sum_{i'\in S,k\in S_\omega}\hat{x}_{i'jk}\le\sum_{i'\in S,k\in S_\omega}x^S_{i'jk}\le 1,$$

where the last inequality follows from feasibility of $x^S_{i'jk}$ in (4.1) for $A = S$. An identical argument establishes feasibility of $\hat{x}_{ijk}$ for the second constraint of (4.1) when $A = S \cup \{i\}$ as well. This completes the proof of feasibility and thus also the proof of $\sum_{j \in D_\omega} \beta^*_{ij} + \sum_{k \in S_w} \beta^*_{ik} + f_\omega(S) \leq f_\omega(S \cup \{i\})$, i.e., $\sum_{j \in D_\omega} \beta^*_{ij} + \sum_{k \in S_w} \beta^*_{ik} \leq \rho_i(S)$.

Then, since $\alpha^{S \cup T}_j$ and $\alpha^{S \cup T}_k$ are feasible for (4.3) for $A = S \cup T$, by weak duality we have

$$
\begin{aligned}
f(S \cup T) &\leq \sum_{j \in D_\omega} \alpha^{S \cup T}_j + \sum_{k \in D_\omega} \alpha^{S \cup T}_k, \\
&= \sum_{j \in D_\omega} \alpha^S_j + \sum_{k \in D_\omega} \alpha^S_k + \sum_{i \in T \setminus S} \left(\sum_{j \in D_\omega} \beta^*_j + \sum_{k \in D_\omega} \beta^*_k \right), \\
&\leq 2f(S) + \sum_{i \in T \setminus S} \rho_i(S).
\end{aligned}
$$

□

With these results, the proofs of Theorems 4.1 and 4.2 follow immediately from Propositions 4.2 and 4.3, respectively.

Proof of Theorem 4.1. By Lemma 4.4 and Proposition 4.4, the objective function of (K-match) satisfies Condition 4.1, and therefore the result follows from Proposition 4.2. □

Proof of Theorem 4.2. By Lemma 4.4 and Proposition 4.4, the objective function of (K-match) satisfies Condition 4.2, and therefore the result follows from Proposition 4.3. □

4.4.2 Tight Example for Greedy K-match

In this Section, we present an example showing the bound in Theorem 4.2 is tight for the greedy algorithm, even for the deterministic version of the problem where $|\Omega| = 1$.

Example 4.2. *Consider an instance of* (K-match) *with* $|\Omega| = 1$*, so all input is deterministic, and we drop the subscript* ω*. We represent the matching profits using an undirected graph* $G(U \cup D \cup S, E)$*. If there is an edge between* j *and* i *and an edge between* k *and* i*, then* $p_{ijk} = 1$*, otherwise* $p_{ijk} = 0$*. We say a location node* i *is "connected" to a demand node* j*, when there is an edge between* i *and* j *in the graph (similarly for a supply node* k*).*

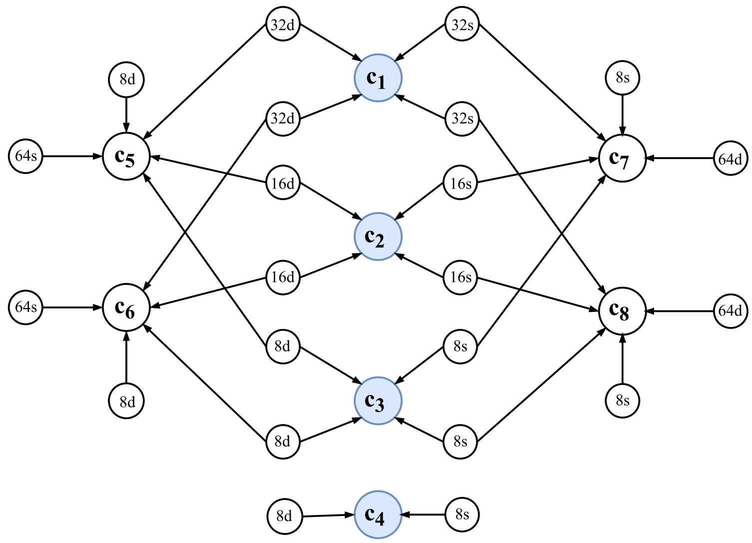

Figure 4.2: Example 4.2 for $K = 4$

Given an even integer K, let there be $2K$ location nodes in U. For $1 \leq i \leq K$, location node i is connected to m_i demand nodes and m_i supply nodes that are not connected to any other location nodes i' with $1 \leq i' \leq K$, where $m_i = (K-2)^{i-1}K^{K-i}$.

For $K+1 \leq i \leq K+K/2$, location node i is connected to K^{K-1} supply nodes and $(K-2)^{K-1}$ demand nodes that are not connected to any other location nodes, and overlaps with the demand nodes connected to the first $K-1$ location nodes in the following way. For $K+1 \leq i \leq K+K/2$ and $1 \leq i' \leq K-1$, location node i is connected to $2m_{i'}/K$ of the demand nodes that location node i' is connected to, that are not connected to any other location nodes $\tilde{i}$ with $K+1 \leq \tilde{i} \leq K+K/2$.

Similarly, for $K+K/2+1 \leq i \leq 2K$, location node i is connected to $K^K - 1$ demand nodes and $(K-2)^{K-1}$ supply nodes that are not connected to any other location nodes, and overlaps with the supply nodes connected to the first $K-1$ location nodes in the following way. For $K+K/2+1 \leq i \leq 2K$ and $1 \leq i' \leq K-1$, location node i is connected to $2m_{i'}/K$ of the supply nodes that location node i' is connected to, that are not connected to any other location nodes $\tilde{i}$ with $K+K/2+1 \leq \tilde{i} \leq 2K$.

Example 4.2 is illustrated in Figure 4.2 for the case when $K = 4$, where, for example, a circle labeled $32d$ represents 32 demand nodes, each with the connections depicted. Blue shaded location nodes are those selected by greedy.

Proposition 4.5. *For the instance in Example 4.2 for given K, the ratio of optimal profit to the profit obtained by Algorithm 7 is at least*

$$\frac{2}{1-(1-\frac{2}{K})^K}.$$

Proof. Let $U^d = \{i|K+1 \leq i \leq K+K/2\}$ represent the set of location nodes that overlap with the demand nodes connected to the first $K-1$ location nodes, and $U^s = \{i|K+K/2+1 \leq i \leq 2K\}$ represent the set of location nodes that overlap with the supply nodes connected to the first $K-1$ location nodes. For location node i, let d_i^t (s_i^t) represent the number of unmatched demand (supply) nodes connected to location node i at the beginning of iteration t of the greedy algorithm. Let $m_i^t = \min(d_i^t, s_i^t)$, which represents the number of matches that can be added by selecting location node i in iteration t of the greedy algorithm. Then we claim that greedy selects location node $i = t$ at iteration t, and that $d_i^t \leq (K-2)^{t-1}K^{K-t}$ for each $i \in U^d$ and $s_i^t \leq (K-2)^{t-1}K^{K-t}$ for each $i \in U^s$, which we prove by induction.

First we demonstrate that location node $i \in U^d$ ($i \in U^s$) is connected to K^{K-1} demand (supply) nodes at the start of the algorithm. This follows from summing the counts of all the demand (supply) nodes which location node i is connected to:

$$\begin{aligned}(K-2)^{K-1} + \sum_{i'=1}^{K-1} \frac{2m_{i'}}{K} &= (K-2)^{K-1} + \frac{2}{K}\sum_{i'=1}^{K-1}(K-2)^{i-1}K^{K-i},\\ &= (K-2)^{K-1} + \frac{2K^{K-1}}{K}\sum_{i'=1}^{K-1}(1-\frac{2}{K})^{i-1},\\ &= (K-2)^{K-1} + \frac{2K^{K-1}}{K}\left(\frac{1-(1-\frac{2}{K})^{K-1}}{\frac{2}{K}}\right),\\ &= (K-2)^{K-1} + K^{K-1}\left(1-(1-\frac{2}{K})^{K-1}\right),\\ &= K^{K-1}.\end{aligned}$$

Then, for $t = 1$, it is clear that all location nodes $i \in U^d \cup U^s$ have $d_i^1 = s_i^1 = K^{K-1} = (K-2)^{t-1}K^{K-t}$. Further, all location nodes $i \leq K$ have $m_i^t \leq K^{K-1}$, while location node one has $m_1^t = K^{K-1}$, so location node one is an optimal greedy selection. Now assume the inductive hypothesis up to iteration t. Then, a given location node $i \in U^d$ is connected to $2m_t/K$ of the demand nodes that are matched on location node t that is selected in iteration t by the induction hypothesis. Thus, also by the induction hypothesis, we have

$$\begin{aligned} d_i^{t+1} &= d_i^t - \frac{2m_t}{K}, \\ &\leq (K-2)^{t-1}K^{K-t} - \frac{2}{K}(K-2)^{t-1}K^{K-t}, \\ &= (K-2)^t K^{K-t-1}. \end{aligned}$$

For $i \in U^s$, the same argument applied to supply nodes implies $s_i^{t+1} \leq (K-2)^t K^{K-t-1}$. Thus, in iteration $t+1$, any location node $i \in U^d \cup U^s$ cannot add more than $(K-2)^t K^{K-t-1}$ matches. Further, any currently unselected location node $i \leq K$ has $m_i^{t+1} \leq (K-2)^t K^{K-t-1}$, while location node $t+1$ has $m_i^{t+1} = (K-2)^t K^{K-t-1}$, and so location node $t+1$ is an optimal greedy selection in iteration $t+1$. This completes the induction proof, and it follows that greedy selects the first K location nodes, which achieves the following number of matches

$$\frac{K^K \left(1 - (1 - \frac{2}{K})^K\right)}{2}$$

Note that choosing the last K location nodes would each give K^{K-1} matches, for a total of K^K matches, which proves the claim. □

4.5 Conclusion

In this chapter we introduce the K-match problem, which models the selection of a limited number of locations in order to maximize the value of matches made between stochastic supply and demand. We show that the problem generalizes the classic K-cover problem, and therefore is NP-hard to approximate better than a factor $1/(1-e^{-1}) \approx 1.58$. Then we consider two algorithms for approximating K-match, exchange and greedy, and provide constant factor approximations for each. For the exchange algorithm we prove an approximation factor of $3 - 2/K + \epsilon$ for any $\epsilon > 0$, while showing the greedy algorithm offers a better factor of $2/(1-e^{-2}) \approx 2.31$, which we show is tight through a family of examples. In order to establish these results we generalize the notion of a submodular set function with a new definition of γ-cover modularity, which we show implies general approximation guarantees for the algorithms we consider.

Chapter 5

Understanding the Value of Fulfillment Flexibility in an Online Retailing Environment

5.1 Introduction

In recent years, online retail sales have soared to unprecedented levels. In 2017, online sales totaled over 452 billion dollars in the U.S. (US Department of Commerce, 2018), and over 5.48 trillion yuan in China (National Bureau of Statistics of China, 2018). Along with rapid growth, online retailing has also experienced the emergence of fierce competition, with e-retailers competing in price, service, and marketing. This market environment offers a fertile ground for leveraging strategic operational improvements, as better service capabilities may not only lead to decreased costs, but also increased revenues through attracting new customers.

In this chapter, we illustrate how to use fulfillment flexibility to improve online retailers' distribution systems through an industrial collaboration. Our industrial partner is one of the largest e-retailers in China with over 300 million active users. A core strategy of the e-retailer is its 211 program, where it promises to deliver orders placed before 11am within the same day, and to deliver orders placed between 11am to 11pm before 3pm of the following day. To support its 211 program, the e-retailer operates under a two-tier distribution system, which consists of regional-distribution-centers (RDCs) and front-distribution-centers (FDCs). Each FDC is responsible for fulfilling demand from its local district, while an RDC is responsible for fulfilling demand from its local district, as well as allocating inventories to a region including multiple FDCs. Under this distribution system, customers in the districts covered by an RDC or an FDC are offered the 211 program.

In practice, the 211 program may not be offered to customers in certain districts at certain times because of stockouts. When a stock keeping unit (SKU) in a district's distribution center (DC) is stocked out, the e-retailer can still offer the SKU to that district with longer delivery time, if the district can be served by another DC. The operation of fulfilling demand of one DC using the inventories of another DC is known as *spillover*. In our industrial partner's current distribution system, spillover is only allowed from the RDC to other FDCs. While the current RDC to FDCs structure is simple to implement, a significant amount of lost sales may occur due to limited inventory and demand fluctuations. Moreover, the fulfillment costs of fulfilling an FDC's demand from the RDC may also be substantial. Therefore, in this chapter, we propose to reduce the combined lost sales and fulfillment costs by adding demand fulfillment flexibilities at the FDCs, such that the demand from an individual FDC can be spilled over to some of the other FDCs.

To achieve our objective of understanding the value of fulfillment flexibility, we propose a model which takes real world data from our industrial partner as inputs, and use the model to evaluate the benefits of additional fulfillment flexibility. To highlight the importance of fulfillment policies under flexibility, we use our model to demonstrate that a greedy fulfillment policy may increase the total supply chain costs even compared to the system with less fulfillment flexibility. While finding an effective fulfillment policy is crucial for understanding the benefit of fulfillment flexibilities, it poses several significant challenges. First, finding the optimal fulfillment policy leads to solving a complex dynamic program, which is computationally intractable due to the curse of dimensionality. Second, to maintain customer satisfaction, the e-retailer is always committed to local fulfillment when orders arrive at a local DC and that DC has inventory, but this adds additional constraints to our fulfillment policy. Third, because the fulfillment decision is made on the fly by inventory managers, the policy needs to be intuitive and straightforward.

These challenges motivate our development of a policy that maintains greedy fulfillment for local demand, while placing threshold limits on spillover demand, which we call the *spillover limit* policy. The spillover limit policy is shown to have good theoretical properties and perform well in simulations. We then propose an inventory allocation scheme and use both the allocation scheme and the spillover limit policy to evaluate the value of fulfillment flexibility for the e-retailer.

The rest of the chapter is organized as follows. In the rest of Section 5.1, we introduce some general notations, define our model, and provide the literature review. In Section 5.2, we define the greedy fulfillment policy, discuss its limitations, and use it to motivate the spillover limit policy. In Section 5.3, we show that the spillover limit policy is asymptotically optimal under a general setting. In Section 5.4, we describe the motivation and the procedure of our inventory allocation policies. In Section 5.5, we use the data from our industrial partner as well as machine learning tools to estimate parameters for our model, use numerical simulations to identify an effective spillover limit policy, and finally evaluate the benefits of fulfillment flexibility. Proofs of mathematical results are provided in Section 5.6. We conclude this chapter in Section 5.7.

5.1.1 General Notations

This chapter adopts the following set of notations. A vector/matrix and its scalar components are distinguished using bold letters and unbold letters respectively, e.g., given $\mathbf{f} \in \mathbb{R}^{n^2}$, the (j, j')-th entry of $\mathbf{f}$ is denoted using $f_{jj'}$. For clarity, we often distinguish between a random variable and its realization using capital and lowercase letters, respectively. For a random variable X, $\mathbb{E}[X]$ denotes the expectation of X. Letters $\mathbb{R}$ and $\mathbb{R}_+$ are used to denote the set of reals and nonnegative reals, respectively. The function $(y)^+$ denotes the positive part of y, which is equivalent to $\max(y, 0)$. Moreover, the chapter adopts the Big O/Small O notations, where $f(T) = O(g(T))$ implies that there exists some M such that $f(T) \leq Mg(T)$ for sufficient large T, while $f(T) = o(g(T))$ implies that $f(T) < Mg(T)$ for sufficient large T and all $M > 0$.

5.1.2 Model Framework

We propose a continuous-time dynamic model driven by the data and practice of our industrial partner. In our model, the e-retailer considers its SKUs separately. The e-retailer decides the allocation of its inventories at time 0, while demand arrives during the interval $[0, T]$. We assume T to be a positive integer, which can be naturally interpreted as T days; and for any positive integer $t \leq T$, we define the interval $[t-1, t]$ as the duration of day t. There are $N+1$ distribution centers (DCs) and $N+1$ districts, indexed by $0, 1, \ldots, N$, where DC j is primarily responsible for the orders arriving at district j. For succinctness, we often use demand for DC j to describe the demand of district j. Orders arrive at DCs sequentially, and the demand distribution for DC j during day t is represented by D_j^t, which is assumed to be i.i.d. across t unless stated otherwise. For concreteness, we assume that for each day t, the demand arrivals occur uniformly at random over the interval $[t-1, t]$. The daily uniform arrival assumption is actually not needed for our analysis, as the performances of our policies only depend only on the amount and sequence of order arrivals to each DC. The notion of "daily" demand in our model is important in practice, as it is much easier to forecast the daily demand instead of the highly nonstationary demand arrival process each day. Finally, we let $D_j(T) = \sum_{t=1}^{T} D_j^t$ represent the total demand arriving at DC j across T days.

At time zero, the e-retailer allocates inventories to different DCs. After the allocation is finished at time 0, a fulfillment decision will be made each time a demand arrives at a DC. In general, the e-retailer has three types of fulfillment options for a demand arrival: local fulfillment - satisfying the demand from its origination DC, spillover fulfillment - satisfying the demand from a DC different from its origination; and lost sale - allowing the demand to be lost. In addition to inventory availability, the ability to perform spillover fulfillment also depends on whether DC j has the flexibility to satisfy demand at DC j'. The flexibility structure of inter-DC fulfillment is described by a set of ordered pairs denoted by A, where $(j, j') \in A$ if and only if DC j has the flexibility to satisfy demand at DC j'. Spillover fulfillment typically lowers demand due to longer delivery time. To model this, we assume the customer may abandon its intended purchase when demand is not filled locally. In particular, each unit of demand at DC j that is not filled locally will be abandoned by the customer independently with probability α_j and become a lost sale.

The unit fulfillment cost of satisfying demand at DC j' from DC j (allowing $j' = j$) is denoted by $c_{jj'}$, while the unit lost sales cost at DC j is denoted by p_j. The e-retailer's goal is to minimize the sum of total fulfillment and lost sales cost. Throughout the chapter, we use x_j and r_j to denote the inventory held at DC j before and after the inventory allocation, respectively, and make the following mild assumption about $\mathbf{c}$ and $\mathbf{p}$.

Assumption 5.1.

$$(j, j) \in A,\, p_j \geq c_{jj} > 0,\ \forall\, 0 \leq j \leq N, \tag{5.1}$$

$$c_{j'j'} \leq c_{jj'},\ \forall (j, j') \in A, \tag{5.2}$$

$$p_{j'} + c_{jj} \leq p_j + c_{jj'},\ \forall (j, j') \in A. \tag{5.3}$$

In Assumption 5.1, Equation (5.1) implies that the cost of lost sales and local fulfillment are always positive and a lost sale is more costly than local fulfillment, Equation (5.2) implies that it is always cheaper to fulfill an order locally instead of spillover, and Equation (5.3) implies that the firm is always willing to trade a lost sale at j and a spillover fulfillment at j' for a local fulfillment at j and a lost sale at j'. Intuitively, Assumption 5.1 indicates that the e-retailer prefers local fulfillment over lost sales and spillover fulfillment at any DC.

The model proposed in this chapter was developed in conjunction with expert practitioners in order to capture the most relevant factors affecting the fulfillment problem faced by e-retailers. While the model makes some simplifying assumptions, experts from our industrial partner agree that the model is detailed enough to capture the first order effect of fulfillment flexibilities, and therefore is useful for assessing important strategic decisions. For example, the model treats SKUs separately, ignoring the fact that an order of multiple SKUs may be bundled to reduce shipping costs. Also by looking at a model which ends at time T, we are effectively focusing on the allocation and fulfillment decisions during one replenishment cycle, without considering how leftover inventories will affect the next replenishment cycle. While practically important, each of these complicating model features was judged to not have a critical impact on the value of fulfillment flexibility, and so are not considered in our model.

Our model has the distinctive feature that the e-retailer is committed to local fulfillment, i.e., the retailer will always fulfill an order with the local DC if that DC has inventory, and unavailability of local fulfillment may cause order abandonment. To the best of our knowledge, this feature has not appeared in the online retailing literature. While this feature is incorporated in our model to align with our industrial partner's business practice, other online retailers may exhibit a similar preference towards local fulfillment. Thus, while our numerical simulation uses data from our partner, our proposed policies and methodologies have the potential to be applied to other e-businesses.

5.1.3 Related Literature

Our work is related to the recent efforts in studying online order fulfillment. In general, finding the exact optimal fulfillment decisions for even the single SKU problem is difficult (Acimovic and Graves, 2015), and researchers have resorted to various innovative fulfillment policies. Acimovic and Graves (2015) consider the problem of minimizing shipping costs for a single SKU, and propose an LP-based policy that incorporates the forecast of future orders. Xu et al. (2009) consider the problem of minimizing the number of bundled shipments with multiple SKUs, and propose a policy to reevaluate and reassign orders through solutions of integer programs periodically. Jasin and Sinha (2015) consider a problem similar to Xu et al. (2009) that incorporates different costs of bundled shipments. Finally, Lei et al. (2018) consider fulfillment and pricing jointly. Our model differs from the aforementioned papers in two significant ways. First, the retailer in our model is committed to "local fulfillment." Second, while substantial parts of this chapter focus on fulfillment policy, we also apply our policies to study the strategic decision of adding partial fulfillment flexibilities, whereas papers in the existing literature assume that the retailer has complete fulfillment flexibilities.

The consideration of designing a flexible fulfillment network is related to the concept of process flexibility, which dates back to Jordan and Graves (1995a). Because of the vast literature in that area, we only mention the papers that are related to online retailing. Asadpour et al. (2018) consider online resource allocation under the classical chaining structure. Xu et al. (2018) extend Asadpour et al. (2018) to general flexibility structures, and provide methods for designing flexible systems under the online allocation setting. Unlike our model, the decision makers in Asadpour et al. (2018) and Xu et al. (2018) are interested in minimizing just lost sales and do not consider the additional costs of spillover fulfillments.

The fulfillment policies we study are also generally related to capacity control in the network revenue management literature. While our fulfillment problem differs significantly from problems in revenue management, the structure of our spillover limit policy is motivated by the protection levels in revenue management originated from Littlewood (1972), and is closely connected with the inventory rationing literature started by Topkis (1968) and Kaplan (1969). Moreover, our analysis of the spillover limit policy is inspired by the asymptotic analysis of Gallego and Van Ryzin (1997) and Cooper (2002) using the solution of a linear program (LP). We refer interested readers to Talluri and Van Ryzin (2006) for details for protection levels and LP-based analysis for network revenue management.

We also briefly survey other relevant literature in online retailing other than online order fulfillment. Acimovic and Graves (2017) proposed policies for inventory replenishment under a greedy fulfillment policy with replenishment leadtime. Chen (2017) studies both the placement (which SKUs to hold at which DCs) and the replenishment decision at DCs. Govindarajan et al. (2018) studies the joint allocation and fulfillment problem for an omnichannel retail network with both brick-and-mortar stores and an online channel. While Govindarajan et al. (2018) also propose a policy with threshold limits on fulfillments, their threshold limits the quantity of fulfillment for all online orders where online orders are fulfilled in batches, whereas the policy proposed in this chapter limits the number of spillovers and our model makes fulfillment decisions immediately after each order arrives. Finally, Wei et al. (2017) considers dynamic order fulfillment with delivery deadlines and expedited shipment options.

5.2 Fulfillment Policies

To understand the value of fulfillment flexibility, we first find effective fulfillment policies. In general, finding an optimal fulfillment policy via dynamic programming is intractable computationally due to the curse of dimensionality. In the special case where there is no abandonment and uniform fulfillment costs across DCs, the optimal fulfillment problem reduces to the online stochastic b-matching problem, which has received considerable attention in academia due to its application in online advertising display. Even for the b-matching problem, the optimal policy is difficult to find, and researchers have focused on developing policies with approximation factors and competitive ratios, e.g., Karp et al. (1990) and Feldman et al. (2009), Manshadi et al. (2012), and Jaillet and Lu (2013). As a result, we will focus on finding simple policies that can be easily implemented by online retailers, provide good theoretical guarantees, and perform well in numerical studies. First, we discuss the greedy fulfillment policy.

5.2.1 Greedy fulfillment policy

Perhaps the simplest policy to describe and implement is the greedy policy, which fulfills each arriving demand from the least expensive DC with available inventory at the time of arrival (subject to any random abandonment that may occur in the process of spillover fulfillment). This policy is attractive for its simplicity and transparency, as well as the relatively low amount of information it requires to implement: a manager need only keep an updated list of the inventory level at each DC. Currently, our industrial partner uses the greedy policy for its fulfillment and similar greedy fulfillment policies are also adopted by other e-retailers (Acimovic and Graves, 2017).

However, greedy fulfillment can lead to sub-optimal performance, using inventory that should have been saved for local fulfillment to fulfill demand at another DC while paying a premium for spillover fulfillment. To illustrate this phenomenon, we consider a simple example.

Example 5.1. *Consider a fulfillment network with 3 DCs ($N = 2$). Recall that each DC is responsible for orders from its own district, but may also be able to fulfill orders in other districts depending on the flexibility structure. There are T days. Orders arrive at each district as independent Poisson processes, where DCs 1 and 2 have a rate of 1 per day, while DC 0 has a rate of 0. We assume DC 0 starts with an infinite amount of inventory, DC 1 starts with T inventory, and DC 2 starts with 0 inventory. Motivated by our industry collaboration, we consider two flexibility structures: one is* *called "dedicated" which only allows spillover from DC 0 to either DC 1 or 2, while the "flexible" structure augments the dedicated structure with full flexibility between DCs 1 and 2, i.e., it adds arcs* $(1,2)$ *and* $(2,1)$ *(see Figure 5.1). The abandonment rates are set to* $\alpha_j = 0$*, i.e., customers do not abandon orders upon spillover.*

For each flexibility structure, we analyze the expected cost of the greedy fulfillment policy. Note that no lost sales will occur in the system under a greedy policy, as DC 0 always has enough inventory to fulfill all demand. Suppose that local fulfillment cost is defined as $c_{jj} = c_l$ *for all* j*, the DC 0 to DC 1 or 2 spillover costs are* $c_{01} = c_{02} = c_r$*, and the spillover costs between DCs 1 and 2 are* $c_{12} = c_{21} = c_f \leq c_r$*. As* T *grows large, the total cost for the "dedicated" structure is equal to the local fulfillment cost, plus the spillover cost of fulfilling all demand at DC 2 using DC 0, and plus the spillover cost of fulfilling demand at DC 1 using DC 0 when DC 1 runs out of inventory. Therefore, the total cost for the "dedicated" structure is*

$$c_l \mathbb{E}[\min(D_1(T), T)] + c_r \mathbb{E}[(D_1(T) - T)^+] + c_r \mathbb{E}[D_2(T)] = c_l T + c_r T + O(\sqrt{T}).$$

To compute the total costs for the "flexible" structure, note that DC 1 will on average use $\mathbb{E}[\min(D_1(T) + D_2(T), T)]$ *amount of inventories, because demand from both DC 1 and DC 2 is first fulfilled with inventories at DC 1. Because the demand for DC 1 and DC 2 have the same independent Poisson process, the average amount of DC 1 inventories consumed,* $\mathbb{E}[\min(D_1(T)+D_2(T), T)]$*, will be split equally by DC 1 and DC 2. Note that* $\mathbb{E}[\min(D_1(T) + D_2(T), T)] = T - O(1)$*, thus, the total cost for the "flexible" structure, is equal to*

$$0.5(c_f + c_l)(T - O(1)) + c_r \mathbb{E}[D_1(T) + D_2(T)] - c_r(T - O(1)) = 0.5 c_f T + 0.5 c_l T + c_r T + O(1).$$

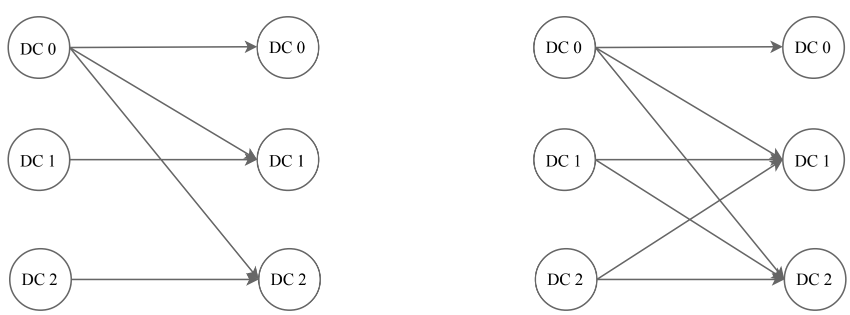

Figure 5.1: Systems in Example 5.1: Left is dedicated, right is flexible

Example 5.1 shows that adding flexibility while maintaining the greedy fulfillment policy can actually increase costs by as much as 50% if c_r/c_l is large and $c_f \approx c_r$. Therefore, we next turn our attention to a new class of fulfillment policies.

5.2.2 Spillover limit policy

The reason greedy performed poorly in Example 5.1 is because DC 1 used its inventories to fulfill the demand of DC 2, even though it should have kept its inventory for its local orders in the future. Motivated by this observation, we propose a class of spillover limit policies, which fulfill local demand greedily, while placing threshold limits on spillover fulfillment. Intuitively, the threshold limits protect some inventories for just local fulfillment, and therefore avoid too many spillovers in situations similar to Example 1.

A specific spillover limit policy sets two type of limits. For $j \neq j'$, the policy sets a limit $\theta_{j'j}$ on the amount of spillover demand at DC j that may be filled with inventory from DC j'. Further, for each j, the policy also sets a limit θ_j^{DC} on the aggregate DC j fulfillment of spillover demand from any other DC $j' \neq j$. The policy then makes fulfillment decisions as follows. When a demand arrives at DC j, the demand is fulfilled locally if DC j has inventory. Otherwise, the policy fills the demand from the lowest cost flexible DC j' which has inventory, has not filled more than $\theta_{j'j}$ units of spillover demand from DC j already, and has not filled more than $\theta_{j'}^{DC}$ units of spillover demand from any DC (subject to any random abandonment that may occur in the process of spillover fulfillment). If no such flexible DC exists, the demand is lost.

The spillover limit policy is attractive for multiple reasons. First, it puts practical limits on the amount of spillover that can occur, allowing the e-retailer to avoid excessive spillover costs. Second, the policy improves over greedy with minimal added complexity; the only additional parameters are the spillover limits, and the only additional informational requirements are recording the amount of spillover on each non-local arc in A. As a result, the policy allows each DC to decide whether to fulfill in a decentralized manner, making it easy to implement in practice. Second, the policy satisfies the e-retailer's practical constraint that it may not withhold inventory from local fulfillment due to customer delivery expectations or service level requirements. Lastly, the performance of the policy only depends on the sequence and the total number of order arrivals at each DC, making the policy robust to the timing of order arrivals throughout the day, which is a highly non-stationary stochastic process for our industrial partner (as is common in the context of e-retailing).

In addition, note that the greedy policy is a special case of a spillover limit policy, with $\theta_{j'j}$ and θ_j^{DC} set to infinity for all j, j'. Much discussion in this chapter focuses on identifying a spillover limit policy that improves upon the greedy policy. More specifically, in Section 5.3, we prove that a spillover limit policy is superior to greedy in asymptotic settings, and in Section 5.5.2, we identify an effective spillover limit policy that consistently outperforms greedy in numerical simulations.

5.3 Analysis of Fulfillment Policies

In this section, we analyze and compare the greedy and spillover limit policies. Recall that r_j represents the inventory held at DC j after allocation, and D_j^t, the demand at DC j during day t, is i.i.d. across $t = 1, \ldots, T$. We will analyze greedy and spillover limit policies under an asymptotic regime where both T (the number of days) and r_j are scaled to infinity at the same rate. This scaling regime is commonly used in network revenue management to identify asymptotically optimal policies (e.g., Gallego and Van Ryzin (1997), Cooper (2002)). Throughout this section, for $j = 0, 1, \ldots, N$, we let $r_j = Tb_j$ for some base inventory level b_j, and let the demand for DC j on each day have mean μ_j and finite variance σ_j^2. Throughout this section, we also assume that Assumption 5.1 introduced in Section 5.1.2 holds, which intuitively, suggests that the e-retailer prefers local fulfillment over lost sales and spillover at each DC j.

Under fixed flexibility structure and daily demand distributions, we will analyze the expected cost of the optimal and greedy policy over T days, which are denoted by $C^*(T)$ and $C^{Greedy}(T)$, respectively. Because the demand distribution at each DC is i.i.d. across days, and the cost of both lost sales and fulfillment are strictly positive, $C^*(T)$ increases linearly with T. As a result, we say that a policy is *asymptotically optimal* if its expected cost is upper-bounded by $C^*(T) + o(T)$. First, we describe a setting where the greedy policy is asymptotically optimal.

Lemma 5.1. *Suppose that $b_j \geq \mu_j$ for all $0 \leq j \leq N$, then*

$$\frac{C^{Greedy}(T)}{C^*(T)} \leq 1 + O\big(T^{-\frac{1}{2}}\big).$$

Lemma 5.1 shows that the greedy policy should perform well when the condition $b_j \geq \mu_j$ for all $0 \leq j \leq N$ holds, and we indeed observe this in Section 5.5.3, when numerical simulations are performed using data from the real world. However, the requirement of $b_j \geq \mu_j$ for all $0 \leq j \leq N$

is somewhat stringent, and when the condition fails, greedy is neither asymptotically optimal, nor does it perform well in simulations. For example, the greedy policy for the more flexible structure in Example 5.1 is not asymptotically optimal, as $\lim_{T\to\infty} C^{Greedy}(T)/C^*(T)$ is at least $(0.5(c_f + c_l) + c_r)/(c_l + c_r)$.

Next, we will show that there is a spillover limit policy achieving asymptotic optimality under more general settings. We start by stating the condition implying the asymptotic optimality of the spillover limit policy.

Assumption 5.2. *For any $(j, j'), (j', j'') \in A$, we have $(j, j'') \in A$ and*

$$c_{jj''} + c_{j'j'} \leq c_{jj'} + c_{j'j''}. \tag{5.4}$$

In other words, Assumption 5.2 implies that in the virtual scenario where there are simultaneously two orders, one at DC j' one at DC j'' and DC j' has inventory, then DC j' should always prioritize satisfying its local order first. Unlike the condition in Lemma 5.1, Assumption 5.2 depends only on the costs of fulfillment, but not on the initial inventory $\mathbf{r}$. This is advantageous, as it implies that Assumption 5.2 can be checked independently of inventory allocation policies.

There are many systems satisfying Assumption 5.2. In particular, both systems in Example 5.1 satisfy Equation (5.4). Moreover, Assumption 5.2 is also satisfied by the flexible systems we investigate in Section 5.5. Next, we formally state the asymptotic optimality of a spillover limit policy under Assumption 5.2. We use π to denote a spillover limit policy, and $C^\pi(T)$ to denote the policy's expected cost.

Theorem 5.1. *Under Assumption 5.2, there exists a spillover limit policy π such that*

$$\frac{C^\pi(T)}{C^*(T)} \leq 1 + O\big(T^{-\frac{1}{2}}\big).$$

Theorem 5.1 provides theoretical justification for using a spillover limit policy. While finding an effective spillover limit policy requires effort, it is significantly easier compared to searching over the space of all dynamic fulfillment policies. In Section 5.5, we demonstrate that there exists a spillover limit policy which is very effective for our proposed flexible system in non-asymptotic regimes under the data from our industrial project.

The intuition behind the proof of Theorem 5.1 is similar in spirit to the asymptotic analysis in network revenue management (NRM), e.g., (Gallego and Van Ryzin, 1997; Cooper, 2002; Talluri and Van Ryzin, 2006). However, the analysis is significantly more complex because of the fact that the spillover limit policy is locally greedy, i.e., each DC must fulfill its local order when inventories are available. In fact, the locally greedy constraint makes Assumption 5.2 crucial. In Section 5.6.4, we provide a system that violates Assumption 5.2, and any locally greedy policy for that system fails to be asymptotically optimal.

We describe the roadmap for establishing Theorem 5.1. First, we consider the following optimization problem.

$$
\begin{aligned}
\min_{\mathbf{f},\mathbf{l}} \quad & \sum_{j} \left(p_j l_j + \sum_{j'} c_{jj'} f_{jj'} \right) \\
\text{s.t.} \quad & l_j + \sum_{0 \le j' \le N} f_{j'j} \ge T\mu_j, \ \forall\ 0 \le j \le N, \\
& \sum_{j' \ne j} f_{j'j} \le (1-\alpha_j)(T\mu_j - f_{jj}), \forall\ \ 0 \le j \le N, \\
& \sum_{0 \le j' \le N} f_{jj'} \le r_j, \ \forall\ \ 0 \le j \le N, \\
& f_{jj'} = 0, \ \forall\ (j, j') \notin A, \\
& f_{jj'} \ge 0, \ \forall\ (j, j') \in A,\ l_j \ge 0, \forall\ 0 \le j \le N.
\end{aligned}
\tag{5.5}
$$

We name the above optimization problem the deterministic (demand) linear program (DLP) for fulfillment and use C^{DLP} to denote its objective value, as the optimization solves the optimal fulfillment problem with deterministic demands and constant abandonment rate for spillovers. Like analysis for NRM problems, our proof establishes C^{DLP} to be a lower bound for $C^*(T)$, uses the optimal solution of the DLP to propose a spillover limit policy, and finally, shows that the expected cost of that policy is bounded by $C^{DLP} + O(\sqrt{T})$. There are two technical hurdles for Theorem 5.1 that were not present in the NRM analysis. First, in addition to demand uncertainty, our model also incorporates order abandonment, which directly depends on the fulfillment policy, and complicates establishing both lower and upper bounds on the cost of the online problem. To establish that C^{DLP} is a lower bound for $C^*(T)$, we need to carefully unfold the information structure, relax only some of the information and apply Jensen's inequality at different stages of the decision tree. To establish an upper bound on the expected lost cost under random abandonment, we require a careful treatment of conditional expectations. Second, because DCs cannot hold back their inventories for their local demand, we need to use Assumption 5.2 to derive a condition on the optimal solution of the DLP, which in turn allows us to bound the deviation of the spillover limit policy compared to the optimal solution of the DLP.

5.4 Allocation Policies

In this section, we discuss policies for inventory allocation at time 0, before any order arrives. We will restrict our attention to two simple allocation policies. To model the allocation constraints, we assume that the inventory vector after allocation, $\mathbf{r}$, must be contained in some feasible set $\mathcal{R}$. Throughout the chapter, we assume that $\mathcal{R}$ is polyhedral.

Deterministic Demand Allocation Policy. The first allocation policy we consider is to allocate the inventories assuming that the retailer will see a deterministic demand equal to the expected value of the demand distribution, and a constant fraction α_j of the spillover demand at DC j. Under the deterministic demand allocation policy, we solve the following optimization problem:

$$\min_{\mathbf{r}\in\mathcal{R}} f(\mathbf{r}), \text{where } f(\mathbf{r}) \text{ is defined as:} \tag{5.6}$$

$$\begin{aligned} f(\mathbf{r}) = \min_{\mathbf{f},\mathbf{l}} \quad & \sum_{j} \left(p_j l_j + \sum_{j'} c_{jj'} f_{jj'} \right) \\ \text{s.t.} \quad & l_j + \sum_{0\le j'\le N} f_{j'j} \ge T\mu_j, \ \forall\ 0 \le j \le N, \\ & \sum_{j'\ne j} f_{j'j} \le (1-\alpha_j)\big(T\mu_j - f_{jj}\big), \forall\ \ 0 \le j \le N, \\ & \sum_{0\le j'\le N} f_{jj'} \le r_j, \ \forall\ \ 0 \le j \le N, \\ & f_{jj'} = 0, \ \forall\ (j,j') \notin A, \\ & \mathbf{f}, \mathbf{l} \ge \mathbf{0}. \end{aligned} \tag{5.7}$$

The optimization problem defined by (5.6) can be solved as a linear program (LP). In the optimization problem, variable $\mathbf{r}$ models inventory allocation, while variables $\mathbf{f}$ and $\mathbf{l}$ model fulfillment and lost sales under deterministic demand.

The intuition behind the deterministic demand allocation policy is that $f(\mathbf{r})$ should be a reasonable approximation for the expected lost sales and fulfillment cost, with inventory position $\mathbf{r}$ under an effective inventory policy. Indeed, Lemma 5.4 in Section 5.6.1 shows that $f(\mathbf{r})$ is a lower bound on the expected lost sales and fulfillment cost under any inventory policy. If costs satisfy Assumption 5.2, the deterministic demand allocation policy is asymptotically optimal in the following sense: Suppose that $\mathbf{r}^*$ is the optimal soution for formulation (5.6), by Theorem 5.1, there exists a spillover limit policy achieving an expected cost of at most $f(\mathbf{r}^*) + O(T^{-1/2})$.

An important drawback of the deterministic demand allocation policy is that it will try to match inventories with the expected demand at each DC, and completely ignores the possibility that each DC may stockout or hold leftover inventories due to demand variability. As a result, the inventory allocation policy may not work well if the variabilities of demand are relatively large. This motivates our next allocation policy.

Stochastic Programming Allocation Policy. To specify our stochastic programming allocation policy, let $d_{j\omega}(T)$ denote a realization of $D_j(T)$, the total DC j demand over the time horizon, for a demand scenario $\omega \in \Omega$, which occurs with probability μ_ω. Consider the following stochastic optimization model.

$$\min_{\mathbf{r}\in\mathcal{R}} \tilde{f}(\mathbf{r}), \text{ where } \tilde{f}(\mathbf{r}) \text{ is defined as:} \tag{5.8}$$

$$\begin{aligned}
\tilde{f}(\mathbf{r}) = \min_{\mathbf{r},\mathbf{f},\mathbf{l}} \quad & \sum_{j,\omega} \mu_\omega \Big(p_j l_{j\omega} + \sum_{j'} c_{jj'} f_{jj'\omega} \Big) \\
\text{s.t.} \quad & l_{j\omega} + \sum_{0\le j'\le N} f_{j'j\omega} \ge d_{j\omega}(T), \ \forall\ 0 \le j \le N, \omega, \\
& \sum_{j'\ne j} f_{j'j\omega} \le (1-\alpha_j)\big(d_{j\omega}(T) - f_{jj\omega}\big), \forall\ \ 0 \le j \le N, \omega, \\
& \sum_{0\le j'\le N} f_{jj'\omega} \le r_j, \ \forall\ \ 0 \le j \le N, \omega, \\
& f_{jj'\omega} = 0, \ \forall\ (j,j') \notin A, \\
& \mathbf{f}, \mathbf{l} \ge \mathbf{0}.
\end{aligned} \tag{5.9}$$

The stochastic programming allocation policy allocates the inventories as if the fulfillment decisions are made offline, after all demands are realized. Under the stochastic programming allocation policy, variables $\mathbf{f}$ and $\mathbf{l}$ are indexed by demand scenarios, therefore accounting for different optimal offline fulfillment decisions subject to demand uncertainty.

Like the deterministic demand allocation policy, our intuition behind the stochastic programming allocation policy is to use $\tilde{f}(\mathbf{r})$ as an approximation for the expected cost, with inventory position $\mathbf{r}$ under an effective dynamic fulfillment policy. By Lemma 5.3 in Section 5.6.1, we have that $\tilde{f}(\mathbf{r})$

is a lower bound on the expected lost sales and fulfillment cost under any inventory policy, and furthermore, by Jensen's Inequality, we have $\tilde{f}(\mathbf{r}) \geq f(\mathbf{r})$, implying that $\tilde{f}(\mathbf{r})$ approximtes the the expected cost better than $f(\mathbf{r})$. In our numerical studies in Section 5.5.2, we see that the stochastic programming allocation policy indeed performs better than the determinstic demand allocation policy, often by a significant margin.

5.5 Assessing the Value of Flexibility

In this section we use data from our industrial partner, a large e-retailer, to assess the impact of adding flexibility to their fulfillment network in a region of China. We begin with an overview of the e-retailer's current distribution infrastructure and fulfillment operations, and our proposed implementation of additional flexibility. To protect our industrial partner from revealing sensitive information to its competition, the numerical values presented in this section are masked versions of the real data.

5.5.1 Specific Modeling Features from Our Industrial Collaboration

Following the e-retailer's current practice of treating regions independently, we focus our analysis on a specific region of its fulfillment network in central China. The region has a total of 6 DCs: one centrally located regional-distribution-center (RDC) and 5 front-distribution-centers (FDCs). The relative geographic positioning of the DCs is depicted in Figure 5.2. In our model we denote the RDC with index 0, and denote the 5 FDCs with indices 1 through 5, as depicted in Figure 5.2.

Flexibility. The e-retailer's current fulfillment practice allows spillover fulfillment only from the RDC to any FDC. We call this the "dedicated" structure, reflecting the fact that each FDC is dedicated to serving only its local demand (depicted in the left panel of Figure 5.2). Furthermore, it fulfills its orders greedily, without holding back its inventories at any DC. Our purpose is to augment this dedicated structure with more flexibility and better fulfillment operations, and a few practical constraints must be taken into consideration when doing so. First, the e-retailer's supply chain strategy views the RDC as providing centralized support for the FDCs, and hence the e-retailer's logistics team does not want to allow spillover from an FDC back to the RDC. Further, resource constraints and geographical considerations dictate that not all FDCs be flexible with each other; i.e., it is practical to add just a few FDC to FDC flexibility links to the system.

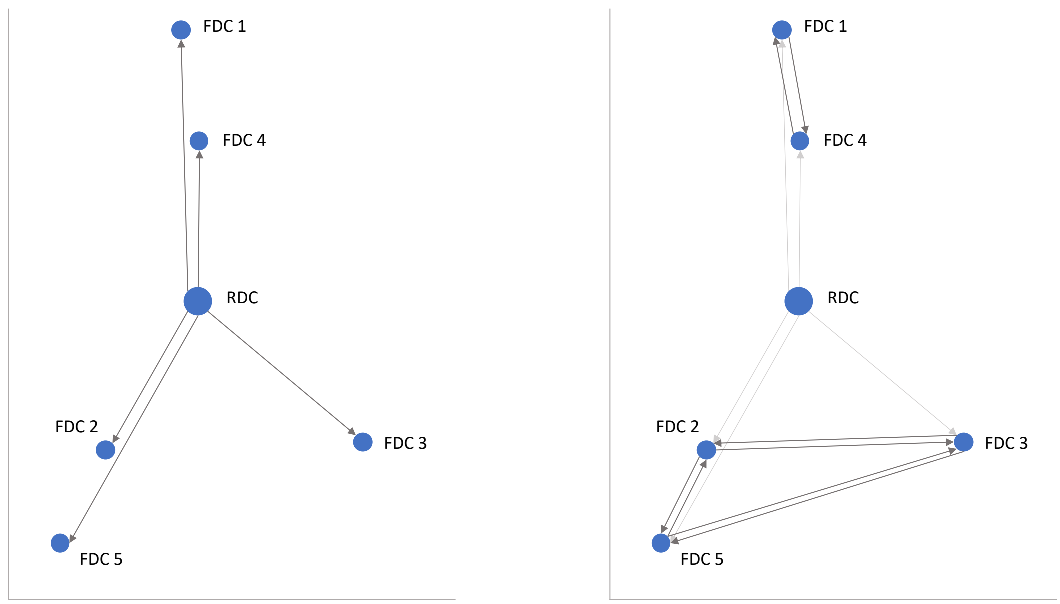

Figure 5.2: DC flexibility within a region: current (left) and proposed (right)

With these constraints in mind, we consider the e-retailer's current flexibility structure depicted in the left panel of Figure 5.2. There is a natural grouping of FDCs 1 and 4 to the North of the RDC and FDCs 2, 3, and 5 to the South, and so we propose adding flexibility among these two disjoint sets of DCs. Since there are a relatively small number of FDCs in each group, we propose adding full flexibility among each set of DCs. Specifically, we add eight flexibility links (out of a possible twenty) which are $(1, 4)$ and $(4, 1)$ in the North, and $(2, 3)$, $(3, 2)$, $(2, 5)$, $(5, 2)$, $(3, 5)$, and $(5, 3)$ in the South. These additional links are depicted in the right panel of Figure 5.2. We call the augmented structure the "flexible" structure. Our industrial collaborators agree with us that the flexible structure is reasonable and interesting, as it does not add too many flexibility links, and does not allow the southern FDCs to spillover to northern FDCs and vice versa, which is considered as wasteful given the central location of the RDC.

Next, we describe how the different parameters of our model are obtained. Some parameters come directly from the data provided by the e-retailer, while some are estimated through statistical and machine learning tools.

Starting Inventory and Replenishment Time. Using the data from our collaboration, we compile a cross-sectional dataset for 982 SKUs, which include a starting inventory position, and days until the next replenishment on a randomly selected day during the non-peak selling period in spring of

2018. This dataset is fed directly to our model to determine the inventory positions before allocation and T, the day until the next replenishment.

Costs. The cost of fulfilling a unit of DC j demand from DC j' is estimated by the e-retailer's logistics department according to the e-retailer's historical shipping cost data. The estimated cost of fulfillment is a function of the weight of the SKU and the distance between the DCs. Specifically, the cost is calculated as $c_{j'j} = a_{j'j} + b_{j'j}(w - d_{j'j})^+$, where w is the weight of the SKU in kilograms, $a_{j'j}$ is the flat rate for the first $d_{j'j}$ kilograms, and $b_{j'j}$ is the variable rate per kilogram above the first $d_{j'j}$ kilograms. Both $a_{j'j}$, $b_{j'j}$ and $d_{j'j}$ are determined based on a zone shipping system where the DCs are located. In other words, $a_{j'j}$ is determined via lookup in a flat cost reference table based on which zone DC j and j' are located in, and similarly for $b_{j'j}$ and $d_{j'j}$. As it turns out, the region we consider covers only two distinct vectors (a, b, d), for in-province fulfillment and inter-province fulfillments. For two DCs that are located in the same province, although the spillover fulfillment costs between them are no higher than their local fulfillment costs, the e-retailer would still prefer local fulfillment, as spillover fulfillment takes longer and can cause customer abandonment. The weights for all 982 SKUs are obtained in a SKU description dataset. The weights vary greatly, ranging from 0.002 kilograms to 25.58 kilograms, thus, the spillover and fulfillment costs across SKUs are in general very different. We checked the fulfillment costs and the flexibility structures we consider and verified that Assumptions 5.1 and 5.2 are satisfied for each SKU.

We estimate the unit cost of incurring a lost sale by the gross profit of a given SKU, which is equal to the selling price minus the purchasing price. Typically the e-retailer offers the same price for a given SKU across all DCs in a region, and so the lost sales cost in our simulation is the same for each DC, i.e., $p_j = p$ for all j, where p is the gross profit of the SKU. The gross profit for SKUs varies greatly; excluding the negative values due to clearance or temporary promotions, the gross profit ranges from 0.02 to 1222.

Estimating Demand. Demand distributions are estimated using one year of historical sales data during a non-peak selling period. To estimate the uncensored demand, historical sales data observations are included only for days when inventory is not stocked out. To cover a range of distribution shapes, the normal, gamma, Poisson, and negative binomial distributions are all considered as possible underlying parametric distribution types. The parameter values for each distribution type are fit using maximum likelihood estimation with a SKU's historical sales, and the best distribution is chosen based on the Kolmogorov-Smirnov goodness-of-fit statistic.

Estimating Abandonment. The abandoment rate is the most difficult to estimate among all parameters in our model. The intuitive method of taking the ratio of average demand on days with inventory to the average demand spilled over to other DCs on days without inventory is problematic, as there is a positive correlation between the busy days to the probability of inventory stockouts. To better estimate the abandonment rate, we apply a standard machine learning method known as Nonnegative Matrix Factorization (NMF) (see Lee and Seung (2001) for a reference) to recover the unseen demand, using one year of historical sales data.

We briefly describe the NMF method we used to recover the unseen demand. Let n be the number of SKUs, m be the number of days (we use one year of data), and p be the number of latent variables. The NMF method factorizes the SKU demand matrix $\mathbf{V}$ with dimension $n \times m$ into two matrices, namely, matrix $\mathbf{W}$ with dimension $n \times p$ and matrix $\mathbf{H}$ with dimension $p \times m$. The parameter p is used to control the complexity of the model. We use Kullback-Leibler divergence as the objective to train the NMF model and apply L_1 and L_2 regularization to deal with overfitting. We then use the trained model to fill in the missing demand for those days when inventory is out-of-stock. We note that when inventory is positive in the beginning of the day but zero at the end of the day, we still use the NMF method to fill in the demand as if there is no inventory in the beginning of the day.

Now we have an estimated demand d'_{it} for sku i on day t when the inventory is out-of-stock at the end of the day. We then calculate $\alpha_{it} = (d'_{it} - d_{it})^+/d'_{it}$ to estimate the percentage of unseen demand for SKU i on day t. As there are only few out-of-stock days for a single sku, we aggregate α for these days across all SKUs to draw a statistically significant conclusion. We believe this serves as a reasonable estimate on the abandoment rate, and more accurate estimation on the abandoment rate for different SKUs is an interesting research topic by itself.

5.5.2 Policies

In this section we perform preliminary tests on several fulfillment and allocation policies using the e-retailer's data in order to identify an effective policy for taking advantage of additional flexibility in the fulfillment network. To perform these tests on a broad range of policies, we choose a sample of SKUs from the 982 SKUs in our dataset. The reason for selecting a sample of SKUs on which

Table 5.1: Starting inventory and mean demand for 16 SKUs

Inv.	Mean Daily Dem.	Days (T)	Inv./Dem.
164	20.6	28	0.28
41	10.0	8	0.51
113	20.6	9	0.61
50	3.9	19	0.67
108	8.3	16	0.81
49	8.2	7	0.86
190	7.2	30	0.88
587	37.1	17	0.93
26	7.9	3	1.09
104	9.5	10	1.10
105	10.0	9	1.16
43	4.5	8	1.21
91	10.8	7	1.21
151	16.7	6	1.51
560	8.7	30	2.13
2645	40.3	14	4.69

to perform these tests is twofold. First, testing many policies across all 982 SKUs would lead to prohibitively long run times. Second, many of the SKUs provided in the data have inventory positions or demand distributions that make it hard to assess the impact of different policies. For example, SKUs with a very high starting inventory position relative to expected demand until the next replenishment allow almost all demand to be filled locally, and so the impact of different allocation and fulfillment policies is muted. Likewise, SKUs with a very low starting inventory position relative to expected demand lead to mostly lost sales, regardless of the allocation and fulfillment policy. Thus, in this section, we test several policies on a sample of SKUs in order to identify the most effective policies, and we will then test these policies on all 982 SKUs in Section 5.5.3.

With this goal in mind, we choose a sample of 16 SKUs in conjunction with our industrial partner in order to assess the policies' effectiveness over a reasonable range of starting inventory positions and demand distributions. The starting inventory position and average daily demand, both aggregated across all six DCs, are recorded in columns 1 and 2 of Table 5.1 for each of the 16 SKUs. The number of days until the next replenishment for each SKU is recorded in the third column, which determines the time horizon for the simulation, T. The last column of the table records the ratio of total starting inventory to total expected demand over the horizon, i.e., the first column divided by the product of the second and third columns. The SKUs are selected in order to cover a wide (but reasonable) range of this ratio; with total inventory ranging from about one fourth of total expected demand to about four times total expected demand.

While these SKUs may not be representative of all SKUs for the purpose of estimating the overall magnitude of the impact provided by additional flexibility, that is not our purpose in this section. Our present goal is to assess the relative effectiveness of alternative allocation and fulfillment policies, and this sample of SKUs provides a meaningful set of scenarios for which to make these comparisons. We then use the insights gathered from this analysis to assess the overall impact of adding flexibility for all SKUs in Section 5.5.3.

In this section and the next, we assess the effectiveness of various policies using simulation. For each SKU we generate 1,000 sample paths and record the cost incurred by each policy along each sample path. In order to track the fulfillment decisions of each policy, we also record the total units of local fulfillment, total units of FDC to FDC spillover, total units of RDC to FDC spillover, and total units of lost sales for each policy along each sample path. The costs that we report for each policy are the average cost of these 1,000 sample paths. Next we describe the specific policies we consider.

Fulfillment. We first describe the fulfillment policies we test.

Greedy. The greedy policy, discussed in Section 5.2.1, fulfills each unit of arriving demand from the lowest cost DC with available inventory. We test this policy on both the dedicated structure, and the flexible structure we propose. The dedicated structure with the greedy fulfillment policy represents our industrual partner's current practice, and so serves as a base case against which we measure other policies. We test the flexible structure with the greedy fulfillment policy to consider the effect of adding flexibility without changing the fulfillment policy, and to demonstrate what further improvements can be made with a better fulfillment strategy. In the computational tests below we will denote the greedy policy in the dedicated structure and flexible structure as D-GRD and F-GRD, respectively.

Besides greedy, we test four additional fulfillment policies on the flexible structure, in order to determine the best policy for taking advantage of the additional flexibility. The first three are variations of the spillover limit policy introduced in Section 5.2.2, which we discuss next.

DLP Spillover Limit. The first variant of spillover limit policy we consider is based on the proof of Theorem 5.1 and is constructed using an optimal solution to the LP formulated as (5.5) from

Section 5.3. In particular, letting $(\mathbf{f}^*, \mathbf{l}^*)$ denote an optimal solution to (5.5), the policy sets spillover limits $\theta_{jj'} = f^*_{jj'}$ for $j' \neq j$ and $\theta^{DC}_j = \sum_{j' \neq j} f^*_{jj'}$ for each j. In our computational tests, we denote the DLP spillover limit policy as DLP. Intuitively, the policy limits the level of spillover fulfillment to the optimal level for average demand from the DLP. This policy is our starting point for testing the spillover limit policy, since we show it achieves asymptotic optimality as T grows large in the proof of Theorem 5.1. However, we observe in our numerical simulations that, compared to greedy, the DLP spillover limit policy generally has too much lost sales. This is because DLP optimistically solves an LP with deterministic demand, resulting in setting the limit on many spillover fulfillment links to zero. Thus, we consider two alternative spillover limit policies.

Mean Protecting Spillover Limit. An observation on why DLP performs poorly is that the limit on many spillover fulfillment links is set too low. This motivates us to propose the mean protecting spillover limit policy, which relaxes DLP by setting $\theta^{DC}_j = (r_j - T\mu_j)^+$, and sets $\theta_{jj'} = \infty$, i.e., we do not set a limit on the amount of spillover allowed between specific DC pairs, rather, we only set a limit on the aggregate spillover sourced from a single DC. Intuitively, the mean protecting spillover limit policy reserves enough inventory to fulfill the expected demand locally; making any excess available for flexible fulfillment. In our computational tests we denote this policy by MP.

Littlewood Spillover Limit. Next, we propose a spillover limit policy that takes both cost and demand distributional information into account when deciding what amount of inventory to reserve for local fulfillment (note that both DLP and MP use only the expected demand). A intuitve way to incorporate cost can be achieved through an adaptation of Littlewood's two-class model for revenue management (Littlewood, 1972). Littlewood's model considers a finite capacity C that may serve two classes of stochastic demand that arrive sequentially: a sale to the first class provides profit p_L, while a sale to the second class provides a higher profit, $p_H \geq p_L$. The problem is to decide how much capacity to make available to the first (low profit) demand class in order to maximize expected profit, and is solved by Littlewood's rule, $(C - Q_H(1 - p_L/p_H))^+$, where Q_H is the quantile function (or inverse CDF in the case of continuous demand) for the second (high profit) demand class. For a random variable X, the quantile function of a probability $p \in (0, 1)$ is defined as $Q(p) = \inf\{x | p \leq \mathbb{P}(X \leq x)\}$. This simple rule is straightforward to implement as the quantile function can be easily approximated through sampling. In our numerical simulation, when p is greater than 0.99 or less than 0.01, we set $Q(p)$ to infinity or zero, respectively, to avoid the significant sampling errors of estimating extreme percentiles.

To adapt Littlewood's rule to our fulfillment context, let us first consider setting the spillover limit $\theta_{jj'}$ for some $(j, j') \in A$. Naturally, the inventory at DC j corresponds to capacity, while the local demand at DC j and spillover demand from DC j' (assuming DC j' is out of inventory) correspond to the two classes of demand. Since Littlewood's rule is stated in terms of profits, we compute the "profit" estimates for local and spillover demand by subtracting fulfillment costs from the lost sales cost, i.e., $p - c_{jj}$ for local fulfillment and $p - c_{jj'}$ for spillover, because these quantities are the cost reductions we receive if we choose to fulfill the order instead of losing the order. Now, we have a Littlewood model with $p_L = p - c_{jj'}$ and $p_H = p - c_{jj}$, and we set the amount of DC j inventory to make available for spillover fulfillment of demand at DC j' as

$$\theta_{jj'} = \left(r_j - Q_j\left(\frac{c_{jj'} - c_{jj}}{p - c_{jj}}\right)\right)^+,$$

where Q_j is the quantile function of the demand distribution at DC j.

We use a similar method to set an aggregate flexibility limit for DC j, θ_j^{DC}. Since there are multiple spillover costs associated with using DC j inventory for spillover fulfillment, i.e., $c_{jj'}$ for all $(j, j') \in A$, we apply an extension of Littlewood's rule known as the EMSR-a heuristic (Belobaba, 1989; Talluri and Van Ryzin, 2006). In particular, let $c_j^s = \min\{c_{jj'} | j' \neq j, (j, j') \in A\}$, the EMSR-a heuristic suggests a potection limit applying Littlewood's rule with $p_L = p - c_j^s$ and $p_H = p - c_{jj}$. Therefore, the Littlewood spillover limit policy sets the aggregate spillover limit at DC j as

$$\theta_j^{DC} = \left(r_j - Q_j\left(\frac{c_j^s - c_{jj}}{p - c_{jj}}\right)\right)^+.$$

In our computational tests, we denote the Littlewood spillover limit policy as LTL. We note that LTL is a heuristic because our fulfillment model does not fully match the original model studied in by Littlewood (1972), which has spillover demand requests arrive before local fulfillment requests. Nethertheless, we observe that LTL performs very well in numerical simulations, likely because of the robustness of Littlewood's rule.

Dual Cost. The final fulfillment policy we consider is adapted from a linear programming heuristic developed by Acimovic and Graves (2015), and provides a point of reference for comparing against a leading heuristic from the literature on e-retail fulfillment. The basic idea of the policy in Acimovic and Graves (2015) is to solve an allocation LP using expected demand, then make fulfillment decisions using costs adjusted by the optimal dual values associated with the inventory constraints in the LP. The intuition of this approach is that the dual variables provide a linear approximation to the dynamic programming value function, and thus allow the fulfillment decision to incorporate an estimate of the expected value of the resulting inventory position.

We adapt this type of dual cost framework to our fulfillment setting. Since our setting dictates that demand must be filled locally if there is available inventory, we modify Acimovic and Graves (2015) by maintaining a greedy local fulfillment policy, while adopting a dual cost policy for spillover fulfillment. In particular, we solve the LP defined by (5.7) to obtain optimal dual variables, ν_j^* for each j, associated with the inventory constraints $\sum_{j'} f_{jj'} \leq r_j$. The dual variable ν_j^* encodes the reduced cost in the optimal objective value resulting from a change in the inventory level r_j. Thus, when considering using a unit of inventory at DC j to fill demand, we may estimate the increase in future cost resulting from the reduction of inventory as $-\nu_j^*$. We note that the dual variable is non-positive, $\nu_j^* \leq 0$ (since increasing the value of r_j in the original program cannot increase the optimal objective value), and hence $-\nu_j^*$ represents an estimate of the cost increase from using a unit of DC j inventory. Therefore, we may estimate the impact of using a unit of DC j inventory to fill demand at DC j' as $c_{jj'} - \nu_j^*$, where $c_{jj'}$ is the immediate fulfillment cost and $-\nu_j^*$ is the estimate of future costs caused by reducing the inventory level at DC j.

Thus, when faced with a spillover fulfillment decision at DC j, the dual cost policy chooses the lowest cost fulfillment option based on these dual adjusted costs. Formally, at a given point in the course of the time horizon, let U_j denote the set of DCs that may provide spillover fulfillment to DC j in the flexibility structure, and that have available inventory. Then the dual cost policy fills from the minimizer of

$$\min\Big(\min_{j' \in U_j} \{c_{j'j} - \nu_{j'}^*\}, p_j \Big),$$

where if p_j is the minimizer, then the policy chooses to let the demand become a lost sale. As mentioned above, to satisfy the business constraints that local demand is filled greedily, the dual costs are only considered for spillover fulfillment decisions. In our computational tests, the dual cost policy is denoted as DUAL.

Inventory Allocation Policies. In addition to fulfillment, we also wish to identify an effective inventory allocation policy. For each of the fulfillment policies described above, we test the two inventory allocation policies introduced in Section 5.4: the deterministic demand allocation and the stochastic programming allocation, described by (5.6) and (5.8), respectively. As mentioned above, our industrial partner's supply chain strategy treats the RDC as a central support for the FDCs,

Table 5.2: Average costs for 16 SKUs by fulfillment and allocation policy

Policy	Det. Alloc.	SP Alloc.
D-GRD	3,393.7	3,150.6
F-GRD	3,056.6	2,916.3
DLP	3,325.3	3,248.2
MP	3,186.3	3,057.3
LTL	3,044.5	2,900.9
DUAL	3,102.8	2,983.3

and thus only allows inventory reallocation from the RDC to each FDC. These constraints define the feasible set of inventory allocations, given the initial inventory position $\mathbf{x}$, as

$$\mathcal{R} = \left\{ \mathbf{r} \in \mathbb{R}_+ \middle| \sum_{0 \le j \le N} r_j = \sum_{0 \le j \le N} x_j,\ r_j \ge x_j\ \forall 1 \le j \le N \right\}.$$

For both allocation policies, we use a rounding scheme to identify an integer solution maintaining feasibility for $\mathcal{R}$. In particular, we round down the optimal values for $r_1, r_2, \dots r_N$ into integers, while rounding up r_0 so that $\sum_{0 \le j \le N} r_j = \sum_{0 \le j \le N} x_j$, and use those integer values to allocate inventories.

For each of the 16 SKUs, we perform both the deterministic demand and stochastic programming allocations, then simulate the performance of each fulfillment policy. For the stochastic programming allocation, we solve the stochastic program (5.8) using 1,000 demand samples. In Table 5.2, we report the average expected cost across the 16 SKUs for each combination of allocation and fulfillment policy. For each fulfillment policy, the stochastic programming allocation policy provides lower expected cost than the deterministic demand allocation policy. Intuitively, the stochastic programming allocation policy is better because $\tilde{f}(\mathbf{r})$ approximates the expected fulfillment costs much more accurately compared $f(\mathbf{r})$. Thus, for the remainder of our computational tests, we focus on the stochastic programming allocation policy.

Resolving Heuristics. Next, we consider the benefit of updating the fulfillment policy periodically throughout the time horizon. This idea of "resolving" the heuristic policy is common in the revenue management and online resource allocation literature (see e.g., Cooper (2002) and Jasin and Kumar (2012)), and typically leads to improved performance since the policy can adapt to the demand realizations as they are realized. In the context of our industrial collaborator's fulfillment decisions, because the daily demand distribution at each DC is provided, it is feasible to consider resolving daily, with updated policy parameters computed at the beginning of each day. In particular, for

Table 5.3: Average costs for 16 SKUs for static and resolved policies

Policy	Avg. Cost Static	Avg. Cost Resolved	Avg. Cost Ratio Reslv./D-GRD
D-GRD	3,150.6	-	100.0%
F-GRD	2,916.3	-	95.4%
DLP	3,248.2	3,120.0	106.8%
MP	3,057.3	2,932.5	95.7%
LTL	2,900.9	2,900.0	94.6%
DUAL	2,983.3	2,914.3	94.9%

the DLP, MP, LTL and DUAL policies, we compare the "static" heuristic described above, with a resolved version that simply recomputes the policy parameters at the beginning of each day given the starting inventory position on that day, and demand distribution over the remaining days in the horizon. The greedy heuristics are inherently static, since there are no policy parameters to update, and so we do not consider a resolved version for these heuristics. As mentioned above, we consider each of these heuristics in conjunction with the stochastic programming allocation.

The average expected costs across the 16 SKUs for each of these heuristics are reported in the first two columns of Table 5.3. For each heuristic tested, the resolved version provides lower cost than the static version on average. Although the improvement from resolving is small for some heuristics, we note that resolving is a generally beneficial strategy as it allows increased flexibility to adjust to demand as it is realized. Therefore, the remainder of our discussion and computational tests will focus on the resolved DLP, MP, LTL and DUAL heuristics (as well as the static greedy heuristics).

In Table 5.3, we also report the average cost ratio for the re-solved version of our policies in column three. These values are calculated by dividing the cost of each resolved policy (or static in the case of F-GRD) by the cost of D-GRD for each SKU, then taking the average of this ratio across all 16 SKUs. The cost ratio statistic provides a normalization across SKUs to compare the relative cost change provided by each policy, reducing the variation across SKUs so that a single high cost SKU does not dominate the average.

From Table 5.3 we see that the LTL policy is the best in terms of both average cost and average cost ratio, with the next closest competitors being DUAL and F-GRD. We also note that the average cost for LTL did not change significantly from "Static" to "Resolved". This shows that LTL policy is robust with respect to the need for resolving, which is advantageous in practice. In contrast, a heuristic like DUAL may suffer significant performance loss without frequent resolving.

While the difference between the average cost ratio and the relative average cost between LTL, DUAL and F-GRD is not large, this difference is important to our industrial partner. For our industrial partner, an extra 0.1% decrease in the total lost sales and fulfillment costs equates to a financial gain on the order of millions of U.S. dollars. To ensure that these differences between LTL, DUAL and F-GRD are not caused by variations in the sampled demand, we present some statistical significance test for each SKU. Since the simulations were run on the same 1,000 sample paths for each policy, we perform a paired T-test for each SKU to test the significance of the difference in pairwise cost between LTL and both F-GRD, and DUAL. Out of the 16 SKUs, for 10 SKUs the cost reduction from F-GRD to LTL is statistically significant, and for the rest of the SKUs, F-GRD and LTL costs are either tied or do not have statistically significant differences. Similarly, for 10 SKUs the cost reduction from DUAL to LTL is statistically significant, for 3 SKUs there is a statisticallly significant cost increase from DUAL to LTL, and for the rest of the SKUs, F-GRD and LTL are either tied or do not have statistically significant differences. These tests suggests that the difference between LTL and both DUAL and F-GRD are not caused by the randomness of our simulation instances. Moreover, for the 3 SKUs with statistically significant cost increase from DUAL to LTL, we note that LTL cost is at most 0.35% larger than that of the DUAL cost, whereas the largest cost reduction from DUAL to LTL (among all 16 SKUs) is almost an order of magnitude larger at 3.14%. Thus, we conclude that the LTL policy provides very close to the best performance among all policies tested for each SKU.

Finally, we note some further advantages that the LTL policy maintains over DUAL, the next best alternative in terms of average cost. Intuitively, the DUAL policy is limited by the fact that it only uses the mean of the demand distribution when solving the LP to obtain the dual costs, which is why this policy makes a more significant gain when moving from the static to resolved version (Table 5.3). Meanwhile, the LTL policy leverages more demand information since it calculates a quantile of the distribution function, and thus is able to provide better performance, even for the static version. Further, from an implementation perspective we note that the DUAL policy requires more coordinated computational effort to resolve, since a central entity needs to solve the LP each day given the current inventory position and expected future demand of all DCs. On the other hand, the LTL policy can be carried out in a more decentralized fashion, since each DC can calculate its spillover limit given its own distribution of future demand. Therefore, we recommend the resolved LTL policy to e-retailers seeking to implement flexibility in their fulfillment network, and this is the

Table 5.4: Average costs and fulfillment metrics across 866 SKUs

Policy	Avg. Cost	Local	FDC-FDC	RDC-FDC	Lost Sale
D-GRD	2,816.1	277.4	0.00	2.25	31.62
F-GRD	2,797.2	275.7	4.26	1.34	29.99
LTL	2,794.3	276.7	3.76	0.73	30.10

main policy which we will test on all SKUs in Section 5.5.3.

5.5.3 Impact of Flexibility

In this section we assess the impact of adding flexibility to our industrial partner's fulfillment network by simulating demand for all 982 SKUs from our data. Based on the preliminary tests of Section 5.5.2, the main policy we test is the resolved Littlewood Spillover Limit, denoted LTL. We will measure this policy against the e-retailer's status quo fulfillment strategy, the greedy policy with the dedicated structure, denoted D-GRD. For comparison purposes we will also test the greedy policy on the flexible structure, denoted F-GRD, since it is our industrial partner's status quo fulfillment policy and would be straightforward to implement (and it achieved close to the second lowest average cost in our preliminary tests of Section 5.5.2, see Table 5.3).

We implement these tests on the SKUs from our data, with a few exclusions. First, of the 982 SKUs from our data, 43 were excluded because they had non-positive lost sales cost due to promotional status. A further 73 SKUs were excluded because they had no on-hand inventory at any DC on the randomly chosen day used to create our dataset, and so assessing fulfillment policies on these SKUs would be meaningless. We tested our policies on the remaining 866 SKUs that had positive lost sales cost and a positive starting inventory level at any DC.

In Table 5.4, the first column records the average cost for the three policies tested across all 866 SKUs, while the remaining columns contain the average of various fulfillment metrics. In particular, column two records the average units of local fulfillment, column three the average FDC to FDC spillover fulfillment, column four the average RDC to FDC spillover fulfillment, and column five the average lost sales. The LTL policy achieves the lowest average cost, which is a 0.77% reduction from the status quo D-GRD policy. Moreover, the fulfillment metrics of Table 5.4 provide us some insights to explain the usefulness of the flexible fulfillment structure and the effectiveness of LTL. First, the flexible structure decreases the average amount of lost sales by about 5% for both the LTL policy and the F-GRD policy (the F-GRD policy reduces lost sales by 0.36% more than the

Table 5.5: SKU level cost ratio of F-GRD and LTL to D-GRD across 866 SKUs

Cost Ratio	Average	Worst Case	# SKUs > 1
F-GRD/D-GRD	99.11%	101.82%	79
LTL/D-GRD	99.00%	100.18%	12

Table 5.6: SKU level cost ratio by inventory to demand ratio

Inv./Dem.	# SKUs	LTL/D-GRD	LTL/F-GRD
≤ 1	234	98.09%	99.60%
> 1	632	99.33%	99.99%
Overall	866	99.00%	99.88%

LTL policy, which is small compared to the 5% lost sales reduction compared to the D-GRD policy). The reason LTL provides lower average cost than F-GRD is because LTL significantly decreases the amount of both FDC-FDC and RDC-FDC spillover fulfillment compared to F-GRD.

Furthermore, Table 5.5 highlights that the LTL policy provides better control of cost on a SKU level than F-GRD. For each SKU we calculate the ratio of F-GRD to D-GRD cost, and the ratio of LTL to D-GRD cost. Table 5.5 records the average and worst case of these ratios, as well as the count of SKUs for which the ratio is larger than 1. The LTL policy provides an average cost reduction over D-GRD of 1.00% for each SKU, while the average cost reduction of F-GRD over D-GRD is 0.89%. Further, we find that F-GRD has higher cost than D-GRD for 79 out of 866 SKUs, with the worst case cost being 1.82% larger. In contrast, LTL has higher cost than D-GRD for only 12 SKUs, and the worst case cost is only 0.18% higher. This illustrates the main point of Example 5.1, that a greedy policy can use too much flexibility and lead to higher costs. The LTL policy does a good job of controlling the use of flexibility to achieve lower costs across SKUs.

In Table 5.6 we report the cost ratio of the LTL policy to D-GRD and F-GRD, for SKUs with different inventory to average demand ratios. For SKUs with less inventory than average demand, LTL provides an average of 1.91% improvement over D-GRD, and 0.4% improvement over F-GRD. For SKUs with more inventory than average demand, LTL still provides an average of 0.67% improvement over D-GRD, but is almost equal on average to the cost of F-GRD. This illustrates the intuition of Lemma 5.1, that with enough inventory greedy will perform well. Finally, across all SKUs, we see that LTL provides a 1.00% improvement over D-GRD, and a 0.12% improvement over F-GRD. Thus, adding flexibility to the fulfillment network provides about a 1% improvement in costs, with the LTL policy providing about 10% of that 1% improvement.

For completeness, we also consider implementing the LTL policy for our industrial partner's current dedicated structure. Intuitively, this policy does not have much room for improvement over greedy since the only flexibility in the system is from the RDC to the FDCs, and so putting limits on spillover can only improve the cost at the RDC. We simulate LTL on the dedicated structure and find that the average cost is 0.09% lower than the average cost of D-GRD. This supports our conclusion that the proposed additional flexibility is a major driver of the cost improvement we observe in our simulations.

In summary, compared to the status quo D-GRD policy, the LTL policy reduces lost sales by about 5% while keeping local fulfillment at a similar level, and reduces overall cost by about 0.77%, with an average of about 1% cost reduction for each SKU. From these simulations we estimate that additional flexibility can offer the e-retailer around a 1% improvement in lost sales and fulfillment costs when implemented with the LTL fulfillment policy. While we do not provide the exact financial information of our industrial partner in this chapter, we note that 0.1% improvement is estimated to increase the profit on the order of millions of dollars.

5.6 Policy Analysis

In this section, we provide formal proofs of the various mathematical claims made throughout this chapter.

5.6.1 Technical Proof Components

We first present several technical lemmas that were not stated formally in the chapter. These lemmas will serve as useful building blocks for our proof of Theorem 5.1, as well as justifications for our allocation policy.

Lemma 5.2. *For T i.i.d. random variables $(X_i)_{1\leq i\leq T}$ with finite mean μ and finite variance σ^2 we have*

$$\mathbb{E}[(\sum_{1\leq i\leq T} X_i - T\mu)^+] \leq \frac{1}{2}\sqrt{T}\sigma,$$

Proof. From the identity $(z)^+ = (|z| + z)/2$, we have

$$\begin{aligned}\mathbb{E}\Big[\Big(\sum_{1\le i\le T} X_i - T\mu\Big)^+\Big] &= \frac{1}{2}\mathbb{E}\Big[\Big|\sum_{1\le i\le T} X_i - T\mu\Big|\Big],\\ &\le \frac{1}{2}\sqrt{\mathbb{E}\Big[\Big(\sum_{1\le i\le T} X_i - T\mu\Big)^2\Big]},\\ &= \frac{1}{2}\sqrt{T}\sigma\end{aligned}$$

where the inequality follows from Jensen's inequality. □

Lemma 5.3. *For a given initial inventory allocation* $\mathbf{r}$*, the optimal value of the stochastic program* (5.9)*,* $\tilde{f}(\mathbf{r})$*, provides a lower bound on the expected cost of the optimal fulfillment policy.*

Proof. We note that there are two types of randomness in the online fulfillment problem, randomness in demand, and randomness in abandonment due to spillover. We will prove the result using three successive relaxations of these sources of randomness in the online problem.

The first relaxation we make to the online problem is to assume that at the beginning of each day, it is revealed how many demand units will arrive on that day in total for all DCs, where this value is drawn according to the aggregate distribution of demand at all DCs. The demand and abandonment will still occur randomly in an online fashion throughout the day, and the decision maker still needs to make decisions online. This is a relaxation because the decision maker has a strictly larger information set at each point in time than she did in the original problem. Under this relaxation, consider some time τ during a day, and consider the sequence and timing of demand arrivals up to time τ on that day. We observe that the decision maker's conditional distribution over the sequence of remaining arrivals on that day after time τ is only dependent upon the sequence of arrivals up to time τ, and not dependent upon the specific timing of these arrivals. This is because the decision maker knows how many arrivals will occur under the relaxation, and because the arrivals occur uniformly at random throughout the day. Thus, since costs are not discounted in time also, without loss of generality we exclude time from the state variable under this relaxation, and only consider the total number of arrivals on a day (revealed at the beginning of the day) and the sequence of arrivals as it unfolds throughout a day.

Therefore, we formulate this relaxed problem with a scenario tree. The scenario tree evolves from a root node at level zero (denoting the initial state before any randomness is realized) through the uncertainty realized on each day. On each day, the first level of "arrival" nodes encode how many arrivals will occur on that day. From these nodes, the sub-tree on each day is comprised of alternating levels of demand nodes, denoted by ν, and abandonment nodes, denoted by ξ. In particular, from an arrival node, the tree branches into the first level of demand nodes, from which each node branches into the first level of abandonment nodes, from which each node branches into the second level of demand nodes, and so on. When the sub-tree has progressed through the total number of arrivals for that day, it will branch into a new set of arrival nodes for the next day.

Associated with each demand node ν is a vector of binary values $(d_{0\nu}, \ldots, d_{n\nu})$ denoting the demand arriving for DC j at that node. Reflecting the fact that demand units arrive one by one, for each demand node ν, only one DC j has $d_{j\nu} = 1$, while $d_{j'\nu} = 0$ for the remaining DCs $j' \neq j$. Associated with each abandonment node ξ is a vector of binary values $(a_{0\xi}, \ldots, a_{n\xi})$, with $a_{j\xi} = 0$ denoting that demand will be abandoned upon spillover, and $a_{j\xi} = 1$ denoting that demand will not be abandoned upon spillover.

Each demand (abandonment) level of the tree has a probability measure over its nodes ν (ξ) denoted by μ_ν (μ_ξ) that obeys the laws of conditional probability in relation to the previous level of the tree (i.e., a node's probability should equal the sum of probabilities of its children). The demand node probabilities are governed by the daily demand process specified in Section 5.1.2 and are independent of the abandonment node probabilities. For a demand node ν with demand at DC j (i.e., $d_{j\nu} = 1$), there are two children abandonment nodes, $\xi_0(\nu)$ and $\xi_1(\nu)$, with $a_{j\xi_0(\nu)} = 0$, occurring with probability $\mu_{\xi_0(\nu)} = \alpha_j \mu_\nu$ and $a_{j\xi_1(\nu)} = 1$ occurring with probability $\mu_{\xi_1(\nu)} = (1 - \alpha_j)\mu_\nu$. We may assume $a_{j\xi_0(\nu)} = a_{j\xi_1(\nu)} = 0$ for all other DCs $j' \neq j$ at each abandonment node $\xi_0(\nu)$ and $\xi_1(\nu)$.

Let $\mathfrak{R}(\nu)$ and $\mathfrak{R}(\xi)$ denote the nodes along the path from the root of the tree to demand node ν and abandonment node ξ, respectively, and let u denote an arbitrary element of $\mathfrak{R}(\nu)$ or $\mathfrak{R}(\xi)$. For an abandonment node ξ, let $\nu(\xi)$ denote the parent demand node which it follows in the tree.

Associated with each demand node ν is a local fulfillment decision, denoted $f_{jj\nu}$, which represents the amount of arriving DC j demand to fill locally. Naturally, this quantity should be non-negative and no larger than the arriving demand, and so the decision satisfies the constraints

$$0 \leq f_{jj\nu} \leq d_{j\nu}, \; \forall j, \nu. \tag{5.10}$$

Further, associated with each abandonment node ξ is a spillover fulfillment decision, denoted by the variables $f_{j'j\xi}$ for $j' \neq j$, representing the amount of spillover DC j demand to fill with inventory from DC j', and $l_{j\xi}$, representing the amount of DC j lost sales. The total spillover demand filled should be zero if the demand is abandoned (i.e., if $a_{j\xi} = 0$) and should also be less than the remaining demand after the local fulfillment decision at demand node $\nu(\xi)$, giving the following spillover constraints

$$\begin{aligned} &\sum_{j' \neq j} f_{j'j\xi} \leq a_{j\xi}, \ \forall j, \xi, \\ &\sum_{j' \neq j} f_{j'j\xi} \leq d_{j\nu(\xi)} - f_{jj\nu(\xi)}, \ \forall j, \xi, \end{aligned} \tag{5.11}$$

Each unit of demand must either be filled locally, filled by spillover, or lost which is captured by the following constraints on each abandonment node ξ

$$l_{j\xi} + \sum_{j' \neq j} f_{j'j\xi} + f_{jj\nu(\xi)} \geq d_{j\nu(\xi)}, \ \forall j, \xi. \tag{5.12}$$

Finally, given initial inventory levels r_j, the total fulfillment along any sample path cannot exceed the initial inventory. Let $\mathcal{T}$ denote the set of terminal abandonment nodes, then we have the following constraints

$$\sum_{u \in \mathfrak{R}(\xi)} \sum_{j'} f_{jj'u} \leq r_j, \ \forall j, \xi \in \mathcal{T}. \tag{5.13}$$

The relaxed problem is

$$\begin{aligned} \min_{\mathbf{f}, \mathbf{l}} \quad & \sum_{j,\nu} \mu_\nu c_{jj} f_{jj\nu} + \sum_{j,\xi} \mu_\xi \left(p_j l_{j\xi} + \sum_{j' \neq j} c_{jj'} f_{jj'\xi} \right) \\ \text{s.t.} \quad & \mathbf{f}, \mathbf{l} \geq \mathbf{0}, \ \text{and satisfy (5.10), (5.11), (5.12), (5.13).} \end{aligned} \tag{5.14}$$

Next we relax this problem further, so that the decisions and randomness of the abandonment node ξ are collapsed into its parent demand node $\nu(\xi)$. In particular, the spillover fulfillment and lost sales decisions are made concurrently with the local fulfillment decision at the demand node, denoted by $f_{j'j\nu}, l_{j\nu}$, and abandonment occurs at a deterministic rate α_j for every unit of demand not filled locally. Thus, constraints (5.11) and (5.12) are replaced by

$$\sum_{j' \neq j} f_{j'j\nu} \leq (1 - \alpha_j)(d_{j\nu} - f_{jj\nu}), \ \forall j, \nu, \tag{5.15}$$

$$l_{j\nu} + \sum_{j'} f_{j'j\nu} \geq d_{j\nu}, \ \forall j, \nu. \tag{5.16}$$

Letting $\mathcal{T}'$ denote the set of terminal demand nodes, constraint (5.13) is replaced by

$$\sum_{u \in \Re(\nu)} \sum_{j'} f_{jj'u} \leq r_j, \ \forall j, \nu \in \mathcal{T}'. \tag{5.17}$$

Thus, the relaxed problem is

$$\begin{aligned} \min_{\mathbf{f},\mathbf{l}} \quad & \sum_{j,\nu} \mu_\nu \left(p_j l_{j\nu} + \sum_{j'} c_{jj'} f_{jj'\nu} \right) \\ \text{s.t.} \quad & \mathbf{f}, \mathbf{l} \geq \mathbf{0}, \text{ and satisfy } (5.10), (5.15), (5.16), (5.17). \end{aligned} \tag{5.18}$$

Now we claim that the optimal objective value of (5.18) is less than the optimal objective value of the first relaxed problem (5.14). We show this by demonstrating that any feasible solution to (5.14) can be used to construct a feasible solution to (5.18) with the same objective value.

Consider a feasible solution to (5.14), $f_{jj\nu}, f_{j'j\xi}, l_{j\xi}$. Construct a solution to (5.18) as follows. Let $f'_{jj\nu} = f_{jj\nu}$. Further, let the spillover fulfillment and lost sales at demand node ν be the average of the corresponding quantities on abandonment nodes $\xi_0(\nu)$ and $\xi_1(\nu)$, i.e., for $j' \neq j$, let $f'_{j'j\nu} = \alpha_j f_{j'j\xi_0(\nu)} + (1-\alpha_j) f_{j'j\xi_1(\nu)}$, and $l'_{j\nu} = \alpha_j l_{j\xi_0(\nu)} + (1-\alpha_j) l_{j\xi_1(\nu)}$. This solution is feasible to (5.15) since

$$\begin{aligned} \sum_{j' \neq j} f'_{j'j\nu} &= \alpha_j \sum_{j' \neq j} f_{j'j\xi_0(\nu)} + (1-\alpha_j) \sum_{j' \neq j} f_{j'j\xi_1(\nu)}, \\ &\leq (1-\alpha_j)(d_{j\nu} - f'_{jj\nu}), \end{aligned}$$

where the inequality follows from the first constraint in (5.11) for $\xi_0(\nu)$, $\sum_{j' \neq j} f_{j'j\xi_0(\nu)} \leq a_{j\xi_0(\nu)} = 0$ and the second constraint in (5.11) for $\xi_1(\nu)$, $\sum_{j' \neq j} f_{j'j\xi_1(\nu)} \leq d_{j\nu} - f_{jj\nu} = d_{j\nu} - f'_{jj\nu}$. Feasibility for constraints (5.16) and (5.17) follow directly from constraints (5.12) and (5.13), respectively, by taking the convex combination of abandonment scenarios $\xi_0(\nu)$ and $\xi_1(\nu)$. Thus, the constructed solution is feasible to (5.18).

The constructed solution has objective value

$$\sum_{j,\nu} \mu_\nu \left(p_j l'_{j\nu} + \sum_{j'} c_{jj'} f'_{jj'\nu} \right),$$

$$= \sum_{j,\nu} \mu_\nu \left(\alpha_j \left(p_j l_{j\xi_0(\nu)} + \sum_{j' \neq j} c_{jj'} f_{jj'\xi_0(\nu)} \right) + (1 - \alpha_j) \left(p_j l_{j\xi_1(\nu)} + \sum_{j' \neq j} c_{jj'} f_{jj'\xi_1(\nu)} \right) \right)$$

$$+ \sum_{j,\nu} \mu_\nu c_{jj} f_{jj\nu},$$

$$= \sum_{j,\nu} \mu_\nu c_{jj} f_{jj\nu} + \sum_{j,\xi} \mu_\xi \left(p_j l_{j\xi} + \sum_{j' \neq j} c_{jj'} f_{jj'\xi} \right),$$

which is the same objective value as the original solution for (5.14), completing the proof that (5.18) provides a lower bound.

Finally, it is clear that (5.9) is a relaxation of (5.18), since the former problem allows the fulfillment decisions to be made after all demand has arrived, where the first and second constraints of (5.9) correspond to the aggregation of constraints (5.16) and (5.15) over sample paths, while the third constraint of (5.9) corresponds to the constraint (5.17). □

Lemma 5.4. *For a given initial inventory allocation* $\mathbf{r}$*, the optimal value of the deterministic program* (5.7)*,* $f(\mathbf{r})$*, provides a lower bound on the optimal value of* (5.9)*,* $\tilde{f}(\mathbf{r})$*, and thus also provides a lower bound on the expected cost of the optimal fulfillment policy.*

Proof. Because (5.9) is a linear minimization problem, its optimal objective value is convex in the right hand side constraint values. Therefore $f(\mathbf{r}) \leq \tilde{f}(\mathbf{r})$ follows by Jensen's inequality, the remainder follows from Lemma 5.3. □

Lemma 5.5. *Consider the following linear programming problem:*

$$
\begin{aligned}
\min_{\mathbf{f},\mathbf{l}} \quad & \sum_{j} \left(p_j l_j + \sum_{j'} c_{jj'} f_{jj'} \right) \\
s.t. \quad & l_j + \sum_{0 \le j' \le N} f_{j'j} \ge T\mu_j, \ \forall\ 0 \le j \le N, \\
& \sum_{j' \ne j} f_{j'j} \le (1-\alpha_j)(T\mu_j - f_{jj}), \forall \ \ 0 \le j \le N, \\
& \sum_{0 \le j' \le N} f_{jj'} \le r_j, \ \forall \ \ 0 \le j \le N, \\
& f_{jj'} = 0, \ \forall\ (j, j') \notin A, \\
& f_{jj'} \ge 0, \ \forall\ (j, j') \in A,\ l_j \ge 0, \forall\ 0 \le j \le N.
\end{aligned}
\tag{5.19}
$$

Suppose that Assumption 5.2 holds, then there exists an optimal solution $\mathbf{f}^*$ *to* (5.19)*, such that if* $f^*_{jj'} > 0$ *for* $j' \ne j$ *then* $f^*_{jj} = T\mu_j$*.*

Proof. We prove the lemma by showing that given any optimal solution $\mathbf{f}$, if $f_{jj'} > 0$ and $f_{jj} < T\mu_j$ for some $(j, j') \in A$, then we can always find another optimal solution $\mathbf{f}'$ such that either $f'_{jj'} = 0$ or $f'_{jj} = T\mu_j$. This is proved by considering two cases, either $f_{j''j} > 0$ for some $(j'', j) \in A$, or else $f_{j''j} = 0$ for all $(j'', j) \in A$.

Case 1. Given $f_{j''j} > 0$, construct a new solution as follows. For some $\epsilon > 0$ let $f'_{jj'} = f_{jj'} - \epsilon$, $f'_{jj} = f_{jj} + \epsilon$, $f'_{j''j} = f_{j''j} - \epsilon$, $f'_{j''j'} = f_{j''j'} + \epsilon$. We claim this solution remains feasible for small enough $\epsilon > 0$. To see this, first note all variables that are reduced by ϵ are strictly positive, so they remain non-negative for small enough ϵ. Next, note that we only change variables for DC's j, j' and j'', and so we consider each of these DCs and its associated constraints. For DC j, the first constraint of (5.19) remains satisfied since we increase f_{jj} and decrease $f_{j''j}$ by the same amount, the second constraint remains satisfied since we decrease $f_{j''j}$ on the left hand side by $(1-\alpha_j)\epsilon$, while we effectively decrease the right hand side by a smaller amount, $(1-\alpha_j)\epsilon$, when we increase f_{jj} by ϵ, and the third constraint remains satisfied since we increase f_{jj} and decrease $f_{jj'}$ by the same amount. For DC j' the first and second constraint of (5.19) remain satisfied because we increase $f_{j''j'}$ and decrease $f_{jj'}$ by the same amount, while the third constraint remains satisfied because we do not change any flow originating from DC j'. For DC j'' we do not change any flows whose destination is DC j', so the first two constraints are satisfied, while the last constraint is satisfied because we increase $f_{j''j'}$ and decrease $f_{j''j}$ by the same amount.

Thus, the constructed solution is feasible, and gives a change in cost of $\epsilon(c_{jj}+c_{j''j'}-c_{jj'}-c_{j''j})$. This change is non-positive by Assumption 5.2, and so the new solution remains optimal. The largest ϵ that maintains feasibility is $\min(T\mu_j - f_{jj}, f_{jj'}, f_{j''j})$. If the minimal value is obtained by either of the first two arguments, then we have constructed an optimal solution $\mathbf{f}'$ such that either $f'_{jj'} = 0$ or $f'_{jj} = T\mu_j$. Otherwise, if $f_{j''j}$ is the minimizer, then the new optimal solution $\mathbf{f}'$ will have $f'_{j''j} = 0$, and we can repeat the procedure described in Case 1 until we either find the satisfactory optimal solution, or end up in Case 2 below.

Case 2. If $f_{j''j} = 0$ for all $(j'', j) \in A$, then construct a new solution as follows. For some $\epsilon > 0$ let $f'_{jj'} = f_{jj'} - \epsilon$, $f'_{jj} = f_{jj} + \epsilon$, $l'_j = l_j - \epsilon$, and $l'_{j'} = l_{j'} + \epsilon$. By an argument similar to Case 1, it is straightforward to see that this solution remains feasible for small enough $\epsilon > 0$. The constructed solution gives a change in cost of $\epsilon(c_{jj} + p_{j'} - c_{jj'} - p_j)$, which is non-positive by Assumption 5.1, and the new solution remains optimal. Since $f_{j''j} = 0$ for all $(j'', j) \in A$, by the first constraint of (5.19) we have $l_j \geq T\mu_j - f_{jj}$, and so the largest ϵ that maintains feasibility is $\min(T\mu_j - f_{jj}, f_{jj'})$, which will make either either $f'_{jj'} = 0$ or $f'_{jj} = T\mu_j$.

Combining Case 1 and Case 2, we have that for any optimal solution $\mathbf{f}$, if $f_{jj'} > 0$ and $f_{jj} < T\mu_j$ for some $(j, j') \in A$, we can always find another optimal solution $\mathbf{f}'$ such that either $f'_{jj'} = 0$ or $f'_{jj} = T\mu_j$. Because A has a finite number of links, we will always find an optimal solution $\mathbf{f}^*$ satisfying the condition in Lemma 5.5. □

5.6.2 Proof of Lemma 5.1

Proof of Lemma 5.1. We first observe that every unit of demand that arrives must be either be filled locally, filled by spillover or lost. By Assumption 5.1, the lowest cost option among these options is local fulfillment, and thus the cost incurred by each unit of DC j demand is at least $c_{jj} > 0$. Therefore, the expected cost of any online policy must be larger than the expected cost of filling all demand locally, so we have

$$T \sum_j c_{jj}\mu_j \leq C^*(T).$$

The expected amount of local fulfillment at DC j is $\mathbb{E}[\min(Tb_j, D_j(T))] \leq T\mu_j$. Further, note that the maximum cost that can be incurred by the greedy policy for a unit of DC j demand that is not filled locally is p_j (since the lost sale option is always available). Then the greedy policy incurs expected cost at most

$$\sum_j c_{jj} T\mu_j + \sum_j p_j \mathbb{E}[(D_j(T) - Tb_j)^+] \leq T \sum_j c_{jj}\mu_j + \sum_j p_j \mathbb{E}[(D_j(T) - T\mu_j)^+],$$

$$\leq T \sum_j c_{jj}\mu_j + O(\sqrt{T}),$$

where the first inequality follows from $b_j \geq \mu_j$, and the second from Lemma 5.2. Therefore we have

$$\frac{C^{Greedy}(T)}{C^*(T)} \leq 1 + O\big(T^{-\frac{1}{2}}\big).$$

□

5.6.3 Proof of Theorem 5.1

Proof of Theorem 5.1. Consider the optimization problem specified in (5.19). Let C^{DLP} denote the optimal objective of (5.19), and note that C^{DLP} is equal to $f(\mathbf{r})$ with $r_j = Tb_j$. By Lemma 5.4, we have that $C^{DLP} \leq C^*(T)$. By Assumption 5.2, there exists an optimal solution to (5.19) satisfying the condition in Lemma 5.5; let $(\mathbf{f}^*, \mathbf{l}^*)$ denote this optimal solution. Construct a spillover limit policy π using $\mathbf{f}^*$ as follows: let $\theta_{jj'} = f^*_{jj'}$ for all $j \neq j'$, and let $\theta^{DC}_j = \sum_{j \neq j'} f^*_{jj}$. Because the limits are set based on a linear program with the expected demand, we call π the DLP spillover limit policy. We will prove Theorem 5.1 by bounding the expected costs of π in three parts: local fulfillment, spillover fulfillment, and lost sales, compared to the costs of (5.19).

First, we prove a small consequence of the condition in Lemma 5.5: if $f^*_{jj} < T\mu_j$, then $f^*_{jj} = r^*_j$. To see this, first note that since $f^*_{jj} < T\mu_j$, we must have $f^*_{jj'} = 0$ for all $j' \neq j$ by Lemma 5.5. Suppose contrary to the claim that $f^*_{jj} < r^*_j$. Then, it is feasible to increase f^*_{jj} to $\min(r^*_j, T\mu_j)$, while decreasing the values of variable l_j and variables $f^*_{j'j}$ for $j \neq j$ such that $l_j + \sum_{j \neq j} f^*_{j'j}$ decrease by $\min(r^*_j, T\mu_j) - f^*_{jj}$. Because $c_{jj} \leq p_j$ and $c_{jj} \leq c_{j'j}$ for any $j' \neq j$, we have that the new solution is also optimal. Thus, we can assume without loss of optimality that $f^*_{jj} = r^*_j$ if $f^*_{jj} < T\mu_j$.

Local Fulfillment. We claim that under the DLP spillover limit policy, π, the expected amount of local fulfillment (i.e., DC j demand that is filled by DC j inventory) is less than f^*_{jj}. To see this, consider two cases. First, if $f^*_{jj} = T\mu_j$, then since local fulfillment must be no more than the total demand on any sample path, we have expected local fulfillment bound by $T\mu_j = f^*_{jj}$. Otherwise, if $f^*_{jj} < T\mu_j$, then $f^*_{jj} = r^*_j$. But local fulfillment must be no more than r^*_j on any sample path, implying that f^*_{jj} is again an upper bound on the expected local fulfillment.

Spillover Fulfillment. By definition of π, on any sample path the policy can not fill more than $f^*_{j'j}$ of DC j demand from DC j' inventory, thus the expected amount of this spillover fulfillment is less than $f^*_{j'j}$.

Lost Sales. We claim that under π, the expected amount of lost sales at DC j is bound by

$$l^*_j + O(\sqrt{T}).$$

To show this, we note that there are two types of lost sales that can arise under π. The first type occurs when DC j stocks out of inventory and an arriving demand is lost due to customer abandonment. The second type occurs when DC j stocks out of inventory and an arriving demand is not lost due to customer abandonment, but there are no available flexible inventories to fill the demand. To bound the first type of lost sales, note that by definition π, there is at least $r_j - \sum_{j' \neq j} f^*_{jj'} \geq f^*_{jj}$ DC j inventory available to fill DC j demand. Thus, given total DC j demand $D_j(T)$, at most $(D_j(T) - f^*_{jj})^+$ of it will arrive when DC j is stocked out. Each unit of this demand will become a lost sale due to abandonment independently with probability α_j, and so the expected amount of type one lost sales given total demand $D_j(T)$, is at most $\alpha_j(D_j(T) - f^*_{jj})^+$. Thus, the expected amount of type one lost sales is less than

$$\mathbb{E}[\alpha_j(D_j(T) - f^*_{jj})^+] \leq \alpha_j \mathbb{E}[(D_j(T) - T\mu_j)^+] + \alpha_j(T\mu_j - f^*_{jj}),$$

where the inequality follows from the triangle inequality and $f^*_{jj} \leq T\mu_j$.

Next we will bound type two lost sales. On a given sample path, let $\hat{D}_j$ denote the amount of DC j demand that arrives when DC j is stocked out, and that is not lost due to customer abandonment. This is the demand that may spillover to other flexible DCs with available inventory. For each j' such that $f^*_{j'j} > 0$, under policy π, DC j' has at least $r_j - \sum_{k \neq j, j'} f^*_{j'k}$ units of inventory that can be only used for either local fulfillment at j' or spillover fulfillment at DC j. Because $r_j - \sum_{k \neq j, j'} f^*_{j'k} \geq f^*_{j'j'} + f^*_{j'j}$ by the feasibility of $\mathbf{f}^*$, on any sample path, we have that there is at least

$$\min(f^*_{j'j}, f^*_{j'j} + f^*_{j'j'} - D_{j'}) = f^*_{j'j} - (D_{j'}(T) - f^*_{j'j'})^+$$

amount of inventories stored at DC j' that can be used to fulfill demand at DC j, where $D_{j'}(T)$ may change depending on the sample path. Therefore, on any sample path, the amount of type two lost sales is less than

$$
\begin{aligned}
&\big(\hat{D}_j - \sum_{j'\neq j} f^*_{j'j} + \sum_{j'\neq j} \mathbb{1}_{\{f^*_{j'j}>0\}}(D_{j'}(T) - f^*_{j'j'})^+\big)^+, \\
&\leq \big(\hat{D}_j - \sum_{j'\neq j} f^*_{j'j}\big)^+ + \sum_{j'\neq j} \mathbb{1}_{\{f^*_{j'j}>0\}}(D_{j'}(T) - f^*_{j'j'})^+, \\
&= \big(\hat{D}_j - \sum_{j'\neq j} f^*_{j'j}\big)^+ + \sum_{j'\neq j} \mathbb{1}_{\{f^*_{j'j}>0\}}(D_{j'}(T) - T\mu_{j'})^+,
\end{aligned} \tag{5.20}
$$

where the second line follows from the triangle inequality of $(\cdot)^+$, and the third line follows from Lemma 5.5. To bound the expected value of the first component in (5.20), we first take the expected value conditional on $D_j(T)$, (i.e., the expectation is taken over the random abandonment)

$$
\begin{aligned}
&\mathbb{E}\big[\big(\hat{D}_j - \sum_{j'\neq j} f^*_{j'j}\big)^+ | D_j(T)\big] \\
&\leq \mathbb{E}\big[\big(\hat{D}_j - \mathbb{E}[\hat{D}_j]\big)^+ | D_j(T)\big] + \mathbb{E}\big[\big(\mathbb{E}[\hat{D}_j] - \sum_{j'\neq j} f^*_{j'j}\big)^+ | D_j(T)\big], \\
&\leq \frac{1}{2}\sqrt{\alpha_j(1-\alpha_j)(D_j(T) - f^*_{jj})^+} + \big((1-\alpha_j)(D_j(T) - f^*_{jj})^+ - \sum_{j'\neq j} f^*_{j'j}\big)^+,
\end{aligned} \tag{5.21}
$$

where the first line is again the triangle inequality of $(\cdot)^+$, and the second line follows from Lemma 5.2 and the fact that, given $D_j(T)$, $\hat{D}_j$ is a binomial random variable with probability of success $1-\alpha_j$ and number of trials $(D_j(T) - f^*_{jj})^+$. For the first term in (5.21), taking expectations we have

$$
\mathbb{E}\Big[\frac{1}{2}\sqrt{\alpha_j(1-\alpha_j)(D_j(T) - f^*_{jj})^+}\Big] \leq \frac{1}{2}\sqrt{\alpha_j(1-\alpha_j)T\mu_j} = O(\sqrt{T}),
$$

where the inequality follows from $(D_j(T) - f^*_{jj})^+ \leq D_j(T)$ and Jensen's inequality. For the second term, taking expectations we have

$$
\begin{aligned}
&\mathbb{E}\Big[\big((1-\alpha_j)(D_j(T) - f^*_{jj})^+ - \sum_{j'\neq j} f^*_{j'j}\big)^+\Big] \\
&\leq (1-\alpha_j)\mathbb{E}\big[(D_j(T) - T\mu_j)^+\big] + \big((1-\alpha_j)(T\mu_j - f^*_{jj})^+ - \sum_{j'\neq j} f^*_{j'j}\big)^+, \\
&= (1-\alpha_j)\mathbb{E}\big[(D_j(T) - T\mu_j)^+\big] + (1-\alpha_j)(T\mu_j - f^*_{jj}) - \sum_{j'\neq j} f^*_{j'j}.
\end{aligned}
$$

where the first line follows from the triangle inequality of $(\cdot)^+$, and the second from the second constraint of (5.19). Thus, combining the bounds on the first and second term of (5.21), we get

$$\mathbb{E}\big[(\hat{D}_j - \sum_{j'\neq j} f^*_{j'j})^+ | D_j(T)\big]$$

$$\leq (1-\alpha_j)\mathbb{E}\big[(D_j(T) - T\mu_j)^+\big] + (1-\alpha_j)(T\mu_j - f^*_{jj}) - \sum_{j'\neq j} f^*_{j'j} + O(\sqrt{T}).$$

Now substitute the above into (5.20), and add type one lost sales, we can upper-bound the total expected lost sales as

$$T\mu_j - f^*_{jj} - \sum_{j'\neq j} f^*_{j'j} + \mathbb{E}\big[(D_j(T) - T\mu_j)^+\big] +$$

$$O(\sqrt{T}) + \sum_{j'\neq j} \mathbb{1}_{\{f^*_{j'j}>0\}} \mathbb{E}\big[(D_{j'}(T) - T\mu_{j'})^+\big]$$

$$\leq l^*_j + O(\sqrt{T}),$$

where the inequality follows from the first constraint of (5.19) and the $\sqrt{T}$ bound of Lemma 5.2 for each $\mathbb{E}\big[(D_{j'}(T) - T\mu_{j'})^+\big]$ term.

With these bounds on expected fulfillment, we have shown that the expected cost of π is less than

$$\sum_j \Big(p_j l^*_j + \sum_{j'} c_{jj'} f^*_{jj'}\Big) + O(\sqrt{T}) = C^{DLP} + O(\sqrt{T}) \leq C^*(T) + O(\sqrt{T}).$$

Finally, note that C^{DLP} scales linearly with T because $c_{jj'} > 0$ for all $(j, j') \in A$ and $p_j > 0$ for all j. Thus, we have

$$\frac{C^\pi(T)}{C^*(T)} \leq \frac{C^*(T) + O(\sqrt{T})}{C^*(T)} \leq 1 + O\big(T^{-\frac{1}{2}}\big)$$

□

5.6.4 Non-asymptotic Optimality Example Without Assumption 5.2

Example 5.2. *Consider a fulfillment network with 3 DCs ($N = 2$). Recall that each DC is responsible for orders from its own district, and we assume that DC 0 can fulfill orders for DC 1, while DC 1 can fulfill orders for DC 2 (see Figure 5.3). There are T days, orders arrive at district 1 and 2 as independent Poisson processes at a rate of 1 per day, while district 0 has no orders. We assume that both DC 0 and DC 1 starts with T inventory, while DC 2 starts with 0 inventory. The*

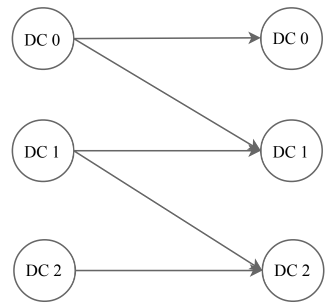

Figure 5.3: System considered in Example 5.2

abandonment rates are set to $\alpha_j = 0$*, i.e., customers do not abandon orders upon spillover. Note that this fulfillment network violates Assumption 5.2, as* $(0,1),(1,2) \in A$*, but* $(0,2) \notin A$*.*

Suppose that local fulfillment costs and lost sales costs are equal to c_l *and* p *at all DCs, respectively. Also, suppose that the spillover costs are uniformly equal to* c_f*. We assume that* $p \gg c_f > c_l$*, so it is the retailer's best interest to avoid lost sales as much as possible. First, consider policy* π_1 *where it uses only DC 0 to fulfill orders at DC 1, and DC 1 to fulfill orders at DC 2. In this case, the total lost sales costs under* π_1 *is equal to*

$$p\mathbb{E}[(D_2(T) - T)^+] = p\sqrt{T}.$$

On the other hand, consider a policy in which DC 1 is locally greedy in fulfilling demand in its own district. In that case, note that the lost sales at DC 2 are minimized if the policy does not hold back any inventories at DC 1 for orders at DC 2. Let π_2 *be such a policy. Under* π_2*, DC 1 will on average use* $\mathbb{E}[\min(D_1(T) + D_2(T), T)]$ *amount of inventories, because demand from both DC 1 and DC 2 is first fulfilled with inventories at DC 1. Because the demand for DC 1 and DC 2 have the same independent Poisson process, the average amount of inventories consumed at DC 1,* $\mathbb{E}[\min(D_1(T) + D_2(T), T)]$*, will be split equally to both DC 1 and DC 2. This implies that the lost sales cost under* π_2 *is at least*

$$p(\mathbb{E}[D_2(T)] - \mathbb{E}[\min(D_1(T) + D_2(T), T)]/2) = 0.5pT + O(1).$$

Therefore, if $p \gg c_f > c_l$*, any policy that does not hold back the inventories for local fulfillment cannot be asymptotically optimal, as* T *approaches infinity.*

5.7 Conclusion

In conclusion, we collaborate with a large e-retailer to apply flexibility design, a classical operations management strategy, in the context of the rapidly growing e-retailing industry. Motivated by the actual practice of the e-retailer and its data, we propose a model including key features such as commitment to local orders and customer abandonment due to spillover, which is new to the online resource allocation literature. We demonstrate both the effectiveness and limitation of the naive greedy policy, and use it to motivate the more powerful spillover limit policy. The spillover limit policy has a simple form, is easy to implement in practice, and is asymptotically optimal under a general condition. Finally, we use the spillover limit policy to estimate the benefit of adding fulfillment flexibility into a distribution system. We estimate that our newly proposed flexibility structure will decrease the lost sales plus fulfillment costs by approximately 1%, which will lead to a significant increase in profit for our industrial partner.

Chapter 6

Conclusion

The research in this dissertation has considered practical and approximately optimal solutions for decision problems arising in complex operational settings with uncertain demand. We have consider three classes of problems, assemble-to-order problems, matching problems on a network, and fulfillment flexibility in e-retailing. For each of these problems, we have shown how to use the theory of optimization and applied probability to design a policy that is guaranteed to perform well on a wide range of problem instances. These results provide insight for managers in how to think about solving these problems, as well as implementable algorithms for making the decisions. It is my hope that the techniques developed in this dissertation will find new and extended uses in a wide range of applications.

Bibliography

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.

Adelman, Daniel. 2007. Dynamic bid prices in revenue management. *Operations Research* **55**(4) 647–661.

Agrawal, Shipra, Zizhuo Wang, Yinyu Ye. 2014. A dynamic near-optimal algorithm for online linear programming. *Operations Research* **62**(4) 876–890.

Alijani, Reza, Siddhartha Banerjee, Sreenivas Gollapudi, Kostas Kollias, Kamesh Munagala. 2017. Two-sided facility location. *arXiv preprint arXiv:1711.11392* .

Ansari, Sardar, Laura Albert McLay, Maria E Mayorga. 2017. A maximum expected covering problem for district design. *Transportation Science* **51**(1) 376–390.

Asadpour, A., X. Wang, J. Zhang. 2018. Online resource allocation with limited flexibility. Working Paper, New York University, NY.

Baker, Kenneth R, Michael J Magazine, Henry LW Nuttle. 1986. The effect of commonality on safety stock in a simple inventory model. *Management Science* **32**(8) 982–988.

Ball, Michael O, Maurice Queyranne. 2009. Toward robust revenue management: Competitive analysis of online booking. *Operations Research* **57**(4) 950–963.

Belobaba, Peter P. 1989. Or practice—application of a probabilistic decision model to airline seat inventory control. *Operations Research* **37**(2) 183–197.

Bernstein, Fernando, Gregory A DeCroix, Yulan Wang. 2011. The impact of demand aggregation through delayed component allocation in an assemble-to-order system. *Management Science* **57**(6) 1154–1171.

Blume, Lawrence E, David Easley, Jon Kleinberg, Eva Tardos. 2009. Trading networks with price-setting agents. *Games and Economic Behavior* **67**(1) 36–50.

Caro, Felipe, Kumar Rajaram, Jens Wollenweber. 2012. Process location and product distribution with uncertain yields. *Operations research* **60**(5) 1050–1063.

Chen, Annie I-An. 2017. Large-scale optimization in online-retail inventory management. Ph.D. thesis, Massachusetts Institute of Technology.

Chen, Xi, Jiawei Zhang, Yuan Zhou. 2015. Optimal sparse designs for process flexibility via probabilistic expanders. *Operations Research* **63**(5) 1159–1176.

Cho, Soo-Haeng, Hoon Jang, Taesik Lee, John Turner. 2014. Simultaneous location of trauma centers and helicopters for emergency medical service planning. *Operations Research* **62**(4) 751–771.

Cooper, William L. 2002. Asymptotic behavior of an allocation policy for revenue management. *Operations Research* **50**(4) 720–727.

Corless, Robert M, Gaston H Gonnet, David EG Hare, David J Jeffrey, Donald E Knuth. 1996. On the Lambert W function. *Advances in Computational Mathematics* **5**(1) 329–359.

Cornuejols, Gerard, Marshall L Fisher, George L Nemhauser. 1977. Location of bank accounts to optimize float: An analytic study of exact and approximate algorithms. *Management science* **23**(8) 789–810.

de Almeida Correia, Gonçalo Homem, António Pais Antunes. 2012. Optimization approach to depot location and trip selection in one-way carsharing systems. *Transportation Research Part E: Logistics and Transportation Review* **48**(1) 233–247.

DeCroix, Gregory A, Jing-Sheng Song, Paul H Zipkin. 2009. Managing an assemble-to-order system with returns. *Manufacturing & Service Operations Management* **11**(1) 144–159.

DeValve, Levi, Saša Pekeč, Yehua Wei. 2019a. Matching supply and demand in a resource constrained service network. *Working Paper, Duke University.*

DeValve, Levi, Saša Pekeč, Yehua Wei. 2019b. A primal-dual approach to analyzing ATO systems. *Working Paper, Duke University.*

DeValve, Levi, Yehua Wei, Di Wu, Rong Yuan. 2019c. Understanding the value of fulfillment flexibility in an online retailing environment. *Working Paper, Duke University.*

Doğru, 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.

Doğru, 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.

Durrett, Rick. 2013. *Probability: Theory and Examples*. 4th ed.

Dye, Shane, Leen Stougie, Asgeir Tomasgard. 2003. The stochastic single resource service-provision problem. *Naval Research Logistics* **50**(8) 869–887.

Feige, Uriel. 1998. A threshold of ln(n) for approximating set cover. *Journal of the ACM (JACM)* **45**(4) 634–652.

Feige, Uriel, Vahab S Mirrokni, Jan Vondrak. 2011. Maximizing non-monotone submodular functions. *SIAM Journal on Computing* **40**(4) 1133–1153.

Feldman, Jon, Aranyak Mehta, Vahab Mirrokni, S Muthukrishnan. 2009. Online stochastic matching: Beating 1-1/e. *Foundations of Computer Science, 2009. FOCS'09. 50th Annual IEEE Symposium on*. IEEE, 117–126.

Fisher, Marshall L, George L Nemhauser, Laurence A Wolsey. 1978. An analysis of approximations for maximizing submodular set functions II. *Polyhedral combinatorics*. Springer, 73–87.

Freund, Daniel, Shane G Henderson, David B Shmoys. 2017. Minimizing multimodular functions and allocating capacity in bike-sharing systems. *arXiv preprint arXiv:1611.09304* .

Gallego, Guillermo, Garrett Van Ryzin. 1997. A multiproduct dynamic pricing problem and its applications to network yield management. *Operations research* **45**(1) 24–41.

Gerchak, Yigal, Michael J Magazine, A Bruce Gamble. 1988. Component commonality with service level requirements. *Management Science* **34**(6) 753–760.

Glasserman, Paul, Yashan Wang. 1998. Leadtime-inventory trade-offs in assemble-to-order systems. *Operations Research* **46**(6) 858–871.

Gourdin, Éric, Martine Labbé, Gilbert Laporte. 2000. The uncapacitated facility location problem with client matching. *Operations Research* **48**(5) 671–685.

Govindarajan, Aravind, Amitabh Sinha, Joline Uichanco. 2018. Joint inventory and fulfillment decisions for omnichannel retail networks. Working Paper, University of Michigan.

Gupta, Anupam, R Ravi, Amitabh Sinha. 2007. LP rounding approximation algorithms for stochastic network design. *Mathematics of Operations Research* **32**(2) 345–364.

Hagerup, Torben, Christine Rüb. 1990. A guided tour of Chernoff bounds. *Information Processing Letters* **33**(6) 305–308.

Halman, Nir, James B Orlin, David Simchi-Levi. 2012. Approximating the nonlinear newsvendor and single-item stochastic lot-sizing problems when data is given by an oracle. *Operations Research* **60**(2) 429–446.

He, Long, Ho-Yin Mak, Ying Rong, Zuo-Jun Max Shen. 2017. Service region design for urban electric vehicle sharing systems. *Manufacturing & Service Operations Management* **19**(2) 309–327.

Huh, Woonghee Tim, Ganesh Janakiraman, John A Muckstadt, Paat Rusmevichientong. 2009. Asymptotic optimality of order-up-to policies in lost sales inventory systems. *Management Science* **55**(3) 404–420.

Jaillet, Patrick, Xin Lu. 2013. Online stochastic matching: New algorithms with better bounds. *Mathematics of Operations Research* **39**(3) 624–646.

Jain, Kamal, Vijay V Vazirani. 2001. Approximation algorithms for metric facility location and k-median problems using the primal-dual schema and Lagrangian relaxation. *Journal of the ACM* **48**(2) 274–296.

Jasin, Stefanus, Sunil Kumar. 2012. A re-solving heuristic with bounded revenue loss for network revenue management with customer choice. *Mathematics of Operations Research* **37**(2) 313–345.

Jasin, Stefanus, Amitabh Sinha. 2015. An lp-based correlated rounding scheme for multi-item ecommerce order fulfillment. *Operations Research* **63**(6) 1336–1351.

Jauvion, Grégoire, Nicolas Grislain, Pascal Dkengne Sielenou, Aurélien Garivier, Sébastien Gerchinovitz. 2018. Optimization of a ssp's header bidding strategy using thompson sampling. *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. ACM, 425–432.

Jordan, W., S. Graves. 1995a. Principles on the benefits of manufacturing process flexibility. *Management Science* **41**(4) 577–594.

Jordan, William C, Stephen C Graves. 1995b. Principles on the benefits of manufacturing process flexibility. *Management Science* **41**(4) 577–594.

Kaplan, Alan. 1969. Stock rationing. *Management Science* **15**(5) 260–267.

Kapuscinski, Roman, Rachel Q Zhang, Paul Carbonneau, Robert Moore, Bill Reeves. 2004. Inventory decisions in Dell's supply chain. *Interfaces* **34**(3) 191–205.

Karp, Richard M, Umesh V Vazirani, Vijay V Vazirani. 1990. An optimal algorithm for on-line bipartite matching. *Proceedings of the twenty-second annual ACM symposium on Theory of computing*. ACM, 352–358.

Kleywegt, Anton J, Alexander Shapiro, Tito Homem-de Mello. 2002. The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization* **12**(2) 479–502.

Kolliopoulos, Stavros G. 2003. Approximating covering integer programs with multiplicity constraints. *Discrete Applied Mathematics* **129**(2-3) 461–473.

Kolliopoulos, Stavros G, Neal E Young. 2001. Tight approximation results for general covering integer programs. *Foundations of Computer Science, 2001. Proceedings. 42nd IEEE Symposium on*. IEEE, 522–528.

Lee, Daniel D, H Sebastian Seung. 2001. Algorithms for non-negative matrix factorization. *Advances in neural information processing systems*. 556–562.

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.

Levi, Retsef, Martin Pál, Robin O Roundy, David B Shmoys. 2007. Approximation algorithms for stochastic inventory control models. *Mathematics of Operations Research* **32**(2) 284–302.

Levi, Retsef, Robin Roundy, David Shmoys, Maxim Sviridenko. 2008. A constant approximation algorithm for the one-warehouse multiretailer problem. *Management Science* **54**(4) 763–776.

Levi, Retsef, Robin O Roundy, David B Shmoys. 2006. Primal-dual algorithms for deterministic inventory problems. *Mathematics of Operations Research* **31**(2) 267–284.

Littlewood, Kenneth. 1972. Forecasting and control of passenger bookings. *Airline Group International Federation of Operational Research Societies Proceedings, 1972* **12** 95–117.

Louveaux, François V, Dominique Peeters. 1992. A dual-based procedure for stochastic facility location. *Operations research* **40**(3) 564–573.

Lu, Lijian, Jing-Sheng Song, Hanqin Zhang. 2015. Optimal and asymptotically optimal policies for assemble-to-order N- and W-systems. *Naval Research Logistics* **62**(8) 617–645.

Lu, Yingdong, Jing-Sheng Song. 2005. Order-based cost optimization in assemble-to-order systems. *Operations Research* **53**(1) 151–169.

Lu, Yingdong, Jing-Sheng Song, David D. Yao. 2005. Backorder minimization in multiproduct assemble-to-order systems. *IIE Transactions* **37**(8) 763–774.

Lu, Yingdong, Jing-Sheng Song, Yao Zhao. 2010. No-holdback allocation rules for continuous-time assemble-to-order systems. *Operations Research* **58**(3) 691–705.

Lubell, David. 1966. A short proof of Sperner's lemma. *Journal of Combinatorial Theory* **1**(2) 299.

Ma, Will, David Simchi-Levi. 2017. Online resource allocation under arbitrary arrivals: Optimal algorithms and tight competitive ratios. *Working Paper, MIT*.

Mangasarian, O.L. 1979. Uniqueness of solution in linear programming. *Linear Algebra and its Applications* **25** 151–162.

Manshadi, Vahideh H, Shayan Oveis Gharan, Amin Saberi. 2012. Online stochastic matching: Online actions based on offline statistics. *Mathematics of Operations Research* **37**(4) 559–573.

Moshkovitz, Dana. 2015. The projection games conjecture and the NP-hardness of ln n-approximating set-cover. *Theory of Computing* **11**(1) 221–235.

Muller, Joann. 2010. BMW's push for made-to-order cars. *Forbes* (September 9). `https://www.forbes.com/forbes/2010/0927/companies-bmw-general-motors-cars-bespoke-auto.html`.

Nadar, Emre, Mustafa Akan, Alan Scheller-Wolf. 2014. Optimal structural results for assemble-to-order generalized M-systems. *Operations Research* **62**(3) 571–579.

Nagarajan, Mahesh, Yehuda Bassok. 2008. A bargaining framework in supply chains: The assembly problem. *Management Science* **54**(8) 1482–1496.

Nagarajan, Mahesh, Greys Sošić. 2009. Coalition stability in assembly models. *Operations Research* **57**(1) 131–145.

Nagarajan, Viswanath, Cong Shi. 2016. Approximation algorithms for inventory problems with submodular or routing costs. *Mathematical Programming* **160**(1-2) 225–244.

National Bureau of Statistics of China. 2018. Total retail sales of consumer goods in december 2017. URL `http://www.stats.gov.cn/english/pressrelease/201801/t20180126_1577681.html`.

Nemhauser, George L, Laurence A Wolsey, Marshall L Fisher. 1978. An analysis of approximations for maximizing submodular set functions I. *Mathematical Programming* **14**(1) 265–294.

Nemirovski, Arkadi, Anatoli Juditsky, Guanghui Lan, Alexander Shapiro. 2009. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization* **19**(4) 1574–1609.

Nguyen, Thành. 2015. Coalitional bargaining in networks. *Operations Research* **63**(3) 501–511.

Nguyen, Thanh, Rakesh Vohra. 2017. Near feasible stable matchings with couples. *Working Paper, Purdue*.

O'Mahony, Eoin, David B Shmoys. 2015. Data analysis and optimization for (citi) bike sharing. *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*. 687–694.

Papadimitriou, Christos H, Kenneth Steiglitz. 1998. *Combinatorial optimization: algorithms and complexity*. Courier Corporation.

Plambeck, Erica L, Amy R Ward. 2006. Optimal control of a high-volume assemble-to-order system. *Mathematics of Operations Research* **31**(3) 453–477.

Plambeck, Erica L, Amy R Ward. 2007. Note: A separation principle for a class of assemble-to-order systems with expediting. *Operations Research* **55**(3) 603–609.

Prebid.org. 2018. How to optimize header bidding setup. Software Documentation. URL `http://prebid.org/overview/how-many-bidders-for-header-bidding.html`.

Press, William H, Saul A Teukolsky, William T Vetterling, Brian P Flannery. 2002. *Numerical Recipes in C: The Art of Scientific Computing*. 2nd ed. Cambridge University Press.

Qin, Rui, Yong Yuan, Fei-Yue Wang. 2017. Optimizing the revenue for ad exchanges in header bidding advertising markets. *Systems, Man, and Cybernetics (SMC), 2017 IEEE International Conference on*. IEEE, 432–437.

Ravi, R, Amitabh Sinha. 2006. Hedging uncertainty: Approximation algorithms for stochastic optimization problems. *Mathematical Programming* **108**(1) 97–114.

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.

Robbins, Herbert. 1948. The asymptotic distribution of the sum of a random number of random variables. *Bulletin of the American Mathematical Society* **54**(12) 1151–1161.

Sayedi, Amin. 2018. Real-time bidding in online display advertising. *Marketing Science* **37**(4) 553–568.

Schulz, Andreas S, Martin Skutella. 2002. Scheduling unrelated machines by randomized rounding. *SIAM Journal on Discrete Mathematics* **15**(4) 450–469.

Shapiro, Alexander, Darinka Dentcheva, Andrzej Ruszczyński. 2009. *Lectures on Stochastic Programming: Modeling and Theory*. SIAM.

Shmoys, David B, Chaitanya Swamy. 2006. An approximation scheme for stochastic linear programming and its application to stochastic integer programs. *Journal of the ACM* **53**(6) 978–1012.

Shmoys, David B, Éva Tardos, Karen Aardal. 1997. Approximation algorithms for facility location problems. *Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing*. ACM, 265–274.

Shu, Jia, Mabel C Chou, Qizhang Liu, Chung-Piaw Teo, I-Lin Wang. 2013. Models for effective deployment and redistribution of bicycles within public bicycle-sharing systems. *Operations Research* **61**(6) 1346–1359.

Simchi-Levi, David, Yehua Wei. 2012. Understanding the performance of the long chain and sparse designs in process flexibility. *Operations research* **60**(5) 1125–1141.

Song, Jing-Sheng. 2002. Order-based backorders and their implications in multi-item inventory systems. *Management Science* **48**(4) 499–516.

Song, Jing-Sheng, Yao Zhao. 2009. The value of component commonality in a dynamic inventory system with lead times. *Manufacturing & Service Operations Management* **11**(3) 493–508.

Song, Jing-Sheng, Paul Zipkin. 2003. Supply chain operations: Assemble-to-order systems. *Supply Chain Management: Design, Coordination and Operation*, *Handbooks in Operations Research and Management Science*, vol. 11. Elsevier, 561–596.

Swamy, Chaitanya, David B Shmoys. 2012. Sampling-based approximation algorithms for multistage stochastic optimization. *SIAM Journal on Computing* **41**(4) 975–1004.

Talluri, Kalyan T, Garrett J Van Ryzin. 2006. *The theory and practice of revenue management*, vol. 68. Springer Science & Business Media.

Tamir, Arie. 1976. On totally unimodular matrices. *Networks* **6**(4) 373–382.

Teo, Chung-Piaw, Dimitris Bertsimas. 2001. Multistage lot sizing problems via randomized rounding. *Operations Research* **49**(4) 599–608.

Topkis, Donald M. 1968. Optimal ordering and rationing policies in a nonstationary dynamic inventory model with n demand classes. *Management Science* **15**(3) 160–176.

US Department of Commerce. 2018. Quarterly retail e-commerce sales 4th quarter 2017. Press Release. URL `https://www2.census.gov/retail/releases/historical/ecomm/17q4.pdf`.

van Jaarsveld, Willem, Alan Scheller-Wolf. 2015. Optimization of industrial-scale assemble-to-order systems. *INFORMS Journal on Computing* **27**(3) 544–560.

Vazirani, Vijay V. 2013. *Approximation Algorithms*. Springer Science & Business Media.

Vidakovic, Ratko. 2017. The beginner's guide to header bidding. *AdProfs* (March 30). `https://adprofs.co/beginners-guide-to-header-bidding/`.

Wei, Lai, Stefanus Jasin, Roman Kapuscinski. 2017. Shipping consolidation with delivery deadline and expedited shipment options. Working Paper.

Williamson, David P, Michel X Goemans, Milena Mihail, Vijay V Vazirani. 1995. A primal-dual approximation algorithm for generalized Steiner network problems. *Combinatorica* **15**(3) 435–454.

Xin, Linwei, David A Goldberg. 2016. Optimality gap of constant-order policies decays exponentially in the lead time for lost sales models. *Operations Research* **64**(6) 1556–1565.

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 Zhang. 2018. Online demand fulfillment under limited flexibility. Working Paper.

Zhu, Zhanguo, Feng Chu, Linyan Sun. 2010. The capacitated plant location problem with customers and suppliers matching. *Transportation Research Part E: Logistics and Transportation Review* **46**(3) 469–480.

Zipkin, Paul. 2016. Some specially structured assemble-to-order systems. *Operations Research Letters* **44**(1) 136–142.