EMERGING TOPICS IN SUPPLY CHAIN MANAGEMENT: PRODUCT SUBSTITUTION, DEMAND AMBIGUITY, AND ENVIRONMENTAL AND SOCIAL RESPONSIBILITY

A Dissertation

Submitted to the Faculty

of

Purdue University

by

Chengzhang Li

In Partial Fulfillment of the

Requirements for the Degree

of

Doctor of Philosophy

August 2019

Purdue University

West Lafayette, Indiana

THE PURDUE UNIVERSITY GRADUATE SCHOOL
STATEMENT OF DISSERTATION APPROVAL

Dr. Qi Annabelle Feng, Co-Chair
Krannert School of Management
Dr. Mengshi Lu, Co-Chair
Krannert School of Management
Dr. Gökçe Esenduran
Krannert School of Management
Dr. Zheng Justin Jia
Krannert School of Management
Dr. J. George Shanthikumar
Krannert School of Management

Approved by:

Dr. Yanjun Li
Head of the School Graduate Program

Dedicated to my family

ACKNOWLEDGMENTS

First and foremost, I would like to express my sincere gratitude to my advisors and dissertation committee co-chairs, Professor Qi Annabelle Feng and Professor Mengshi Lu. They are not only great research advisors but also great mentors. During the past six years, they have provided me with sufficient academic guidance, have contributed a significant amount of time, and have been tolerant of the mistakes I have made along my Ph.D. journey.

I am also deeply grateful to Professor J. George Shanthikumar for serving on my dissertation committee, teaching Stochastic Models and Stochastic Dynamic Programming, and his guidance on my research. It has been my great pleasure working with him during my Ph.D. study. His brilliant ideas in the class and our research discussion always enlighten me, and his dedication and mastery of knowledge have always been a great inspiration to me.

I would also like to express my great appreciation to my committee members, Professor Gökçe Esenduran and Professor Zheng Justin Jia, who have raised critical questions which helped me to improve the quality of the dissertation significantly.

I would also like to thank the other Operations Management faculty members, especially, Professor Gemma Berenguer, Professor Susan Lu, Professor Zhan Pang, and Professor Pengyi Shi for advising me on the Ph.D. study and the job search. I am also very thankful to Dr. Olga Senicheva for her help with my teaching.

I would also like to thank my best friends who have always supported me on the way including Xing Fang, Taotao He, Tao Jiang, Yutong Xue, and many others. A very special appreciation to Zhongjie Ma and Yuanchen Li for their help and accompany during the past years.

Finally, I am very grateful to my family. I appreciate the unconditional love and support from my parents, Mr. Hao Li and Ms. Defen Zhang.

TABLE OF CONTENTS

Page

LIST OF TABLES

LIST OF FIGURES

Figure | Page

ABSTRACT

Li, Chengzhang Ph.D., Purdue University, August 2019. Emerging Topics in Supply Chain Management: Product Substitution, Demand Ambiguity, and Environmental and Social Responsibility. Major Professors: Qi Annabelle Feng, Mengshi Lu.

This study examines several emerging topics in supply chain management including the dynamic product substitution, the joint optimization of price and order quantity with demand ambiguity, and the implementation of the environmental and social responsibility (ESR) programs. We provide below a brief description of the results obtained for the specific problems considered in this study.

In the first problem discussed in Chapter 2, we present a dynamic model, in which the firm replenishes product inventories from uncertain sources and dynamically allocates available products to meet the uncertain demands with the flexibility of substitution. To address the analytical challenge associated with multi-product management, we develop an approximation algorithm that leverages the value of substitution, while allowing separability of the future profit among the products. Through extensive numerical analysis, we demonstrate that our approximation yields good performance measured by the percentage profit gap against an upper bound problem. We show that substitution can generate significant benefits when the supply capacities are moderate, the supply and demand uncertainties are high, or the replenishment cycle is short.

In the second problem discussed in Chapter 3, we study the problem of jointly optimizing the price and order quantity for a perishable product in the presence of demand ambiguity. We employ the minimax regret decision criterion to minimize the worst-case regret, which is defined as the difference between the optimal profit that could be obtained with perfect information and the realized profit using the

decision made with ambiguous demand information. We characterize the optimal pricing and ordering decisions under the minimax regret criterion and compare their properties with those in the classical models that seek to maximize the expected profit. We compare the minimax regret approach with two other approaches that are commonly used under demand ambiguity, namely the max-min robust approach and the regression-based data-driven approach. In the demand ambiguity setting, we show that the minimax regret approach avoids the high degree of conservativeness that is often incurred in the max-min approach. In the data-driven setting, we show via a numerical study that the minimax regret approach outperforms the classical regression-based approach when data is scarce, when the demand has high volatility, or when the demand model is misspecified.

In the third problem discussed in Chapter 4, we focus on the problem of administering ESR programs throughout a complex supply network. We apply a bilateral bargaining framework to analyze to what extent an ESR initiator should directly engage higher-tier suppliers, as opposed to delegating the assurance of ESR compliance to its first-tier suppliers. We show that the eventual structure of negotiation relationships can be derived by finding a shortest path tree in the supply network with the arc cost defined as a monotone function of the negotiating parties' relative bargaining power. We find that the ESR initiator tends to delegate ESR compliance negotiation to a supplier that is strong in negotiations with higher-tier suppliers. When the supply network is complex (i.e., wide and deep), directly engaging all suppliers for ESR compliance can lead to a larger gain by the initiator than fully delegating the negotiations with higher-tier suppliers to the first-tier ones.

1. INTRODUCTION

This study examines several emerging issues in supply chain management. First, while the uncertainties from both supply and demand sides increase the complexity of managing a mix of supplies and demands, with the advancement of information technology, production substitution is proposed as an operational strategy to mitigate the mismatch between the supplies and demands. However, the design of the substitution policy can be challenging and the coordination between the substitution and replenishment decisions can be complex. We study the dynamic substitution policy in the presence of supply and demand uncertainties with a dynamic model. With the notion of stochastic linearity and an efficient approximation algorithm, we investigate the value of the dynamic substitution.

Second, the coordination of the inventory and the pricing decisions plays an important role in the marketing-operations interfaces. While the conventional approaches assume that the firms have the complete information of the customer demand, the exact distributional information of the random demand is unknown or ambiguous in practice. To facilitate the decision making with the demand ambiguity, we employ the minimax regret framework to coordinate the price and inventory decisions. We characterize the structural properties of the optimal decisions. We also compare the minimax regret criterion with other two widely-used approaches, namely the max-min robust approach and the regression-based data-driven approach.

Third, with increasing customers' awareness of the environmental and social impact of the supply chains, the compliance of ESR policy has become a core of the business. Specifically, responsible sourcing has become an important strategy for the firms to ensure the compliance of the supply chains with the ESR requirements. With the increasing complexity in the supply networks, a common problem faced by the brand retailers and manufactures is, how should the initiator implement the ESR program and ensure the compliance of all members involved? We apply the multi-unit bilateral bargaining framework to investigate to what extent an ESR initiator should directly engage higher-tier suppliers, or as opposed to delegating the assurance of ESR compliance to its first-tier suppliers.

A brief introduction of this dissertation is provided below.

1.1 Dynamic Product Substitution

In recent years, the advancement of e-commerce and information technology has enabled retailers to fulfill consumer demands with unprecedented flexibility. Take online grocery delivery as an example. Due to the highly uncertain supply processes, it is common that the originally ordered items may be out-of-stock or may not have the desired qualities. To address the supply uncertainty, Instacart—one of the leading start-up companies in online grocery delivery—asks customers to select replacement options when placing the orders, which include substituting with customer-specified items or allowing Instacart's employees to choose what they consider to be the best alternatives (Instacart, 2018). Another example arises in omnichannel retailing. More often than not, traditional brick-and-mortar stores are involved in online sales to increase revenue, and in the meantime use their physical outlets to enhance customer experiences (Kumar et al., 2018). While omnichannel selling allows customers to buy online and pick up at store, it is generally believed that fulfilling online orders through central warehouses allows for demand pooling and therefore reduces inventory costs. Interestingly, recent studies by Pulse Commerce (2018) find that "some of the most successful [omnichannel retailers] fulfill an increasing proportion of online orders by shipping from stores, instead of from warehouses." Compared with shipping solely from warehouses, shipping from stores can provide faster and cheaper delivery, increase inventory turnover when shipping from overstocked stores, and avoid lost sales when an item is out of stock online (Pulse Commerce, 2018).

The common strategy in both examples is the dynamic remixing of the supplies (for different products or from different locations) using substitution to match the mix of demands. Such dynamic substitution has become popular in many industries. However, due to the nature of uncertain and likely dependent product demands, the design of efficient substitution policy is challenging. As a result, the existing studies on product substitution focus heavily on the case of upgrade (or downward substitution), which greatly limits the applicability of the results developed. In the aforementioned examples, the products or locations may not have a rank order of their values and the existing development does not address such situations.

Moreover, the supply planning for a multi-product system can also be complex, imposing added difficulties in coordinating replenishment decisions with substitution decisions. It is nevertheless believed that the replenishment and substitution can be operated separately in the sense that if the firm can make appropriate replenishment decisions to obtain the right supply mix, it does not have much need for product substitution (see, e.g., Shumsky and Zhang, 2009; Yu et al., 2015; Yao et al., 2016). However, this conclusion is drawn under condition where possible complications in the supply process are ignored. It is unclear how the supply conditions can affect the synergy between replenishment and substitution decisions.

In Chapter 2, we formulate a dynamic model, in which the firm replenishes product inventories from uncertain sources and dynamically allocates available products to meet the uncertain demands with the flexibility of substitution. To address the analytical challenge associated with multi-product management, we develop an approximation algorithm that leverages the value of substitution, while allowing separability of the future profit among the products. This approximation algorithm iteratively solves a transportation problem in a network with the Monge property. The application of the Monge property allows for dealing with general substitution structures, which generalizes the commonly studied downward substitution models. Through extensive numerical analysis, we demonstrate that our approximation yields good performance measured by the percentage profit gap against an upper bound problem. We also show that substitution can generate significant benefit when the supply capacities are moderate, when the supply and demand uncertainties are high, or when the replenishment cycle is short.

1.2 Demand Ambiguity

Coordinating pricing and inventory decisions plays an important role in managing the marketing-operations interface. Pricing strategies can effectually shape customer demand as well as demand uncertainty, which is the main consideration of inventory management; inventory decisions can improve operational efficiency and further enhance the effectiveness of pricing. Joint price-inventory management problems have been widely studied in the operations research literature. Among various joint price-inventory models, a fundamental one is the pricing newsvendor problem where the firm needs to determine the price and order quantity of a perishable product before the selling period. Despite its parsimonious setting, the pricing newsvendor problem forms the building block of many operations management models involving pricing and inventory decisions.

With the advancement of technology, firms can now collect real-time information of the customer demand and adjust the selling price based on the information. The conventional modeling approach for pricing and inventory management problem requires the complete distributional information of the random demand. However, in practice, the exact distributional form is typically unknown or ambiguous. In other words, there is a gap between the information available in practice and the modeling approach used in research. To bridge this gap, inspired by recent development in inventory management and pricing with demand ambiguity (Perakis and Roels, 2008; Caldentey et al., 2016), we develop a robust and tractable approach in the context of the pricing newsvendor problem, which can tackle demand ambiguity and is viable for data integration.

Specifically, we consider a widely used demand model where the demand is a function of the unit selling price and an uncertain factor. To allow direct comparison with the conventional approach, we consider the case where the form of the demand function is known, but the distributional information of the uncertain factor remains ambiguous. We assume that the firm only knows the support of the uncertain factor, e.g., the interval within which the uncertain factor lies with high confidence. With such demand ambiguity, the firm cannot maximize the expected profit using the traditional approach. Therefore, we adopt the minimax regret decision criterion to determine the price and order quantity that minimizes the worst-case regret. The regret is defined as the gap between the optimal profit that the firm could obtain with perfect demand information and the realized profit using decisions made with ambiguous demand information. The minimax regret criterion is an important alternative to maximizing expected payoff in decision theory. It has been adopted in inventory management or pricing to tackle ambiguity and generate new insights (see, Perakis and Roels, 2008; Caldentey et al., 2016).

In Chapter 3, we employ the minimax regret decision criterion to minimize the worst-case regret, which is defined as the difference between the optimal profit that could be obtained with perfect information and the realized profit using the decision made with ambiguous demand information. First, we characterize the optimal pricing and ordering decisions under the minimax regret criterion and compare their properties with those in the classical models that seek to maximize the expected profit. Specifically, we explore the impact of inventory risk by comparing the optimal price and the risk-free price, and study comparative statics with respect to the degree of demand ambiguity and the unit ordering cost. Second, we compare the minimax regret approach with two other approaches that are commonly used under demand ambiguity, namely the max-min robust approach and the regression-based data-driven approach. In the demand ambiguity setting, we show that the minimax regret approach avoids the high degree of conservativeness that is often incurred in the max-min approach. In the data-driven setting, we show via a numerical study that the minimax regret approach outperforms the classical regression-based approach when data is scarce, when the demand has high volatility, or when the demand model is misspecified.

1.3 Environmental and Social Responsibility

Consumers have become increasingly conscious about the environmental and social impact of the supply chains that generate the products or services they consume (Agrawal and Lee, 2016). Responsible sourcing has become an important part of corporate responsibility and many firms have started working with their suppliers and investing on environmental and social responsibility (ESR) initiatives. According to the International Association of Contract and Commercial Managers, nearly three-quarters of companies include a sustainability clause in their procurement contracts (Ecovadis, 2018). Based on 1,409 interviews with CEOs in 83 countries, PwC (2016) finds that 64% of CEOs consider ESR to be core to their business rather than a stand-alone program.

Years ago, many firms would treat ESR effort as a pure cost. Increasingly, firms realize that well implemented ESR measures can induce significant increase in customer loyalty, thereby generating long-term value to the firm. A survey by Nielsen reports that 55 percent of global online customers across 60 countries express their willingness to pay more for products and services provided by companies that are committed to positive social and environmental impact (Nielsen, 2014). Cone Communications also confirms that 91 percent of global consumers are likely to switch brands to one associated with a good cause, given comparable price and quality (Cone Communications, 2013).

A non-compliance instance, even if occurring at an upstream supplier with no direct relationship, can greatly damage the brand image of the downstream firm. For example, the entire organic eggs industry in Germany suffered significantly from a loss of trust among consumers due to an instance of dioxin contamination discovery in 2010. The source of contamination was found to be the maize meal imported from Ukraine (Wiese and Toporowski, 2013). With the relentless move of globalization, ensuring ESR throughout a large scale supply network is a challenging task. A recent study by Sedex, one of the largest non-profit organizations for promoting responsible sourcing, concerns ten companies with collectively 3,922 supplier relationships (Sedex, 2013). This study suggests that the average number of ESR non-compliance by tier-two and tier-three suppliers are 18% and 27% higher than that by tier-one suppliers, concluding that "[the] greatest and most critical [ESR] risks are found deeper down the supply chain."

Aware of the impact of far upstream suppliers on ESR success, many companies have gone deep into their supply network to emphasize the importance of full compliance for their ESR campaigns. Starbucks commits itself to 100 percent ethically sourced coffee from more than 170,000 farmers. H&M, a Swedish fashion firm, aims toward full traceability of their supply process involving 1.6 million workers across more than 800 suppliers (Chen, 2017). In many cases, a narrowly focused ESR effort that intends to solve issues in one part of the supply network but overlooks others would backfire. An instance occurred at PUMA, a German sportswear manufacturer. Shortly after PUMA's corrective measures to address the worker safety and fairness issues at one of its Chinese suppliers, Taiway Sports Inc., devastating labor conditions are reported at another supplier, Dongguan Surpassing Shoe Co. Ltd. The report quickly prompted significant criticisms and questions on PUMA's commitment to social responsibility.

The assurance of full compliance of an ESR initiative throughout the supply chain can be achieved with different ways of implementation. A recent study by Ecovadis (2018) shows that most companies *delegate* the responsibility of ensuring higher-tier compliance to their immediate tier-one suppliers. For example, Kroger, the largest supermarket chain in the U.S., requires its tier-one suppliers to comply with its ESR policies including anti child labor, anti discrimination, and anti harassment. In addition, suppliers must ensure that any of their contractors, distributors, and suppliers also comply with the same code of conduct.[1] In contrast, Walmart is known for *direct engagement* with its entire supply network. For example, in the program for improving labor conditions in Mexico produce supply chain, Walmart directly works with growers in Mexico. In the program for improving workplace safety in Bangladesh apparel supply chain, Walmart closely engages the local garment factories.

In Chapter 4, we apply a bilateral bargaining framework to analyze to what extent an ESR initiator should directly engage higher-tier suppliers, as opposed to delegating the assurance of ESR compliance to its first-tier suppliers. Our bargaining framework not only generalizes the conventional Shapley value approach by allowing the flexibility of modeling imbalanced power distribution among the firms, but also provides an explicit way of implementing the resulting gain sharing among the firms through negotiated contract terms. We show that the eventual structure of negotiation relationships can be derived by finding a shortest path tree in the supply network with the arc cost defined as a monotone function of the negotiating parties' relative bargaining power. These developments allow us to analyze ESR implementation in generally extended supply networks. We find that the ESR initiator tends to delegate ESR compliance negotiation to a supplier that is strong in negotiations with higher-tier suppliers. When the supply network is complex (i.e., wide and deep), directly engaging all suppliers for ESR compliance can lead to a larger gain by the initiator than fully delegating the negotiations with higher-tier suppliers to the first-tier ones.

1.4 Organization of the Dissertation

In Chapter 2, we formulate a dynamic model, in which the firm replenishes product inventories from uncertain sources and dynamically allocates available products to meet the uncertain demands with the flexibility of substitution. Through extensive numerical analysis, we demonstrate the performance of the approximation algorithm and investigate the benefit of the dynamic product substitution.

[1] https://www.thekrogerco.com/wp-content/uploads/2017/09/code-of-conduct.pdf

In Chapter 3, we study the problem of jointly optimizing the price and order quantity for a perishable product in the presence of demand ambiguity. We characterize the optimal pricing and ordering decisions under the minimax regret criterion and compare their properties with those in the classical models that seek to maximize the expected profit. We also compare the minimax regret approach with two other approaches that are commonly used under demand ambiguity, namely the max-min robust approach and the regression-based data-driven approach.

In Chapter 4, we apply a bilateral bargaining framework to analyze to what extent an ESR initiator should directly engage higher-tier suppliers, as opposed to delegating the assurance of ESR compliance to its first-tier suppliers. We show that how to derive the eventual structure of negotiation relationships. With such development, we investigate the effect of the supply networks on the ESR implementations.

Chapter 5 concludes the dissertation and provides suggestions for future research.

Chapter 2 is based on Feng et al. (2019a). Chapter 3 is based on Feng et al. (2019b). Chapter 4 is based on Feng et al. (2019c). I would like to express my sincere appreciation to my co-authors, Professors Qi Annabelle Feng, Mengshi Lu, and J. George Shanthikumar for their invaluable contributions.

2. DYNAMIC SUBSTITUTION FOR SELLING MULTIPLE PRODUCTS UNDER SUPPLY AND DEMAND UNCERTAINTIES

2.1 Synopsis

To derive a general substitution policy under both supply and demand uncertainties, we study the problem of replenishment for multiple products and inventory-demand allocation among those products. At the beginning of a replenishment cycle, the firm determines the supply orders issued for the products. The amount delivered for each product is a stochastic function of the order placed for that product. In each period within a replenishment cycle, the firm decides how to allocate the available inventories to meet the demand observed. The value of using one product to substitute for another product is less than the value of using that product to meet its own demand.

To understand the general substitution structure (that is not restricted to downward substitution), we first examine the single-period version of the problem. We show that downward substitution is a special case of a general substitution structure under which the benefits of substitution between products reveal the reverse Monge property (Monge, 1781). The product allocation decision in our model is a solution of the transportation problem, for which it is well-known that the problem can be solved efficiently when the transportation cost (benefit) reveals the (reverse) Monge property. For our product allocation problem, we construct a network of the corresponding transportation problem by inserting a fictitious demand node for each actual demand. The flow from a supply node to an actual demand node represents the product allocation decision, while the flow from a supply node to a fictitious demand node represents the leftover stock after the allocation. We show that the reverse Monge property is preserved with the addition of the fictitious demands, and thus the network problem can be solved efficiently with the algorithm developed by Vaidyanathan (2013).

When there are multiple periods of demand occurrence, the problem becomes significantly complex because the allocation of products depends not only on the substitution benefits, but also on the future values of the products. The future value of one product, in view of possible future substitutions, depends heavily on the inventory levels on all other products. Consequently, the time complexity of the dynamic substitution problem grows exponentially with the number of products due to the curse of dimensionality. Given that the optimal policy is difficult to compute, we propose an approximation that allows the firm, in addition to allocating on-hand inventories to the demands in the current period, to "convert" the leftover inventories among the products, while forgoing the opportunities of substitution in any future periods. By converting one product to another, the former product is then reserved to meet the future demand for the latter one. Though this approach excludes the possibility of future substitutions, the product conversion makes up for the loss by remixing the leftover inventories in anticipation of possible future demand scenarios. The immediate advantage of this approach is the separability of the future profit function in the remixed product inventories. With separable future profit functions, we can apply the development by Hochbaum and Shanthikumar (1990) to solve the product allocation problem through a series of linear programs, and the number of linear programs is of the polynomial order of the number of products. Furthermore, we show that the solution of each linear program can be derived by constructing a network in the way similar to that for the single-period setting. The main difference is that we now iteratively insert appropriate fictitious demands to represent the approximate marginal profit at different levels of leftover inventories. The constructed networks preserve the Monge property, with which the computational time can be significantly reduced.

With the substitution policy developed, we then examine the firm's stocking decisions at the beginning of each replenishment cycle. Applying the notion of stochastic linearity in midpoint (Feng and Shanthikumar, 2018), we show that the stocking problem can be converted to a concave maximization problem for general stochastic supply functions. Thus, the stocking decisions can be computed efficiently.

Because the optimal solution of the problem is difficult to compute, we evaluate the performance of our approximated solution against an upper bound problem. The upper bound of the firm's profit is computed by assuming that the firm can observe all future demands at the beginning of a replenishment cycle. In this case, the firm can simply aggregate the demands during the cycle by product, and replenish the aggregate quantities accordingly. Our analysis suggests that the approximated solution results in a profit gap from the upper bound problem that is increasing in both supply and demand uncertainties, while decreasing in the replenishment cycle. We further examine the value of substitution generated from the approximated solution against a lower bound problem, in which no substitution is ever allowed. We show that the characteristics of the supply function can have a major impact on the effectiveness of substitution. While an increased supply uncertainty generally increases the benefit generated from product substitution, the limits on supply capacities can increase or decrease the value of substitution. These observations highlight the importance of modeling the supply process in studying substitution policies.

The remainder of the chapter is organized as follows. We summarize the related literature and spell out our contribution in the next section. In Section 2.3, we lay out the dynamic model and establish the concavity of the profit functions. We treat the special case of the single-period model in Section 2.4, which becomes a building block of the approximation developed for the multi-period setting in Section 2.5. Section 2.6 presents the evaluation of the approximation algorithm and analyzes the value of substitution. Section 2.7 concludes our study. Proofs of all formal results are relegated to the appendix.

2.2 Literature Review

There is a vast literature studying product substitution or resource allocation in the context of inventory management. Our work develops a framework for dynamic product allocation under general substitution structure and uncertain supplies.

Pasternack and Drezner (1991) is one of the earliest papers that study the inventory control problem with product substitution under special substitution structures. They consider a single-period inventory model for two products that can be used as substitutes for each other, though at different revenue levels. They find that the total optimal stocking levels with substitution may be more or less than those without substitution. Bassok et al. (1999) extend the model of Pasternack and Drezner (1991) to a multi-product case. They consider multiple demand classes with downward substitution and propose a greedy policy that derives the optimal solution. Through numerical analysis, they show that the benefit of substitution is higher when demand variability is higher, substitution cost is lower, profit margins are lower, salvage values are higher, and product prices and costs are similar. Hsu and Bassok (1999) further extend this study by introducing random supply yields. Under a more general setting, Kalagnanam et al. (2000) analyze the problem of matching items in an order book with available surplus inventory, for which they present an efficient network-flow-based heuristic. Netessine et al. (2002) study the effect of demand correlation on the optimal capacity decisions with upgrading. When the demands follow the multivariate normal distribution, the change of demand correlation affects the optimal capacity for adjacent resources in opposite directions. In contrast to all these studies that focus on the single-period case, we consider dynamic product substitution under uncertain supply and demand processes.

Our work is also closely related to the stream of literature on revenue management with upgrading. Shumsky and Zhang (2009) consider a dynamic capacity management problem with downward substitution for products in adjacent classes, i.e., the unmet demand in a class can be fulfilled only by the products from the next higher class. The optimal policy calls for allocating the products for its class first, and then satisfy the lower-end unmet demand before the leftover stock reaches a certain threshold. Xu et al. (2011) consider a similar setting with two substitutable products and allow the customers to decide whether or not to accept the product substitution. They characterize an optimal substitution policy similar to the one in Shumsky and Zhang (2009). Yu et al. (2015) extend the setting in Shumsky and Zhang (2009) to allow general downward substitution that can go beyond adjacent products. The results derived from these papers rely heavily on the downward substitutable relationships among the products. Our development, in contrast, applies to general substitutions under which the products may not have a rank order of their values.

Another distinct feature of our model is the consideration of supply uncertainty in the context of dynamic product substitution, which is a key driver of the mismatch between supply and demand in reality. Feng et al. (2018) presents a thorough review of the literature in multi-sourcing under supply uncertainties, which focus on how to diversify the supplier base by trading off between suppliers' cost and reliability (see, e.g., Feng and Shi, 2012; Chen et al., 2013; Feng et al., 2018). Different from the multi-sourcing problem, our work considers replenishment decisions for multiple products under uncertain supplies with product substitution. Moreover, we do not restrict the form of supply uncertainty (e.g., proportional random yield).

Our framework can also be applied to the reactive transshipment problem. Specifically, inventories at different locations can be regarded as inventories for multiple products and the transportation costs between different locations correspond to the substitution costs in our setting. The majority of the literature focuses on analyzing the effect of transshipment on the stocking decisions with specific transshipment networks and designing efficient heuristics for the transshipment decisions (see, Krishnan and Rao, 1965; Karmarkar, 1981; Robinson, 1990; Archibald et al., 1997; Axsäter, 2003, etc.). All these papers consider the multi-period setting, and replenishment is allowed in each period. In our model, replenishment is only allowed at the beginning of each replenishment cycle but not allowed within the replenishment cycle, which increases the complexity of the problem. Yao et al. (2016) studies the two-location transshipment problem with fixed initial inventory levels, where the transshipment decisions are made without future demand information or additional replenishment opportunities. They show that with proper initial inventory levels, the relative benefit of dynamic transshipment diminishes as the length of the planning horizon increases. We obtain similar observations under general substitution structures through numerical analysis. Hu et al. (2008) and Chen et al. (2015) study the two-location transshipment problem with uncertain capacities. We consider a general class of supply functions, which includes the uncertain capacities in Hu et al. (2008) and Chen et al. (2015) as special cases.

Our work is also related to network revenue management (see, e.g., Talluri and Van Ryzin, 2006), where the firm needs to decide the allocation of limited resources to demands arriving over multiple periods. Due to computational intractability, most existing literature focus on designing heuristics for resource allocation (see, e.g., Cooper, 2002; Reiman and Wang, 2008; Jasin and Kumar, 2012). Whereas these studies investigate the resource allocation policy upon each demand arrival, we consider the allocation of products/resources to meet batch demands. Another stream of literature studies network revenue management with discrete choice models where customers choose from the assortment of products offered by the firm (see, e.g., Liu and Van Ryzin, 2008; Kunnumkal and Topaloglu, 2010; Zhang and Cooper, 2005). Different from this stream of study, we are concerned with the firm-driven scheme where the firm decides the quantities of products substituted for meeting demands.

In terms of methodology, to determine the substitution policy, the key challenge is to solve the allocation problem efficiently in each period, which reduces to the stochastic transportation problem (see Williams, 1963). With convex objective and linear constraints, most existing study apply nonlinear optimization techniques and show the convergence of the algorithm (see., e.g., Shetty, 1959; Holmberg, 1995). However, the existing approaches are restricted to continuous demands though the firms typically deal with discrete ones in reality. In addition, the time complexity of those proposed algorithms remains unknown. Different from these works, we propose a polynomial-time algorithm based on the framework developed in Hochbaum and Shanthikumar (1990) that requires solving a series of linear programs. By constructing appropriate networks and exploiting the Monge property of the cost matrix, we can solve each linear program efficiently with the algorithm developed in Vaidyanathan (2013). Unlike most convex optimization algorithms which require approximating discrete products with continuous inventory levels, our approach can handle discrete supply and demand, which are common in multi-product replenishment and substitution, and results in the exact optimal solution.

2.3 The Model

2.3.1 Problem Formulation

The firm under consideration manages multiple products to satisfy customers' demands over a planning horizon of T periods. The products are indexed by $j \in N = \{1, 2, \ldots, n\}$. The demand for product i in period t, denoted by a random variable $D_{t,i}$, is independent of other products' demand. Also, the demands in period t, denoted by $\mathbf{D}_t = (D_{t,1}, D_{t,2}, \ldots, D_{t,n})$, is independent of those in other periods. The unit selling price for product j is p_j. The firm incurs a per-unit holding cost of h_j for each unsold product and a per-unit good-will loss of g_j for each unmet demand in a period. The firm may substitute one product with another in meeting the demands. We use $k_{i,j}$ to denote the per-unit cost of using product i to meet the demand for product j, while allowing the possibility of $k_{i,j} = \infty$. We define the *per-unit net gain* of using product i to meet demand j as

$$r_{i,j} = p_i + h_i + g_j - k_{i,j}. \tag{2.1}$$

It is natural to assume that $r_{i,i} \geq r_{i,j}$ and $r_{i,i} \geq r_{j,i}$ for $j \neq i$ so that matching a unit with its own demand generates a higher value than substituting.

The firm faces a replenishment cycle of C periods, i.e., procurement orders can only be placed in periods $\mathcal{C} = \{1, 1+C, 1+2C, \ldots\}$. The replenishment cycle C is often highly product and firm specific. For example, it is typical in apparel retailing that stocks are replenished from overseas suppliers with a time window of three months, while it is often the case that fast-moving grocery items are replenished daily in large retail stores. We consider uncertain supply streams and assume zero delivery lead time. Specifically, when ordering $q_{t,j}$ units of product j, the firm receives a random amount of $S_{t,j}(q_{t,j})$. Let $\mathbf{q}_t = (q_{t,1}, q_{t,2}, \ldots, q_{t,n})$ and $\mathbf{S}_t(\mathbf{q}_t) = (S_{t,1}(q_{t,1}), S_{t,2}(q_{t,2}), \ldots, S_{t,n}(q_{t,n}))$ denote the vectors of order quantities and delivery quantities for period $t \in \mathcal{C}$, respectively. We assume that the supply process of one product is independent of those of other products and the supply process in one period is independent of those in other periods.

At the beginning of period t, $t = 1, 2, \ldots, T$, the firm reviews the on-hand inventory $\mathbf{x}_t = (x_{t,1}, x_{t,2}, \ldots, x_{t,n})$ and, if $t \in \mathcal{C}$, determines the order quantities $\mathbf{q}_t$. After observing the demands $\mathbf{d}_t = (d_{t,1}, d_{t,2}, \ldots, d_{t,n})$ and receiving the replenishment quantities $\mathbf{s}_t = (s_{t,1}, s_{t,2}, \ldots, s_{t,n})$ if $t \in \mathcal{C}$, the firm decides an allocation policy $\mathbf{y}_t = \{y_t^{i,j} : i, j \in N\}$, where $y_t^{i,j}$ is the amount of product i allocated to meet the demand for product j. Thus, the dynamics of the on-hand inventory levels are given by

$$\mathbf{x}_{t+1} = \begin{cases} \mathbf{x}_t - \mathbf{y}_t \mathbf{1} & \text{if } t \notin \mathcal{C}, \\ \mathbf{x}_t - \mathbf{y}_t \mathbf{1} + \mathbf{s}_t & \text{if } t \in \mathcal{C}. \end{cases}$$

The sequence of events is depicted in Figure 2.1. Let $V_t(\mathbf{x}_t)$ denote the firm's optimal expected profit at the beginning of period t when the on-hand inventories are $\mathbf{x}_t$. Then the dynamic programming equation can be written as

$$V_t(\mathbf{x}_t) = \begin{cases} \max_{\mathbf{q}_t \geq \mathbf{0}} \{\mathbb{E}[W_t(\mathbf{x}_t + \mathbf{S}(\mathbf{q}_t), \mathbf{D}_t) - \mathbf{g}^\top \mathbf{D}_t - \mathbf{h}^\top (\mathbf{x}_t + \mathbf{S}(\mathbf{q}_t)) - \mathbf{c}^\top \mathbf{S}(\mathbf{q}_t)]\} & \text{if } t \in \mathcal{C}, \\ \mathbb{E}[W_t(\mathbf{x}_t, \mathbf{D}_t) - \mathbf{g}^\top \mathbf{D}_t] - \mathbf{h}^\top \mathbf{x}_t & \text{if } t \notin \mathcal{C}, \end{cases} \tag{2.2}$$

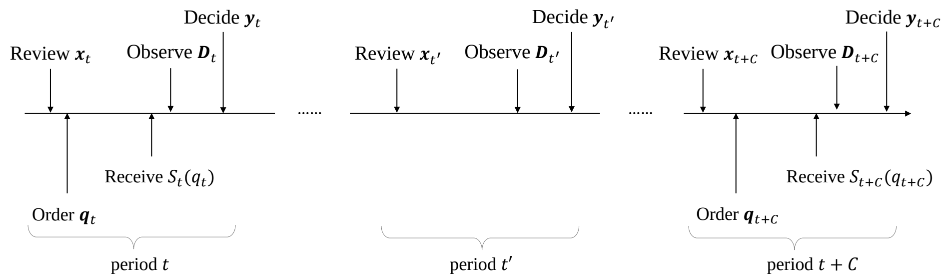

Fig. 2.1.: Sequence of events

where

$$W_t(\mathbf{z}_t, \mathbf{d}_t) = \max_{\mathbf{y}_t \in \mathbb{R}_+^{n \times n}} \quad \mathbf{R} \odot \mathbf{y}_t + V_{t+1}(\mathbf{z}_t - \mathbf{y}_t \mathbf{1}) \tag{2.3}$$

$$\text{s.t.} \quad \mathbf{z}_t - \mathbf{y}_t \mathbf{1} \geq \mathbf{0}, \tag{2.4}$$

$$\mathbf{d}_t - \mathbf{y}_t^\top \mathbf{1} \geq \mathbf{0}. \tag{2.5}$$

Here, $\mathbf{R} = (r_{i,j})_{i,j \in N}$ and $\mathbf{R} \odot \mathbf{y}_t = \sum_{i,j \in N} r_{i,j} y_t^{i,j}$. The function $W_t(\mathbf{z}_t, \mathbf{d}_t)$ computes the firm's optimal expected profit in period t when the available inventories are $\mathbf{z}_t$ and the realized demands are $\mathbf{d}_t$. The terminal condition is

$$V_{T+1}(\mathbf{x}_{T+1}) = \mathbf{v}^\top \mathbf{x}_{T+1}, \tag{2.6}$$

where $\mathbf{v} = (v_1, v_2, \ldots, v_n)$ is the vector of salvage values of unsold inventories.

2.3.2 Preliminaries: The Concavity of the Profit Function

In this subsection, we investigate the structural property of the multi-period replenishment and allocation problem defined in (2.2)–(2.5). As discussed in §2.2, most existing literature rely on specific substitution structures (see, e.g., Shumsky and Zhang, 2009; Chen et al., 2015; Yu et al., 2015) or random supply functions (see, e.g., Hsu and Bassok, 1999). We do not impose these restrictions. Feng and Shanthikumar (2018) develop the theory of *stochastic linearity in midpoint*, which enables the analysis of general demand and supply functions in operations management problems.

Equipped with the notion of stochastic linearity in midpoint, we can transform the original optimization problem in (2.2)–(2.5) into an equivalent concave problem for a general class of stochastic supply functions. This notion is defined based on the concave ordering. Specifically, a random variable Y_1 is said to be smaller than another random variable Y_2 in the concave order, written as $Y_1 \leq_{cv} Y_2$, if $\mathbb{E}[\phi(Y_1)] \leq \mathbb{E}[\phi(Y_2)]$ for all concave functions (see., Shaked and Shanthikumar, 2007).

Definition 2.3.1 (Feng and Shanthikumar, 2018) *A function* $\{Y(x), x \in \mathcal{X}\}$ *for some convex* $\mathcal{X}$ *is stochastically linear in midpoint, written as* $\{Y(x), x \in \mathcal{X}\} \in SL(mp)$*, if for any* $x_1, x_2 \in \mathcal{X}$*, there exist* $\hat{Y}(x_1)$ *and* $\hat{Y}(x_2)$ *defined on a common probability space such that*

(i) $\hat{Y}(x_i) =^d Y(x_i),\ i = 1, 2,$ *and*

(ii) $\frac{\hat{Y}(x_1)+\hat{Y}(x_2)}{2} \leq_{cv} Y(\frac{x_1+x_2}{2})$.

Let $\mu_{t,i}(q) = \mathbb{E}[S_{t,i}(q)]$, which we assume to be continuous and non-decreasing. Let $\bar{\mu}_{t,i} = \lim_{q\to+\infty} \mathbb{E}[S_{t,i}(q)]$, defined on the extended real numbers, and $\bar{\boldsymbol{\mu}}_t = (\bar{\mu}_{t,1}, \bar{\mu}_{t,2}, \ldots, \bar{\mu}_{t,n})$. Let $q_{t,i}(\mu) = \inf\{q : \mu_{t,i}(q) \geq \mu\}$, which is the inverse of $\mu_{t,i}(q)$, and $\boldsymbol{\mu}_t = (\mu_{t,1}, \mu_{t,2}, \ldots, \mu_{t,n})$. Define $\hat{S}_{t,i}(\mu) =^d S_{t,i}(q_{t,i}(\mu))$ and $\hat{\mathbf{S}}_t(\boldsymbol{\mu}_t) = (\hat{S}_{t,1}(\mu_{t,1}), \hat{S}_{t,2}(\mu_{t,2}), \ldots, \hat{S}_{t,n}(\mu_{t,n}))$. The value function $V_t(\mathbf{x}_t)$ in (2.2) for $t \in \mathcal{C}$ can be written as,

$$V_t(\mathbf{x}_t) = \max_{\mathbf{0}\leq\boldsymbol{\mu}_t\leq\bar{\boldsymbol{\mu}}} \left\{\hat{V}_t(\mathbf{x}_t, \boldsymbol{\mu}_t) := \mathbb{E}[W_t(\mathbf{x}_t + \hat{\mathbf{S}}_t(\boldsymbol{\mu}_t), \mathbf{D}_t)] - \mathbf{g}^\top \mathbb{E}[\mathbf{D}_t] - \mathbf{h}^\top(\mathbf{x}_t + \boldsymbol{\mu}_t) - \mathbf{c}^\top \boldsymbol{\mu}_t\right\}. \quad (2.7)$$

The following theorem establishes the concavity of the problem in (2.2)–(2.5) when the transformed supply functions $\hat{S}_{t,i},\ i = 1, 2, \ldots, n$, are stochastically linear in midpoint.

Theorem 2.3.1 *If* $\{\hat{S}_{t,i}(\mu), 0 \leq \mu \leq \bar{\mu}_{t,i}\} \in SL(mp)$ *for* $i = 1, 2, \ldots, n$*, then* $\hat{V}_t(\mathbf{x}_t, \boldsymbol{\mu}_t)$ *is concave in* $(\mathbf{x}_t, \boldsymbol{\mu}_t)$ *for* $t \in \mathcal{C}$*;* $V_t(\mathbf{x}_t)$ *is concave in* $\mathbf{x}_t$ *and* $W_t(\mathbf{z}_t, \mathbf{d}_t)$ *is concave in* $(\mathbf{z}_t, \mathbf{d}_t)$*, for* $t = 1, 2, \ldots, T$*.*

As pointed out by Feng and Shanthikumar (2018), there is a rich class of random supply functions that are stochastically linear in midpoint, including the widely used proportional random yield model and the random capacity model. Theorem 2.3.1 enables us to analyze inventory management with product substitution under general uncertain supply processes.

2.4 The Single-Period Problem

In this section, we focus on the single-period version of the model formulated in (2.2)–(2.5). Specifically, we examine the following two-stage stochastic program:

$$\max_{\mathbf{q}\in\mathbb{R}^n_+} \left\{\mathbb{E}[W(\mathbf{S}(\mathbf{q}),\mathbf{D}) - \mathbf{c}^\top \mathbf{S}(\mathbf{q})]\right\} - \mathbf{g}^\top \mathbb{E}[\mathbf{D}], \tag{2.8}$$

where

$$W(\mathbf{x},\mathbf{d}) = \left\{\max_{\mathbf{y}\in\mathbb{R}^{n\times n}_+} \quad \mathbf{R}\odot\mathbf{y} : \mathbf{x} - \mathbf{y}\mathbf{1} \geq \mathbf{0}, \mathbf{d} - \mathbf{y}^\top\mathbf{1} \geq \mathbf{0}\right\}. \tag{2.9}$$

Several versions of this problem have been studied in the literature. Pasternack and Drezner (1991) formulate a two-product model in which demand of one can be substituted by the other in the event of stockout of the former. Bassok et al. (1999) extends the analysis to multi-product substitution, while imposing the restriction of downward substitution (i.e., only a high-value product can be used to meet the shortage of a low-value product). Netessine et al. (2002) study a similar model with a focus of understanding the effect of demand correlation on the value of substitution. Hsu and Bassok (1999) extends the downward substitution model to consider the case where the supply of the products comes from a production process with random yields proportional to the input. In our formulation, the restriction of downward substitution translates into the requirement that $r_{i,j} \leq 0$ whenever $p_i > p_j$. We do not impose such a restriction in our model. Moreover, we allow the supply processes of the products follow general stochastic input-output relationships.

2.4.1 The Monge Property

The allocation problem defined in (2.9) is known as a transportation problem. The fastest min-cost flow algorithm designed by Orlin (1993) solves this problem with a run-time of $O(n \log n(m + n \log n))$, where n is the total number of nodes (i.e., supplies and demands) and m is the number of arcs (i.e., feasible substitution relations). Monge (1781) first observes that an optimal solution of this problem can be computed faster if the total supply equals the total demand (i.e., the network is balanced) and the matrix $\mathbf{R}$ satisfies a certain property, which Hoffman (1963) later names after Monge.

Definition 2.4.1 A $n_1 \times n_2$ matrix $\mathbf{C} = \{c_{i,j} : i = 1, 2, ..., n_1, j = 1, 2, ..., n_2\}$ is a *Monge matrix* (or Monge array) if $c_{i_1,j_1} + c_{i_2,j_2} \leq c_{i_1,j_2} + c_{i_2,j_1}$ for $1 \leq i_1 < i_2 \leq n_1$ and $1 \leq j_1 < j_2 \leq n_2$.

To test the Monge property in a matrix, it is sufficient to check the adjacent rows and columns, which results in a complexity of $O(n_1 n_2)$. In fact, Hoffman (1963) described a slightly more general property by introducing the Monge sequence.

Definition 2.4.2 (Hoffman, 1963) Given a $n_1 \times n_2$ matrix $\mathbf{C} = \{c_{i,j} : i = 1, 2, ..., n_1, j = 1, 2, ..., n_2\}$, let $\pi_M = (z_1, ...z_{n_1 n_2})$ be a permutation of $S_z = \{(i, j) : i = 1, ..., m; j = 1, ..., n\}$ and let $n(i, j)$ be the position of (i, j) in the permutation π_M such that $z_{n(i,j)} = (i, j)$. We say π_M is a *Monge sequence* if $\min\{(n(i_1, j_1), n(i_2, j_2)\} < \min\{n(i_1, j_2), n(i_2, j_1)\}$ implies $c_{i_1,j_1} + c_{i_2,j_2} \leq c_{i_1,j_2} + c_{i_2,j_1}$ for all possible i_1, i_2, j_1, j_2.

Clearly, a Monge sequence exists in a Monge matrix. For a general matrix C, the complexity of judging the existence and constructing a Monge sequence is $O(n_1^2 n_2 \log n_2)$; see Alon et al. (1989). It is well known that for a balanced transportation problem with a Monge cost matrix, a greedy allocation along the Monge sequence is optimal. Burkard (2007) describes several examples of Monge matrices that often appear in practice:

- $c_{i,j} = a_i + b_j$ for any real $a_i, 1 \leq i \leq n_1$ and $b_j, 1 \leq j \leq n_2$.
- $c_{i,j} = a_i b_j$ for any increasing $a_i, 1 \leq i \leq n_1$ and decreasing $b_j, 1 \leq i \leq n_2$.
- $c_{i,j} = \min\{a_i, b_j\}$ for any increasing $a_i, 1 \leq i \leq n_1$ and decreasing $b_j, 1 \leq i \leq n_2$.
- $c_{i,j} = d(P_i, Q_j)$ where $P_1, P_2, \ldots, P_{n_1}$ and $Q_1, Q_2, \ldots, Q_{n_2}$ are points on disjoint paths P and Q on a convex polygon, and $d(\cdot, \cdot)$ is the Euclidean distance.

Because our problem involves profit maximization instead of cost minimization, we focus on the *reverse* Monge property (i.e., $r_{i_1,j_1} + r_{i_2,j_2} \geq r_{i_1,j_2} + r_{i_2,j_1}$ for $1 \leq i_1 \leq i_2 \leq n$ and $1 \leq j_1 \leq j_2 \leq n$) as benefits can be treated as negative costs. In the first example described above, the additive cost (reward) often arises in product substitution and two-sided market matching. Specifically, in the product substitution context, the substitution reward is determined by the benefit of meeting demand for type j and the cost of using product of type i. For example, the practice of free upgrading in airline industry or auto rental industry naturally results in the linear additive form (Shumsky and Zhang, 2009; Yu et al., 2015). In two-sided market matching, for example, Hu and Zhou (2016) describe that the matching reward in the carpooling platforms generally has two additive components: the first one is the dis-utility associated with the distance traveled by the driver to pick up the passenger, and the other is the reward related to the distance between the pickup and drop-off locations of the passenger. Another example is the centralized medical residency assignment that matches the medical residents to residency programs in order to achieve a stable matching (Agarwal, 2015). The utility of a pair of a medical resident and a residency program depends on the preferences of two sides on each other. Both the additive form and the multiplicative form are commonly used to define the reward of a pair, which are special cases of the Monge property (i.e., supermodular reward functions). The last example described above is the discrete version of Monge's original observation (i.e., transportation problems). Specifically, in the transportation problem that aims to minimize the total cost of shipping goods from one set of locations to another, the transportation costs are related to the distances between those locations. Therefore, when applying our framework to the transshipment problem, the substitution costs can be regarded as the transportation costs between different locations, which would result in Monge matrices.

Hoffman (1963) shows that when a (reverse) Monge sequence exists in the cost (benefit) matrix, it is optimal to follow a *greedy allocation rule*. Specifically, in the context of our problem, the allocation is prioritized according to the sequence $(z_1, z_2, \ldots, z_{n_1 n_2})$. At the kth step with $z_k = (i_k, j_k)$, one allocates the available supply (i.e., excluding the amount assigned in previous steps) of i_k to meet the demand of j_k to the extent possible. Such a greedy algorithm significantly reduces the computational complexity.

The result developed by Hoffman (1963), however, is not directly applicable to solve the allocation problem in (2.9) when the total supply does not necessarily match the total demand (i.e., $\sum_{i=1}^{n_1} x_{t,i} \neq \sum_{i=1}^{n_2} d_{t,i}$), which is likely the situation in practice. We focus on the case when the total supply exceeds the total demand. This is a more relevant case when we try to extend the results for the multi-period setting in §2.5. The case with the total demand exceeding the total supply can be treated similarly.

2.4.2 The Network Construction

For a transportation problem with over supply, one would naturally add a fictitious demand that may consume a quantity of $\sum_{i=1}^{n} d_i - \sum_{i=1}^{n} x_i \geq 0$. Any supply to this node results in zero benefit. While such an approach makes the network balanced, the added zero benefit vector destroys the Monge property in the extended matrix. As a result, there is no guarantee that the greedy allocation along the original Monge sequence is optimal.

Many studies attempt to address the unbalanced transportation problem with a Monge cost matrix. The latest algorithm is developed by Vaidyanathan (2013). He constructs a network with a source node connecting all the supply nodes and the supply nodes connecting to the demand nodes according to the cost matrix R. His algorithm iteratively augments flows from the source to each demand node along the shortest path. The complexity of this algorithm is $O(m\log n)$, where m is the number of arcs and n is the number of demand nodes. Although the developments for the transportation problem are usually for cost minimization, their applications to profit maximization problems can be easily done by reversing the signs (i.e., considering the negative of the cost matrix as the benefit matrix).

The algorithm developed by Vaidyanathan (2013) certainly solves our single-period allocation problem defined in (2.9), while not offering too much insights into the specific substitution problem. Moreover, it does not help when we extend the problem to multiple periods. Therefore, we propose an alternative network construction. Specifically, define an $n \times 2n$ matrix $\hat{\mathbf{R}}$ with

$$\hat{r}_{i,j} = \begin{cases} r_{i,j} & \text{for } j = 2k-1, \\ r_{i,j} - \Delta_j & \text{for } j = 2k, \end{cases} \quad \text{for } k = 1, 2, \ldots, n,$$

where $\Delta_j \in (0, r_{j,j}]$. In other words, the $(2k-1)$st column of matrix $\hat{R}$ is simply the kth column of matrix R, and the $(2k)$th column of matrix $\hat{R}$ is the kth column of matrix R minus a constant $\Delta_j \in (0, r_{j,j}]$.

Lemma 2.4.1 *If the original benefit matrix* $\mathbf{R}$ *has the reverse Monge property (i.e.,* $r_{i_1,j_1} + r_{i_2,j_2} \geq r_{i_1,j_2} + r_{i_2,j_1}$ *for* $1 \leq i_1 < i_2 \leq n$ *and* $1 \leq j_1 < j_2 \leq n$*), then so does the extended matrix* $\hat{\mathbf{R}}$.

We construct the network G based on matrix $\hat{\mathbf{R}}$ as follows (see also the graphical illustration in Figure 2.2). The source nodes, which are also the supply nodes, are indexed by $i = 1, 2, \ldots, n$. Node $2n+1$ is the sink. The demand nodes are indexed by $i, i = n+1, n+2, \ldots, 2n$, and the capacity on arcs $(n+i, 2n+1)$ corresponds to the demand for product i. Moreover, we replicate each demand i as i', which corresponds to the virtual demand node that can consume left-over inventories. The excess flows at each node are defined as

$$e(i) = \begin{cases} -\sum_{i=1}^{n} x_i & \text{if } i = 2n+1, \\ x_i & \text{if } i = 1, 2, \ldots, n, \\ 0 & \text{otherwise.} \end{cases}$$

The arc benefits and arc capacities are, respectively,

$$\lambda_{i,j} = \begin{cases} r_{i,j-n} & \text{if } i = 1, 2, \ldots, n, j = n+1, n+2, \ldots, 2n, \\ r_{i,j-n} - \Delta_{j-n} & \text{if } i = 1, 2, \ldots, n, j = (n+1)', (n+2)', \ldots, (2n)', \\ 0 & \text{if } j = 2n+1. \end{cases}$$

$$u_{i,j} = \begin{cases} d_{i-n} & \text{if } i = n+1, n+2, \ldots, 2n, j = 2n+1, \\ \infty & \text{otherwise,} \end{cases}$$

where, with a slight abuse of notation, we take $(n+j)' - n = j$ for $j = 1, 2, \ldots, n$. In other words, the matrix $\hat{\mathbf{R}}$ defines the benefits for arcs that connecting the (virtual) demand and the supply. The benefit of all other arcs are set to be zero. For the arc connecting a demand node $n+j$ to the sink node $2n+1$, which is called as *demand arc*, its capacity is the demand amount d_j. All other arcs have infinite capacities. The arcs that connect supply nodes and the demand nodes (including the virtual demand nodes) are *transportation arcs*. The next result suggests that a solution to the transportation problem defined in graph G also solves our substitution problem.

Theorem 2.4.1 *If the original benefit matrix* $\mathbf{R}$ *has the reverse Monge property, then*

i) the max-benefit flow solution for G *corresponds to an optimal allocation for problem (2.9) when* $\Delta_i = r_{i,i}$ *for* $i = 1, 2, \ldots, n$*;*

ii) the algorithm developed by Vaidyanathan (2013) finds the optimal solution in network G *with a run time of* $O(n^2 \text{logn})$.

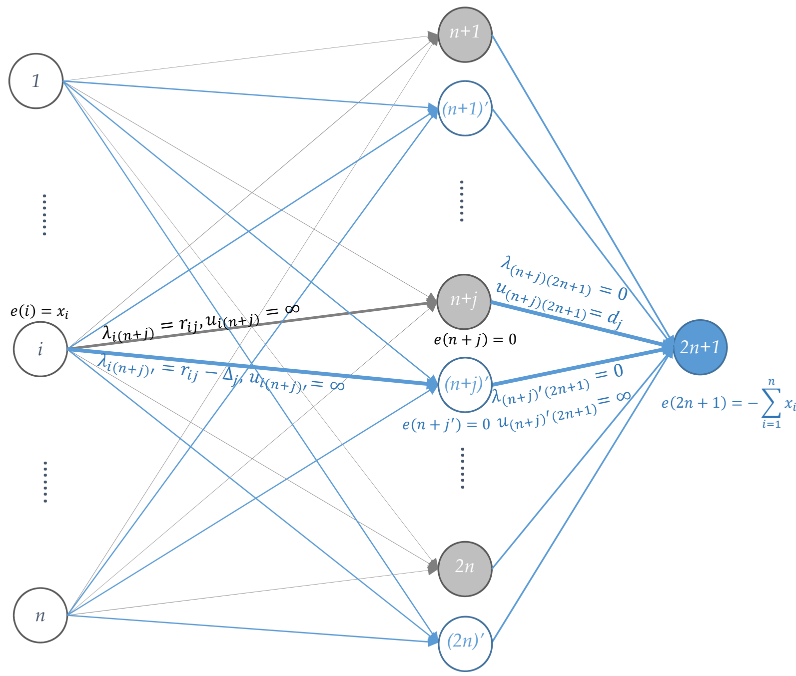

Fig. 2.2.: Network G

2.4.3 Interpretation Using the Special Case of Downward Substitution

To understand the allocation policy established in Theorem 2.4.1, we consider examples of downward substitution, the most studied situation in the literature. The products are ordered based on their values with a higher value product given a smaller index. Downward substitution happens when the supply of a higher value product is used to meet the demand of a lower value product.

In the example in Figure 2.3-(a), substitution is only possible between adjacent products (i.e., $r_{i,j} \geq 0$ if $j-1 \leq i \leq j$ and $r_{i,j} = -\epsilon$ for any $\epsilon > 0$ otherwise). An allocation sequence is marked using arrows between nonnegative entries of $r_{i,j}$ in the figure. According to this sequence, we use the supply for product 1 to meet the demand for product 1, the supply for product 2 to the demand for product 2, the remaining supply for product 1 to the demand for product 2,..., and finally the remaining supply for product 3 to the remaining demand for product 4. This sequence is a Monge sequence under three conditions. Specifically, using the product supply to meet its own demand is more profitable than either downgrading the supply to meet a lower valued demand (i.e., $r_{i,i} \geq r_{i,i+1}$) or meeting the demand by substituting with a higher value supply (i.e., $r_{i,i} \geq r_{i-1,i}$). Moreover, it is not profitable to use a product's supply to meet a lower value demand, while meeting that product's demand using a higher value supply (i.e., $r_{i,i} \geq r_{i,i+1} + r_{i-1,i}$). With these conditions, the optimal allocation along the Monge sequence suggests one should first use the supply to meet its own demand. If a product has excess supply and the adjacent lower value product has unmet demand, then allocate the former to the latter to the extent possible.

A special case of the example in Figure 2.3-(a) is analyzed by Shumsky and Zhang (2009). They assume that the substitution cost is independent of the demand being substituted, i.e., $k_{i,j} = \bar{k}_i$ for all $j = 1, 2, \ldots, n$. For $j - 1 \leq i \leq j$, $p_i + g_i \geq p_j + g_j$ and $\bar{k}_i \geq \bar{k}_j$. It is easy to verify that these conditions imply the conditions needed for the Monge sequence identified in Figure 2.3-(a).

In Figure 2.3-(b), general downward substitutions are allowed (i.e., $r_{i,j} \geq 0$ if $i \leq j$, and $r_{i,j} = -\epsilon$ for any $\epsilon > 0$ otherwise). For the sequence identified in the figure to be a Monge sequence, one needs to make sure that substitution between products with a closer value is more profitable (i.e., $r_{i,j} \geq r_{i,j+1}$ and $r_{i,j} \geq r_{i-1,j}$ for $i \leq j$). Moreover, it is more profitable to assign the high (low) value supply to the low (high) value demand than to match the supplies and demands with the similar value ranks (i.e., $r_{i,j} + r_{i-1,j+1} \geq r_{i-1,j} + r_{i,j+1}$ for $i \leq j$). According to this Monge sequence, the greedy allocation calls for prioritizing the demands according to their values. The supply of a product is allocated to its own demand. In the event of shortage, the available supply with a closer value is used to substitute before one with a more different value is.

Special cases of the setting presented in Figure 2.3-(b) have been studied by several authors. Bassok et al. (1999) and Netessine et al. (2002) assume $p_i + g_i \geq p_j + g_j$,

$\bar{k}_i \geq \bar{k}_j$ and $r_{i,j} \geq 0$ for $i \leq j$.[1] Yu et al. (2015) impose conditions $p_i > p_j$, $g_i \geq g_j$ and $\bar{k}_i \geq \bar{k}_j$ for $i \leq j$. It is easy to see that assumptions imposed in these studies imply the conditions in Figure 2.3-(b).

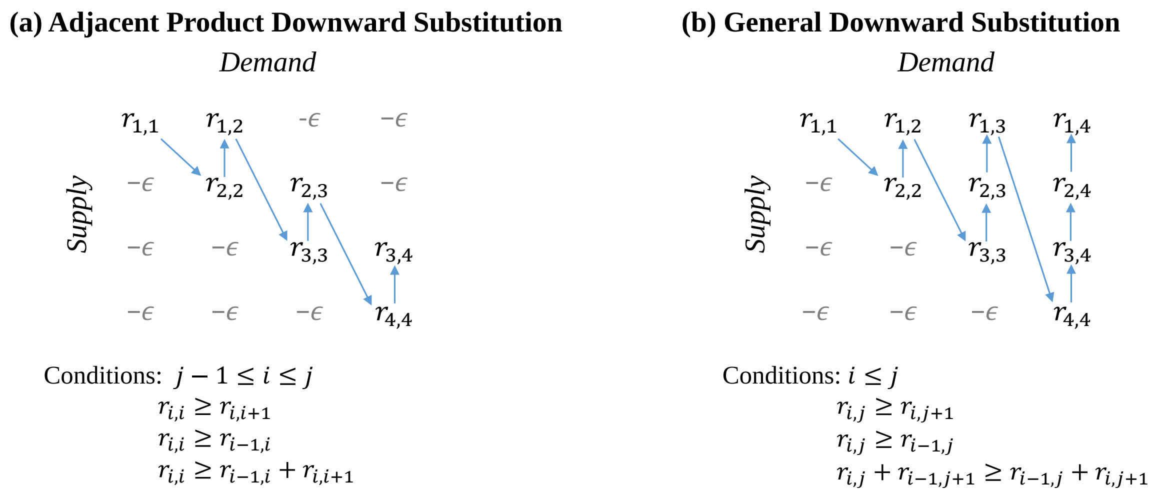

Fig. 2.3.: Benefit matrix for downward substitution

2.4.4 The Stocking Decisions

Recall that Theorem 2.3.1 establishes the concavity of the multi-period replenishment and product substitution problem. Applying the concavity result to the single-period problem, we can show that the expected profit is concave in the mean supplies. In specific, using the notations defined in §3.3 and suppressing the subscripts for time periods, problem (2.8) can be transformed to the following problem,

$$\max_{\mathbf{0}\leq\boldsymbol{\mu}\leq\bar{\boldsymbol{\mu}}} \left\{ \hat{V}(\boldsymbol{\mu}) := \mathbb{E}[W(\hat{\mathbf{S}}(\boldsymbol{\mu}), \mathbf{D})] - \mathbf{g}^\top \mathbb{E}[\mathbf{D}] - \mathbf{c}^\top \boldsymbol{\mu} \right\}. \tag{2.10}$$

By Theorem 2.3.1, $\hat{V}(\boldsymbol{\mu})$ is concave in $\boldsymbol{\mu}$ for a general class of supply functions that are stochastically linear in midpoint. In addition, the development in §§2.4.2 enables

[1] Bassok et al. (1999) also consider a salvage value v_j of unsold items and assume $v_i \geq v_j$ for $i < j$. Note that one can transform their model into ours by redefine unit price as $p_j - v_j$ and procurement cost as $c_j - v_j$.

us to evaluate $W(\mathbf{x}, \mathbf{d})$ efficiently for given available inventory $\mathbf{x}$ and demand realization $\mathbf{d}$. Therefore, we can apply standard approaches for concave maximization with the transformed objective function and find the optimal mean supply vector $\boldsymbol{\mu}$. The optimal vector of stocking decisions $\mathbf{q}$ can be recovered using its one-to-one correspondence to $\boldsymbol{\mu}$.

2.5 The Multi-Period Problem

For the multi-period problem, the optimal product allocation policy based on the Monge sequence no longer applies because of the nonlinear value-to-go function. Furthermore, even though Theorem 2.3.1 guarantees the concavity of the problem for general uncertain supply processes, the time complexity of the dynamic programming increases exponentially in the number of products, known as the curse of dimensionality, making the problem computationally intractable for practical applications. In view of these, the focus of this section is to develop close-to-optimal solutions that can be computed efficiently. The algorithm developed below has three essential elements, namely, the additively separable value function approximation, the efficient concave separable optimization, and the efficient optimal allocation based on the Monge property. In §2.6, we demonstrate the performance of this algorithm as well as discussing the insights generated.

2.5.1 An Approximation of the Value Function

To solve the allocation problem defined in (2.3)–(2.5) efficiently, the crux of the computational challenge lies in evaluating the value-to-go function $V_{t+1}(\mathbf{x}_{t+1})$, which is the optimal expected profit from period $t+1$ onward given available inventory $\mathbf{x}_{t+1}$. Instead of evaluating it exactly, we approximate $V_{t+1}(\mathbf{x}_{t+1})$ with the expected profit under the following simple policy.

Suppose that we are in period t, and from period $t+1$ onwards, we do not allow any product substitution. Then, it is easy to see that the future expected profit becomes separable in $\mathbf{x}_{t+1}$. Specifically, let $\tilde{V}_{t+1,i}(x_i), x_i \geq 0$ denote the expected profit from product i when the beginning inventory for that product in period $t+1$ is x_i. Then, we can compute $\tilde{V}_{t+1,i}(\cdot)$ recursively using the following expressions:

$$\tilde{V}_{k,i}(x) = \begin{cases} \max_{q \geq 0} \mathbb{E}\Big[W_i(x + S_i(q), D_{k,i}) - c_i S_i(q) + \tilde{V}_{k+1,i}\big((x + S_i(q) - D_{k,i})^+\big)\Big], \\ \qquad\qquad\qquad\qquad\qquad\qquad\qquad \text{if } k \in \{t+1, \ldots, T\} \cap \mathcal{C}, \\ \mathbb{E}\Big[W_i(x, D_{k,i}) + \tilde{V}_{k+1,i}\big((x - D_{k,i})^+\big)\Big], \ \text{if } k \in \{t+1, \ldots, T\} \setminus \mathcal{C}, \end{cases} \tag{2.11}$$

$$W_i(z, d) = r_{i,i} \min\{z, d\} - g_i d - h_i z, \tag{2.12}$$

$$\tilde{V}_{T+1,i}(x) = v_i x. \tag{2.13}$$

Applying Theorem 2.3.1 to the special case with a single product, we have that if $\{\hat{S}_i(\mu), 0 \leq \mu \leq \bar{\mu}_i\}$ is stochastically linear in midpoint, then $\tilde{V}_{t+1,i}(x)$ is concave in x for all $i \in N$ and $t = 1, \ldots, T$. Therefore, the single-dimensional function $\tilde{V}_{t+1,i}(x)$ can be evaluated efficiently.

The restriction of no future substitution certainly reduces the computational burden. However, it also induces the loss of value from substitution. To make up for that, we allow the firm to "convert" certain amount of product i's inventory into additional product j's inventory at the end of period t. Through such conversions, the firm can re-balance its stocks based on its anticipation of future demands. Recall the omnichannel example in §4.1. Product "conversion" in this case can be viewed as proactive transshipment between the stores. The per-unit conversion cost from production i to product j is

$$m_{i,j} := (p_j - p_i) + (h_j - h_i) + k_{i,j} = r_{j,j} - r_{i,j},$$

which consists of the potential loss of revenue, $(p_j - p_i)$, the potential increase in inventory holding cost, $(h_j - h_i)$, and the substitution cost $k_{i,j}$.

Let $\tilde{y}_t^{i,j}$ denote the amount of product i's inventory "converted" to product j's inventory at the end of period t, and $\tilde{x}_{t+1,i}$ the amount of product i's inventory at the beginning of period $t+1$ after the "conversion." Let $\mathbf{M} = (m_{i,j})_{i,j \in N}$, $\tilde{\mathbf{y}}_t = (\tilde{y}_t^{i,j})_{i,j \in N}$, and $\tilde{\mathbf{x}}_{t+1} = (\tilde{x}_{t+1,i})_{i \in N}$. Incorporating the product conversion decisions in the original allocation problem (2.3)–(2.5) and replacing the value-to-go function $V_{t+1}(\mathbf{x}_{t+1})$ with the sum of the approximation functions $\tilde{V}_{t+1,i}(x_{t+1,i})$ defined in (2.11)–(2.13), we have the following optimization problem,

$$\begin{aligned}
\tilde{W}(\mathbf{z}_t, \mathbf{d}_t) = \max \quad & \mathbf{R} \odot \mathbf{y}_t - \mathbf{M} \odot \tilde{\mathbf{y}}_t + \sum_{i \in N} \tilde{V}_{t+1,i}(\tilde{x}_{t+1,i}) && (2.14) \\
\text{s.t.} \quad & \mathbf{y}_t \mathbf{1} + \tilde{\mathbf{y}}_t \mathbf{1} = \mathbf{z}_t && (2.15) \\
& \mathbf{y}_t^\top \mathbf{1} \leq \mathbf{d}_t && (2.16) \\
& \tilde{\mathbf{y}}_t^\top \mathbf{1} = \tilde{\mathbf{x}}_{t+1} && (2.17) \\
& \mathbf{y}_t, \tilde{\mathbf{y}}_t \in \mathbb{R}_+^{n \times n}, \tilde{\mathbf{x}}_{t+1} \in \mathbb{R}_+^n. && (2.18)
\end{aligned}$$

Solving the above problem optimally, we obtain an allocation decision $\mathbf{y}_t$ given on-hand inventories $\mathbf{z}_t$ and realized demands $\mathbf{d}_t$. Note that the "conversion" decision $\tilde{\mathbf{y}}_t$ is only used to approximate the value function but not actually implemented. Whereas other approximation methods may also apply, the additive separability of the above approximation enables us to further reduce the allocation problem in (2.14)–(2.18) to a series of linear programs (Hochbaum and Shanthikumar, 1990), which we present in the next subsection.

2.5.2 Efficient Optimization with Separable Concave Value Functions

The problem defined in (2.14)–(2.18) is a variant of the *stochastic transportation problem*. Holmberg (1995) reviews efficient solution methods for this problem and develops an approach combining separable programming with piecewise linear approximation and mean value cross decomposition, which is so far the most efficient method for this problem. However, the separable programming approach in Holmberg (1995) does not provide any guarantee of the accuracy of the approximation. In this subsection, we present an efficient algorithm based on the framework developed by Hochbaum and Shanthikumar (1990). Our approach not only guarantees that the approximation achieves an arbitrary level of accuracy, but also limits the size of the linear program solved in each iteration.

A direct application of Hochbaum and Shanthikumar (1990)'s method to our problem suggests solving a total of $\lceil \log_2 (\frac{2n}{\epsilon} \|(\mathbf{z}_t, \mathbf{d}_t)\|_\infty) \rceil$ linear programs to achieve ϵ-optimality (i.e., the difference between the approximate solution and the true optimal solution has a supremum norm less than or equal to ϵ) and the number of variables of each linear program is at most $16n^3 + 8n^2$. These complexities are established for the case when the supplies and demands take continuous values, a common assumption in the product substitution literature. In reality, however, the supplies and demands involved in multi-product substitution are often discrete (recall the examples discussed in §4.1). The continuous approximation often provides easier analysis on the marginal benefits or allows for deriving the structural properties of the policy. For our purpose of computation, however, we can in fact achieve a reduced computational complexity and find the exact optimal solution with discrete supplies and demands.

We note that $\tilde{V}_{t+1,i}(x_i)$ is a piecewise linear function when the demand is discrete. We can further approximate this function by a set of piecewise linear functions defined base on a scaling factor s:

$$V^s_{t+1,i}(x_i) = \left(\left\lfloor \tfrac{x_i}{s} + 1 \right\rfloor - \tfrac{x_i}{s} \right) \tilde{V}_{t+1,i}\left(s \left\lfloor \tfrac{x_i}{s} \right\rfloor \right) + \left(\tfrac{x_i}{s} - \left\lfloor \tfrac{x_i}{s} \right\rfloor \right) \tilde{V}_{t+1,i}\left(s \left\lfloor \tfrac{x_i}{s} + 1 \right\rfloor \right), \; i \in N. \quad (2.19)$$

It is clear that s is the length of the segments of the piecewise linear function and $V^1_{t+1,i}(\cdot) = \tilde{V}_{t+1,i}(\cdot)$.

The main idea is to solve the problem with the approximated function $V^s_{t+1,i}(x_i)$ iteratively by gradually reducing the scaling factor s. Specifically, let the scaling factor in the l-th iteration be $s_l = 2^{l_0 - l}, l = 1, 2, \ldots, l_0$, where l_0 is determined based on problem parameters. Then, we solve the following optimization problem for the allocation decisions:

$$
\begin{aligned}
(\mathbf{IP}{-}s_l): \qquad \tilde{W}_t^{s_l}(\mathbf{z}_t, \mathbf{d}_t) = \max \quad & \mathbf{R} \odot (s_l \mathbf{y}_t) - \mathbf{M} \odot (s_l \tilde{\mathbf{y}}_t) + \sum_{i \in N} V_{t+1,i}^{s_l}\big(s_l \tilde{x}_{t+1,i}\big) \\
s.t. \quad & \mathbf{y}_t \mathbf{1} + \tilde{\mathbf{y}}_t \mathbf{1} = \lfloor \mathbf{z}_t / s_l \rfloor \\
& (\mathbf{y}_t)^\top \mathbf{1} \le \lfloor \mathbf{d}_t / s_l \rfloor \\
& (\tilde{\mathbf{y}}_t)^\top \mathbf{1} = \tilde{\mathbf{x}}_{t+1} \\
& \underline{\mathbf{x}}_{t+1}^{(l)} \le \tilde{\mathbf{x}}_{t+1} \le \bar{\mathbf{x}}_{t+1}^{(l)} \\
& \mathbf{y}_t, \tilde{\mathbf{y}}_t \in \mathbb{Z}_+^{n \times n}, \ \tilde{\mathbf{x}}_{t+1} \in \mathbb{Z}_+^n
\end{aligned}
$$

We shall note that the linear relaxation of the integer constraints in the above formulation does not change the optimal solution as it is well known for transportation problems. Let $\big(s_l \mathbf{y}_t^{(l)}, s_l \tilde{\mathbf{y}}_t^{(l)}, s_l \tilde{\mathbf{x}}_{t+1}^{(l)}\big)$ denote the optimal solution of the above linear program. As shown by Hochbaum and Shanthikumar (1990), the following bound on the difference between the exact and approximate solutions holds

$$
\Big\| \big(\mathbf{y}_t^*, \tilde{\mathbf{y}}_t^*, \tilde{\mathbf{x}}_{t+1}^*\big) - s_l \big(\mathbf{y}_t^{(l)}, \tilde{\mathbf{y}}_t^{(l)}, \tilde{\mathbf{x}}_{t+1}^{(l)}\big) \Big\|_\infty \le n(2n+1)s_l. \tag{2.20}
$$

Therefore, we set $\underline{\mathbf{x}}_{t+1}^{(l)} = 2\tilde{\mathbf{x}}_{t+1}^{(l-1)} - (4n^2 + 2n)\mathbf{1}$ and $\bar{\mathbf{x}}_{t+1}^{(l)} = 2\tilde{\mathbf{x}}_{t+1}^{(l-1)} + (4n^2 + 2n)\mathbf{1}$.

With the above formulation, we can solve the allocation problem (2.14)–(2.18) exactly using Algorithm 1. The total number of linear programs that are solved in Algorithm 1 is $\big\lceil \log_2\big(\frac{2}{2n+1} \left\|(\mathbf{z}_t, \mathbf{d}_t)\right\|_\infty\big)\big\rceil$ and the number of variables of each linear program is $8n^3 + 4n^2$. Whereas these linear programs can be solved with any standard off-the-shelf solver, in the next subsection, we present a more efficient algorithm based on network optimization.

Algorithm 1 Optimal Allocation for Discrete Demand (Hochbaum and Shanthikumar, 1990).

1: *Initialization*

2: $l_0 \leftarrow \left\lceil \log_2\left(\frac{2}{2n+1}\left\|(\mathbf{z}_t, \mathbf{d}_t)\right\|_\infty\right)\right\rceil$, $\tilde{\mathbf{x}}_{t+1}^{(0)} \leftarrow \mathbf{0}$

3: *Optimization*

4: **for** $l \leftarrow 1, 2, \ldots, l_0$ **do**

5: $\quad s_l \leftarrow 2^{l_0-l}$, $\underline{\mathbf{x}}_{t+1}^{(l)} \leftarrow 2\tilde{\mathbf{x}}_{t+1}^{(l-1)} - (4n^2+2n)\mathbf{1}$, $\bar{\mathbf{x}}_{t+1}^{(l)} \leftarrow 2\tilde{\mathbf{x}}_{t+1}^{(l-1)} + (4n^2+2n)\mathbf{1}$

6: $\quad$ Solve $(\mathrm{IP} - s_l)$

7: **end for**

8: $\mathbf{y}_t \leftarrow \mathbf{y}_t^{(l_0)}$

2.5.3 Improvement on Allocation Using Network Optimization

In this subsection, we discuss a network representation of the linear program (IP-s_l) formulated in the previous subsection. By exploring the reverse Monge property in the network, we devise an algorithm that solves the allocation problem more efficiently than the usual linear programming algorithm. Similar to the algorithm for the one-period allocation problem, we first construct a network with nodes representing initial on-hand inventory, fulfilled demand, or leftover inventory, and arcs representing inventory allocation. We then show that the max-benefit flow in the constructed network corresponds to the optimal solution to the original problem. Moreover, the benefit matrix of the constructed network satisfies the reverse Monge property, and therefore, the max-benefit flow in the network can be found efficiently with a greedy algorithm similar to the one in §§2.4.2.

Let G^{s_l} denote the constructed network for $(\mathrm{IP} - s_l)$ (see the graphical illustration in Figure 2.4). There are n supply nodes corresponding to the on-hand inventory of the n products. The amount of product i's supply equals to the scaled on-hand inventory $\lfloor z_{t,i}/s_l \rfloor$. There are n actual demand nodes, and the actual demand volume for product i is equal to the scaled demand $\lfloor d_{t,i}/s_l \rfloor$. In addition to actual demands, we create virtual demand nodes to represent inventory conversion. The virtual demands of product i correspond to the segments of the piecewise linear function $V^{s_l}_{t+1,i}(\tilde{x}_{t+1,i})$. Recall that $\underline{x}^{(l-1)}_{t+1,i}$ and $\bar{x}^{(l-1)}_{t+1,i}$ are lower and upper limits of $\tilde{x}_{t+1,i}$, and s_l is the length of the linear segments. Therefore, the number of virtual demands introduced for each product is equal to $K+1$, where $K := \left(\bar{x}^{(l-1)}_{t+1,i} - \underline{x}^{(l-1)}_{t+1,i}\right)/s_l$. For each product i, its virtual demands are indexed from 0 to K. Virtual demand 0 of product i represents the linear segment from 0 to $s_l \underline{x}^{(l)}_{t+1,i}$ (recall that $\underline{x}^{(l)}_{t+1,i}$ is the lower limit of $x_{t+1,i}$). Accordingly, its demand volume is set to $\underline{x}^{(l)}_{t+1,i}$ units of the scaled inventory $\lfloor z_{t,i}/s_l \rfloor$. Virtual demand k, for $k = 1, 2, \ldots, K$, represents the linear segment from $s_l\big(\underline{x}^{(l)}_{t+1,i} + k - 1\big)$ to $s_l\big(\underline{x}^{(l)}_{t+1,i} + k\big)$ and has demand volume equal to one unit of the scaled inventory. In total, there are $n(K+2)$ demand nodes (either actual or virtual). The demand nodes are indexed in such a way that the actual demand of product i has index $j = (i-1)(K+1) + i$ and virtual demand k of product i has index $j = (i-1)(K+1)+i+1+k$, for $k = 0, 1, 2, \ldots, K$. We use the $n(K+2)$-dimensional vector $\mathbf{d}^{(l)}_t = \left(d^{(l)}_{t,1}, d^{(l)}_{t,2}, \ldots, d^{(l)}_{t,n(K+2)}\right)$ to represent the amount of actual and virtual demand, where $d^{(l)}_{t,j}$ is defined as

$$d^{(l)}_{t,j} = \begin{cases} \lfloor d_{t,i}/s_l \rfloor, & \text{if } j = (i-1)(K+1) + i, \\ \underline{x}^{(l)}_{t+1,i}, & \text{if } j = (i-1)(K+1) + i + 1, \\ 1, & \text{if } j = (i-1)(K+1) + i + 1 + k,\ \forall k = 1, 2, \ldots, K, \end{cases}$$

for some $i = 1, 2, \ldots, n$.

Let $r^{(l)}_{i,j}$ denote the unit benefit of allocation from supply i to demand j. First, consider allocation from supply i to demand $j = (j'-1)(K+1) + j'$, that is, the actual demand of product j'. The benefit per unit is s_l times the original per unit net gain, that is, $r^{(l)}_{i,j} = s_l r_{i,j'}$. Next, consider allocation from supply i to demand $j = (j'-1)(K+1)+j'+1+k$, that is, virtual demand k of product j', for $k = 0, \ldots, K$. In this case, allocation means converting product i's inventory into that of product j'. The benefit corresponds to the increase in the expected value of product (j')'s inventory minus the loss incurred in the conversion. We have $r_{i,j}^{(l)} = s_l r_{i,j'} - \Delta_{j',k}^{(l)}$, where

$$\Delta_{j,k}^{(l)} = \begin{cases} s_l r_{j,j} - \Big(V_{t+1,j}^{s_l}\big(s_l \underline{x}_{t+1,j}^{(l)}\big) - V_{t+1,j}^{s_l}(0)\Big)/\underline{x}_{t+1,j}^{(l)}, & \text{if } k = 0, \\ s_l r_{j,j} - \Big(V_{t+1,j}^{s_l}\Big(s_l\big(\underline{x}_{t+1,j}^{(l)} + k\big)\Big) - V_{t+1,j}^{s_l}\Big(s_l\big(\underline{x}_{t+1,j}^{(l)} + k - 1\big)\Big)\Big), & \text{if } k = 1, 2, \ldots, K. \end{cases}$$

Let $\mathbf{R}^{(l)}$ denote the $n \times n(K+2)$ matrix with element $r_{i,j}^{(l)}$, which we refer to as the extended benefit matrix. The following lemma establishes that the extended benefit matrix preserves the reverse Monge property of the original benefit matrix.

Lemma 2.5.1 *If the original benefit matrix* $\mathbf{R}$ *has the reverse Monge property (i.e.,* $r_{i_1,j_1} + r_{i_2,j_2} \geq r_{i_1,j_2} + r_{i_2,j_1}$ *for* $1 \leq i_1 < i_2 \leq n$ *and* $1 \leq j_1 < j_2 \leq n$*), then so does the extended matrix* $\mathbf{R}^{(l)}$*.*

In network G^{s_l} as depicted in Figure 2.4, the supply of product i is represented by the source node i, $i = 1, 2, \ldots, n$. The actual demand of product i is represented by node $n + i$, $i = 1, 2, \ldots, n$. Virtual demand k of product i is represented by node $(n + i)_k$, $k = 0, 1, \ldots, K$, $i = 1, 2, \ldots, n$. Node $2n + 1$ is the sink node. The amount of excess flow at each node is defined as

$$e(i) = \begin{cases} \lfloor z_{t,i}/s_l \rfloor, & \text{if } i = 1, 2, \ldots, n, \quad \text{(a supply node)} \\ -\sum_{i=1}^{n} \lfloor z_{t,i}/s_l \rfloor, & \text{if } i = 2n + 1, \quad \text{(the sink node)} \\ 0, & \text{otherwise} \quad \text{(a demand node).} \end{cases}$$

There are two types of arcs. Arcs from the source nodes to the demand nodes represent inventory allocation or conversion, and the weight $\lambda_{i,j}$ of these arcs correspond to benefits in the extended benefit matrix $\mathbf{R}^{(l)}$. These arcs have unlimited capacities. Arcs from the demand nodes to the sink node have limited capacities $u_{i,j}$ which correspond to the scaled demand vector $\mathbf{d}_t^{(l)}$. These arcs have zero benefit. The arc benefits and capacities are summarized as follows.

$$\lambda_{i,j} = \begin{cases} s_l r_{i,j-n}, & \text{if } i = 1,\ldots,n,\ j = n+1,\ldots,2n, \\ s_l r_{i,j-n} - \Delta_{j-n,k}^{(l)}, & \text{if } i = 1,\ldots,n, \\ & j = (n+j')_k,\ \text{for some } j' = 1,\ldots,n, k = 0,\ldots,K, \\ 0, & \text{otherwise.} \end{cases}$$

$$u_{i,j} = \begin{cases} \lfloor d_{t,i-n}/s_l \rfloor, & \text{if } i = n+j',\ \text{for some } j' = 1,2,\ldots,n,\ j = 2n+1, \\ \underline{x}_{t+1,i-n}^{(l)}, & \text{if } i = (n+j')_0,\ \text{for some } j' = 1,2,\ldots,n,\ j = 2n+1, \\ 1, & \text{if } i = (n+j')_k,\ \text{for some } j' = 1,2,\ldots,n, k = 1,\ldots,K,\ j = 2n+1, \\ \infty, & \text{otherwise,} \end{cases}$$

where we take $(n+i)_k - n = i$.

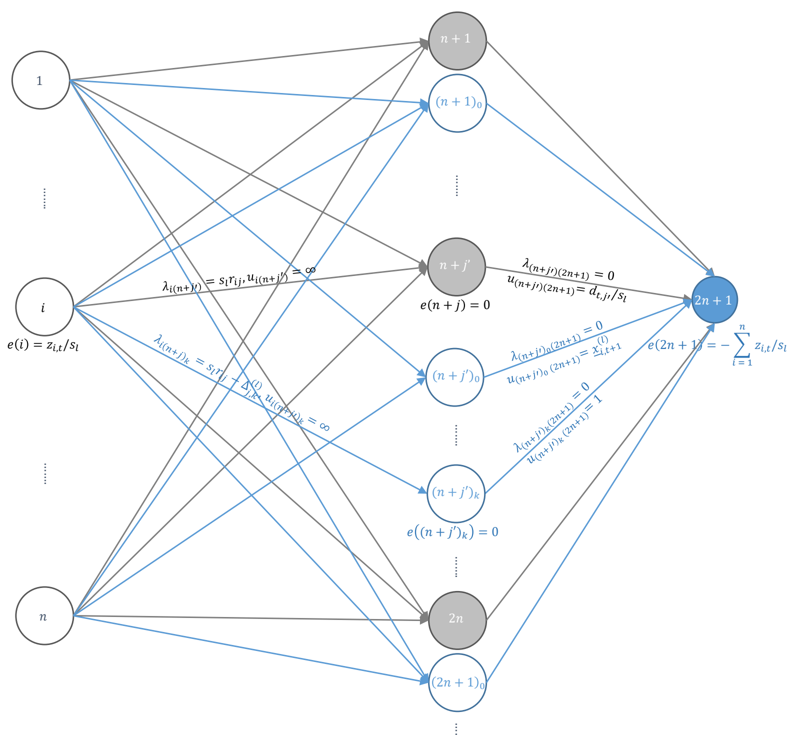

Fig. 2.4.: Network G^{s_l}

Lemma 2.5.2 *The max-benefit flow in network G^{s_l} corresponds to an optimal allocation for problem* $(\text{IP} - s_l)$*. Moreover, if the original benefit matrix* $\mathbf{R}$ *satisfies the reserve Monge property, then the algorithm developed by Vaidyanathan (2013) finds the max-benefit flow in network G^{s_l} with a run time of $O(n^4 \log n)$.*

Lemma 2.5.2 suggests that the linear program $(\text{IP}-s_l)$ can be solved efficiently as a network optimization problem when the benefit matrix of our problem has the reverse Monge property. Combining this result with the development by Hochbaum and Shanthikumar (1990), we conclude the optimality and time complexity of Algorithm 1.

Theorem 2.5.1 *If the original benefit matrix* $\mathbf{R}$ *satisfies the reserve Monge property, then Algorithm 1 finds an optimal solution for problem (2.14)-(2.18) with a run time of $O\Big(n^4 \log n \log\big(\frac{2}{2n+1}\left\|(\mathbf{z}_t, \mathbf{d}_t)\right\|_\infty\big)\Big)$.*

2.6 Numerical Analysis and Observations

In this section, we evaluate the performance of the approximate algorithm developed in the previous section, as well as examining the firms' substitution policies. For these purposes, we define two benchmark models, which give a lower bound and an upper bound of the optimal expected profit.

Consider the special case in which the demands over the entire planning horizon T are observed at the beginning of period 1. In this case, the problem reduces to the single-period planning described in §2.4 for a demand mix of $\sum_{t=1}^{T} \mathbf{d}_t$. The resulting optimal profit can be easily obtained as

$$\overline{V} = \max_{\mathbf{q} \in \mathbb{R}_+^n} \left\{ \mathbb{E}\left[W\left(\mathbf{S}(\mathbf{q}), \sum_{t=1}^{T} \mathbf{d}_t \right) - \mathbf{c}^\top \mathbf{S}(\mathbf{q}) \right] \right\} - \mathbf{g}^\top \left(\sum_{t=1}^{T} \mathbb{E}[\mathbf{D}_t] \right), \qquad (2.21)$$

where $W(\cdot, \cdot)$ is defined in (2.9). In view of the concavity of the profit functions (recall Theorem 2.3.1), we can easily deduce that $\overline{V}$ is an *upper bound* of the optimal profit for our problem defined in (2.2)–(2.5).

Consider another special case in which there is no benefit of substitution, i.e., $r_{i,j} = 0$ whenever $i \neq j$. In this case, the problem reduces to one of managing replenishment for each product separately. The resulting optimal profit can be computed as

$$\underline{V} = \sum_{i=1}^{n} \tilde{V}_{1,i}(\mathbf{0}), \tag{2.22}$$

where $\tilde{V}_{t,i}(\cdot)$ is defined in (2.11). Clearly, the restriction on substitution leads to suboptimality. We have the following relationship:

$$\underline{V} \leq \tilde{V} \leq V^* \leq \overline{V}, \tag{2.23}$$

where $\tilde{V}$ is the profit under the heuristic developed in §2.5 and $V^* = V_1(\mathbf{0})$ is the optimal profit.

In the numerical analysis below, we consider $n = 4$ products with discretized log-normal demands in each period (i.e., $\log(\mathbf{D}_t)$ follows a joint normal distribution). The mean demands are $\mu = \mathbf{9}$ and the coefficient of variation of the demand $\mathbf{cv}$ varying from 0.5 to 3. We consider positively correlated, uncorrelated, and negatively correlated scenarios with the correlation matrices of the underlying normal distribution are given by, respectively,

$$\rho_+ = \begin{bmatrix} 1 & 0.5 & 0.2 & 0 \\ 0.5 & 1 & 0.5 & 0.2 \\ 0.2 & 0.5 & 1 & 0.5 \\ 0 & 0.2 & 0.5 & 1 \end{bmatrix}, \qquad \rho_0 = I, \qquad \rho_- = \begin{bmatrix} 1 & -0.5 & -0.2 & 0 \\ -0.5 & 1 & -0.5 & -0.2 \\ -0.2 & -0.5 & 1 & -0.5 \\ 0 & -0.2 & -0.5 & 1 \end{bmatrix}.$$

We set $\mathbf{h} = \mathbf{v} = \mathbf{0}$, $\mathbf{g} = (0.4, 0.3, 0.2, 0.1)$ and

$$\mathbf{R} = \begin{bmatrix} 3.0 & 2.0 & 1.3 & 0.7 \\ 1.8 & 2.5 & 1.8 & 1.2 \\ 0.8 & 1.5 & 2.0 & 1.4 \\ 0.1 & 0.8 & 1.3 & 1.8 \end{bmatrix}.$$

We vary the procurement costs $\mathbf{c}$ within $[0.4, 1.6]$ and the planning horizon T within $\langle 5, 20 \rangle$. Only one replenishment is allowed at the beginning of the planning horizon. The supply function is of the form $S_i(q_i) = \min\{q_i, K_i\}$, where K_i is the nonnegative random capacity for product i. We take $\Pr\{K_i = \underline{K}\} = 1 - \Pr\{K_i = \overline{K}\} = 0.2$.

2.6.1 Performance of the Approximation Algorithm

Given that it is difficult to compute the optimal solution to our problem, we measure the performance of the approximate algorithm against the upper bound problem instead. That is, we compute the percentage profit gap

$$\frac{\overline{V} - \tilde{V}}{\overline{V}} \times 100\%.$$

The actual optimality gap, measured against V^* as opposed to $\overline{V}$, is certainly smaller than the above in view of (2.23).

In the instances reported in Table 2.1, the performance gap is within the range of [5%, 11.57%]. The instances of high performance gaps occur when the optimal profit is extremely small. We observe that the relative performance gap reduces when the procurement cost decreases, the demand variability decreases, the supply uncertainty decreases, or the length of the replenishment cycle increases.

Intuitively, demand correlation can have a major effect on the product substitutions, and our approximation, assuming substitution only in the current period, would lead to a large profit loss with negatively correlated demands. Interestingly, the performance gap reported in Table 2.1 is not very sensitive to the demand correlation. This suggests that once product "conversion" allows for capturing most of the benefit from product substitution. This observation further justifies our approximation approach.

2.6.2 Benefit of Dynamic Product Substitution

To measure the benefit of substitution using the policy generated by our approximate algorithm, we measure the profit gap against the lower bound profit $\underline{V}$ (recall that $\underline{V}$ is computed in (2.22) by assuming no substitution is allowed). That is,

$$\frac{\tilde{V} - \underline{V}}{\underline{V}} \times 100\%.$$

From Table 2.2, we observe that the value of product substitution increases when the product procurement cost increases. A higher procurement cost induces a lower

Table 2.1.: Performance of the approximation algorithm measured by $(\bar{V} - \tilde{V})/\bar{V} \times 100\%$

Capacities	*Deterministic*		*Uncertain*			
	Unlimited	Limited	Limited			Unlimited
			(Low)	(Med)	(High)	
$c = 0.4$	0.18	0.53	0.54	0.78	1.26	1.02
0.8	0.60	0.74	0.80	1.21	2.00	2.42
1.2	1.50	1.58	1.63	2.24	5.02	6.34
1.6	4.40	4.30	4.56	5.61	8.57	11.57
$T = 5$	2.10	2.17	2.20	2.67	4.07	5.45
10	1.50	1.58	1.63	2.24	5.02	6.34
15	1.23	1.26	1.36	2.03	3.35	4.50
20	1.04	1.06	1.16	1.88	3.10	4.14
$c_{\mathrm{v}} = 0.5$	0.83	0.89	0.96	1.65	2.89	3.97
1	1.50	1.58	1.63	2.24	5.02	6.34
2	2.73	2.87	2.88	3.37	4.78	6.26
3	3.75	3.83	3.86	4.24	5.69	7.32
$\rho = \rho_+$	1.23	1.18	1.28	2.02	3.44	4.76
$\boldsymbol{\rho_0}$	1.50	1.58	1.63	2.24	5.02	6.34
ρ_-	1.43	1.74	1.74	2.31	3.62	4.90

Notes. The base setting is highlighted with boldface values in the first column. $\log(\mathbf{D}_t) \sim N(\mu, \Sigma)$ with $\mu_t = 9$. $\Pr\{K = \underline{k}\} = 1 - \Pr\{K = \overline{k}\} = 0.2$. We set $\underline{k} = \overline{k} = 9T$ under *unlimited deterministic supplies*, $\underline{k} = \overline{k} = \infty$ under *limited and deterministic supplies*, and $\underline{k} = 0$ and $\overline{k} = \infty$ under *unlimited random supplies.* For *limited random supplies*, we choose $\underline{k} \in \{0.8, 0.5, 0.2\}$ with $\mu_K = 0.2\underline{k} + 0.8\overline{k}$ fixed to represent low, medium, and high supply variabilities.

replenishment quantity and thus increases the chance of stockout. Consequently, product substitution becomes more useful.

The benefit derived from substitution decreases when the replenishment cycle becomes longer or the demand variability becomes smaller. This is simply because the chance of using a product's own supply to meet its own demand becomes higher. We further note that product substitution leads to a larger value when the demands are negatively correlated than when they are positively correlated, as intuition may suggest.

Table 2.2.: Performance of the approximation algorithm measured by $(\tilde{V} - V)/V \times 100\%$

Capacities	*Deterministic*		*Uncertain*			
	Unlimited	Limited	Limited			Unlimited
			(Low)	(Med)	(High)	
$c = 0.4$	0.18	0.53	0.54	0.78	1.26	1.02
0.8	0.60	0.74	0.80	1.21	2.00	2.42
1.2	1.50	1.58	1.63	2.24	5.02	6.34
1.6	4.40	4.30	4.56	5.61	8.57	11.57
$T = 5$	2.10	2.17	2.20	2.67	4.07	5.45
10	1.50	1.58	1.63	2.24	5.02	6.34
15	1.23	1.26	1.36	2.03	3.35	4.50
20	1.04	1.06	1.16	1.88	3.10	4.14
$c_{\mathrm{v}} = 0.5$	0.83	0.89	0.96	1.65	2.89	3.97
1	1.50	1.58	1.63	2.24	5.02	6.34
2	2.73	2.87	2.88	3.37	4.78	6.26
3	3.75	3.83	3.86	4.24	5.69	7.32
$\rho = \rho_+$	1.23	1.18	1.28	2.02	3.44	4.76
$\boldsymbol{\rho_0}$	1.50	1.58	1.63	2.24	5.02	6.34
ρ_-	1.43	1.74	1.74	2.31	3.62	4.90

Notes. The base setting is highlighted with boldface values in the first column. $\log(\mathbf{D}_t) \sim N(\mu, \Sigma)$ with $\mu_t = 9$. $\Pr\{K = \underline{k}\} = 1 - \Pr\{K = \overline{k}\} = 0.2$. We set $\underline{k} = \overline{k} = 9T$ under *unlimited deterministic supplies*, $\underline{k} = \overline{k} = \infty$ under *limited and deterministic supplies*, and $\underline{k} = 0$ and $\overline{k} = \infty$ under *unlimited random supplies.* For *limited random supplies*, we choose $\underline{k} \in \{0.8, 0.5, 0.2\}$ with $\mu_K = 0.2\underline{k} + 0.8\overline{k}$ fixed to represent low, medium, and high supply variabilities.

In general, a higher level of supply uncertainty leads to a larger benefit derived from product substitution, because there is a greater chance of the mismatch between the supply mix and the demand mix. The presence of supply limitation, however, can increase or decrease the value of product substitution, as we observe from the instances with deterministic capacities. On the one hand, extremely limited capacities induce a high likelihood of stockout across all products in the early periods during a replenishment cycle, leaving little overstock of any product to be used for other demands. On the other hand, extremely ample capacities allow the firm to freely choose the right replenishment quantities to match the demand mix, reducing the need for product substitution. Therefore, product substitution is most valuable when the capacities are moderate, with which a careful cost and benefit trade-off can lead to close supply and demand levels. As demand uncertainties materialize after replenishment, there is a good chance of understock for certain products and overstock for other products. The scarcity or ampleness of capacities is relative to the procurement cost—a given capacity level is ample (scarce) when the procurement cost is low (high), as suggested by Table 2.2.

In all the instances reported in Table 2.2, the value of product substitution under our approximate algorithm is significant, between 4.30% and 30.88%. This makes an interesting contrast to a common observation from the literature (see, e.g., Shumsky and Zhang, 2009; Yu et al., 2015; Yao et al., 2016) that dynamic product substitution does not generate much value unless the replenishment decisions are made suboptimally. This contrast is due to the fact that most previous studies do not model the supply processes (i.e., assuming any replenishment quantities are fully delivered). Our model, taking into account general supply functions, is more realistic and provides a strong justification for the wide adoption of substitution policies. The significance of the values reported here also suggests that our approximation appropriately explores the relationships among the products in generating the allocation solutions.

2.7 Concluding Remarks

We study the problem of coordinating product replenishment and substitution decisions in a dynamic environment with both supply and demand uncertainties. There are two key features of our model. The first is the general substitution structure with the benefit matrix revealing the reverse Monge property. The reverse Monge property allows us to generalize the existing studies on downward product substitution, yet develop an efficient approach to derive the decisions of multi-product allocation to meet the demands. The second feature of our model is the consideration of general supply functions that are stochastically linear in midpoint. We show that the value of product substitution is greatly affected by both the limits and the variabilities of the supplies.

3. DATA-DRIVEN INVENTORY AND PRICING MANAGEMENT WITH ADVERSARIAL MODELS

3.1 Synopsis

To study the problem of joint optimizing the price and the inventory decisions with ambiguous demand information, we consider a widely used demand model where the demand is a function of the unit selling price and an uncertain factor. To allow direct comparison with the conventional approach, we consider the case where the form of the demand function is known, but the distributional information of the uncertain factor remains ambiguous. We assume that the firm only knows the support of the uncertain factor, e.g., the interval within which the uncertain factor lies with high confidence. With such demand ambiguity, the firm cannot maximize the expected profit using the traditional approach. Therefore, we adopt the minimax regret decision criterion to determine the price and order quantity that minimizes the worst-case regret. The regret is defined as the gap between the optimal profit that the firm could obtain with perfect demand information and the realized profit using decisions made with ambiguous demand information. The minimax regret criterion is an important alternative to maximizing expected payoff in decision theory. It has been adopted in inventory management or pricing to tackle ambiguity and generate new insights (see, Perakis and Roels, 2008; Caldentey et al., 2016).

Our contributions are summarized as follows:

1. We propose a robust and tractable approach for joint price and inventory optimization under demand ambiguity to minimize the worst-case regret. We explore the properties of the optimization problem and characterize the optimal solutions, which can be computed efficiently.

2. We study the properties of the optimal minimax regret decisions, such as the impact of inventory risk and monotone comparative statics with respect to the degree of demand ambiguity and unit ordering cost. Comparison of these properties with those of the classical models unveils both similarities and distinctions between the minimax regret model and the classical models.

3. We compare the minimax regret approach with two other widely used approaches that can tackle demand ambiguity. Compared with the max-min robust approach that maximizes the worst-case profit, the minimax regret approach avoids extreme conservativeness and provides robust performance that does not depend on the choice of uncertainty sets. Compared with the regression-based data-driven approach, the minimax regret approach leads to higher realized profit when data is scarce, when demand is highly volatile, or when the demand model is misspecified.

The rest of this chapter is organized as follows. Section 3.2 reviews related literature. Section 3.3 presents the model for pricing newsvendor problem under minimax regret. Section 3.4 characterizes the optimal ordering and pricing decision under the minimax regret framework. Section 3.5 compares the properties of the minimax regret decisions with the optimal solutions of the classical models. Section 3.6 describes the implementation of the minimax regret approach in a data-driven setting and compares its performance with other approaches.

3.2 Literature Review

There exists a vast literature on coordinating pricing and inventory decisions (Elmaghraby and Keskinocak, 2003; Chan et al., 2004; Yano and Gilbert, 2005; Chen and Simchi-Levi, 2012). The price-demand relationship is typically modeled by a certain parametric form, and the distribution of the random factor is given. Petruzzi and Dada (1999) presented a unified framework for the additive demand (Mills, 1959), the multiplicative demand (Karlin and Carr, 1962), and the more general multiplicative-additive demand (Young, 1978). More general conditions that can guarantee the unimodality of the expected profit function were provided by Yao et al. (2006), Lu and Simchi-Levi (2013), Roels (2013), and Luo et al. (2016). Kocabiyikoglu and Popescu (2011) took a different approach with a new demand model based on lost-sales rate (LSR) elasticity and show that structural properties of the problem can be fully characterized by conditions regarding LSR elasticity. Feng and Shanthikumar (2018) considered general supply and demand functions in pricing and inventory models. They introduced the notion of stochastic linearity in mid-point and showed that it can guarantee the concavity of the expected profit function for a much more general class of demand and supply functions. Different from studies on the standard pricing newsvendor problem, Chou et al. (2012) considered the pricing and inventory problem in a two-sided market given the cross-side network effect, where an intermediary coordinates the price and the order quantity of platforms sold to consumers as well as the royalty charged to content developers for software. An interesting research question that has been extensively discussed for the pricing newsvendor problem is what effect demand uncertainty and inventory risk have on the optimal price decision. Mills (1959) found that, with additive demand, the optimal price is smaller than the risk-free price, i.e., the price that maximizes the product of price and expected demand. Karlin and Carr (1962) found that the optimal price is larger than the risk-free price with multiplicative demand. Petruzzi and Dada (1999) summarized these results in the previous literature with a unified framework. Lu and Simchi-Levi (2013) and Roels (2013) also studied the effect of additive-multiplicative demand model on the relationship between the optimal price and the risk-free price. Instead of assuming a risk-neutral objective, Agrawal and Seshadri (2000) and Chen et al. (2009) considered the effect of risk-aversion in the pricing newsvendor problem.

All the aforementioned papers assume that the demand function and the distribution of the random factor are known. These approaches may have some limitations for application when the complete demand information is not available. Therefore, robust optimization and data-driven approaches were proposed to overcome these limitations under demand ambiguity. For the standard newsvendor problem where the firm only determines the order quantity, Scarf (1958) first presented the distributionally-robust optimization approach under ambiguous demand specified by the first and second moments. A detailed discussion on this modeling framework was given in Gallego and Moon (1993). Natarajan et al. (2017) further incorporated distribution asymmetry to improve the performance of the robust solution. Another stream of literature devises data-driven approaches to determine the ordering quantities based on historical data. Levi et al. (2007) proposed sampling-based policies for the newsvendor problem. They solved the sample average approximation counterpart of the problem and proved bounds for the number of samples required such that the expected cost associated with their approximate solutions is close to the true optimal one with high probability. Levi et al. (2015) extended the work of Levi et al. (2007) by identifying the conditions for the demand distribution where the sample-based approach is effective and proves tighter bounds. Godfrey and Powell (2001) approximated the objective cost function with a sequence of piecewise linear functions using a technique called Concave Adaptive Value Estimation (CAVE). Bookbinder and Lordahl (1989) proposed estimating the critical quantile of the demand distribution using the bootstrap method. Liyanage and Shanthikumar (2005) and Chu et al. (2008) introduced a new approach called operational statistics, which integrates parameter estimation and optimization. Jain et al. (2011) introduced operational objective learning for the newsvendor problem with inventory-dependent demand. Rudin and Vahn (2014) studied the newsvendor problem with feature information, which may include the price as a feature. They proposed machine learning methods with or without regularization to determine the optimal order quantity.

There are fewer papers that consider joint pricing-ordering problems using robust optimization or data-driven approaches. Burnetas and Smith (2000) proposed a fully data-driven approach for the pricing newsvendor problem. The pricing problem is modeled as a multi-arm bandit problem where the price is selected from a set of price levels. Given the price level, the optimal order quantity is determined using stochastic approximation with historical observations. This approach can be applied to the setting where the number of feasible prices is small and the observations for each price level are adequate. Chu et al. (2017) applied operational statistics in the pricing newsvendor problem that integrates parameter estimation and profit maximization. They considered the multiplicative demand of which the random factor follows the exponential distribution. The elasticity parameter that determines the price decision and the scale parameter that governs the quantity decisions is unknown. They proposed approaches for both scenarios when the scale parameter is unknown and when neither of the parameters is known. Feng et al. (2013) studied the dynamic pricing and inventory management over multiple periods with demand estimation. They identified the sufficient conditions under which a base-stock list-price policy is optimal and developed the constrained maximum likelihood estimation approach to obtain estimates with the generalized additive model (GAM). Recently, Chen et al. (2017) presented a non-parametric algorithm for the joint pricing and inventory management problem where the demand distribution is unknown and the estimated profit function based on the historical data is not necessarily unimodal. They designed a learning-while-doing algorithm that integrates exploration and exploitation and mitigates estimation biases due to demand censoring. They showed that the convergence rate of the regret of their approach is of order $T^{-\frac{1}{5}}(\log T)^{\frac{1}{4}}$. Lu et al. (2016) presented a data-driven approach for the pricing-newsvendor problem. Based on historical data and domain knowledge, without assuming parametric forms of the demand model, they estimated the conditional quantile path of the demand using parametric programming. They further enhanced the estimates and decisions using smoothing and kernelization as well as additional domain knowledge. Harsha et al. (2016) proposed a data-driven framework that translates the data-driven optimization problem into statistical estimations of mean, quantile, and superquantile of the demand distribution. Additional drivers besides the price can be incorporated into the demand model for statistical accuracy. Fu et al. (2017) considered the profit sharing problem in an ambiguity averse supply chain under price and demand uncertainty. They developed a distributionally robust Stackelberg game model with limited price and demand information. Compared with the wholesale price contract, the profit sharing agreement is shown to be more beneficial to both supplier and retailer.

Our work is different from the aforementioned papers as we adopt the minimax regret framework while considering both pricing and ordering decisions. The regret is defined as the gap between the optimal profit that the firm can acquire with perfect demand information and the realized profit with ambiguous demand information. The minimax regret decision rule was originally proposed by Savage (1951), and has since been adopted in various settings such as operations research, economics, marketing, and computer science. Yue et al. (2006) and Perakis and Roels (2008) adopted this framework where the firm aims to minimize the maximum opportunity cost from not making the optimal decision with moment information of the demand. Levi et al. (2011) solved the minimax regret problem with absolute mean spread information. Zhu et al. (2013) also applied the minimax regret approach in the newsvendor model where the regret is defined as the ratio between the expected cost with partial information and the expected cost with complete information. Jiang et al. (2013) adopted the minimax regret framework for the unit commitment problem in power systems. Bergemann and Schlag (2008) studied the static pricing problem where the firm only knows the interval of the customer valuations. They proposed a randomized pricing algorithm to minimize the maximum regret since the nature will adversely choose the customer valuation from the interval. Caldentey et al. (2016) considered intertemporal pricing under minimax regret. They characterized the optimal price trajectory to minimize the maximum regret associated with ambiguous information regarding customers' valuations and the delay in the purchase of strategic customers. Handel and Misra (2015) studied a two-period dynamic pricing problem with ambiguous demand information under minimax regret framework. Since pricing decisions are made sequentially for two periods and the underlying demand curve must be consistent, the pricing decision for the first selling period will affect the learning outcome and thus further influence the decision in the second period. Given this interaction between the firm and the nature, the optimal pricing decision can be found by solving the dynamic pricing problem backward. Chen et al. (2016) considered a similar setting as in Handel and Misra (2015) but with a fixed initial inventory. The pricing decisions are contingent on the remaining inventory at the beginning of each period. Their computational study shows that the expected profit of using such approach can closely approximate that of the conventional approach that requires complete demand information. These papers considered either inventory or pricing decisions. To the best of our knowledge, the only paper that considers pricing and inventory jointly under minimax regret is Wang et al. (2014). They considered a robust pricing newsvendor problem where only upper and lower limits of the market size and the interval of possible consumer's willingness-to-pay are known. The realized demand is determined by the actual market size and the willingness-to-pay function. Different from Wang et al. (2014), we consider the multiplicative-additive demand model, which is widely used in classical pricing newsvendor models. Our model can be viewed as a robust counterpart of the classical models (e.g., Petruzzi and Dada, 1999) under the minimax regret framework.

3.3 Model

The firm under consideration sells a single type of product in a single period. The demand denoted by $D(p, \theta)$ is a function of the selling price p and an uncertain factor θ. $D(p, \theta)$ is assumed to be decreasing in the unit price p and increasing in θ without loss of generality. We assume that unsold inventory is of zero salvage value at the end of the selling period. The profit from ordering y units of products and setting the selling price as p for a given θ is of the form:

$$\pi(p, y; \theta) = p\min\{D(p, \theta), y\} - cy.$$

In the classical model with complete distributional information, the uncertain factor θ is known to have a cumulative distribution function Φ. In this case, we refer to the uncertain factor as the random factor. The firm seeks to maximize the expected profit, i.e.,

$$\max_{p,y}\ \mathbb{E}_{\Phi}\big[\pi(p,y;\theta)\big]. \tag{3.1}$$

3.3.1 Minimax Regret Model with Demand Ambiguity

Under ambiguous demand, assume the firm only knows the support of the random factor θ denoted by $[\underline{\theta},\bar{\theta}]$. Before the selling period starts, the firm needs to determine the unit selling price p and the order quantity y where the unit procurement cost of the product is c. Let $[\underline{p},\bar{p}]$ denote the range of feasible unit selling prices that the firm can set. For example, the lower bound $\underline{p}$ can be set at the discretion of the firm or simply set as the unit procurement cost c. For the upper bound, it can be set as the maximum price that guarantees the demand to be nonnegative, i.e., $\bar{p}=\sup\{p : D(p,\theta)\geq 0, \forall\theta\in[\underline{\theta},\bar{\theta}], p\geq \underline{p}\}$. Without complete demand information, the firm cannot maximize the expected profit. Therefore, we adopt the minimax regret framework that the firm aims at minimizing the worst-case regret. The regret is defined as the gap between the optimal profit that the firm can generate with perfect demand information and the realized profit using the decisions made with ambiguous demand information. First, we consider a benchmark (hindsight) model where the firm determines order quantity and price decisions after observing the realized value of θ. Since the firm has the perfect demand information, the optimal order quantity is $y^*(\tilde{p})=D(\tilde{p},\theta)$ given the unit selling price $\tilde{p}$. The optimal profit of the firm with perfect demand information is $\phi(\tilde{p},\theta)=(\tilde{p}-c)D(\tilde{p},\theta)$ and the optimal hindsight profit, $\phi^*(\theta)$, is defined as follows:

$$\phi^*(\theta)=\max_{\tilde{p}}(\tilde{p}-c)D(\tilde{p},\theta). \tag{3.2}$$

Given the uncertain factor θ, we have $\phi^*(\theta)\geq\pi(p,y;\theta)$ for any (p,y). The regret of the firm with order quantity y and selling price p is defined as follows:

$$R(p,y;\theta)=\phi^*(\theta)-\pi(p,y;\theta)=\max_{\tilde{p}}(\tilde{p}-c)D(\tilde{p},\theta)-p\min\{D(p,\theta),y\}+cy. \tag{3.3}$$

Due to the demand ambiguity, the firm aims at minimizing the worst-case regret. Specifically, the firm first decides the order quantity y and the selling price p. Then, an adversarial nature can choose the uncertain factor θ from the interval $[\underline{\theta}, \bar{\theta}]$ to maximize the regret given the order quantity y and the selling price p. Therefore, the maximum regret of the firm with (p, y) is of the form:

$$R(p, y) = \max_{\theta \in [\underline{\theta}, \bar{\theta}]} R(p, y; \theta).$$

Then, given the response from the nature, the firm chooses (p, y) to minimize the maximum regret $R(p, y)$. Thus, the optimization problem for the firm is defined as follows:

$$\min_{p,y} \max_{\theta \in [\underline{\theta}, \bar{\theta}]} R(p, y; \theta). \tag{3.4}$$

3.4 Optimal Ordering and Pricing Decisions

In this section, we first consider a general class of demand functions of the additive-multiplicative form. Given the unit price p and the uncertain factor θ, the realized demand is of the form:

$$D(p, \theta) = \mu(p) + \sigma(p)\theta. \tag{3.5}$$

In the classical pricing newsvendor model, the additive-multiplicative demand function is similar to the form in (3.5) and the distribution of the random factor is known. Specifically, when $\mu(p) = 0$, it becomes the multiplicative form and when $\sigma(p) = 1$, it becomes the additive form. Due to the demand ambiguity, we assume that the firm only knows the interval in which θ lie with high probability, denoted by $[\underline{\theta}, \bar{\theta}]$, which is defined as the uncertainty set. Therefore, our definition of the demand function can be regarded as a counterpart of the one in the traditional pricing newsvendor problem under demand ambiguity.

Assumption 1 *We make the following assumptions on the demand function.*

(i) $\phi(p, \theta) = (p - c)D(p, \theta)$ *is concave in* p, $\theta \in \{\underline{\theta}, \bar{\theta}\}$.

(ii) $\sigma(p) \geq 0$ *and* $p\sigma(p)$ *is concave in* p.

Assumption (i) ensures the concavity of the profit function for the extreme cases of θ. Similar assumptions are common in existing literature. For example, Federgruen and Heching (1999) assumes that demand function $D(p, \theta)$ is decreasing and concave in p, which is a sufficient condition for assumption (i). For assumption (ii), the condition of $\sigma(p) \geq 0$ essentially guarantees the monotonicity of demand function with respect to θ. For the second part, the concavity of $p\sigma(p)$ is also common in the literature of pricing-newsvendor model. For example, Kocabiyikoglu and Popescu (2011) and Luo et al. (2016) assume that both $p\mu(p)$ and $p\sigma(p)$ are concave in p. Luo et al. (2016) summarize some classes of functions that satisfy this assumption. Note that when Assumption 1 holds, the objective function of the optimization problem (3.2) is concave and there exists a unique optimal solution given the uncertain factor θ. Let $\underline{p}^*$ denote the optimal solution when $\theta = \underline{\theta}$ and $\bar{p}^*$ the optimal solution when $\theta = \bar{\theta}$. Equivalently, we have $\phi^*(\underline{\theta}) = \phi(\underline{p}^*, \underline{\theta})$ and $\phi^*(\bar{\theta}) = \phi(\bar{p}^*, \bar{\theta})$.

To obtain the optimal decisions for the pricing newsvendor problem under the minimax regret framework, we investigate the properties of the objective function in (3.4). The following proposition establishes the convexity of the regret function and characterizes the maximum regret.

Proposition 3.4.1 *The optimal hindsight profit* $\phi^*(\theta)$ *is convex and increasing in* θ *and the regret function* $R(p, y; \theta)$ *is convex in* θ. *Consequently, the maximum regret is* $R(p, y) = \max\{R(p, y; \underline{\theta}), R(p, y; \underline{\theta})\}$.

Since $\phi(p, \theta)$ is linear in θ and the maximum of convex functions is still convex, the optimal hindsight profit $\phi^*(\theta)$ is convex in θ. Then, we can show that the regret function $R(p, y; \theta)$ is convex in θ as well. To maximize the regret, the nature will set θ to the lower bound or the upper bound of the interval $[\underline{\theta}, \bar{\theta}]$ due to the convexity of the regret function. Therefore, we can rewrite the minimax regret problem as follows:

$$\min_{p,y} \max \left\{ R(p, y; \underline{\theta}), R(p, y; \bar{\theta}) \right\}. \tag{3.6}$$

Instead of deciding p and y simultaneously, we can determine p and y sequentially as follows:

$$\min_p \min_y \max\left\{R(p,y;\underline{\theta}), R(p,y;\bar{\theta})\right\}. \tag{3.7}$$

Consider the inner optimization of (3.7), the firm first determines the order quantity y given the unit selling price p. The following proposition characterizes the optimal order quantity $y^*(p)$ that minimizes the maximum regret and the resulting optimal regret function $R\big(p, y^*(p)\big)$ for a given unit selling price p.

Proposition 3.4.2 *When Assumption 1 holds,*[1] *there exists p_l and p_h such that $\underline{p} \leq p_l < p_h \leq \bar{p}$ and the optimal order quantity:*

$$y^*(p) = \begin{cases} \frac{1}{p}\big(\phi^*(\bar{\theta}) - \phi^*(\underline{\theta})\big) + D(p,\underline{\theta}), & \text{if } p \in [p_l, p_h], \\ D(p,\bar{\theta}), & \text{otherwise.} \end{cases}$$

The maximum regret associated with the optimal order quantity:

$$R\big(p, y^*(p)\big) = \begin{cases} \frac{c}{p}\phi^*(\bar{\theta}) + \big(1 - \frac{c}{p}\big)\phi^*(\underline{\theta}) - (p-c)D(p,\underline{\theta}), & \text{if } p \in [p_l, p_h], \\ \phi^*(\bar{\theta}) - (p-c)D(p,\bar{\theta}), & \text{otherwise.} \end{cases}$$

The regret function $R\big(p, y^*(p)\big)$ is convex on $[\underline{p}, \bar{p}]$ and the optimal selling price $p^* = \arg\min\left\{p : R\big(p, y^*(p)\big), p \in [\underline{p}, \bar{p}]\right\}$ lies in the interval $[p_l, p_h]$.

Proposition 3.4.2 first explicitly characterizes the optimal order quantity given the unit selling price p. When the unit selling price is optimally chosen, which means it lies in $[p_l, p_h]$, then, the role of the ordering decision is to balance between the overage regret and the underage regret, both of which are associated with the inventory risk. Otherwise, when $p \notin [p_l, p_h]$, the regret associated with the realization of $\theta = \bar{\theta}$ dominates the one with $\theta = \underline{\theta}$, and thus, the optimal order quantity will always be $D(p, \bar{\theta})$. Hence, with the optimal ordering decision, we can further characterize the maximum regret as a function of the unit selling price p. Specifically, we find the regret function is convex and identify the interval which always contains the optimal solution. Therefore, the optimal unit selling price and order quantity can be efficiently calculated.

[1] For part (ii) of Assumption 1, Proposition 3.4.2 only requires that $p\sigma(p)$ is unimodal.

3.5 Properties of the Minimax Regret Decisions

As mentioned earlier, the minimax regret pricing newsvendor problem can be viewed as a robust counterpart to the classical model that seeks to maximize the expected profit with complete distributional information. In this section, we compare the properties of the optimal decisions under these two frameworks and discuss the intuition behind both similarities and differences between these properties. Specifically, we first observe the monotone relationship between the optimal order quantity and the unit selling price, which does not necessarily hold in the classical models. Second, we further investigate the effect of inventory risk on the optimal decisions by comparing the optimal price and the risk-free price (i.e., the optimal price that minimizes the regret of the firm when there is no risk of inventory mismatch). Third, we study how the optimal price changes with respect to the degree of demand ambiguity. At last, we show with counterexamples that the monotone relationship between the optimal decision variables and the unit ordering cost in the classical pricing newsvendor model does not necessarily hold under the minimax regret framework. We summarize the comparison results in Table 3.1, and present the details in the following sections.

3.5.1 Monotone Optimal Order Quantity

Corollary 3.5.1 *When Assumption 1 holds, the optimal order quantity $y^*(p)$ is strictly decreasing in the unit selling price p.*

The above corollary is a direct result from Proposition 3.4.2 and it indicates a monotone relationship between the optimal order quantity and the given unit selling price p. However, in the classical pricing newsvendor model, the optimal order quantity for a given price p is defined as: $y^*(p) = \sigma(p)\Phi^{-1}(\frac{p-c}{p})$ in the multiplicative demand and $y^*(p) = \mu(p) + \Phi^{-1}(\frac{p-c}{p})$ in the additive demand where $\Phi(\cdot)$ is the cumulative distribution function of the random factor. The monotonicity of the inventory-price relationship, which depends on the form of the demand function and the distribution of the random factor, does not necessarily hold.

Table 3.1.: Comparison between the minimax regret model and the classical model.

Demand Model	Decision	Properties and Comparative Statics	
		Classical Model	Minimax Regret Model
Additive	Price	Smaller than risk-free price	(i) Linear-additive: smaller than risk-free price
			(ii) General-additive: may be larger than risk-free price
		Decreasing in demand variability[2]	(i) Linear-additive: decreasing in demand ambiguity
			(ii) General-additive: may be non-monotone
		Strictly increasing in unit ordering cost	May be non-monotone
	Order Quantity	$\mu(p) + \Phi^{-1}(1 - c/p)$	$\mu(p) + \underline{\theta} + \big(\phi^*(\bar{\theta}) - \phi^*(\underline{\theta})\big)/p$, if $p \in [p_l, p_h]$
		May be non-monotone in price	Strictly decreasing in price
		Strictly decreasing in unit ordering cost	May be non-monotone
Multiplicative	Price	Larger than risk-free price	Larger than risk-free price
		Increasing in demand variability[3]	Increasing in demand ambiguity
		Strictly increasing in ordering cost	May be non-monotone
	Order Quantity	$\sigma(p)\Phi^{-1}(1 - c/p)$	$\sigma(p)\underline{\theta} + \phi^*(1)\big(\bar{\theta} - \underline{\theta}\big)/p$, if $p \in [p_l, p_h]$
		May be non-monotone in price	Strictly decreasing in price
		Strictly decreasing in unit ordering cost	Strictly decreasing in unit ordering cost

3.5.2 Effect of Inventory Risk

In this section, we study the effect of inventory risk by comparing the optimal price and the risk-free price under the minimax regret framework, which is the price that minimizes the regret without inventory risk. Specifically, we can rewrite the regret function in (3) as follows:

$$R(p, y; \theta) = \underbrace{\phi^*(\theta) - (p - c)D(p, \theta)}_{\text{Regret of Price}} + \underbrace{(p - c)\big(D(p, \theta) - y\big)^+ + c\big(y - D(p, \theta)\big)^+}_{\text{Regret of Inventory}}. \tag{3.8}$$

The total regret can be decomposed into two parts: one is associated with the regret from choosing the selling price indicated by the difference of the first two terms in (3.8), and the second part is the regret associated with the order quantity y represented by the last two terms. Now we consider a scenario where the firm always meets the realized demand perfectly given the price (e.g., the firm does not need to pre-order and always receives ample supply). Equivalently, there is no risk of mismatch between supply and demand. Therefore, the risk-free pricing problem under the minimax regret framework can be formulated as follows:[4]

$$\min_{p} \left\{ \max_{\theta \in [\underline{\theta}, \bar{\theta}]} \phi^*(\theta) - (p - c)D(p, \theta) \right\}. \tag{3.9}$$

Let p_0 be the solution to the above problem, which we refer to as the risk-free price. The following proposition characterizes the risk-free price under the additive-multiplicative demand model.

Proposition 3.5.1 *If Assumption 1 holds, let* $\bar{\bar{p}} = \max\{\underline{p}^*, \bar{p}^*\}$ *and* $\underline{\underline{p}} = \min\{\underline{p}^*, \bar{p}^*\}$*, then* p_0 *is the unique price in* $[\bar{\bar{p}}, \underline{\underline{p}}]$ *that satisfies the following equation:*

$$(p - c)\sigma(p) = \frac{\phi^*(\bar{\theta}) - \phi^*(\underline{\theta})}{\bar{\theta} - \underline{\theta}}. \tag{3.10}$$

Specifically, when the demand function is of the additive form, the risk-free price is $p_0 = c + \frac{\phi^*(\bar{\theta}) - \phi^*(\underline{\theta})}{\bar{\theta} - \underline{\theta}}$*. When the demand function is of the multiplicative form, we have* $p_0 = \bar{p}^* = \underline{p}^*$*.*

Proposition 3.5.1 first characterizes the solution to (3.9) with the additive-multiplicative demand function. With additional information of the demand form, we can further specify the form of the risk-free price. Next, we consider two widely-used subcases of the general demand model which are the additive demand and the multiplicative demand. We identify the relationship between the optimal selling price p^* and the risk-free price p^0 under the additive demand and the multiplicative demand.

Proposition 3.5.2 *If the demand function is of the form* $D(p, \theta) = \mu(p) + \theta$ *where* $\mu(p)$ *is linearly decreasing in* p*, then the optimal selling price is smaller than the risk-free price,* $p^* \leq p^0$*. If the demand function is of the form* $D(p, \theta) = \sigma(p)\theta$*, then the optimal selling price is larger than the risk-free price,* $p^* \geq p^0$*.*

[4] An alternative setting for the risk-free pricing problem is when the firm has perfect demand information. We consider the case where demand is still ambiguous, but the firm is relieved of the risk of inventory mismatch, i.e., it does not need to pre-order and always receives ample supply. For the classical model with expected profit maximization and additive-multiplicative demand, these two settings are equivalent. In contrast, for the minimax regret model, they are different. Our analysis in Section 3.5.3 can be used to study the alternative setting under minimax regret by setting demand ambiguity to zero. The results parallel those in Section 3.5.2.

To establish that the optimal price is smaller than the risk-free price with the additive demand, we impose the assumption that the demand function is of linear form. If we relax this assumption, the relationship between the optimal price and the risk-free price becomes less clear. Specifically, we provide an example where the risk-free price is smaller than the optimal price in the additive demand case when the demand function is not of linear form.

Example 3.5.1 *Consider the following demand function, which satisfies Assumption 1:*

$$d(p) = \begin{cases} \frac{1}{8}(87 - p + \frac{9975}{p-1}), & 1 < p \le 6, \\ 559.5 - 50p, & p > 6. \end{cases} \tag{3.11}$$

We set the unit purchase cost $c = 1$ *and the uncertainty set* $[\underline{\theta}, \bar{\theta}] = [-10.5, 10.5]$. *We find that the risk-free price* $p_0 = 6$ *and the optimal price* $p^* = 6.019$.[5] *Thus,* $p_0 < p^*$ *with the demand function defined in* (3.11). *Therefore, with additive demand, if the demand function is not of linear form, the optimal price is not guaranteed to be smaller than the risk-free price.*

To further describe the intuition behind Proposition 3.5.2 and the above observation, we discuss the interplay between the regret associated with the pricing decision and that of the inventory. By Proposition 3.4.1, the nature will choose either $\underline{\theta}$ or $\bar{\theta}$ whichever can lead to the maximum regret of the firm. Let $\bar{R}$ ($\underline{R}$) denote the total regret when $\theta = \bar{\theta}$ ($\underline{\theta}$) and we use subscript p (y) to indicate the regrets that is associated with the pricing decision (inventory risk). By Proposition 3.4.2, we can substitute the formula of the optimal order quantity in (3.8) for $\theta \in \{\underline{\theta}, \bar{\theta}\}$:

$$\begin{aligned} \bar{R} &= \underbrace{\phi^*(\bar{\theta}) - (p-c)D(p,\bar{\theta})}_{\bar{R}_p} + \underbrace{(p-c)\big(D(p,\bar{\theta}) - D(p,\underline{\theta})\big) - (1-c/p)\big(\phi^*(\bar{\theta}) - \phi^*(\underline{\theta})\big)}_{\bar{R}_y}. \\ \underline{R} &= \underbrace{\phi^*(\underline{\theta}) - (p-c)D(p,\underline{\theta})}_{\underline{R}_p} + \underbrace{c/p\big(\phi^*(\bar{\theta}) - \phi^*(\underline{\theta})\big)}_{\underline{R}_y}. \end{aligned}$$

[5]For more details, please refer to the online appendix.

We first consider the additive demand and discuss how these regrets change with respect to p in the neighborhood of the risk-free price p_0. For the regrets associated with the pricing decision, it is easy to verify that $\bar{R}_p$ is decreasing in p and $\underline{R}_p$ is increasing in p in the neighborhood of p_0. Because p_0 is the price that perfectly balances $\bar{R}_p$ and $\underline{R}_p$, and we have $\bar{p}^* > \underline{p}^*$ with additive demand, which is a direct result from Lemma A.0.4 in the appendix. Second, it can be shown that $\bar{R}_y$ is increasing in p while $\underline{R}_y$ is decreasing in p. By Corollary 3.5.1, the optimal order quantity $y^*(p)$ is decreasing in p. In addition, the range of realized demand does not depend on p. Thus, the regret $\bar{R}_y$, which is associated with the shortage of inventory, is increasing in p while $\underline{R}_y$, which is associated with the overage of inventory, is decreasing in p. Finally we note that the role of the optimal price p^* is to balance $\bar{R}_p$ ($\underline{R}_p$) and $\bar{R}_y$ ($\underline{R}_y$). Specifically, when the additive demand is of linear form, we prove that the effect of p on $\underline{R}_p$ is larger than the one on $\underline{R}_y$ in the neighborhood of p_0. Consequently, the derivative of the total regret with respect to the selling price p in the neighborhood of p_0 is positive. Therefore, the optimal price is smaller than the risk-free price in the linear additive demand case. However, in general, the direction at which p^* should move from p_0 depends on the magnitudes of the effects of the unit selling price p in the neighborhood of p_0 on both R_p and R_y. In the above example, we find a setting in which the effect of the unit selling price p on R_p in the neighborhood of p_0 is smaller than the one on R_y and consequently we have $p^* > p_0$ for this specific demand function.

For the multiplicative demand, we observe that both $\bar{R}_p$ and $\underline{R}_p$ are increasing in p and both $\bar{R}_y$ and $\underline{R}_y$ are decreasing in p. As shown by Proposition 3.5.1, we have $p_0 = \bar{p}^* = \underline{p}^*$, thus, any deviation from the risk-free price will lead to increase in the regrets associated with the pricing decision but with relatively small magnitude. However, since the range of realized demand can be significantly reduced by increasing the unit selling price p, for both cases with $\bar{\theta}$ and $\underline{\theta}$, the regrets associated with inventory risk are decreasing in p. Therefore, it is straight forward to conclude that

p^* is larger than the risk-free price p_0 as the effect of the unit selling price p on $\bar{R}_y$ ($\underline{R}_y$) dominates the one on $\bar{R}_p$ ($\underline{R}_p$).

Petruzzi and Dada (1999) summarize the relationship between the optimal price and the risk-free price in the classical pricing newsvendor model. Specifically, they show that the optimal price is larger than the risk-free price in the multiplicative demand case and is smaller than the risk-free price in the additive demand case. They argue that such difference can be explained by the monotonicity of variance and coefficient of variation of the demand. Under the minimax regret framework, the firm chooses the order quantity and the selling price to balance between the regret with pricing decision and the regret with inventory risk. We can summarize the relationship between the optimal price and the risk-free price in a similar fashion as in Petruzzi and Dada (1999). The effect of the selling price on the regret associated with inventory depends on the demand form. In other words, the range of demand realization does not change as the price differs with the additive demand while the range of demand realization is decreasing in the price with the multiplicative demand. Such difference in demand form will consequently affect the interplay between regrets and thus, may drive the optimal price from the risk-free price towards opposite directions.

3.5.3 Effect of Demand Ambiguity

In this section, we investigate the effect of degree of demand ambiguity on the optimal price. Specifically, we consider the relationship between the optimal pricing decision p^* and the range of the uncertainty set, which is defined as the degree of demand ambiguity. When the demand is of additive form, the uncertainty set is constructed as $[-\delta, \delta]$. When the demand is of multiplicative form, the uncertainty set is constructed as $[\theta_0 - \delta, \theta_0 + \delta]$ where θ_0 is a constant and $0 < \delta < \theta_0$. In both cases, δ represents the degree of demand ambiguity. The following proposition characterizes the relationship between the degree of demand ambiguity and the optimal price p^*.

Proposition 3.5.3 *When the demand is of additive form and $\mu(p)$ is linearly decreasing in p, let the uncertainty set be denoted as $[-\delta, \delta]$ where $\delta > 0$, the optimal price p^* is decreasing in δ. When Assumption 1 holds and the demand is of multiplicative form and the uncertainty set is of the form $[\theta_0 - \delta, \theta_0 + \delta]$ where $0 < \delta < \theta_0$, then the optimal price p^* is increasing in δ.*

The intuition for the (opposite) effects of demand ambiguity with different demand models is similar to the discussion in Section 3.5.2 on the relationship between the optimal price and the risk-free price. However, for additive demand, if the demand function is not of linear-additive form, we show with the following counterexample that the optimal price is not necessarily decreasing in the degree of demand ambiguity.

Example 3.5.2 *Consider the following demand function,which satisfies Assumption 1:*

$$d(p) = \begin{cases} 55 - 5p, & 1 < p \leq 6, \\ 0.55 - 0.05p + \frac{123.75}{p-1}, & p > 6. \end{cases} \tag{3.12}$$

We set the unit purchase cost $c = 1$ and the uncertainty set $[\underline{\theta}, \bar{\theta}] = [-\delta, \delta]$ where δ is ranging from 0 to 5. Optimal prices under the minimax regret framework with the demand function defined in (3.12) under different degree of demand ambiguity are presented in Figure 3.1. As can be seen, the optimal price first decreases and then increases in demand ambiguity.

Demand ambiguity is analogous to demand variability under complete distributional information. Li and Atkins (2005) studied monotone optimal price with respect to demand variability ordered by the mean-preserving transformation. Xu et al. (2010) generalized the result of Li and Atkins (2005) in the notion of the convex stochastic order under one additional condition. An alternative variability stochastic order that is more general than the mean-preserving transformation family is the excess wealth order (Shaked and Shanthikumar, 1998, 2007). It can be shown that the optimal price in the classical pricing newsvendor model is decreasing (increasing) in demand variability in the notion of excess wealth order for the additive (multiplicative) case

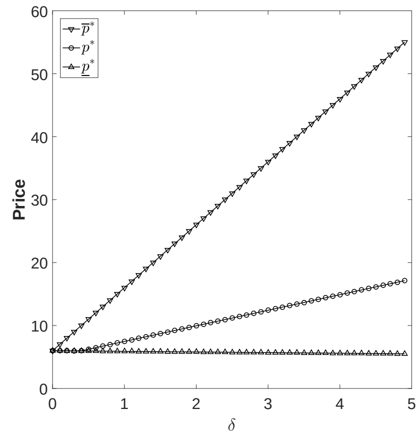

Fig. 3.1.: Optimal price with different demand ambiguity levels under a general additive demand model.

with no additional condition. Our results uncover both similarities and distinctions between the minimax regret and the classical models in how the optimal price changes with respect to demand uncertainty.

3.5.4 Effect of Unit Ordering Cost

In this section, we study the comparative statics of the optimal pricing decision p^* with different unit procurement costs. In classical pricing newsvendor models (i.e., Yao et al., 2006), it has been shown that the optimal price p^*_{NV} is monotonically increasing in the unit procurement cost c with mild assumptions. However, we show with numerical examples that the optimal price p^* may not have a monotone relationship with the unit procurement cost c. For the additive demand, as pointed out by Proposition 3.4.2, the optimal price must lie in the interval $[p_l, p_h]$. Specifically, we find that p_l is decreasing in c and p^* is equal to p_l when c is small, which are demonstrated in the left panel of Figure 3.2. Consequently, the optimal price is not necessarily increasing with respect to c. In the multiplicative demand, the optimal price does not have the monotone relationship with c either.

In terms of the optimal order quantity, the drivers behind the change of the optimal order quantity with different unit ordering costs can be decomposed into two parts. On one hand, the higher cost directly leads to a lower optimal order quantity. On the other hand, Corollary 3.5.1 points out the monotonicity between the optimal order quantity and the unit selling price. Thus, the change of the unit ordering cost can indirectly affect the optimal order quantity through the optimal selling price. Specifically, for the additive demand, as shown in Figure 3.2, the change of the optimal order quantity is mainly driven by the optimal price, which is not monotonically increasing in c. For the multiplicative demand, although the optimal price starts to decrease in c when c is relatively large, we prove in Proposition 3.5.4 that the direct effect of the unit ordering cost on y^* is larger than the indirect effect of the unit ordering cost c through the optimal price. Consequently, in the multiplicative demand case, the optimal order quantity is decreasing in the unit ordering cost, which is formally summarized in the following proposition.

Proposition 3.5.4 *When Assumption 1 holds and the demand function is of the multiplicative form, the optimal order quantity $y^*(p^*)$ is decreasing in c.*

3.5.5 Comparison with Max-min Robust Optimization

To tackle demand ambiguity, one alternative is the widely-used max-min robust optimization approach. The firm aims to maximize the worst-case profit by choosing optimally the unit selling price and the order quantity. Specifically, after the firm determines the unit selling price and the order quantity, an adversarial nature will choose $\theta \in [\underline{\theta}, \bar{\theta}]$ that minimizes the firm's profit. The max-min robust optimization problem in the context of pricing newsvendor problem is formulated as follows:

$$\max_{p,y} \left\{ \min_{\theta \in [\underline{\theta}, \bar{\theta}]} p \min\{D(p, \theta), y\} - cy \right\}.$$

For the inner minimization problem, the nature will choose $\underline{\theta}$ to minimize the firm's profit since the objective function is increasing in θ when the firm's price and inventory

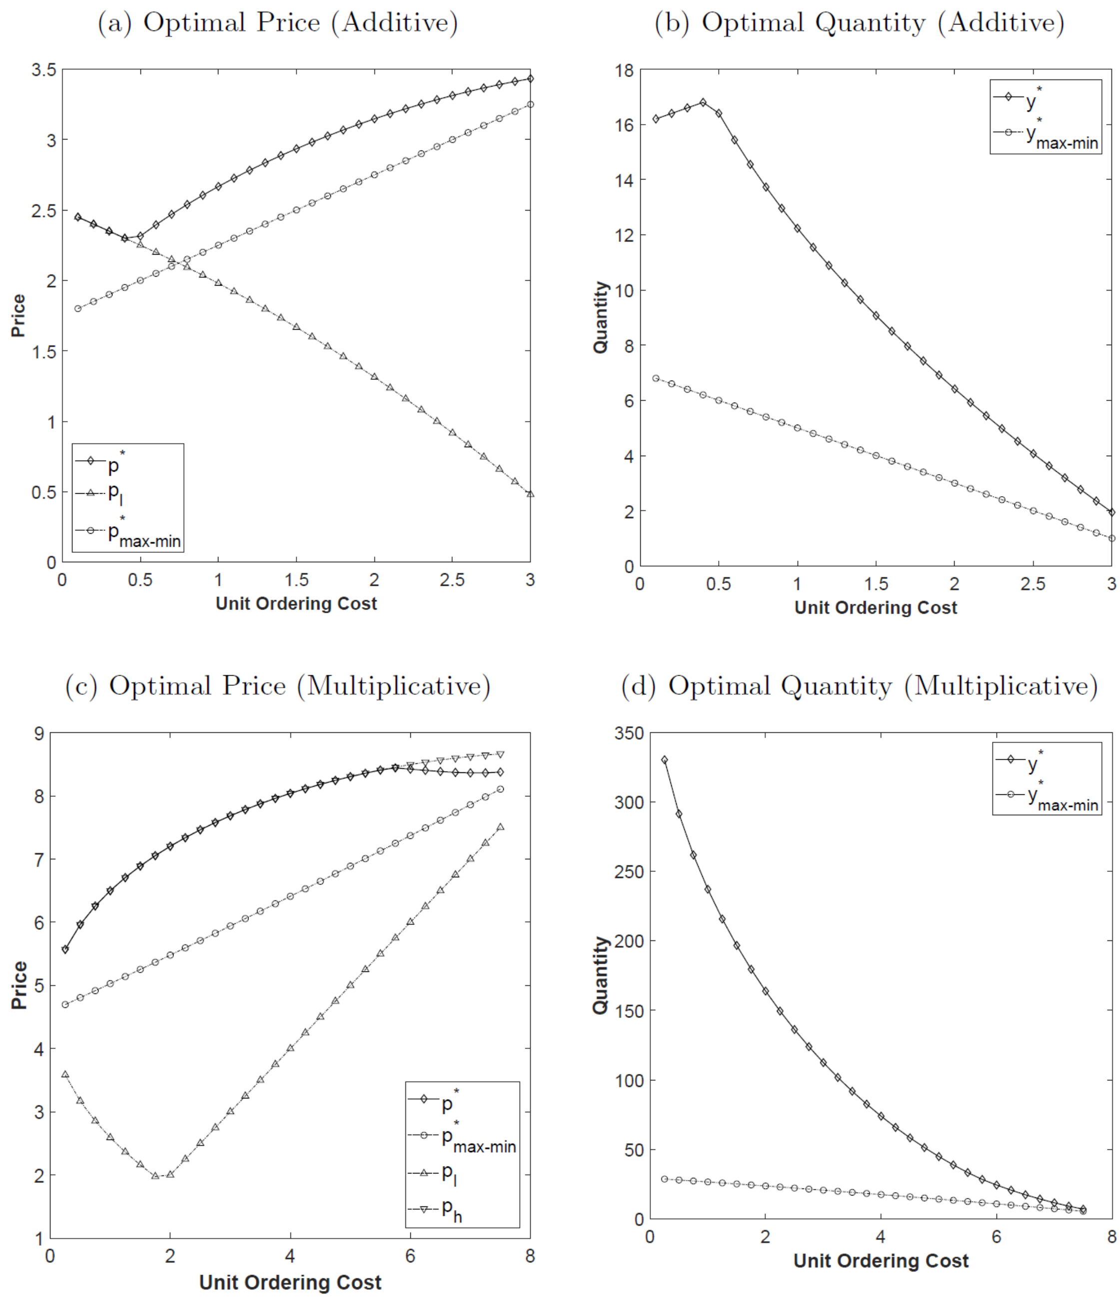

Fig. 3.2.: Optimal price and order quantity with different unit ordering cost.

decisions are fixed. Based on the response of the adversarial nature, the firm will choose $y^*(p) = D(p, \underline{\theta})$ given the unit selling price p to avoid the overage cost. Thus, the firm's optimization problem will be of the following form:

$$\max_p (p-c)D(p, \underline{\theta}).$$

The following proposition formally states the relationship between the optimal prices under two robust optimization frameworks, where $p^*_{\max-\min}$ denotes the optimal price derived from the max-min robust optimization and p^* is the optimal price derived by the minimax regret approach.

Proposition 3.5.5 *When Assumption* 1 *holds, the price derived from the minimax regret framework is higher than the price derived from the max-min robust optimization, that is,* $p^* > p^*_{\max-\min}$.

Proposition 3.5.5 indicates that the optimal price derived from the max-min robust optimization is smaller than that of the minimax regret approach as the former is known to be more conservative (see, Perakis and Roels, 2008; Caldentey et al., 2016). We conduct experiments to compare the performance of the two approaches and present results in the appendix. The numerical results show that while the expected profits using decisions made with the minimax regret approach is increasing in the mean of the random factor, the framework that adopts the max-min robust optimization cannot benefit from such changes in the underlying distribution. Moreover, we also find that the performance of the max-min approach is robust with different choices of uncertainty sets while that of the max-min robust optimization highly depends on such choices. We refer readers to the appendix for a detailed performance comparison between the two approaches.

3.6 Data-Driven Implementation

In this section, we consider a data-driven setting where the firm decides the order quantity and the selling price based on the limited historical data. The conventional approach that requires specifications of the distributional forms of the random demand is not applicable due to demand ambiguity in this data-driven setting. Even if the firm obtains the complete information of the random demand, the conventional approach may need to impose certain conditions for the tractability of the optimization problem. For example, some of classical models assume that the random factor has generalized increasing failure rate (GIFR) so that the expected profit function is unimodal or quasi-concave (i.e., Yao et al., 2006). We show with a numerical example in the appendix that the minimax regret framework does not require complete information of the random factor and can provide robust performance.

Next, we will demonstrate how to make robust decisions using the approach that we developed under the minimax regret framework based on available data. Consider the scenario where the firm observes a set of data $(p_{\mathrm{i}}, d_{\mathrm{i}}), i = 1, 2, \cdots n$, where p_{i} is the unit selling price and d_{i} is the realized demand for the i-th observation. Suppose the firm estimates the demand function using limited observations with linear regression and constructs the estimated empirical distribution of the random factor with the residuals from the linear regression. Then, the firm defines the uncertainty set by setting the upper and lower limits to the $(1 - \alpha/2)\%$ and $(\alpha/2)\%$ quantiles of the estimated empirical distribution of the random factor where $\alpha \in [0, 1]$. Specifically, we choose $\alpha = 0.4$ in this numerical study.[6] With the estimated demand function and the corresponding uncertainty set, the firm decides the order quantity and the unit selling price with the approach that we develop under the minimax regret framework.

In addition, we adopt the traditional approach for the pricing newsvendor problem as a benchmark. The firm can estimate the expected profit function using the estimated demand function and the estimated empirical distribution of the random factor. The traditional approach chooses the order quantity and the selling price that can maximize the estimated expected profit. Finally, with complete demand information, we compare the two approaches by evaluating the expected profit using decisions

[6] In the appendix, we conduct the robustness check for the minimax regret approach using different confidence levels α for constructing the uncertainty sets and the performance of minimax regret approach is robust when α is ranging from 0.2 to 0.6.

obtained by the two approaches. In the following sections, we study the effects of different factors such as sample size, demand variability, and model misspecification on the performance of the minimax regret approach.

3.6.1 Effect of Sample Size

We investigate the effect of sample size by comparing the performance of the two approaches with different sample sizes while other factors, such as demand variability, are fixed. For each sample size, we repeat the experiments for 5000 times. Table 3.2 summarizes the statistics of expected profits obtained with the two approaches. Figure 3.3 compares the boxplots of expected profits generated by the two approaches when the sample size is 4, 12, and 20 and the star indicates the average expected profits. Figure 3.4 compares the 25%, 50%, and 75% quantiles.

Since the firm assumes the correct form of the underlying demand function, more accurate estimations of the demand function and the distribution of the random factor are obtained with larger sample sizes. Therefore, with the increase in the sample size, the performance of the minimax regret approach and the traditional approach will be improved. Second, when the sample size is small to moderate, the performance of the minimax regret framework is better than that of the traditional approach, which means our approach is more robust to uncertainty arising from limited sample sizes. Furthermore, we find that the lower quantiles (e.g., 10% quantiles) of realized profits using decisions obtained by our approach are consistently larger than those associated with the traditional approach, although the traditional approach performs slightly better than the minimax regret approach for the median and higher quantiles. Thus, the robust minimax regret framework can avoid significant loss due to inaccurate estimations and demand ambiguity.

Table 3.2.: Performance comparison with different sample sizes

	Sample Size	4	6	8	10	12	14	16	18	20
Mean	minimax	21.15	22.79	23.65	24.11	24.41	24.55	24.70	24.78	24.84
	traditional	20.26	21.93	23.02	23.66	24.08	24.32	24.51	24.66	24.75
Std.	minimax	7.76	5.62	3.90	2.81	2.10	1.48	1.06	0.72	0.59
	traditional	9.16	7.18	5.33	4.00	3.09	2.37	1.77	1.29	1.06
10th percentile	minimax	16.61	20.90	22.35	23.08	23.54	23.71	23.97	24.10	24.20
	traditional	11.69	18.27	20.76	22.08	22.87	23.30	23.60	23.82	23.98

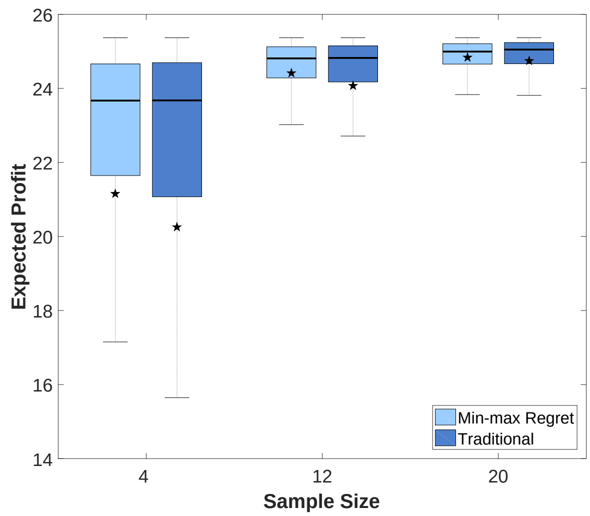

Fig. 3.3.: Performance comparison with different sample sizes

3.6.2 Effect of Demand Variability

We compare the performance of the minimax regret framework and the traditional approach with varying demand variability while the sample size is fixed to 10. The results are presented and summarized in Figure 3.5 and Table 3.3. First, the average expected profits of both approaches are decreasing in demand variability. Because higher demand variability leads to higher inventory costs as well as causes higher variance in the estimates of the demand information. Second, when the demand variability is relatively small, the performance of the minimax regret framework is close to that of the traditional approach. Third, we also observe that when the

(a) Minimax Regret Approach (b) Traditional Approach

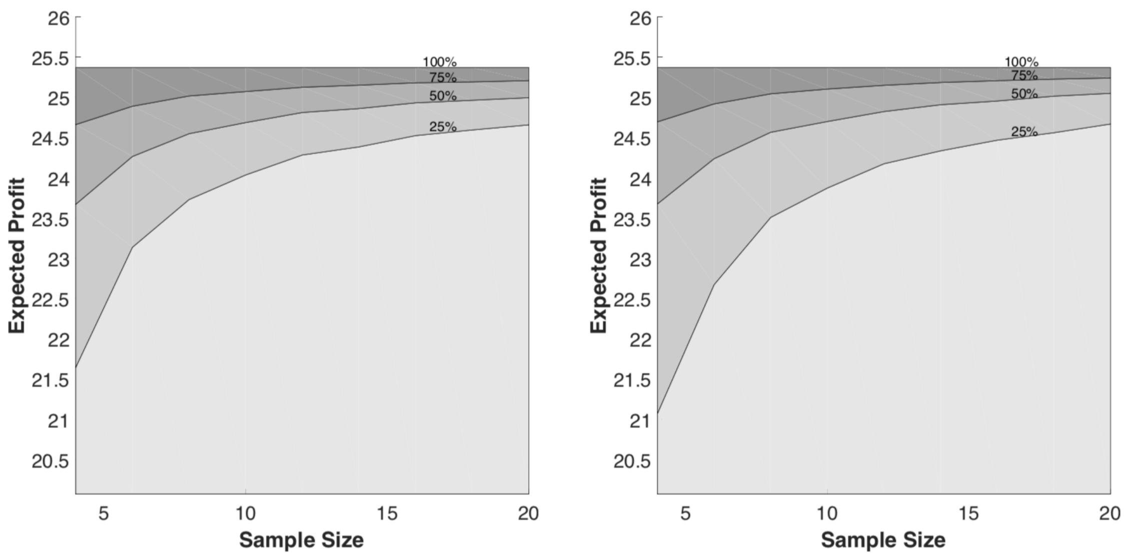

Notes. Consider the linear additive demand function, $d(p) = a - bp$, where $a = 30$, $b = 5$. The unit ordering cost $c = 1$. The random factor $\theta \sim \mathcal{N}(0, \sigma^2)$, where $\sigma = 2$. Sample size ranges from 4 to 20.

Fig. 3.4.: Comparison of 25th, 50th, and 75th percentiles with different sample sizes

demand variability is moderate to large, using the decisions made under the minimax regret framework, the firm can obtain a higher average but smaller standard deviation of the expected profits compared to those of the traditional approach. In addition, the benefit of the minimax regret approach is evident in terms of lower quantiles (e.g., 10% quantiles). This further indicates the robustness of the minimax regret approach as it is less sensitive to randomness from the sampling process.

Table 3.3.: Performance comparison with different demand variance.

Demand Std. Dev.		1	1.5	2	2.5	3	3.5	4	4.5	5	5.5	6
Mean	minimax	29.93	29.26	28.57	27.88	27.17	26.43	25.69	24.92	24.15	23.36	22.58
	traditional	29.94	29.26	28.57	27.86	27.11	26.33	25.51	24.64	23.73	22.79	21.84
Std.	minimax	0.16	0.26	0.39	0.62	0.84	1.21	1.62	2.10	2.61	3.17	3.72
	traditional	0.16	0.27	0.47	0.75	1.12	1.67	2.29	3.04	3.82	4.65	5.44
10th percentile	minimax	29.74	28.97	28.18	27.37	26.54	25.69	24.82	23.96	23.09	22.18	21.29
	traditional	29.76	28.97	28.16	27.32	26.43	25.49	24.49	23.39	22.18	20.85	19.39

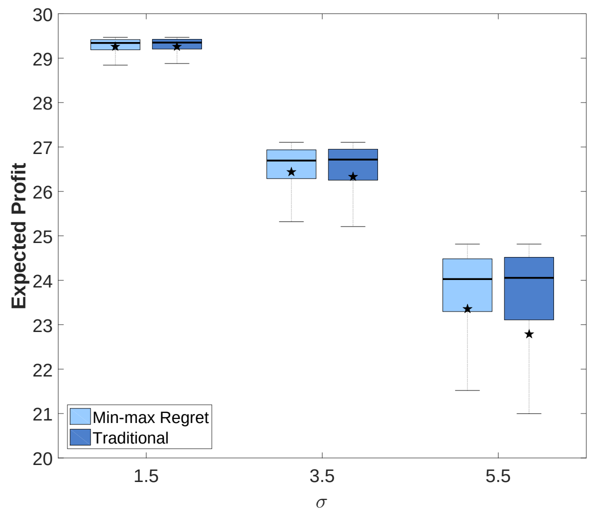

Notes. Consider the additive demand function, $d(p) = a - bp$, where $a = 30$, $b = 5$. The unit ordering cost $c = 1$. Sample size is 10. The random factor $\theta \sim \mathcal{N}(0, \sigma^2)$, where σ ranges from 1 to 6.

Fig. 3.5.: Performance comparison with different demand variability

Based on our previous observation regarding the performance of the two approaches, we propose a numerical experiment that compares the two approaches with different sample sizes and different levels of demand variability. We illustrate the numerical results in Figure 3.6. In summary, if the firm assumes the correct form of the demand function, when the sample size is large or the demand variability is small, the performance of the minimax regret framework is close to that of the traditional approach. When the sample size is small to moderate and the demand variability is moderate to large, the minimax regret framework provides robust solutions compared to the traditional approach of which the performance highly depends on the realized estimates.

3.6.3 Effect of Model Misspecification

In Section 3.6.1 and Section 3.6.2, a critical assumption is imposed that the decision maker assumes the correct demand form, which is not necessarily true in practice. Thus, in this section, we investigate the effect of model misspecification on

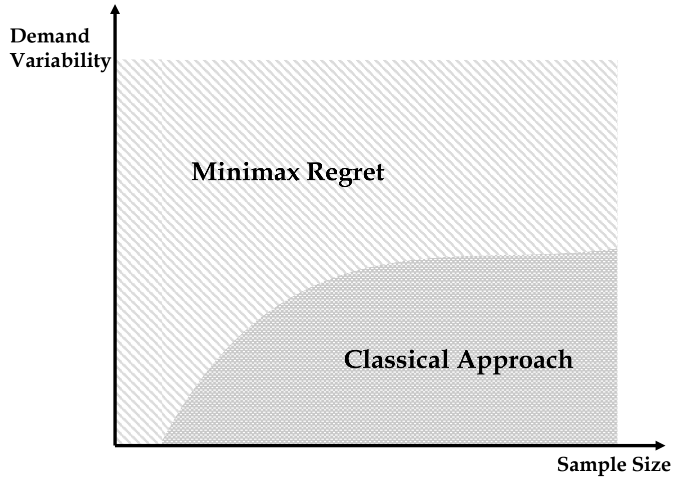

Fig. 3.6.: Performance comparison with different sample sizes and demand variability.

the performance of the minimax regret framework with the traditional approach as a benchmark. Specifically, the firm still assumes that the demand model is of linear additive form but the true demand function is actually of additive-multiplicative form, e.g., $D(p, \theta) = a - bp + \alpha^{-\beta p}\theta$. Model misspecification of this sort cannot easily be identified as the information regarding the demand form is typically unavailable in practice and the number of observations is limited. With model misspecification, we conduct the numerical experiment that compares the performance of the two approaches with different sample sizes and the comparison results are summarized in Table 3.4, Figure 3.7, and Figure 3.8. We first find that the performance of both approaches are indeed affected under model misspecification. In other words, there is a larger gap between the average expected profits of the two approaches and the true optimal profit. Moreover, we observe that the minimax regret approach derives higher average expected profits and results in smaller standard deviations of expected profits than those of the traditional one in the presence of model misspecification. Finally, based on the comparison between the lower quantiles of expected profits, e.g., the 10% quantiles, we find that the minimax regret approach can avoid significant loss due to the combined effects of model misspecification and uncertainty in observed data.

Table 3.4.: Performance comparison under model misspecification.

	Sample Size	7	8	9	10	11	12	13	14	15	16
Mean	minimax	19.69	20.11	20.41	20.89	21.19	21.41	21.52	21.59	21.75	21.89
	traditional	18.26	18.73	19.22	19.75	20.15	20.51	20.74	20.92	21.17	21.35
Std.	minimax	6.87	6.31	5.77	4.79	4.20	3.77	3.60	3.47	3.09	2.68
	traditional	9.39	8.96	8.32	7.42	6.82	6.21	5.81	5.63	5.05	4.62
10th percentile	minimax	17.28	18.18	18.61	19.34	19.86	20.19	20.29	20.37	20.59	20.79
	traditional	9.68	11.72	14.31	15.73	17.16	17.88	18.48	18.98	19.40	19.79

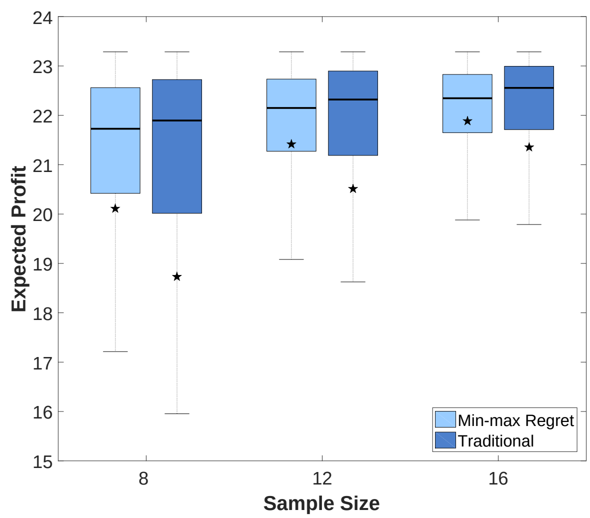

Fig. 3.7.: Performance comparison under model misspecification.

3.7 Concluding Remarks

We adopt the minimax regret framework to study the pricing newsvendor problem with ambiguous demand information. We first identify the optimal order quantity for a given unit price and show that the optimal quantity is decreasing in the price, which is not necessarily true in the classical models. By further exploring the properties of the minimax objective, we characterize the optimal price decision, which can be

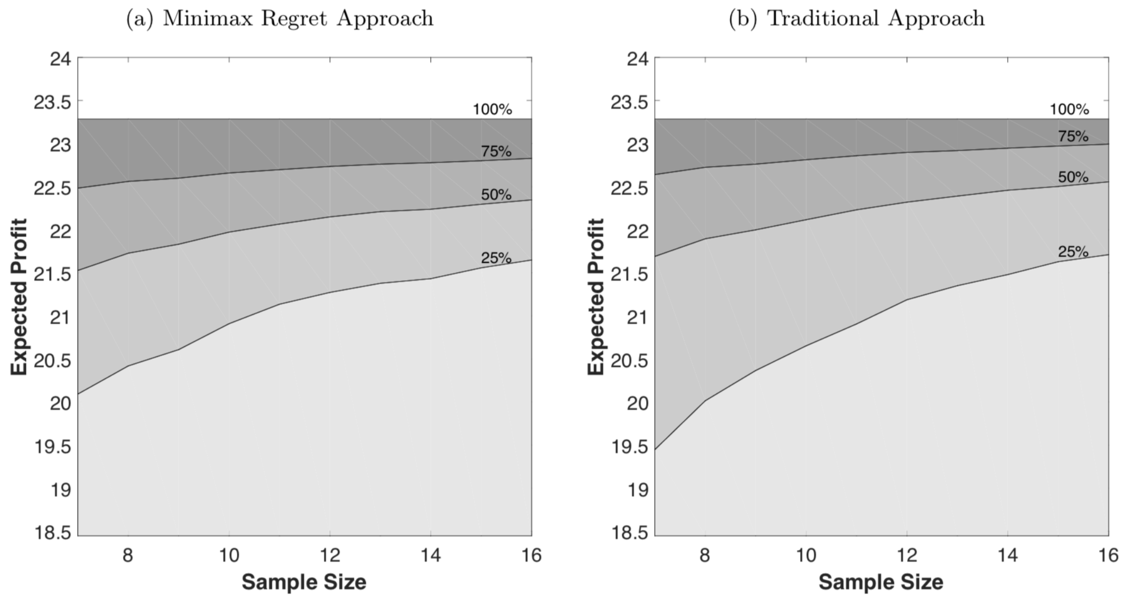

Notes. Consider an additive-multiplicative demand model, $D(p,\theta) = a - bp + \alpha e^{-\beta(p-p_m)}\theta$, where $a = 30, b = 5, p_m = 3.5, \alpha = 2, \beta = 0.1$. The unit ordering cost $c = 1$. The random factor $\theta \sim \mathcal{N}(0, \sigma^2)$ where $\sigma = 3.5$.

Fig. 3.8.: Comparison of 25th, 50th, and 75th percentiles under model misspecification.

computed efficiently. We then conduct a series of analyses to investigate the effects of important factors such as the presence of inventory risk, the unit cost, and the degree of demand ambiguity on the optimal decisions and compare these results with those in classical models. In specific, we show that the optimal price is larger than the risk-free price with multiplicative demand and smaller than the risk-free price with linear-additive demand. For general-additive demand, however, the relationship does not necessarily hold. Similarly, we show that the optimal price is increasing (decreasing) in the degree of demand ambiguity with multiplicative (linear-additive) demand, while monotonicity may not hold for general-additive demand. Further, we show that the optimal order quantity is decreasing in the unit ordering cost with multiplicative demand, but may not be monotone with additive demand. We show with counterexamples that the optimal price is not monotone in the unit ordering cost for either additive or multiplicative demand. In addition, we compare our minimax regret approach with the widely-used max-min robust optimization approach in which the firm aims to maximize the worst-case profit. We first show that the optimal price under the minimax regret framework is larger than the one using the max-min robust optimization.

We implement the minimax regret approach in a data-driven setting where historical data of price and demand are available, and compare its performance with the classical pricing newsvendor model that maximizes the expected profit. First, we find that our approach outperforms the traditional one when the sample size is small to moderate, e.g., when a new product is launched or the firm sells long-tail items. In addition, the lower quantiles of the expected profits using decisions obtained with the minimax regret approach always dominate those of the traditional approach. This suggests our approach may have significant advantage over the traditional one when the firm is risk or loss averse. Second, we find that when the sample size is fixed, the advantage of our approach over the traditional one is larger when demand variability is higher. Third, we study the effect of model misspecification where the demand is generated by an additive-multiplicative model, but misrepresented as a linear-additive model. We find that our approach consistently outperforms the traditional one with different sample sizes, and the lower quantiles of the expected profits of our approach are significantly larger than those of the traditional approach. In summary, our approach can provide robust solutions to avoid significant loss from inaccurate demand estimates and achieve substantial advantage over the traditional approach in a data-scarce or volatile demand environment. We also compare the performance of the minimax regret approach and the max-min robust approach. We find that the performance of the minimax regret approach is more robust and stable with different choices of uncertainty sets compared to the max-min robust optimization approach.

4. IMPLEMENTING ENVIRONMENTAL AND SOCIAL RESPONSIBILITY PROGRAMS IN SUPPLY NETWORKS THROUGH MULTI-UNIT BILATERAL NEGOTIATION

4.1 Synopsis

We study the implementation of an ESR program in a general supply network. The material flows in this network eventually come to a retailer, who owns the brand of the product or services offered to the eventual consumer. When the retailer initiates an ESR campaign, a successful implementation requires all firms' commitment to compliance efforts. Examples of such efforts may include use of biodegradable material for product recyclability, change of production technology for organic produce, installation of devices and procedures to ensure worker safety, and adoption of greenhouse gas emission reduction technologies, depending on the specific goal of the ESR program. In addition to interacting with the immediate tier-one suppliers, the retailer may choose to directly approach the higher-tier suppliers, or delegate the dealing with higher tiers to the tier-one suppliers. In general, we allow any downstream firm to choose between direct engagement and delegation of an upstream firm, with whom the downstream has no direct material exchange but is connected with material flows in the supply network. Thus, to describe possible ESR relationships, we extend the material supply network to the ESR network, in which links are added between firms to reflect the possible ESR relationships.

Modeling of the ESR network allows us to analyze the ESR implementation and the resulting gain allocation. An ESR implementation structure specifies the interactions among the firms that facilitate full compliance of the ESR requirements throughout the network. In particular, each chosen link in the ESR implementation structure connects a downstream firm and an upstream firm, who directly interact with each other on ESR terms. The contract ensures a payment by the downstream firm in exchange for the upstream's ESR compliance. An agreeable payment certainly depends on the relative bargaining power of the two parties. We adopt the multi-unit Nash bargaining framework (Davidson, 1988; Horn and Wolinsky, 1988) to analyze the negotiated outcome for ESR terms. We show that our bargaining solution generalizes the gain allocation derived based on individual firms' Shapley values, a commonly applied method for multi-firm cooperation. Specifically, the outcome based on Shapley value corresponds to the bargaining solution when the bargaining power is equally distributed in any downstream-upstream relationships. Clearly, the bargaining model is flexible to capture imbalanced power distribution in the supply network. More importantly, the bargaining framework specifies implementable contract terms that are highly dependent on the ESR relationships formed in the network, while the Shapley value approach is vacuous about the interactions among the firms.

The equilibrium negotiated payments lead to a gain allocation within each bargaining unit that is proportional to the trading parties' bargaining power. With this observation, the problem of identifying the retailer's most favorable implementation structure is shown to be equivalent to finding a shortest path tree (see, e.g., Lawler, 1976; Ahuja et al., 2017) in the extended ESR network. Exploring the characteristics specific to supply networks, the retailer's best implementation structure can be identified using efficient algorithms within a run time linear to the number of possible relationships. This development allows us to analyze problems with large scale networks.

In a large network, it may not be feasible for the retailer to dictate all the relationships for ESR implementation. Instead, the ESR relationships may be formed sequentially as the downstream firms approach their upstream suppliers and the suppliers in turn approach their upstream suppliers. Interestingly, we show that such a sequential relationship formation leads to an implementation structure that coincides with the retailer's most preferred one. Moreover, the analysis and results are robust when we allow each individual firm to choose among multiple levels of ESR effort, and the chosen effort levels by all firms collectively determine the overall benefit of ESR implementation.

As intuition may suggest, the retailer would intend to delegate the ESR negotiation with a higher-tier supplier when her bargaining power is low, and she would directly engage a higher-tier supplier when her bargaining power is high. In general, the choice between delegation and direct control also highly depends on the depth (i.e., number of tiers) and width (i.e., number of immediate suppliers) of the material supply network. We observe that with a fixed number of firms, the retailer tends to work directly with a larger percentage of higher-tier firms when the network is deeper and narrower, while delegating more ESR assurance responsibilities to the first-tier suppliers when the network is flatter and wider.

We further demonstrate that our model and analysis can be extended to examine the synergy among different ESR initiatives in the supply network. In this case, the retailer can enhance her bargaining position by leveraging other ESR programs in the negotiation for one program. As a result, even if the implementation of individual programs induces an economic loss, all firms in the supply network can enjoy a positive gain when multiple programs are implemented. In other words, the multi-unit bargaining over ESR implementation can induce synergy among multiple ESR programs. We also prove that the gain allocation based on Shapley value in this case corresponds to the equilibrium outcome of the negotiation game with a specific bargaining power distribution in the network.

The remainder of the chapter is organized as follows. We review the related literature and spell out our contributions in the next section. In Section 4.3, we introduce the model and present the analysis using Shapley value as a benchmark. In Section 4.4, we analyze the problem of ESR implementation using the multi-unit bargaining framework and identify the retailer's most preferred implementation structure from the ESR network. We also analyze how the ESR implementation and gain allocation depend on the key characteristics of the supply network. Several extensions of our model and analysis are discussed in Section 4.5. Section 4.6 concludes our study. Proofs of all formal results are relegated to the appendix.

4.2 Literature Review

An increasingly growing amount of research has been devoted to understanding environmental and social responsibility (ESR) in supply chains. A recent survey by Atasu (2016) provides a thorough overview of the developments on a variety of topics in this domain. Our work, in terms of application, is at the intersection of responsible sourcing and structure design for environmentally and socially responsible supply chains.

Many authors have contributed to developing sourcing and procurement strategies with ESR awareness. The existing studies have taken various angles, including supplier selection based on their compliance risks and the level of consumer awareness (Guo et al., 2016), incentive contracts to mitigate suppliers' adulteration (e.g., Babich and Tang, 2012; Chen and Lee, 2017), supplier compliance auditing (Plambeck and Taylor, 2016; Chen and Lee, 2017; Caro et al., 2018; Chen et al., 2015; Fang and Cho, 2015), policy design for supplier compliance and performance improvement (e.g., Corbett and DeCroix, 2001; Karaer et al., 2017; Agrawal and Lee, 2017; Cho et al., 2017; Nguyen et al., 2018), information sharing among suppliers (e.g., Karaer et al., 2017), and management of reverse material flows to suppliers (e.g., Ata et al., 2012). Different from this body of literature, we do not model detailed issues in the sourcing processes. Instead, we focus on the formation of ESR relationships among an existing supply base through extending the physical sourcing network.

Several studies analyze supply chain structures with the concern of ESR. For example, Guo et al. (2016) investigate the impact of supply chain structures on firms' responsible sourcing behaviors. They compare two supply chains with dedicated and shared suppliers, respectively, and find that firms have greater incentives for responsible sourcing when faced with shared suppliers. Letizia and Hendrikse (2016) analyze possible alliance between two suppliers when trading with a common downstream. They take the perspective of property ownership to examine firms' investments for ESR under different alliance structures. Chen et al. (2018) consider two buyers sourcing from a common supplier, as well as from their respective dedicated suppliers. They find that buyer coalition can mitigate the issue of reduced common supplier audit induced by buyer competition. Gui et al. (2018) study a network of firms that independently design their products for recycling. They compare the collective system, in which firms cooperate in return flows and waste processing, against the independently managed recycle process. Esenduran and Kemahlıoğlu-Ziya (2015) study a similar problem and compare collective and individual compliance schemes. They also investigate the impact of regulation on the firm's design choice and the possibility of increasing collection rates. Orsdemir et al. (2016) analyze a supply chain consisting of two competing buyers and their respective dedicated suppliers. They analyze the impact of buyer's backward integration and horizontal sourcing on ESR compliance incentive. Different from these studies, the ESR implementation structure in our model arises endogenously as decisions made by the network members. The ESR implementation structure distributes the gain from the downstream and shares the costs incurred by the upstream throughout a complex supply network.

More importantly, our approach, in terms of methodology, allows for analyzing general power distribution in the network structures. The existing studies usually examine firm interactions in supply chains by assuming exogenously given contracts (e.g., Corbett and DeCroix, 2001), Stackelberg or principle-agent setting, or cooperative game. In bilateral relationships modeled using the Stackelberg or principle-agent approach (e.g., Chen and Lee, 2017; Guo et al., 2017), one party is extremely powerful to dictate the trading terms. A cooperative game (e.g., Letizia and Hendrikse, 2016; Gui et al., 2016), in contrast, typically grants equal power to all parties involved. Our bargaining framework, in contrast, enables the flexibility of modeling varying bargaining power and enriches the understanding on how the shifts of supply chain power impact the formation of ESR implementation structure. Feng and Lu

(2013b) provide a comparison between the Stackelberg game and bargaining game in a two-tier supply chain network, pointing out the fundamental difference between the two games, while our analysis suggests that the cooperative game is a special case of the bargaining game in our context. The solution concept we adopt for our network bargaining problem is the so called Nash-Nash solution developed by Davidson (1988) and Horn and Wolinsky (1988). Several authors have applied this solution concept to supply chain settings (see, e.g., Dukes et al., 2006; Feng and Lu, 2013a; Chen et al., 2016; Chu et al., 2017).

Restricted by tractability, all the aforementioned studies in ESR, as well as most studies on cost or gain sharing in supply chains, assume two or three-tier supply chains. This is because in a complex network, identifying the core or computing the Shapley value often involve combinatorial analysis and the number of possible coalitions grows rapidly with the number of firms involved in the network (see, e.g., Granot and Sošić, 2003; Nagarajan and Bassok, 2008; Chen and Yin, 2010; Gui et al., 2016). Our multi-unit bilateral bargaining model, in contrast, overcomes the analytical difficulty. Our analysis suggests that the equilibrium negotiation relationships correspond to a shortest path tree problem (see, e.g., Lawler, 1976; Ahuja et al., 2017), which can be solved efficiently in large supply networks. More importantly, unlike the cooperative game approach, which is silent on how the cost or gain sharing can be implemented, the bargaining solution describes explicit ways of firm interactions to pass the gain and cost across the entire network.

4.3 The Problem

We consider a general supply chain led by one retailer sourcing from a network of suppliers. The retailer has already established direct or indirect trading relationships with the suppliers in the network through contracts that specify monetary exchange for goods or services among the firms within this network. The material flows and the associated contractual relationships are assumed to be given in our study. We, instead, focus on understanding how the retailer may push for the implementation of an ESR program throughout the entire supply network.

The network of material flows. All suppliers, providing materials or services that collectively lead to the eventual offering by the retailer to the end consumers, are connected with the retailer through the physical material flows. In this network, each node $i \in N = \{0, 1, 2, \ldots, n\}$ represents a firm, with node 0 being the retailer. A direct trading relationship of the physical material flow (i.e., an exchange of products or services) between firm i and firm j is represented through a directed arc, (i, j), in the network, which suggests a purchase by firm i from firm j. In other words, firm i is the downstream and firm j is the upstream in this trade. We use A to denote the set of all arcs in this network. Then, the graph (N, A) describes the material network of this supply chain. Because the material flows in the supply network should lead to the eventual product or service offered by the retailer, all arcs should belong to some directed paths originated from the retailer, node 0, and all nodes should be connected with the retailer through some directed paths.

An example of a material network is provided in Figure 4.1(i). In general, we allow multiple firms to have a common upstream or downstream. We also do not exclude the situation where firms in the same tier trade with one another (e.g., allowing firm 2 to purchase from firm 1 in Figure 4.1(i)). As our focus is on the ESR program implementation, the trading agreements along the material network concerning material exchanges are assumed to be given.

The ESR program. The retailer, who enjoys an enhanced brand image through the ESR program, acts as the initiator of the program. The success of the program requires participation of all firms in the network[1]. As we exemplified in §4.1, lack of compliance by even a single upstream supplier can jeopardize the program, preventing the supply chain from deriving the value from the ESR initiative. By participating in

[1]Other entities such as brand manufactures can also act as the initiator of an ESR program that requires the participation of all firms in their supply networks. Without loss of generality, we assume that the retailer is the initiator.

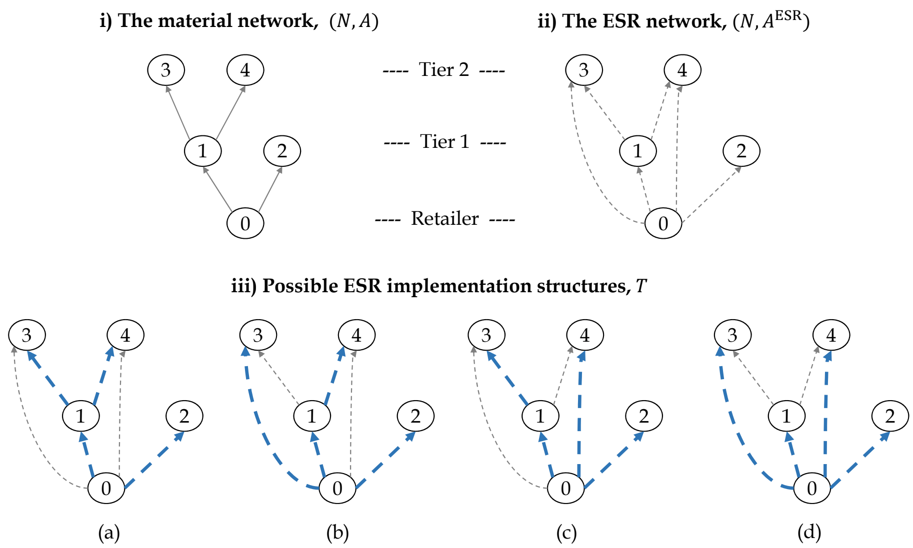

Fig. 4.1.: The material network, ESR network and ESR implementation structures.

the program, firm $i \in N$ agrees to comply with the ESR requirements through making an investment of c_i. This investment is needed for capability building or certification acquisition (Chen and Deng, 2013). If all firms commit to investing, the retailer can materialize an increased revenue, R, as a result of consumers' positive reception of the program. We shall note that, while ESR implementation requires sincere effort, the addition of revenue generated can be significant. Recent studies commissioned by Verizon and Campbell Soup show that well-implemented ESR programs can increase revenue by as much as 20 percent, command price premiums up to 20 percent, and increase customer commitment by as much as 60 percent (Hardcastle, 2015).

We use $\mathbf{x} = \{x_i : i \in N\}$ to denote the vector of investment decisions, where $x_i = 1$ indicates firm i's commitment to complying with the ESR requirements and $x_i = 0$ suggests the absence of firm i's participation. Then, the gain of the ESR program can be defined as

$$V(\mathbf{x}) = R \prod_{i \in N} x_i - \sum_{i \in N} c_i x_i. \tag{4.1}$$

It is easy to see that $V(\mathbf{x}) \leq 0$ if any $x_i = 0, i \in N$. It is natural to assume that $R > \sum_{i \in N} c_i$ so that the gain of the ESR program for the supply chain is positive to avoid the trivial case.

The ESR negotiation network. For a successful implementation of the ESR program, the gain $\Pi \equiv R - \sum_{i \in N} c_i$ should be appropriately allocated among the supply chain members to sustain the outcome. As we mentioned in §4.2, the existing literature focuses on identifying a fair allocation without paying much attention to how such an allocation can be attained in implementation. In reality, the firms involved usually go through extensive communications and negotiations to reach consensus. An agreement states the contractual parties' commitment to invest and the associated transfer payment. The parties engaged in an ESR negotiation may or may not have direct material transfers between them. Recall our examples in §1.3, Kroger interacts mostly with its immediate tier-one suppliers, from whom they directly purchase products. These suppliers, in turn, take the responsibility to ensure the compliance of their upstream suppliers. In contrast, Walmart works with every member along the supply chain to ensure ESR compliance, even though it may not purchase directly from some of the suppliers. Thus, the network of ESR implementation may not be the same as the network of physical material flows, as two firms without direct material exchange may work together and agree on an ESR contract. However, the negotiation on ESR terms would not take place between two firms that are not connected in the physical material network. For instance, a wholesaler of fresh produce would not work with a cotton grower for ESR compliance because the wholesaler does not trade the grower's outputs or products (e.g., garment) made from the grower's outputs. In the modeling language, let $P_{(i,j)}(A)$ denote the set of directed paths connecting firms i and j in the material network (N, A). Then, negotiation is allowed between i and j if and only if $P_{(i,j)}(A) \neq \emptyset$.

We extend the material flow network (N, A) to the ESR negotiation network, denoted by (N, A^{ESR}), by adding additional directed arcs, along which negotiations of ESR terms can be conducted. Specifically, $A^{\text{ESR}} = \{(i, j) : P_{(i,j)}(A) \neq \emptyset, i, j \in$

$N\} \supset A$. In other words, if and only if firms have a direct or indirect material exchange, they may work together on ESR compliance. In Figure 4.1(ii), we present the ESR network corresponding to the material network given in this example.

The ESR implementation structure. In the process of implementing the ESR program, multiple bilateral negotiations are conducted. Our choice of multi-unit bilateral bargaining, instead of collusion bargaining, is consistent with the practice in which bilateral negotiation is most commonly observed. Within each bargaining unit, the negotiation outcome depends on the parties' relative bargaining power. Specifically, in a negotiation between firm i and firm j, the bargaining power of the former is $\theta_{i,j} \in [0, 1]$, while that of the latter is $\theta_{j,i} = 1 - \theta_{i,j}$.

To ensure participation of all firms in the network, each firm must be engaged in at least one negotiation. A feasible ESR implementation structure is described by a selection of negotiation relations (represented by arcs) that connects all the firms (represented by nodes) in the network. In other words, the ESR program is implemented through a rooted tree that spans over all the nodes in the ESR network (N, A^{ESR}). We use $t_{i,j}$ to indicate whether or not the tree contains arc (i, j). In other words, $t_{i,j} = 1$ if firms i and j engage in negotiation for ESR compliance and $t_{i,j} = 0$ otherwise. The feasible set of implementation structures can be defined as

$$\mathcal{T}(A^{\text{ESR}}) = \Big\{T : t_{i,j} = 1 \text{ for } (i, j) \in T \text{ and } t_{i,j} = 0 \text{ for } (i, j) \in A^{\text{ESR}} \backslash T; \textstyle\sum_{i \in N} t_{i,j} = 1,\ j \in N \setminus \{0\}\Big\} \quad (4.2)$$

The above expression suggests that each supplier firm is working with exactly one downstream firm for ESR negotiation. It is easy to verify that this condition leads to a tree in the ESR network (N, A^{ESR}) that spans over all the nodes in the network. Moreover, in this tree, there is exactly one directed path from the retailer node reaching every supplier firm. Figure 4.1(iii) gives all possible implementation structures for the given example. The darkened arcs in each graph form a tree T in (N, A^{ESR}), which specifies the negotiation relationships for ESR implementation. In this example, structure (a) reflects the strategy of Kroger, who delegates the negotiations with the second-tier suppliers to its first-tier ones. Structure (d) represents Walmart's approach of working with every member in the network. Structures (b) and (c) consist of a combination of full delegation and direct engagement.

The sequence of events. An ESR program is implemented through the following sequence of activities.

1. (Relationship Formation) The implementation structure T of the ESR program, which consists of multiple pairs of firms engaging in bilateral bargaining, is formed in network (N, A^{ESR}) based on gain maximizing choices made by the firms involved. This structure specifies the interactions between the firms in ESR contract negotiations.

2. (Contract Negotiations) The retailer initiates the ESR program by committing to investing c_0, provided that all firms in the network participate. The negotiations of all bargaining units take place. Within each negotiation, a contract specifies the upstream's commitment of complying with ESR requirements in exchange for a transfer payment from the downstream. If all bargaining units reach agreements, an industry-wide alliance is formed.

3. (Program Execution) Given the formation of an alliance, firm i invests c_i to ensure the compliance of ESR requirements and the retailer realizes a revenue increase of R. Firms make or receive transfer payments based on the negotiated contracts.

Before analyzing our model, we provide a brief discussion on the commonly used approach that applies the concept of Shapley value. The discussion in the next subsection allows us to provide a clear comparison and demonstrate the advantage of the multi-unit bilateral bargaining framework.

4.3.1 Benchmark: The Shapley Value Based Approach

In studies of gain sharing among the multiple supply chain members, a common approach is to allocate the benefit based on the Shapley value (see., e.g. Granot and Sošić, 2003; Leng and Parlar, 2008; Letizia and Hendrikse, 2016). In this subsection, we briefly derive the solution of our problem using this approach, which serves the purpose of a comparison with our bargaining-based approach discussed in §4.4. Specifically for our model, let $\mathscr{V}(S)$ denote the gain generated by the subset $S \subset N$ of firms who agree to collude in the ESR program. Because the ESR program can realize its benefit only if a grand coalition (i.e., an industry-wide alliance) is formed, we must have

$$\mathscr{V}(S) = \begin{cases} R - \sum_{i \in N} c_i, & \text{if } S = N, \\ 0, & \text{otherwise.} \end{cases} \tag{4.3}$$

Now suppose firms in a set C have agreed to collude and firm i is not part of it. The marginal contribution of firm i joining the set C is simply

$$\mathscr{V}(C \cup \{i\}) - \mathscr{V}(C) = \begin{cases} R - \sum_{i \in N} c_i, & \text{if } C = N \setminus \{i\}, \\ 0, & \text{otherwise.} \end{cases} \tag{4.4}$$

Then, the Shapley value of firm i in some coalition S is thus computed as

$$\phi_i(S) = \sum_{C \subset S \setminus \{i\}} \frac{|C|!(|S| - |C| - 1)!}{|S|!} \big(\mathscr{V}(C \cup \{i\}) - \mathscr{V}(C)\big).$$

In computing the Shapley value, it is assumed that during the process of forming coalition S, all possible sequences of the firms joining the coalition have equal chance to occur. Thus, we average the marginal contributions of firm i over all possible permutations of firms within S. In our model, $\phi_i(S) = 0$ for any $S \subsetneq N$ and

$$\phi_i(N) = \sum_{C \subset N \setminus \{i\}} \frac{|C|!(n - |C|)!}{(n+1)!} \big(\mathscr{V}(C \cup \{i\}) - \mathscr{V}(C)\big) = \frac{1}{n+1} \Big(R - \sum_{i \in N} c_i\Big). \tag{4.5}$$

In other words, when an industry-wide alliance is formed, the allocation based on the Shapley value distributes the gain of ESR equally among all firms.

Naturally, one may argue that equally sharing the benefit derived from the ESR program may not be reasonable in practice. Some firms in the supply chain may have stronger positions than others because they possess specialized technology or have deep market penetration. Also, some firms may play more crucial roles in a specific

ESR initiative than others. All these considerations can lead to discrepancies in the gains received by participating firms. However, the Shapley value based approach treats all firms equally important, reflected by the fact that the sequences of firms joining the alliance are equally likely. Another aspect missing in the Shapley value is the means of gain distribution. It does not specify how payment transfers take place among the participating firms to achieve the resulting allocation. In our discussion below, we show that a bargaining based framework overcomes these shortcomings.

4.4 ESR Implementation through Bargaining

In this section, we discuss in detail the multi-unit bilateral bargaining framework for ESR implementation. In §§4.4.1 we introduce the Nash–Nash bargaining solution and derive the transfer payments between firms under a given implementation structure. In §§4.4.2, we analyze the ESR implementation structure through the lens of the ESR initiator, the retailer. Implications of the derived solutions are discussed in §§4.4.3.

4.4.1 The Multi-Unit Nash Bargaining Framework

Given an implementation structure, i.e., a tree T in the ESR network (N, A^{ESR}), two firms linked by an arc in the tree engage in bilateral negotiation on ESR implementation. If the negotiation is successful, the agreement specifies a transfer payment from the downstream firm to the upstream firm, as well as a commitment by the upstream on ESR investment. Because the ESR program can only be successful with the participation of all firms, the absence of any firm's commitment suggests a zero gain for all firms. Therefore, the transfer payment should depend on all firms' commitment to ESR investment $\mathbf{x} \equiv (x_0, x_1, \ldots, x_n)$. Let $w_{i,j}(\mathbf{x})$ denote a payment from firm i to j given $\mathbf{x}$ if the arc (i, j) is in tree T, i.e., $t_{i,j} = 1$. A firm in the network may make payments to and receive payments from multiple firms. If all negotiations in T are successful, firm i's net contract receivable is

$$W_i(T,\mathbf{x}) = \sum_{v\in\{u:(u,i)\in T\}} w_{v,i}(\mathbf{x}) - \sum_{v\in\{u:(i,u)\in T\}} w_{i,v}(\mathbf{x}).$$

Then, firm i's gain can be computed as

$$\pi_i(T,\mathbf{x}) = W_i(T,\mathbf{x}) + R\cdot \mathbb{I}_{\{i=0\}}\cdot \prod_{j\in N} x_j - c_i x_i, \quad i\in N. \tag{4.6}$$

Because a breakdown of one negotiation suggests a zero gain for all firms, firm i's disagreement point is $D_i^j = 0$, when negotiating with firm j.

Applying the Nash bargaining solution concept (Nash, 1950), the contract parameters, which consist of the downstream's payment $w_{i,j}(\mathbf{x})$ and the upstream's investment choice x_j, are determined through the following maximization problem: (recall $\theta_{i,j}\in[0,1]$ is firm i's bargaining power vis-à-vis firm j and $\theta_{j,i} = 1-\theta_{i,j}$.)

$$\left\{\max_{x_j\in\{0,1\},w_{i,j}(\mathbf{x})} (\pi_i(T,\mathbf{x}) - D_i^j)^{\theta_{i,j}} (\pi_j(T,\mathbf{x}) - D_j^i)^{\theta_{j,i}} :\ \pi_i(T,\mathbf{x}) \geq D_i^j \text{ and } \pi_j(T,\mathbf{x}) \geq D_j^i\right\}, \tag{4.7}$$

for $(i,j)\in T$. It is easy to check that, if the total gain of firms i and j is positive, an optimal $w_{i,j}(\mathbf{x})$ of the above problem must satisfy

$$\frac{\theta_{i,j}}{\theta_{j,i}} = \frac{W_i(T,\mathbf{x}) + \mathbb{I}_{\{i=0\}} R \prod_{\ell\in N} x_\ell - c_i x_i}{W_j(T,\mathbf{x}) + \mathbb{I}_{\{j=0\}} R \prod_{\ell\in N} x_\ell - c_j x_j} = \frac{\pi_i(T,\mathbf{x})}{\pi_j(T,\mathbf{x})}, \tag{4.8}$$

suggesting that a positive trade gain is proportionally allocated between the trading parties according to their respective bargaining power. Given that the tree T consists of n directed arcs connecting all $(n+1)$ firms, there are n bargaining units negotiating in parallel. The equilibrium contracts is the Nash equilibrium of the n Nash bargaining solutions. This solution is called the Nash-Nash solution (see, e.g., Davidson, 1988; Horn and Wolinsky, 1988; Feng and Lu, 2013b).

Theorem 4.4.1 (The Negotiation Equilibrium) *Given any feasible implementation structure T, the equilibrium contract payments satisfy*

$$w_{i,j}(\mathbf{1}) = \frac{\sum_{v\in\{j\}\cup\{u:P_{(j,u)}(T)\neq\emptyset\}} \prod_{a\in P_v(T)} \rho_a}{1+\sum_{v\in N\setminus\{0\}} \prod_{a\in P_v(T)} \rho_a}\Pi + \sum_{v\in\{j\}\cup\{u:P_{(j,u)}(T)\neq\emptyset\}} c_v,$$

and $w_{i,j}(\mathbf{x}) = 0$, *if* $\mathbf{x} \neq \mathbf{1}$, *for all* i, j *such that* $t_{i,j} = 1$. *In the Pareto-dominant equilibrium, all firms invest to comply with the ESR requirements and share a positive portion of the ESR gain* $\Pi \equiv R - \sum_{i \in N} c_i > 0$.

We use the example described in Figure 4.1 to demonstrate the equilibrium outcome of ESR negotiation. This network consists of five firms and there are four feasible implementation structures as depicted in Figure 4.1. Assume that in any negotiation, the bargaining power of the downstream firm is θ and that of the upstream firm is $(1-\theta)$. Then, we can compute each firm's share of the ESR gain by solving the system of equations (4.8), which leads to the results in Table 4.1. It is easy to analyze the allocations in the table to find that the more the retailer delegates the negotiations with the second-tier suppliers to the first tier, the less (more) the retailer gains from ESR when the downstream is stronger (weaker) than the upstream, i.e., $\theta > (<) 0.5$. Also, a second tier supplier gains more (less) from ESR when negotiating directly with its immediate downstream than with the retailer when the negotiation power is more (less) concentrated toward the downstream, i.e., $\theta > (<) 0.5$.

Table 4.1.: The negotiated gain allocation for the example in Figure 4.1 when $\theta_{i,j} = \theta$ for $(i,j) \in T$

	Full Delegation	Partial Delegation		Direct Engagement
	Tree (a)	*Tree* (b)	*Tree* (c)	*Tree* (d)
$\pi_0(T;\mathbf{1})/\Pi$	$\frac{\theta^2}{2-2\theta+\theta^2}$	$\frac{\theta^2}{1+\theta-\theta^2}$	$\frac{\theta^2}{1+\theta-\theta^2}$	$\frac{\theta}{4-3\theta}$
$\pi_1(T;\mathbf{1})/\Pi$	$\frac{\theta(1-\theta)}{2-2\theta+\theta^2}$	$\frac{\theta(1-\theta)}{1+\theta-\theta^2}$	$\frac{\theta(1-\theta)}{1+\theta-\theta^2}$	$\frac{1-\theta}{4-3\theta}$
$\pi_2(T;\mathbf{1})/\Pi$	$\frac{\theta(1-\theta)}{2-2\theta+\theta^2}$	$\frac{\theta(1-\theta)}{1+\theta-\theta^2}$	$\frac{\theta(1-\theta)}{1+\theta-\theta^2}$	$\frac{1-\theta}{4-3\theta}$
$\pi_3(T;\mathbf{1})/\Pi$	$\frac{(1-\theta)^2}{2-2\theta+\theta^2}$	$\frac{\theta(1-\theta)}{1+\theta-\theta^2}$	$\frac{(1-\theta)^2}{1+\theta-\theta^2}$	$\frac{1-\theta}{4-3\theta}$
$\pi_4(T;\mathbf{1})/\Pi$	$\frac{(1-\theta)^2}{2-2\theta+\theta^2}$	$\frac{(1-\theta)^2}{1+\theta-\theta^2}$	$\frac{\theta(1-\theta)}{1+\theta-\theta^2}$	$\frac{1-\theta}{4-3\theta}$

We note in Table 4.1 that if the bargaining power is balanced throughout the network, i.e., $\theta_{i,j} = \theta = 0.5$ for all (i,j), then all firms obtain an equal share (i.e., 20%) of the ESR gain under any implementation structure. This outcome coincides with that derived under the Shapley value based approach; recall equation (4.5). In the next corollary, we formalize this observation.

Corollary 4.4.1 (Connection to the Shapley Value Based Approach) *If firms are equally powerful in any ESR negotiations (i.e., $\theta_{i,j} = \theta_{j,i} = 0.5$ for any $(i, j) \in A^{\text{ESR}}$), then regardless of the implementation structure T, the equilibrium negotiation outcome from the multi-unit bargaining coincides with the allocation based on the Shapley value.*

As suggested from the above corollary, the model based on the Shapley value is a very special case of the multi-unit bargaining framework. This observation is in line with the conclusion made by Gul (1989), who shows that the subgame perfect equilibrium of a bargaining game with random matching of trading parties in a network converges to the solution based on Shapley value when the time between negotiation rounds gets close to zero. Different from the Shapley value approach, our multi-unit bargaining game with given trading relationships corresponds to the limit of a simultaneous alternating offer process (Davidson, 1988). Compared with the random matching process, the alternating offer process captures possibly imbalanced power distribution.

Another important implication of Corollary 4.4.1 is that the solution based on Shapley value is silent on the implementation structure in terms of how the ESR gain is passed along among the participating firms, whereas the multi-unit bargaining outcome is specific to the implementation structure in general. This is evident from Table 4.1 that the gain allocation among firms is generally different under different implementation structures. Our next task is to determine an implementation structure.

4.4.2 The Retailer Preferred Implementation Structure

In most of the ESR programs, the retailer owns the brand of the product or service provided to the end consumer. In the examples quoted in §4.1, the ESR program is pushed by a powerful retailer throughout its supply chain. It is thus natural that the retailer dictates the implementation structure. In this subsection, we use the lens of the retailer in our model to understand the determination of an implementation structure. Later in §§4.5.2, we discuss alternative means for this choice.

To see how the implementation structure T affects the gain allocation in the network, we define $P_i(T)$ as the path in tree T that starts from the retailer and ends at firm i. Based on the description in §4.3, this path is unique in any feasible T. Suppose there are $k(\geq 0)$ firms in between the retailer and firm i in tree T, then $P_i(T)$ consists of a set of $(k+1)$ directed arcs, $\{(0,i_1),(i_1,i_2),\ldots,(i_k,i)\}$. For $(i,j)\in A^{\mathrm{ESR}}$, let

$$\rho_{i,j} = \theta_{j,i}/\theta_{i,j}$$

denote the upstream's bargaining power relative to the downstream's in the ESR network. From (4.8), this quantity is also the ratio of gain obtained by the upstream to that by the downstream. We can derive

$$\pi_i(T,\mathbf{1}) = \pi_{i_k}(T,\mathbf{1})\cdot\rho_{i_k,i} = \pi_{i_{k-1}}(T,\mathbf{1})\cdot\rho_{i_{k-1},i_k}\rho_{i_k,i} = \cdots = \pi_0(T,\mathbf{1})\prod_{a\in P_i(T)}\rho_a. \quad (4.9)$$

In other words, any firm's profit can be expressed as that of the retailer's based on the equilibrium negotiation outcome. Because the ESR gain of the entire network is $\Pi = R - \sum_{i\in N} c_i$, we deduce

$$\pi_0(T,\mathbf{1}) = \Big(1+\sum_{v\in N\setminus\{0\}}\prod_{a\in P_v(T)}\rho_a\Big)^{-1}\Pi, \quad (4.10)$$

$$\pi_i(T,\mathbf{1}) = \frac{\prod_{a\in P_i(T)}\rho_a}{1+\sum_{v\in N\setminus\{0\}}\prod_{a\in P_v(T)}\rho_a}\Pi,\ i\in N\setminus\{0\}. \quad (4.11)$$

The retailer would prefer an implementation structure T that maximizes $\pi_0(T,\mathbf{1})$, which reduces to minimizing $\sum_{v\in N\setminus\{0\}}\prod_{a\in P_v(T)}\rho_a$ over all possible T within the feasible set defined in (4.2). Thus, the retailer's choice of implementation structure is a spanning tree that solves the following optimization problem:

$$\min_{T \in \mathcal{T}(A^{\text{ESR}})} \left\{ \Gamma(T) \equiv \sum_{v \in N \setminus \{0\}} \prod_{a \in P_v(T)} \rho_a \right\}. \tag{4.12}$$

The objective in this optimization problem reflects the magnitude of the ESR gain against the portion retained by the retailer.

For the example in Figure 4.1, the feasible set $\mathcal{T}(A^{\text{ESR}})$ consists of four trees. Comparing the gain allocations in Table 4.1, it is immediate that the retailer would prefer direct engagement (full delegation) when the downstream is more (less) powerful than the upstream in negotiation, i.e., $\theta > (<) 0.5$.

In general, there can be a large number of feasible implementation structures. If one were to solve the minimization problem defined in (4.12) as a mathematical program, one needs to deal with integer decision variables and a nonlinear objective. We, instead, explore the ESR network and solve the problem using the combinatorial optimization approach. As indicated from (4.12), our goal is to find a tree T that reaches each node in the network. The "cost" of reaching each node is the product of the relative bargaining powers (i.e., $\rho_{i,j}$) along the path from the source (i.e., the retailer) to that node. And the objective is to minimize the total cost of reaching all the nodes.

It turns out that a solution of (4.12) corresponds to a *shortest path tree* (see, e.g., Lawler, 1976; Ahuja et al., 2017) in the network with appropriately defined arc costs. However, the total cost of the shortest path tree does not equal to the objective in (4.12). This is because, unlike in the classical shortest path tree problem in which the objective is the sum of the path costs, the objective in (4.12) involves both multiplicative and additive operations.

Theorem 4.4.2 (The Optimal Implementation Structure) *If* $T^* \in \mathcal{T}(A^{\text{ESR}})$ *is a shortest path tree in network* (N, A^{ESR}) *with arc costs* $d_{i,j} = \log \rho_{i,j}$, *then* T^* *minimizes* $\Gamma(T)$ *over* $T \in \mathcal{T}(A^{\text{ESR}})$ *in (4.12).*

The shortest path tree problem can be solved efficiently for general networks that are acyclic. For example, when firm 2 also sources from firm 1 in Figure 4.1(i), the material network is acyclic but is not a tree. If, in addition, firm 4 sources from firm 2, then a cycle $1 \to 4 \to 2 \to 1$, is formed in the material network. Such a situation is uncommon in supply chains, though efficient algorithms exist to solve such problems as long as the length of the cycle is nonnegative (i.e., $\rho_{1,4}\rho_{4,2}\rho_{2,1} > 1$). The detailed discussion on the shortest path tree problem can be found in Appendix A.1. With the result in Theorem 4.4.2, one can analyze the implementation structure in a large scale supply network.

4.4.3 The Extent of Delegation for ESR Implementation

In this subsection, we discuss the retailer's incentive for delegation, as opposed to direct engagement, in ESR implementation. To derive clear insights, we focus on material supply networks that exhibit tree structures. In this case, each supplier sells material to a unique downstream; recall our discussion in §4.3. Thus, the suppliers can be grouped into tiers (with retailer in tier 0) and the number of tiers represents the *depth* of the supply chain.

From (4.10), the distribution of bargaining power in the ESR network plays a critical role in determining the implementation structure. To obtain some intuition, we first derive some formal results for several special settings and then test the general case using an extensive numerical analysis.

Theorem 4.4.3 *Suppose the supply network* (N, A) *is a tree with one retailer (tier 0) and* d *tiers of suppliers, and a firm in tier* k *has bargaining power* $\theta_k, k \in \{0, 1, \ldots, d-1\}$*, when negotiating with any of its upstream firms.*

i) If there exists a $k \in \{1, 2, \ldots, d-1\}$ *such that* $\rho_k \equiv (1-\theta_k)/\theta_k < 1$*, then the retailer does not interact directly with suppliers in tiers* $\{k+1, k+2, \ldots, d\}$ *for ESR implementation.*

ii) If there exists a $k \in \{1, 2, \ldots, d-1\}$ such that $\rho_{k_0} = (1-\theta_{k_0})/\theta_{k_0} > 1$ for all $k_0 \in \{1, 2, \ldots, k-1\}$, then the retailer directly negotiates with suppliers in tiers $\{1, 2, \ldots, k\}$ for ESR implementation.

Theorem 4.4.3 suggests that the retailer tends to delegate the ESR negotiation with a supplier if there is a powerful firm in between the retailer and that supplier along the path of the material flow. If, however, all firms in between the retailer and that supplier are weak in negotiation, the retailer tends to deal with the supplier directly.

Corollary 4.4.2 *For the network described in Theorem 4.4.3,*

i) if θ_k is decreasing in $k \in \{1, 2, \ldots, d\}$, then there exists a k_D such that the retailer delegates the negotiations with tier-k suppliers to their immediate downstream firms for $k \leq k_D$, and to those in tier $(k_D - 1)$ for $k > k_D$.

ii) if θ_k is increasing in $k \in \{1, 2, \ldots, d\}$, then there exists a k_I such that the retailer negotiates directly with tier-k suppliers for $k \leq k_I$ and delegates the negotiation with tier-k suppliers to their immediate downstream firms in tier $(k-1)$ for $k > k_I$.

Corollary 1 explores the implementation structure under monotonically distributed bargaining power. When the suppliers' bargaining power vis-à-vis their respective downstream firms is decreasing along the tiers, the retailer would only deal with tier-one suppliers directly, while delegating ESR assurance of higher tiers. This delegation, interestingly, is to the immediate downstream firms unless the immediate downstream is weak in negotiation (note that $\rho_k > 1$ for any $k > k_D$ from the proof of the corollary). When the suppliers' bargaining power is increasing along the tiers, the retailer would choose to directly deal with suppliers in lower tiers, while leaving the higher tiers to their immediate downstream firms.

We shall remark that, though the results in Theorem 4.4.3 and Corollary 4.4.2 are stated under the assumption that firms in each tier have the same bargaining power, this assumption can be relaxed. Given that the supply network is a tree, there is a unique path from the retailer to each supplier firm. As long as the conditions hold along a path, the corresponding results hold for that path.

In general, the bargaining power distribution can have many variations in view of the large number of possible relationships in the ESR network. To understand how the characteristics of the supply network affect the ESR implementation with general bargaining power distributions, we perform an extensive simulation analysis. In our numerical test, we consider tree-type supply networks with one retailer and n suppliers. The bargaining power in the ESR negotiation is generated as random samples of a random variable $\Theta \in [0, 1]$, which follows some Beta distribution. By varying the shape parameters of the Beta distribution, we can capture different concentrations of bargaining power. While we report the results for Beta$(2, 2)$, changing the distribution parameters does not alter the main insights obtained. For each metric of interest, we simulate 1000 instances of the problem and report the average value.

In Table 4.2, the variable l indicates the number of suppliers trading materials with each downstream firm. When fixing the number of firms n, an increased l implies an increased number of suppliers in each tier and a decreased number of tiers d. Thus, l measures the width of the supply chain. We observe that the width/depth of the supply network has a major impact on the retailer's gain as well as the implementation structure. When the network becomes deeper (i.e., l becomes smaller), the retailer's share of ESR gain increases rapidly. At the same time, the retailer significantly reduces the number of suppliers that it directly interacts with. This is because, when the supply network becomes deeper, there are more possible paths in the ESR network that lead from the retailer to a specific firm. As a result, the retailer has more choices when determining the best implementation structure, leading to an increased retailer's gain. Interestingly, when the size of the network n becomes larger, the retailer tends to delegate more. However, the retailer's share on ESR gain does not seem to be sensitive to the network size.

Table 4.2.: The effect of width/depth of the supply network

	Retailer's % share of ESR gain						% of higher-tier suppliers engaged by retailer				
	$l=1$	$l=2$	$l=3$	$l=4$	$l=5$	$l=n$	$l=1$	$l=2$	$l=3$	$l=4$	$l=5$
$n=100$	31.93	5.59	2.88	2.09	1.80	0.52	0.98	8.98	19.13	27.35	31.56
$n=200$	32.21	4.60	2.05	1.45	1.14	0.26	0.48	5.32	13.21	19.36	24.66
$n=300$	31.93	4.10	1.68	1.11	0.93	0.17	0.33	3.76	10.67	16.86	20.57
$n=400$	32.28	3.82	1.44	0.95	0.76	0.13	0.25	2.90	9.04	14.39	18.63
$n=500$	31.76	3.48	1.31	0.85	0.65	0.10	0.20	2.46	7.78	12.86	17.24

Notes. The left panel gives the retailer's % gain from ESR, and the right panel provides the average % of $(n-l)$ higher-tier suppliers (i.e., excluding tier-1 suppliers) that are directly negotiated by the retailer under the optimal implementation structure. The width measure l is the number of suppliers of each downstream firm, and therefore, the depth of the network is $d = \min_{d^0}\left\{\sum_{k=1}^{d^0} l^k \geq n\right\}$. The bargaining power is randomly generated from $\Theta \sim \text{Beta}(2,2)$. The reported number is the average of 1000 instances.

As discussed earlier, most ESR programs are implemented through a *full delegation* structure. Under such structures, the retailers tend to engage only the first-tier suppliers while delegating the responsibility of ensuring ESR compliance of upstream firms to their associated downstream firms. Walmart is one of the few firms known for *direct engagement*, working directly with every supplier along the material flows. For example, they went directly to growers in Mexico in their produce supply chain and garment factories in Bangladesh in their apparel supply chain. In Table 4.3, we evaluate the retailer's gain under the optimal implementation structure against those under full delegation and direct engagement. We observe from the upper panel of the table that the optimality gaps (measured by the percentage of profit difference compared with the optimal structure) are larger as the network becomes more complex. Moreover, the optimality gap of direct engagement is more sensitive to the depth than to the width of the network, while that of full delegation is more sensitive to the width than to the depth. These observations suggest the importance for the retailer to carefully plan the implementation structure when sourcing from an extended supply chain.

Comparing between the cases of direct engagement and full delegation, we also observe that full delegation results in a larger optimality gap than direct engagement does when the supply network is wider and deeper (marked by the dividing line in the table). This observation supports Walmart's practices as it manages a very complex supply chain.

Table 4.3.: The benefit of optimizing the implementation structure

	Optimality gap of direct engagement (%)					Optimality gap of full delegation (%)				
	$l=1$	$l=2$	$l=3$	$l=4$	$l=5$	$l=1$	$l=2$	$l=3$	$l=4$	$l=5$
$d=2$	21.72	38.60	45.98	48.54	49.78	12.33	31.48	42.15	51.29	57.00
$d=3$	41.79	64.03	70.28	71.77	73.86	24.76	61.16	77.01	83.74	86.28
$d=4$	54.31	79.08	83.58	85.06	85.89	32.16	79.06	91.56	94.61	96.09
$d=5$	61.66	87.06	90.51	91.69	92.33	40.57	90.95	97.36	98.51	98.85
$d=6$	68.17	92.31	94.60	95.40	95.67	44.20	96.70	99.19	99.55	99.67

Optimal Implementation Structure

	Percentage of direct engagement (%)					Longest implementation path				
	$l=1$	$l=2$	$l=3$	$l=4$	$l=5$	$l=1$	$l=2$	$l=3$	$l=4$	$l=5$
$d=2$	50.09	50.10	49.89	49.93	50.01	1.50	1.92	1.99	2.00	2.00
$d=3$	38.25	33.99	31.07	30.13	29.14	2.06	2.79	2.98	3.00	3.00
$d=4$	27.97	20.65	17.60	16.41	15.50	2.66	3.70	3.98	4.00	4.00
$d=5$	24.00	13.21	9.66	8.41	7.60	3.16	4.63	4.97	5.00	5.00
$d=6$	19.46	7.56	5.00	3.96	3.61	3.73	5.61	5.96	6.00	6.00

Notes. The upper panel computes the percentage profit gap from that under the optimal structure (i.e., $(\pi_0(T^*, \mathbf{1}) - \pi_0(T, \mathbf{1}))/\pi_0(T^*, \mathbf{1}) \times 100$ for T being the structure under direct engagement or full delegation). The dividing line indicates whether direct engagement or full delegation results in the larger optimality gap. The left-bottom panel computes the average % of $(n-l)$ higher-tier suppliers that are directly engaged by the retailer under the optimal structure T^* and the right-bottom panel computes the average length (i.e., number of arcs) of the longest path in the optimal structure T^*. The measure l is the number of suppliers of each downstream firm and d is the depth of the network. Thus, the total number of suppliers (excluding retailer) is $n = \sum_{k=1}^{d} l^k$. The bargaining power is randomly generated from $\Theta \sim \text{Beta}(2,2)$ and the reported number is the average of 1000 instances.

Consistent with the observations from Table 4.2, we find from the left-bottom panel of Table 4.3 that the retailer tends to directly engage less portion of suppliers when the network gets wider. In this case, downstream firms are more likely to engage only their immediate upstreams, as suggested from the right-bottom panel of the table that the average length of the longest paths in the optimal implementation structures becomes longer.

4.5 Extensions

In this section, we discuss extensions and alternatives of the model analyzed in the previous section.

4.5.1 Multiple Effort Levels for ESR Compliance

The model analyzed in §4.4 assumes that there is a fixed investment required from each firm to implement the ESR program. We extend our analysis to the case where the firms may choose different effort levels to participate and the benefit generated by the ESR program depends on the collective efforts. Specifically, when firm i exerts an effort of $x_i \in [0,1]$, it results in an investment of $c_i(x_i)$, which is a nonnegative and increasing convex function. The revenue generated by the ESR program is $R(\mathbf{x})$, where $\mathbf{x} = (x_0, x_1, \ldots, x_n) \in [0,1]^{n+1}$. We assume that $R(\mathbf{x}) = 0$ if $\prod_{i \in N} x_i = 0$ and $R(\mathbf{x}) > 0$ is increasing and concave over $\mathbf{x} \in [0,1]^{n+1}$. Note that it is without loss of generality to assume that the highest feasible effort level is one because we can always scale the effort levels in the investment functions and the revenue function. We shall also assume that there exists a feasible vector of effort levels such that an implementation of the ESR program is beneficial, i.e., $\Pi \equiv \max_{\mathbf{x} \in [0,1]^n} \{R(\mathbf{x}) - \sum_{i \in N} c_i(x_i)\} > 0$. Without loss of generality, we assume that $R(\mathbf{x}) - \sum_{i \in N} c_i(x_i)$ is strictly concave such that there exists a unique optimal vector of investment levels that maximizes the net benefit of the entire ESR program.

Since the retailer is the initiator of the ESR program, when deciding the implementation structure, the retailer also commits its effort level x_0 to implement the ESR program. The subsequent contract negotiations should involve the suppliers' effort levels. Thus, the bargaining problem in (4.7) between firm i and firm j now becomes

$$\left\{ \max_{w_{i,j(\mathbf{x})}, x_j \in [0,1]} (\pi_i(T,\mathbf{x}) - D_i^j)^{\theta_{i,j}} (\pi_j(T,\mathbf{x}) - D_j^i)^{\theta_{j,i}} :\ \pi_i(T,\mathbf{x}) \geq D_i^j \text{ and } \pi_j(T,\mathbf{x}) \geq D_j^i \right\},$$

with $c_i x_i$ replaced by $c_i(x_i)$ and $R \cdot \mathbb{I}_{\{i=0\}} \cdot \prod_{j \in N} x_j$ replaced by $R(\mathbf{x}) \cdot \mathbb{I}_{\{i=0\}}$ in firm i's profit defined in (4.6).

Theorem 4.5.1 (Multiple ESR Compliance Effort Levels) *Given any feasible implementation structure $T \in \mathcal{T}(A^{\mathrm{ESR}})$, the equilibrium contract payments satisfy*

$$w_{i,j}(\mathbf{x}) = \frac{\sum_{v \in \{j\} \cup \{u : P_{(j,u)}(T) \neq \emptyset\}} \prod_{a \in P_v(T)} \rho_a}{1 + \sum_{v \in N \setminus \{0\}} \prod_{a \in P_v(T)} \rho_a} \left(R(\mathbf{x}) - \sum_{v \in N} c_v(x_v) \right) + \sum_{v \in \{j\} \cup \{u : P_{(j,u)}(T) \neq \emptyset\}} c_v(x_v),$$

if $R(\mathbf{x}) - \sum_{v \in N} c_v(x_v) > 0$, and $w_{i,j}(\mathbf{x}) = 0$, otherwise, for all i, j such that $t_{i,j} = 1$. In the Pareto-dominant equilibrium, all firms invest to comply with ESR requirements at the centralized optimal investment level $\mathbf{x}^* = \arg\max\{R(\mathbf{x}) - \sum_{i \in N} c_i(x_i) : \mathbf{x} \in [0,1]^{n+1}\}$ and share a positive portion of the ESR gain $\Pi = R(\mathbf{x}^*) - \sum_{i \in N} c_i(x_i^*)$.

Theorem 4.5.1 suggests that the multi-unit bargaining framework leads to efficient effort exertion by all firms in the sense that the gain of the ESR program is maximized. With the result in Theorem 4.5.1, one can then apply Theorem 4.4.2 to find the optimal implementation structure.

4.5.2 Sequential Formation of ESR Implementation Structure

In the model analyzed in §4.4, the ESR implementation structure is dictated by the retailer. In this subsection, we consider an alternative process of implementation, in which the ESR relationships are formed sequentially. Specifically, the retailer initiates the ESR program by approaching some of the suppliers in the network to partner with. The chosen firms, in turn, take the responsibility to ensure the participation of all their upstream suppliers by appropriately choosing their ESR partners. An implementation structure is formed until all firms in the network are connected through ESR partnerships. From the retailer's perspective, such a sequential process can be attractive, as environmental and social responsibility management in an extensive supply chain requires significant governance (see the discussions in, e.g., Huang et al., 2018).

The sequence of events corresponding to a sequentially formed implementation is elaborated below.

1. (Relationship Formation) The ESR implementation structure is formed through a sequential process.

   1.1 In the 1st round of relationship development, the retailer chooses a subset $N_1 \subset N$ to interact directly on ESR terms.

   1.2 In the kth round of relationship development, firms in set N_{k-1} choose their respective upstream negotiation parties within the set $N \setminus \cup_{\ell=1}^{k-1} N_\ell$. Let N_k denote the set of chosen firms in this round. The formation of ESR relationships within the kth round can take one of the two ways.

      (a) [Simultaneous] All firms in set N_{k-1} simultaneously approach their preferred upstream trading parties. When firm $i \in N_k$ is approached by two or more firms in N_{k-1}, firm i can choose at most one to form a trading relationship. If firm i declines to partner with all, then the ESR program fails and the game ends.

      (b) [Sequential] Firms in set N_{k-1} sequentially approach their preferred upstream parties to form trading relationships. We assume that this sequence is randomly determined (i.e., each of the $|N_{k-1}|!$ permutations occurs with equal chance). If a chosen firm in N_k refuses to form a relationship, then the ESR program fails and the game ends.

   If $N = \cup_{\ell=1}^{k} N_\ell$, an implementation structure is formed; otherwise, the game enters the $(k+1)$st round of relationship development.

2. (Contract Negotiations) The retailer initiates the ESR program by committing to investing c_0, provided that all firms in the network participate. The negotiations of all bargaining units take place. Within each negotiation, a contract

specifies the upstream's commitment of complying with ESR requirements in exchange for a transfer payment from the downstream. If all bargaining units reach agreements, an industry-wide alliance is formed.

3. (Program Execution) Given the formation of an alliance, firm i invests c_i to ensure the compliance of ESR requirements and the retailer realizes a revenue increase of R. Firms make or receive transfer payments based on the negotiated contracts.

It is easy to see that at the end of the kth round of relationship development, a subtree T_k is formed which specifies the current implementation structure. The nodes in N_k are the leaf nodes of the subtree T_k. If all firms have established ESR negotiation relationships at the end of the mth round, then $T_m = T$ is the complete implementation structure, which is a spanning tree of the network (N, A^{ESR}).

Consider the game at the beginning of round k and the currently established structure is T_{k-1}. The firms in set N_{k-1}, which corresponds to the leaves in T_{k-1}, are choosing firms in set $N \setminus \cup_{\ell=1}^{k-1} N_\ell$ to establish trading relationships. The eventual equilibrium structure must be an element of $\mathcal{T}_{k-1}(T_{k-1})$, which is the set of all spanning trees in (N, A^{ESR}) containing the subtree T_{k-1}. For any $T \in \mathcal{T}_{k-1}(T_{k-1})$ and $i \in N$, we have from (4.10)–(4.11)

$$\pi_i(T, \mathbf{1}) = \pi_0(T, \mathbf{1}) \cdot \prod_{a \in P_i(T)} \rho_a.$$

If firm i is a leaf node of the subtree T_{k-1} (i.e., $i \in N_{k-1}$), the path from node 0 to node i at the beginning of the kth round of relationship development stays fixed, i.e., $P_i(T) = P_i(T_k)$, regardless of the eventual structure $T \in \mathcal{T}_{k-1}(T_{k-1})$. Then, the above expression suggests that all firms in the set N_{k-1} have the same incentive as the retailer; that is, these firms would like to form ESR partnerships according to the structure that is most preferred by the retailer. However, if firm i is not part of the subtree T_{k-1} (i.e., $i \in N \setminus \cup_{\ell=1}^{k-1} N_\ell$), the path from node 0 to node i depends on which downstream firms is chosen by firm i to interact within the ESR program.

In this case, firm i's incentive is not aligned with that of the retailer. With these observations, we derive the subgame perfect equilibrium of sequentially formed ESR implementation structure.

Theorem 4.5.2 (Sequential Delegation of Relationship Development) *When the ESR relationships are sequentially formed from the downstream firms to the upstream firms, the retailer preferred implementation structure corresponds to a subgame perfect equilibrium, whether the firms move simultaneously or sequentially in each round of relationship development.*

Theorem 4.5.2 formally establishes that the implementation structure identified in §§4.4.2 corresponds to the equilibrium when the firms sequentially form ESR relationships.

4.5.3 Coordination of Multiple ESR Programs

In this subsection, we consider a setting where revenues can be generated by different ESR programs, each requiring the participation of a subset of suppliers. Take Starbucks' supply network as an example. Starbucks manages suppliers who provide raw materials such as coffee beans and tea leaves, while it also deals with suppliers who produce packing materials like paper cups. These two groups of suppliers work on different product families and operate independently from one another. Starbucks may initiate separate ESR programs with specific groups of suppliers. A commitment for offering organic beverage would rely on the first group of suppliers, while an initiative to ensure 100% recycled materials would definitely concern the second group. To understand the coordination of different ESR programs, we need to extend our model and analysis.

For ease of exposition, we demonstrate a model with two product families and with two first-tier suppliers directly trading with the retailer. Our analysis can be extended to multiple product families and multiple first-tier suppliers. As demonstrated in Figure 4.2, let 1 and 2 index the tier-1 suppliers for product families

1 and 2, respectively. The suppliers for product family $k \in \{1,2\}$ form a subset $N_k = \{k\} \cup \{j \in N : P_{(k,j)}(A) \neq \emptyset, P_{(3-k,j)}(A) = \emptyset\}$ with $N = N_1 \cup N_2 \cup \{0\}$. In other words, a supplier in subgroup N_k is connected with tier-1 supplier $k \in \{1,2\}$, and the sets N_1 and N_2 are disjoint. The firms related to product family k form a subnetwork $(\{0\} \cup N_k, A_k^{\text{ESR}})$, where A_k^{ESR} is the set of arcs indicating possible negotiations associated with program k. It is clear that $A^{\text{ESR}} = A_1^{\text{ESR}} \cup A_2^{\text{ESR}}$. Implementation of the ESR program for product family k in the subnetwork (N_k, A_k^{ESR}) generates a revenue of R_k, $k = 1, 2$. If the retailer is able to roll out both programs in (N, A^{ESR}), a revenue of R is generated. The retailer invests $c_{0,k}$ for implementing only ESR program $k = 1, 2$ and c_0 for implementing both. Let $\Pi = R - C_1 - C_2 - c_0$ and $\Pi_k = R_k - C_k - c_{0,k}, k = 1, 2$, where $C_k = \sum_{j \in N_k} c_j$. In view of the possible synergy between the two programs in enhancing the retailer's brand image, we assume

$$\Pi \geq (\Pi_1)^+ + (\Pi_2)^+.$$

This condition suggests that the gain from ESR programs is super-additive, so that there is a positive synergy between the two programs. In other words, the marginal contribution of an additional ESR program is assumed to be greater than or equal to the corresponding investment cost to avoid trivial cases. Note that we do not require the individual program to be profitable.

Consider a feasible structure T under which both ESR programs are implemented. The disagreement point for each firm in $N_1 \cup N_2$ remains zero, i.e., $D_i^j = 0$ for $i, j \in N_1 \cup N_2$ and $(i,j) \in T$ or $(j,i) \in T$, as is in the model analyzed in §4.4. The retailer, however, may obtain a nonzero value in the event of negotiation breakdown. Specifically, if the retailer and firm $i \in N_k$ fail to reach an agreement, the retailer may still implement the ESR program for product family $(3-k)$. Let $\mathbf{x}^k = (x_j : j \in \{0\} \cup N_k)$ denote the ESR investment commitment of the retailer and the suppliers in N_k. The retailer's disagreement point when negotiating with a firm in N_k is

$$D_0^i(T, \mathbf{x}^{3-k}) = D_0^k(T, \mathbf{x}^{3-k}) \equiv R_{3-k} \cdot \prod_{j \in N_{3-k}} x_j - \sum_{j \in \{u \in N_{3-k} : t_{0,u} = 1\}} w_{0,j}(\mathbf{x}^{3-k}) - c_{0,3-k} \cdot x_0.$$

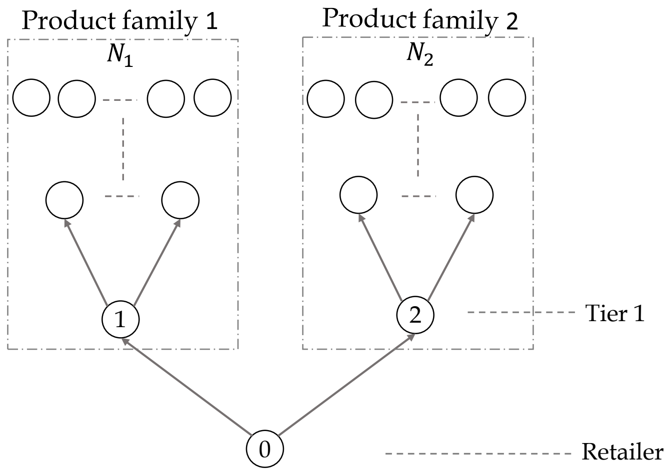

Fig. 4.2.: Multiple ESR programs

Updating the retailer's disagreement point in (4.7), we can solve for the equilibrium negotiated contracts. For any negotiation without the retailer, i.e., $i \neq 0$, the relation in (4.8) continues to hold, i.e., $\rho_{i,j} = \pi_j(T, \mathbf{x})/\pi_i(T, \mathbf{x})$. For negotiations involving the retailer, this relationship is changed to

$$\rho_{0,i} = \frac{\theta_{i,0}}{\theta_{0,i}} = \frac{\pi_i(T, \mathbf{x})}{\pi_0(T, \mathbf{x}) - D_0^k(T, \mathbf{x}^{3-k})}, \quad i \in N_k. \tag{4.13}$$

With these relations, (4.9) is modified to

$$\pi_i(T, \mathbf{1}) = (\pi_0(T, \mathbf{1}) - D_0^k(T, \mathbf{1})) \prod_{a \in P_i(T)} \rho_a, \quad i \in N_k \text{ and } k \in \{1, 2\}. \tag{4.14}$$

Given that the gain of the entire supply chain is Π, we deduce

$$\pi_0(T, \mathbf{1}) + \Gamma_1(T)(\pi_0(T, \mathbf{1}) - D_0^1(T, \mathbf{1})) + \Gamma_2(T)(\pi_0(T, \mathbf{1}) - D_0^2(T, \mathbf{1})) = \Pi, \tag{4.15}$$

where $\Gamma_k(T) = \sum_{v \in N_k} \prod_{a \in P_v(T)} \rho_a$ is the sum of gain ratios of firms related to product family $k \in \{1, 2\}$. The following theorem identifies the retailer's equilibrium gain.

Theorem 4.5.3 (Multiple ESR Programs: The Retailer's Gain) *Given a feasible structure $T \in \mathcal{T}(A^{\mathrm{ESR}})$ with which both ESR programs are implemented, the retailer's gain is*

$$\pi_0(T, \mathbf{1}) = \frac{1 - \Gamma_1(T)\Gamma_2(T)}{(1 + \Gamma_1(T))(1 + \Gamma_2(T))}\Pi + \frac{\Gamma_2(T)}{1 + \Gamma_2(T)}\Pi_1 + \frac{\Gamma_1(T)}{1 + \Gamma_1(T)}\Pi_2,$$

and her disagreement points are

$$D_0^k(T,\mathbf{1}) = \Pi_{3-k} - \frac{\Gamma_{3-k}(T)}{1+\Gamma_{3-k}(T)}(\Pi - \Pi_k),\ k \in \{1,2\}.$$

The supplier i's gain is

$$\pi_i(T,\mathbf{1}) \;=\; \frac{\Pi - \Pi_{3-k}}{1+\Gamma_k(T)} \prod_{a\in P_i(T)} \rho_a,$$

for $i \in N_k, k = 1, 2$.

From Theorem 4.5.3, the retailer may not benefit from an increased gain Π of the entire supply chain (i.e., when $\Gamma_1(T)\Gamma_2(T) > 1$). Such a situation may happen when the upstream firms are very powerful (e.g., the bargaining power of an upstream firm vis-à-vis a downstream one is always above 0.5). An increased value Π_k, $k \in \{1,2\}$, of individual ESR program, however, always leads to an increased gain obtained by the retailer. This is because the retailer enjoys an enhanced bargaining position with an increased disagreement point when negotiating with firms belonging to set N_{3-k}.

If the retailer chooses to implement only program $k \in \{1,2\}$, we can obtain the solution by applying the analysis in §4.4 to the subnetwork (N_k, A_k^{ESR}). From (4.10), the retailer's gain in this case is simply $(\Pi_k)^+(1+\Gamma_k(T))^{-1}$. With Theorem 4.5.3, we can easily derive the retailer's choice of ESR program, as depicted in Figure 4.3. From (4.14) and (4.15), $\Gamma_k(T)$ represents the ratio of gain taken by the suppliers associated with product family k to the trade surplus enjoyed by the retailer. It is intuitive that the retailer would choose to initiate both ESR programs if $\Gamma_1(T)$ and $\Gamma_2(T)$ are sufficiently small. Also if $\Gamma_1(T) < (>)\Gamma_2(T)$, then implementing program 1 is more (less) profitable to the retailer than implementing program 2. It is interesting to note that the retailer would implement both programs when only one of $\Gamma_1(T)$ and $\Gamma_2(T)$ is extremely small. When an extremely small share is taken by the firms in N_1 (N_2), the retailer enjoys not only a large surplus from the trade of product family 1 (2), but also a strengthened bargaining position (i.e., an increased disagreement point) in the negotiations with firms in N_2 (N_1). Working on both ESR programs allows the retailer to reap both benefits.

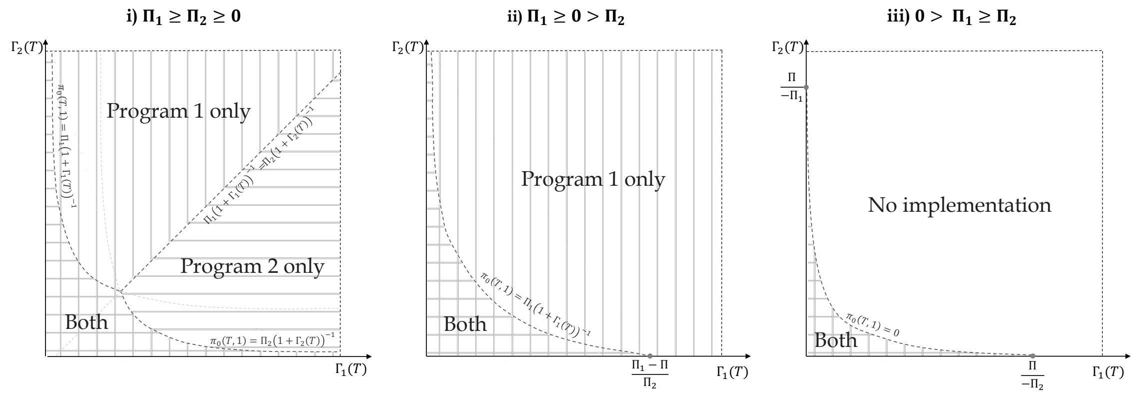

Fig. 4.3.: The retailer's choice of ESR programs

The next theorem suggests that the optimal implementation structure for both ESR programs can be obtained by simply combining the optimal implementation structures for individual ESR programs. Thus, Theorem 4.4.2 can be applied to obtain the optimal implementation structure.

Theorem 4.5.4 (Multiple ESR Programs: Implementation Structure) *If T_k is the retailer preferred structure for implementing program k, $k = 1, 2$, in the sub-network $(N_k, A_k^{\mathrm{ESR}})$, then the retailer preferred structure for implementing both ESR programs is $T = T_1 \cup T_2$.*

When only one ESR program is considered, we have shown that the multi-unit bilateral framework generalizes that based on the Shapley value (recall Theorem 4.4.1). Next, we establish a similar relationship. For the problem with two ESR programs, the characteristic function is defined as

$$\mathscr{V}(S) = \begin{cases} \Pi, & \text{if } S = N, \\ (\Pi_k)^+, & \text{if } S = \{0\} \cup N_k \cup S^0, S^0 \subsetneq N_{3-k}, \\ 0, & \text{otherwise.} \end{cases} \tag{4.16}$$

Following our discussions in §§4.3.1, we can derive the Shapley values of the firms as

$$s_0 = \frac{1}{|N|!}\Big((|N|-1)!\Pi + \sum_{k=1}^{2}\sum_{m=0}^{|N_{3-k}|-1}\binom{|N_{3-k}|}{m}(|N_k|+m)!(\Pi_k)^+\Big), \tag{4.17}$$

$$s_j = \frac{1}{|N|!}\Big((|N|-1)!(\Pi-(\Pi_{3-k})^+) + \sum_{m=0}^{|N_{3-k}|-1}\binom{|N_{3-k}|}{m}(|N_k|-1+m)!(\Pi_k)^+\Big), \tag{4.18}$$

for all $j \in N_k$, $k \in \{1,2\}$. We observe that suppliers within the same product family have the same Shapley values, while those belonging to different product families have different Shapley values in general. With (4.17) and (4.18), we can extend the observation in Corollary 4.4.1.

Corollary 4.5.1 *When a grand coalition is formed, the allocation based on the Shapley value corresponds to that using the multi-unit bilateral bargaining approach under any feasible implementation structure when* $\rho_{0,i} = \frac{s_k}{\Pi-\Pi_{3-k}-|N_k|s_k} > 0, i \in N_k, k \in \{1,2\}$ *and* $\rho_{i,j} = 1$ *for* $i \neq 0$ *and* $(i,j) \in A^{\mathrm{ESR}}$.

According to Corollary 4.5.1, the gain allocation generated using the Shapley value is a special instance of that using the multi-unit bilateral approach.

4.6 Concluding Remarks

In this chapter, we analyze the implementation of ESR initiative through a general supply network. The multi-unit bargaining framework allows us to model various power distributions among the firms in the network, which determine the formation of ESR relationships and the ESR gain allocation across the network. We demonstrate the advantage of our approach over the conventional Shapley value based approach, leading to new understandings of ESR implementation.

5. CONCLUSION AND DIRECTION FOR FUTURE RESEARCH

This chapter concludes the findings of this research. This study dealt with supply chain management and inventory management problems arising in the business world. Several emerging topics are considered, including dynamic product substitution, joint price and inventory decisions under demand ambiguity, and the implementation of ESR program in general supply networks. The first problem concerns the dynamic product substitution policy to meet the mix of demands using the mix of supplies. The second problem studies the robust decisions making for selling perishable products with ambiguous demand information. The third problem examines the implementation of ESR programs in general supply networks under the multi-unit bargaining framework.

For the first problem, we study the coordination of product replenishment and substitution decisions in a dynamic environment with both supply and demand uncertainties. There are two key features of our model. The first is the general substitution structure with the benefit matrix revealing the reverse Monge property. The reverse Monge property allows us to generalize the existing studies on downward product substitution, yet develop an efficient approach to derive the decisions of multi-product allocation to meet the demands. The second feature of our model is the consideration of general supply functions that are stochastically linear in midpoint. We show that the value of product substitution is greatly affected by both the limits and the variabilities of the supplies.

For the second problem, we adopt the minimax regret framework to study the pricing newsvendor problem with ambiguous demand information. We first identify the optimal order quantity for a given unit price and show that the optimal quantity is decreasing in the price, which is not necessarily true in the classical models. By further exploring the properties of the minimax objective, we characterize the optimal price decision, which can be computed efficiently. We then conduct a series of analyses to investigate the effects of important factors such as the presence of inventory risk, the unit cost, and the degree of demand ambiguity on the optimal decisions and compare these results with those in classical models. We also implement the minimax regret approach in a data-driven setting where historical data of price and demand are available, and compare its performance with the classical pricing newsvendor model that maximizes the expectation.

For the third problem, we analyze the implementation of ESR initiative through a general supply network. The multi-unit bargaining framework allows us to model various power distributions among the firms in the network, which determine the formation of ESR relationships and the ESR gain allocation across the network. We demonstrate the advantage of our approach over the conventional Shapley value based approach, leading to new understandings of ESR implementation.

There are several interesting questions that we leave open for future research. In the first problem that considers the supply uncertainty, in addition to product substitution, responsive pricing after observing the supply can also be used to effectively reshape consumer demand. The theory of stochastic linearity in midpoint allows for generalizing our framework to revenue management with both uncertain supply processes and general price-demand relationships.

In the second problem, we have focused on the scenarios where the firm only knows the support information. If the firm obtains additional information, such as the mean or the mode of the random factor, we could further improve the performance of the minimax regret approach by incorporating such additional information. In addition, while our model setting assumes that the parametric form of the demand model is known, a more interesting problem is to model the price-demand relationship with a more general functional form. Specifically, if the firm has partial information regarding the customer valuation distribution, the adversarial nature could choose the functional form of the demand function to maximize the firm's regret. Finally, another valuable question to be investigated is whether a randomized pricing policy could reduce the firm's regret.

In the third problem, the network bargaining framework allows us to analyze complex supply networks, making our model close to reality. There are a number of directions one may take to extend our existing analysis. In our study, we have assumed given material trading terms and focused only on ESR implementation. When the downstream brand firm is in the process of developing a new product or expanding to a new market, it may need to develop the supplier base along with its ESR initiative. In this case, one also needs to model the contracting relationships for the material flow. Incorporating the material exchange into our bargaining framework requires a careful modification of the firms' disagreement points as well as the bargaining power distribution. The trade-off between efficient sourcing and successful ESR implementation is an interesting aspect for future exploration. When leading a new ESR initiative, the downstream brand firm may not be able to implement all measures in one step. Instead, establishing compliance requirements may entail significant capability development as well as supplier education, and the initiative may need to roll out in stages. One may extend the network framework developed in this model to consider sequential roll out of ESR initiatives and possibly diffusion of ESR standards (see, e.g., Castka and Corbett, 2016).

REFERENCES

REFERENCES

Agarwal, N. (2015). An empirical model of the medical match. *The American Economic Review 105*(7), 1939–1978.

Agrawal, V. and D. Lee (2017, October). The effect of sourcing policies on a supplier's sustainable practices. SSRN Scholarly Paper ID 2601974, Social Science Research Network, Rochester, NY.

Agrawal, V. and S. Seshadri (2000). Impact of uncertainty and risk aversion on price and order quantity in the newsvendor problem. *Manufacturing & Service Operations Management 2*(4), 410–423.

Agrawal, V. V. and D. Lee (2016). Responsible sourcing. In *Environmentally Responsible Supply Chains*, pp. 71–86. Springer.

Ahuja, R. K., T. Magnanti, and J. B. Orlin (2017). *Network Flows: Theory, Algorithms, and Applications.* Pearson Education.

Alon, N., S. Cosares, D. S. Hochbaum, and R. Shamir (1989). An algorithm for the detection and construction of monge sequences. *Linear Algebra and Its Applications 114*, 669–680.

Archibald, T. W., S. Sassen, and L. Thomas (1997). An optimal policy for a two depot inventory problem with stock transfer. *Management Science 43*(2), 173–183.

Ata, B., D. Lee, and M. H. Tongarlak (2012). Optimizing organic waste to energy operations. *Manufacturing & Service Operations Management 14*(2), 231–244.

Atasu, A. (2016). *Environmentally Responsible Supply Chains*, Volume 3. Springer.

Axsäter, S. (2003). A new decision rule for lateral transshipments in inventory systems. *Management Science 49*(9), 1168–1179.

Babich, V. and C. S. Tang (2012, February). Managing opportunistic supplier product adulteration: Deferred payments, inspection, and combined mechanisms. *Manufacturing & Service Operations Management 14*(2), 301–314.

Bassok, Y., R. Anupindi, and R. Akella (1999). Single-period multiproduct inventory models with substitution. *Operations Research 47*(4), 632–642.

Bellman, R. (1958). On a routing problem. *Quarterly of Applied Mathematics 16*(1), 87–90.

Bergemann, D. and K. H. Schlag (2008). Pricing without priors. *Journal of the European Economic Association 6*(2-3), 560–569.

Bookbinder, J. H. and A. E. Lordahl (1989). Estimation of inventory re-order levels using the bootstrap statistical procedure. *IIE Transactions 21*(4), 302–312.

Burkard, R. E. (2007). Monge properties, discrete convexity and applications. *European Journal of Operational Research 176*(1), 1–14.

Burnetas, A. N. and C. E. Smith (2000). Adaptive ordering and pricing for perishable products. *Operations Research 48*(3), 436–443.

Caldentey, R., Y. Liu, and I. Lobel (2016). Intertemporal pricing under minimax regret. *Operations Research*.

Caro, F., P. Chintapalli, K. Rajaram, and C. S. Tang (2018, April). Improving supplier compliance through joint and shared audits with collective penalty. *Manufacturing & Service Operations Management 20*(2), 363–380.

Castka, P. and C. Corbett (2016). Adoption and diffusion of environmental and social standards: The effect of stringency, governance, and media coverage. *International Journal of Operations & Production Management 36*(11), 1504–1529.

Chan, L. M., Z. M. Shen, D. Simchi-Levi, and J. L. Swann (2004). Coordination of pricing and inventory decisions: A survey and classification. In *Handbook of Quantitative Supply Chain Analysis*, pp. 335–392. Springer.

Chen, B., X. Chao, and C. Shi (2017). Nonparametric algorithms for joint pricing and inventory control with lost-sales and censored demand. *Working paper*.

Chen, J., A. Qi, and M. Dawande (2018, May). Supplier centrality and auditing priority in socially-responsible supply chains. SSRN Scholarly Paper ID 2889889, Social Science Research Network, Rochester, NY.

Chen, L., S. M. Gilbert, and Y. Xia (2016, June). Product line extensions and technology licensing with a strategic supplier. *Production and Operations Management 25*(6), 1121–1146.

Chen, L. and H. L. Lee (2017, July). Sourcing under supplier responsibility risk: The effects of certification, audit, and contingency payment. *Management Science 63*(9), 2795–2812.

Chen, M. (2017). Why 'Corporate Responsibility' campaigns fail. *The Nation*. `https://www.thenation.com/article/why-corporate-responsibility-campaigns-fail/`, retrieved on Aug 3, 2018.

Chen, R. R. and S. Yin (2010, November). The equivalence of uniform and shapley value-based cost allocations in a specific game. *Operations Research Letters 38*(6), 539–544.

Chen, S., Q. Zhang, and Y.-P. Zhou (2015, March). Impact of supply chain transparency on sustainability under NGO scrutiny. SSRN Scholarly Paper ID 2590152, Social Science Research Network, Rochester, NY.

Chen, W., Q. Feng, and S. Seshadri (2013). Sourcing from suppliers with random yield for price-dependent demand. *Annals of Operations Research 208*(1), 557–579.

Chen, X., X. Gao, and Z. Hu (2015). A new approach to two-location joint inventory and transshipment control via $L^\natural$-convexity. *Operations Research Letters 43*(1), 65–68.

Chen, X. and D. Simchi-Levi (2012). Pricing and inventory management. *The Oxford Handbook of Pricing Management*, 784–822.

Chen, Y., M. Xu, and Z. G. Zhang (2009). Technical note - a risk-averse newsvendor model under the cvar criterion. *Operations Research 57*(4), 1040–1044.

Chen, Y.-J. and M. Deng (2013). Supplier certification and quality investment in supply chains. *Naval Research Logistics (NRL) 60*(3), 175–189.

Chen, Z.-L., N. G. Hall, and H. Kellerer (2016). Dynamic pricing to minimize maximum regret. *Production and Operations Management 26*(1), 47–63.

Cho, S.-H., X. Fang, S. R. Tayur, and Y. Xu (2017, April). Combating child labor: Incentives and information disclosure in global supply chains. SSRN Scholarly Paper ID 2552268, Social Science Research Network, Rochester, NY.

Chou, M. C., C. K. Sim, C.-P. Teo, and H. Zheng (2012). Newsvendor pricing problem in a two-sided market. *Production and Operations Management 21*(1), 204–208.

Chu, L. Y., Y. Rong, and H. Zheng (2017). RFQ, sequencing, and the most favorable bargaining outcome. *Working paper*.

Chu, L. Y., J. G. Shanthikumar, and Z.-J. M. Shen (2008). Solving operational statistics via a bayesian analysis. *Operations Research Letters 36*(1), 110–116.

Chu, L. Y., J. G. Shanthikumar, and Z.-J. M. Shen (2017). Operational learning in pricing and revenue management. *Working paper*.

Cone Communications (2013). Cone Communications/Echo Global CSR study. `http://www.conecomm.com/research-blog/2013-cone-communications-echo-global-csr-study`, retrieved on August 29, 2018.

Cooper, W. L. (2002). Asymptotic behavior of an allocation policy for revenue management. *Operations Research 50*(4), 720–727.

Corbett, C. J. and G. A. DeCroix (2001, July). Shared-savings contracts for indirect materials in supply chains: Channel profits and environmental impacts. *Management Science 47*(7), 881–893.

Davidson, C. (1988). Multiunit bargaining in oligopolistic industries. *Journal of Labor Economics 6*(3), 397–422.

Dukes, A., E. Gal-Or, and K. Srinivasan (2006). Channel bargaining with retailer asymmetry. *Journal of Marketing Research 43*(1), 84–97.

Ecovadis (2018). Sustainability clauses in commercial contracts: The key to corporate responsibility. `https://www.ecovadis.com/library/sustainability-clauses-commercial-contracts-key-corporate-responsibility/`, retrieved on August 3, 2018.

Elmaghraby, W. and P. Keskinocak (2003). Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions. *Management Science 49*(10), 1287–1309.

Esenduran, G. and E. Kemahlıoğlu-Ziya (2015). A comparison of product take-back compliance schemes. *Production and Operations Management 24*(1), 71–88.

Fang, X. and S.-H. Cho (2015, October). Cooperative approaches to managing social responsibility in supply chains: Joint auditing and information sharing. SSRN Scholarly Paper ID 2686018, Social Science Research Network, Rochester, NY.

Federgruen, A. and A. Heching (1999). Combined pricing and inventory control under uncertainty. *Operations Research 47*(3), 454–475.

Feng, Q., J. Jia, and J. G. Shanthikumar (2018). Dynamic multisourcing with dependent supplies. *Management Science*. forthcoming.

Feng, Q., C. Li, M. Lu, and J. G. Shanthikumar (2019a). Dynamic substitution for selling multiple products under supply and demand uncertainties. *Available at SSRN 3309522*.

Feng, Q., C. Li, and M. Lu (2019b). Data-driven inventory and pricing management with adversarial models. *Working paper*.

Feng, Q., C. Li, M. Lu, and J. G. Shanthikumar (2019c). Implementing environmental and social responsibility programs in supply networks through multi-unit bilateral negotiation. *Available at SSRN 3309518*.

Feng, Q. and L. X. Lu (2013a, September). The role of contract negotiation and industry structure in production outsourcing. *Production and Operations Management 22*(5), 1299–1319.

Feng, Q. and L. X. Lu (2013b, May). Supply chain contracting under competition: Bilateral bargaining vs. stackelberg. *Production and Operations Management 22*(3), 661–675.

Feng, Q., S. Luo, and D. Zhang (2013). Dynamic inventory–pricing control under backorder: Demand estimation and policy optimization. *Manufacturing & Service Operations Management 16*(1), 149–160.

Feng, Q. and J. G. Shanthikumar (2017). Supply and demand functions in inventory models.

Feng, Q. and J. G. Shanthikumar (2018). Supply and demand functions in inventory models. *Operations Research 66*(1), 77–91.

Feng, Q. and R. Shi (2012). Sourcing from multiple suppliers for price-dependent demands. *Production and Operations Management 21*(3), 547–563.

Ford, Jr., L. R. (1956). Network flow theory. Technical report, RAND Corporation.

Fu, Q., C. K. Sim, and C.-P. Teo (2017). Profit sharing agreements in decentralized supply chains: a distributionally robust approach. *Operations Research*.

Gallego, G. and I. Moon (1993). The distribution free newsboy problem: review and extensions. *Journal of the Operational Research Society 44*(8), 825–834.

Godfrey, G. A. and W. B. Powell (2001). An adaptive, distribution-free algorithm for the newsvendor problem with censored demands, with applications to inventory and distribution. *Management Science 47*(8), 1101–1112.

Granot, D. and G. Sošić (2003, October). A three-stage model for a decentralized distribution system of retailers. *Operations Research 51*(5), 771–784.

Gui, L., A. Atasu, O. Ergun, and L. B. Toktay (2016). Efficient implementation of collective extended producer responsibility legislation. *Management Science 62*(4), 1098–1123.

Gui, L., A. Atasu, O. Ergun, and L. B. Toktay (2018, January). Design incentives under collective extended producer responsibility: A network perspective. *Management Science*. Forthcoming.

Gul, F. (1989). Bargaining foundations of Shapley value. *Econometrica: Journal of the Econometric Society 51*(1), 81–95.

Guo, R., H. L. Lee, and R. Swinney (2016, December). Responsible sourcing in supply chains. *Management Science 62*(9), 2722–2744.

Guo, X., G. Xiao, and F. Zhang (2017, September). Effect of consumer awareness on corporate social responsibility under asymmetric information. SSRN Scholarly Paper ID 3039862, Social Science Research Network, Rochester, NY.

Handel, B. R. and K. Misra (2015). Robust new product pricing. *Marketing Science 34*(6), 864–881.

Hardcastle, J. L. (2015). CSR programs increase revenue up to 20%, says Verizon, Campbell soup study. `https://www.environmentalleader.com/2015/07/csr-programs-increase-revenue-up-to-20-says-verizon-campbell-soup-study/?amp=1`, retrieved on August 29, 2018.

Harsha, P., R. Natarajan, and D. Subramanian (2016). A data-driven, distribution-free, multivriate approach to the price-setting newsvendor problem. *Working paper*.

Hochbaum, D. S. and J. G. Shanthikumar (1990). Convex separable optimization is not much harder than linear optimization. *Journal of the ACM (JACM) 37*(4), 843–862.

Hoffman, A. J. (1963). On simple transportation problems. *Convexity: Proceedings of Symposia in Pure Mathematics 7*, 317 – 327.

Holmberg, K. (1995, Oct). Efficient decomposition and linearization methods for the stochastic transportation problem. *Computational Optimization and Applications 4*(4), 293–316.

Horn, H. and A. Wolinsky (1988). Bilateral monopolies and incentives for merger. *The RAND Journal of Economics 19*(3), 408–419.

Hsu, A. and Y. Bassok (1999). Random yield and random demand in a production system with downward substitution. *Operations Research 47*(2), 277–290.

Hu, M. and Y. Zhou (2016). Dynamic type matching. *Rotman School of Management Working Paper*.

Hu, X., I. Duenyas, and R. Kapuscinski (2008). Optimal joint inventory and transshipment control under uncertain capacity. *Operations Research 56*(4), 881–897.

Huang, L., J.-S. J. Song, and R. Swinney (2018, October). Managing social responsibility in multitier supply chains. SSRN Scholarly Paper ID 2837332, Social Science Research Network, Rochester, NY.

Instacart (2018). Instacart Help Center–Replacement items. Retrieved on October 29th, 2018.

Jain, A., A. E. B. Lim, J. G. Shanthikumar, and L. Yu (2011). Objective operational learning with application to newsven- dor problem.

Jasin, S. and S. 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.

Jiang, R., J. Wang, M. Zhang, and Y. Guan (2013). Two-stage minimax regret robust unit commitment. *IEEE Transactions on Power Systems 28*(3), 2271–2282.

Kalagnanam, J. R., M. W. Dawande, M. Trumbo, and H. S. Lee (2000). The surplus inventory matching problem in the process industry. *Operations Research 48*(4), 505–516.

Karaer, O., T. Kraft, and J. Khawam (2017). Buyer and nonprofit levers to improve supplier environmental performance. *Production and Operations Management 26*(6), 1163–1190.

Karlin, S. and C. R. Carr (1962). Prices and optimal inventory policy. *Studies in Applied Probability and Management Science 4*, 159–172.

Karmarkar, U. S. (1981). The multiperiod multilocation inventory problem. *Operations Research 29*(2), 215–228.

Kocabiyikoglu, A. and I. Popescu (2011). An elasticity approach to the newsvendor with price-sensitive demand. *Operations research 59*(2), 301–312.

Krishnan, K. and V. Rao (1965). Inventory control in n warehouses. *Journal of Industrial Engineering 16*(3), 212.

Kumar, A., A. Mehra, and S. Kumar (2018). Why do stores drive online sales? Evidence of underlying mechanisms from a multichannel retailer. University of Florida Warrington College of Business Working Paper.

Kunnumkal, S. and H. Topaloglu (2010). A new dynamic programming decomposition method for the network revenue management problem with customer choice behavior. *Production and Operations Management 19*(5), 575–590.

Lawler, E. L. (1976). *Combinatorial Optimization: Networks and Matroids.* Courier Corporation.

Leng, M. and M. Parlar (2008, October). Allocation of cost savings in a three-level supply chain with demand information sharing: A cooperative-game approach. *Operations Research 57*(1), 200–213.

Letizia, P. and G. Hendrikse (2016). Supply chain structure incentives for corporate social responsibility: An incomplete contracting analysis. *Production and Operations Management 25*(11), 1919–1941.

Levi, R., G. Perakis, and J. Uichanco (2011). Regret optimization for stochastic inventory models with spread.

Levi, R., G. Perakis, and J. Uichanco (2015). The data-driven newsvendor problem: new bounds and insights. *Operations Research 63*(6), 1294–1306.

Levi, R., R. O. Roundy, and D. B. Shmoys (2007). Provably near-optimal sampling-based policies for stochastic inventory control models. *Mathematics of Operations Research 32*(4), 821–839.

Li, Q. and D. Atkins (2005). On the effect of demand randomness on a price/quantity setting firm. *IIE Transactions 37*(12), 1143–1153.

Liu, Q. and G. Van Ryzin (2008). On the choice-based linear programming model for network revenue management. *Manufacturing & Service Operations Management 10*(2), 288–310.

Liyanage, L. H. and J. G. Shanthikumar (2005). A practical inventory control policy using operational statistics. *Operations Research Letters 33*(4), 341–348.

Lu, M., Z.-J. M. Shen, and K. Göler (2016). A data-driven approach for the pricing newsvendor problem. *Working paper*.

Lu, Y. and D. Simchi-Levi (2013). On the unimodality of the profit function of the pricing newsvendor. *Production and Operations Management 22*, 615–625.

Luo, S., S. P. Sethi, and R. Shi (2016). On the optimality conditions of a price-setting newsvendor problem. *Operations Research Letters 44*(6), 697–701.

Mills, E. S. (1959). Uncertainty and price theory. *The Quarterly Journal of Economics*, 116–130.

Monge, G. (1781). Déblai et remblai. *Memoires de l'Academie des Sciences*.

Nagarajan, M. and Y. Bassok (2008, August). A bargaining framework in supply chains: The assembly problem. *Management Science 54*(8), 1482–1496.

Nash, J. F. (1950). The bargaining problem. *Econometrica 18*(2), 155–162.

Natarajan, K., M. Sim, and J. Uichanco (2017). Asymmetry and ambiguity in newsvendor models. *Management Science*.

Netessine, S., G. Dobson, and R. A. Shumsky (2002). Flexible service capacity: Optimal investment and the impact of demand correlation. *Operations Research 50*(2), 375–388.

Nguyen, J., K. Donohue, and M. Mehrotra (2018). Closing a supplier's energy efficiency gap through assessment assistance and procurement commitment. *Forthcoming*.

Nielsen (2014). Global consumers are willing to put their money where their heart is when it comes to goods and services from companies committed to social responsibility. http://www.nielsen.com/us/en/press-room/2014/global-consumers-are-willing-to-put-their-money-where-their-heart-is.html, retrieved on August 29, 2018.

Orlin, J. B. (1993). A faster strongly polynomial minimum cost flow algorithm. *Operations Research 41*(2), 338–350.

Orsdemir, A., B. Hu, and V. Deshpande (2016, March). Ensuring corporate social and environmental responsibility through vertical integration and horizontal sourcing. SSRN Scholarly Paper ID 2630733, Social Science Research Network, Rochester, NY.

Pasternack, B. A. and Z. Drezner (1991). Optimal inventory policies for substitutable commodities with stochastic demand. *Naval Research Logistics 38*(2), 221–240.

Perakis, G. and G. Roels (2008). Regret in the newsvendor model with partial information. *Operations Research 56*(1), 188–203.

Petruzzi, N. C. and M. Dada (1999). Pricing and the newsvendor problem: A review with extensions. *Operations Research 47*(2), 183–194.

Plambeck, E. L. and T. A. Taylor (2016, September). Supplier evasion of a buyer's audit: Implications for motivating supplier social and environmental responsibility. *Manufacturing & Service Operations Management 18*(2), 184–197.

Pulse Commerce (2018). Five ways omnichannel retailers increase profits with ship from store programs. Retrieved on October 29th, 2018.

PwC (2016). Redefining business success in a changing world: CEO survey. https://www.pwc.com/gx/en/ceo-survey/2016/landing-page/pwc-19th-annual-global-ceo-survey.pdf, retrieved on August 29, 2018.

Reiman, M. I. and Q. Wang (2008). An asymptotically optimal policy for a quantity-based network revenue management problem. *Mathematics of Operations Research 33*(2), 257–282.

Robinson, L. W. (1990). Optimal and approximate policies in multiperiod, multilocation inventory models with transshipments. *Operations Research 38*(2), 278–295.

Rockafellar, R. T. (2015). *Convex Analysis*. Princeton University Press.

Roels, G. (2013). Risk premiums in the price-setting newsvendor model. *Available at SSRN 2285243*.

Rudin, C. and G.-Y. Vahn (2014). The big data newsvendor: Practical insights from machine learning.

Savage, L. J. (1951). The theory of statistical decision. *Journal of the American Statistical association 46*(253), 55–67.

Scarf, H. (1958). A min-max solution of an inventory problem. In H. Scarf, K. Arrow, and S. Karlin (Eds.), *Studies in the Mathematical Theory of Inventory and Production*, Volume 10, pp. 201–209. Stanford University Press Stanford, CA.

Sedex (2013). Going Deep–The case for multi-tier transparency. `https://cdn.sedexglobal.com/wp-content/uploads/2016/09/Sedex-Transparency-Briefing.pdf`, retrieved on August 3, 2018.

Shaked, M. and J. G. Shanthikumar (1998). Two variability orders. *Probability in the Engineering and Informational Sciences 12*(1), 1–23.

Shaked, M. and J. G. Shanthikumar (2007). *Stochastic orders.* Springer Science & Business Media.

Shetty, C. (1959). A solution to the transportation problem with nonlinear costs. *Operations Research 7*(5), 571–580.

Shumsky, R. A. and F. Zhang (2009). Dynamic capacity management with substitution. *Operations Research 57*(3), 671–684.

Talluri, K. T. and G. J. Van Ryzin (2006). *The Theory and Practice of Revenue Management*, Volume 68. Springer Science & Business Media.

Topkis, D. M. (2011). *Supermodularity and complementarity.* Princeton university press.

Vaidyanathan, B. (2013). Faster strongly polynomial algorithms for the unbalanced transportation problem and assignment problem with monge costs. *Networks 62*(2), 136–148.

Wang, C. X., S. Webster, and S. Zhang (2014). Robust price-setting newsvendor model with interval market size and consumer willingness-to-pay. *International Journal of Production Economics 154*, 100–112.

Wiese, A. and W. Toporowski (2013). CSR failures in food supply chains–An agency perspective. *British Food Journal 115*(1), 92–107.

Williams, A. (1963). A stochastic transportation problem. *Operations Research 11*(5), 759–770.

Xu, H., D. D. Yao, and S. Zheng (2011). Optimal control of replenishment and substitution in an inventory system with nonstationary batch demand. *Production and Operations Management 20*(5), 727–736.

Xu, M., Y. F. Chen, and X. Xu (2010). The effect of demand uncertainty in a price-setting newsvendor model. *European Journal of Operational Research 207*(2), 946–957.

Yano, C. A. and S. M. Gilbert (2005). Coordinated pricing and production/procurement decisions: A review. In *Managing Business Interfaces*, pp. 65–103. Springer.

Yao, D. D., S. X. Zhou, and W. Zhuang (2016). Joint initial stocking and transshipment–Asymptotics and bounds. *Production and Operations Management 25*(2), 273–289.

Yao, L., Y. F. Chen, and H. Yan (2006). The newsvendor problem with pricing: extensions. *International Journal of Management Science and Engineering Management 1*(1), 3–16.

Young, L. (1978). Price, inventory and the structure of uncertain demand. *New Zealand Operations Research 6*(2), 157–177.

Yu, Y., X. Chen, and F. Zhang (2015). Dynamic capacity management with general upgrading. *Operations Research 63*(6), 1372–1389.

Yue, J., B. Chen, and M.-C. Wang (2006). Expected value of distribution information for the newsvendor problem. *Operations Research 54*(6), 1128–1136.

Zhang, D. and W. L. Cooper (2005). Revenue management for parallel flights with customer-choice behavior. *Operations Research 53*(3), 415–431.

Zhu, Z., J. Zhang, and Y. Ye (2013). Newsvendor optimization with limited distribution information. *Optimization Methods and Software 28*(3), 640–667.

APPENDICES

A. PROOFS OF FORMAL RESULTS

Proof of Theorem 2.3.1. The proof of Theorem 2.3.1 requires the following lemma.

Lemma A.0.1 *For* $t = 1, 2, \ldots, T$, *if* $W_t(\mathbf{z}_t, \mathbf{d}_t)$ *is concave in* $(\mathbf{z}_t, \mathbf{d}_t)$ *and* $\{\hat{S}_{t,i}(\mu), 0 \leq \mu \leq \bar{\mu}_{t,i}\} \in SL(mp)$ *for* $i = 1, 2, \ldots, n$, *then* $\hat{V}_t(\mathbf{x}_t, \mu_t)$ *is concave in* $(\mathbf{x}_t, \boldsymbol{\mu}_t)$ *for* $t \in \mathcal{C}$, *and* $V_t(\mathbf{x}_t)$ *is concave in* $\mathbf{x}_t$.

Proof. If $t \notin \mathcal{C}$, then $V_t(\mathbf{x}_t) = \mathbb{E}[W_t(\mathbf{x}_t, \mathbf{D}_t)]$ is concave in $\mathbf{x}_t$ because concavity is preserved under expectation. If $t \in \mathcal{C}$, we show that $\mathbb{E}[W_t(\mathbf{x}_t + \hat{\mathbf{S}}_t(\boldsymbol{\mu}_t), \mathbf{D}_t)]$ is concave in $(\mathbf{x}_t, \boldsymbol{\mu}_t)$. Consider $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{R}_n^+$ and $\boldsymbol{\mu}_1, \boldsymbol{\mu}_2 \in [0, \bar{\boldsymbol{\mu}}_t]$. By the concavity of W_t, given demand realization $\mathbf{d}_t$, we have

$$\tfrac{1}{2}\mathbb{E}\Big[W_t(\mathbf{x}_1 + \hat{\mathbf{S}}_t(\boldsymbol{\mu}_1), \mathbf{d}_t)\Big] + \tfrac{1}{2}\mathbb{E}\Big[W_t(\mathbf{x}_2 + \hat{\mathbf{S}}_t(\boldsymbol{\mu}_2), \mathbf{d}_t)\Big] \leq \mathbb{E}\Big[W_t\Big(\tfrac{1}{2}(\mathbf{x}_1 + \mathbf{x}_2 + \hat{\mathbf{S}}_t(\boldsymbol{\mu}_1) + \hat{\mathbf{S}}_t(\boldsymbol{\mu}_2)), \mathbf{d}_t\Big)\Big] \tag{A.1}$$

Since $\hat{S}_{t,i}(\mu_i) \in SL(mp)$, there exist $\hat{\hat{S}}_{t,i}(\mu_{i,1})$ and $\hat{\hat{S}}_{t,i}(\mu_{i,2})$ defined on a common probability space such that

$$\hat{\hat{S}}_{t,i}(\mu_{i,1}) \quad =^d \quad \hat{S}_{t,i}(\mu_{i,1}), \tag{A.2}$$

$$\hat{\hat{S}}_{t,i}(\mu_{i,2}) \quad =^d \quad \hat{S}_{t,i}(\mu_{i,2}), \tag{A.3}$$

and

$$\frac{1}{2}\hat{\hat{S}}_{t,i}(\mu_{i,1}) + \frac{1}{2}\hat{\hat{S}}_{t,i}(\mu_{i,2}) \leq_{cv} \hat{S}_{t,i}\Big(\frac{\mu_{i,1} + \mu_{i,2}}{2}\Big). \tag{A.4}$$

Since the supply processes $\hat{S}_{t,i}$, $i = 1, \ldots, n$ are independent, and $\hat{\hat{S}}_{t,i}$, $i = 1, \ldots, n$ can be constructed such that they are also independent, by Theorem 7.A.8 in Shaked and Shanthikumar (2007),

$$\frac{1}{2}\hat{\hat{\mathbf{S}}}_t(\boldsymbol{\mu}_1) + \frac{1}{2}\hat{\hat{\mathbf{S}}}_t(\boldsymbol{\mu}_2) \leq_{cv} \hat{\mathbf{S}}_t\Big(\frac{\boldsymbol{\mu}_1 + \boldsymbol{\mu}_2}{2}\Big).$$

Consequently, by the concavity of W_t, we have

$$\mathbb{E}\Big[W_t\Big(\tfrac{1}{2}(\mathbf{x}_1+\mathbf{x}_2)+\tfrac{1}{2}\hat{\mathbf{S}}_t(\boldsymbol{\mu}_1)+\tfrac{1}{2}\hat{\mathbf{S}}_t(\boldsymbol{\mu}_2),\mathbf{d}_t\Big)\Big] \le \mathbb{E}\Big[W_t\Big(\tfrac{1}{2}(\mathbf{x}_1+\mathbf{x}_2)+\hat{\mathbf{S}}_t\Big(\tfrac{\boldsymbol{\mu}_1+\boldsymbol{\mu}_2}{2}\Big),\mathbf{d}_t\Big)\Big]. \quad \text{(A.5)}$$

Combining the above equations,

$$\tfrac{1}{2}\mathbb{E}\Big[W_t\big(\mathbf{x}_1+\hat{\mathbf{S}}_t(\boldsymbol{\mu}_1),\mathbf{d}_t\big)\Big]+\tfrac{1}{2}\mathbb{E}\Big[W_t\big(\mathbf{x}_2+\hat{\mathbf{S}}_t(\boldsymbol{\mu}_2),\mathbf{d}_t\big)\Big] \le \mathbb{E}\Big[W_t\Big(\tfrac{1}{2}(\mathbf{x}_1+\mathbf{x}_2)+\hat{\mathbf{S}}_t\big(\tfrac{\boldsymbol{\mu}_1+\boldsymbol{\mu}_2}{2}\big),\mathbf{d}_t\Big)\Big]. \text{(A.6)}$$

Taking the expectation of both sides with respect to $\mathbf{D}_t$,

$$\tfrac{1}{2}\mathbb{E}\Big[W_t\big(\mathbf{x}_1+\hat{\mathbf{S}}_t(\boldsymbol{\mu}_1),\mathbf{D}_t\big)\Big]+\tfrac{1}{2}\mathbb{E}\Big[W_t\big(\mathbf{x}_2+\hat{\mathbf{S}}_t(\boldsymbol{\mu}_2),\mathbf{D}_t\big)\Big] \le \mathbb{E}\Big[W_t\Big(\tfrac{1}{2}(\mathbf{x}_1+\mathbf{x}_2)+\hat{\mathbf{S}}_t\big(\tfrac{\boldsymbol{\mu}_1+\boldsymbol{\mu}_2}{2}\big),\mathbf{D}_t\Big)\Big]. \text{(A.7)}$$

Since $\mathbb{E}[W_t(\mathbf{x}_t+\hat{\mathbf{S}}(\boldsymbol{\mu}_t),\mathbf{D}_t)]$ is continuous, the above midpoint concavity implies that it is concave in $(\mathbf{x}_t,\boldsymbol{\mu}_t)$. Consequently, $\hat{V}_t(\mathbf{x}_t,\boldsymbol{\mu}_t)$ is jointly concave in $(\mathbf{x}_t,\boldsymbol{\mu}_t)$ since all other terms in (2.7) are linear in $(\mathbf{x}_t,\boldsymbol{\mu}_t)$. Hence, $V_t(\mathbf{x}_t)$ is concave in $\mathbf{x}_t$ when $t\in\mathcal{C}$, as concavity is preserved under maximization. This completes the proof. □

We prove Theorem 2.3.1 by induction. First, consider the last period $t=T$. Since $V_{T+1}(\mathbf{x})$ is linear, $W_T(\mathbf{z}_T,\mathbf{d}_T)$ is concave. By Lemma A.0.1, we can verify that the claim in Theorem 2.3.1 holds for $t=T$. Assume the claim holds for $t=k+1$. Since $W_{k+1}(\mathbf{z}_{k+1},\mathbf{d}_{k+1})$ is concave in $(\mathbf{z}_{k+1},\mathbf{d}_{k+1})$ according to the induction hypothesis, by Lemma A.0.1, $\hat{V}_k(\mathbf{x}_k,\boldsymbol{\mu}_k)$ is concave in $(\mathbf{x}_k,\boldsymbol{\mu}_k)$ if $k\in\mathcal{C}$ and $V_k(\mathbf{x}_k)$ is concave in $\mathbf{x}_k$. Since $V_{k+1}(\mathbf{x}_{k+1})$ is concave in $\mathbf{x}_{k+1}$ by the induction hypothesis, the objective function in (2.3) is concave. Furthermore, since the constraints (2.4) and (2.5) are linear, the optimal profit function $W_k(\mathbf{z}_k,\mathbf{d}_k)$ is concave in $(\mathbf{z}_k,\mathbf{d}_k)$ (see, e.g., Rockafellar, 2015). Therefore, the claim in Theorem 2.3.1 holds for $t=k$. This completes the proof. □

Proof of Lemma 2.4.1. We first claim that given a matrix $\mathbf{R}$ with n rows and $n+k$ columns that has the reverse Monge property, for any column, $R_{\cdot,j}=\{r_{i,j}: i=1,2,\dots,n\}, 1\le j\le n+k$, if we insert a column, $\tilde{R}_{\cdot,j}=R_{\cdot,j}-\Delta\mathbf{1}$, where Δ is a constant, to $\mathbf{R}$ as the $(j+1)$-th column, then the extended matrix $\tilde{\mathbf{R}}$ also has the reserve Monge property. To prove the claim, for $1\le i_1<i_2\le n$ and $1\le j_1<j_2\le n+k$, we consider the following three cases: (i) $j_1\ne j+1$ and

$j_2 \neq j+1$. We have $\tilde{r}_{i_1,j_1} + \tilde{r}_{i_2,j_2} = r_{i_1,j_1} + r_{i_2,j_2} \geq r_{i_1,i_2} + r_{i_2,j_1} = \tilde{r}_{i_1,i_2} + \tilde{r}_{i_2,j_1}$ as the original benefit matrix $\mathbf{R}$ has the reverse Monge property. (ii) $j_1 = j+1$. $\tilde{r}_{i_1,j_1} + \tilde{r}_{i_2,j_2} = r_{i_1,j} - \Delta + r_{i_2,j_2} \geq r_{i_1,j_2} + r_{i_2,j} - \Delta = \tilde{r}_{i_1,j_2} + \tilde{r}_{i_2,j_1}$. The inequality holds because the original benefit matrix has the reverse Monge property. (iii) $j_2 = j+1$. By the similar argument in (ii), we can show that $\tilde{r}_{i_1,j_1} + \tilde{r}_{i_2,j_2} \geq \tilde{r}_{i_1,j_2} + \tilde{r}_{i_2,j_1}$. Therefore, we can conclude that the extended benefit matrix $\tilde{\mathbf{R}}$ has the reverse Monge property.

By applying the above claim for n times, we find that the extended matrix $\hat{\mathbf{R}}$ has the reverse Monge property. $\square$

Proof of Theorem 2.4.1. The proof of Theorem 2.4.1 requires the following lemma.

Lemma A.0.2 *The benefit maximization problem defined in the network G is equivalent to a min-cost problem defined in the network G' (see also the graphical illustration in Figure A.1), where the excess flows at each node are defined as,*

$$e(i) = \begin{cases} \sum_{i=1}^{n} x_i & \text{if } i = 2n+1, \\ 0 & \text{if } i = n+1, (n+1)', n+2, (n+2)', \ldots, 2n, (2n)', \\ -x_i & \text{if } i = 1, 2, \ldots, n, \end{cases}$$

and the arc costs and arc capacities are, respectively,

$$c_{i,j} = \begin{cases} \Delta - r_{j,i-n} & \text{if } i = n+1, n+2, \ldots, 2n, j = 1, 2, \ldots, n, \\ \Delta - r_{j,i-n} + \Delta_j & \text{if } i = (n+1)', (n+2)', \ldots, (2n)', j = 1, 2, \ldots, n, \\ 0 & \text{if } i = 2n+1, \end{cases}$$

$$u_{i,j} = \begin{cases} d_{j-n} & \text{if } i = 2n+1, j = n+1, n+2, \ldots, 2n, \\ \infty & \text{otherwise}, \end{cases}$$

and $\Delta = \max\{r_{i,j} : 1 \leq i, j \leq n\}$.

Proof. Let $\hat{f}_{i,j}$ denote the flow from node i to node j on the transportation arcs in the network G, where $i = 1, 2, \ldots, n$, $j = n+1, n+2, \ldots, 2n, (n+1)', (n+2)', \ldots, (2n)'$. Then, the benefit maximization problem defined in the network G can be written as,

$$\min_{\hat{\mathbf{f}} \in \mathbb{R}_+^{n \times 2n}} \quad (\mathbf{1}\Delta\mathbf{1}^\top - \hat{\mathbf{R}}) \odot \hat{\mathbf{f}} \tag{A.8}$$

$$\text{s.t.} \quad \sum_{j=n+1}^{2n} \hat{f}_{i,j} + \sum_{j=(n+1)'}^{(2n)'} \hat{f}_{i,j} = x_i, i = 1, 2, \ldots, n, \tag{A.9}$$

$$\sum_{i=1}^{n} \hat{f}_{i,j} \leq d_{j-n}, j = n+1, n+2, \ldots, 2n. \tag{A.10}$$

Let $\tilde{f}_{i,j}$ denote the flow from node i to node j on the transportation arcs in the network G', where $i = n+1, n+2, \ldots, 2n, (n+1)', (n+2)', \ldots, (2n)'$, $j = 1, 2, \ldots, n$. Let $\tilde{c}_{i,j} = \Delta - \hat{r}_{j,i}$ and $\tilde{\mathbf{C}} = \{\tilde{c}_{i,j} : i = n+1, n+2, \ldots, 2n, (n+1)', (n+2)', \ldots, (2n)', j = 1, 2, \ldots, n\}$, then the problem in (A.8)–(A.10) can be transformed to the following optimization problem,

$$\min_{\tilde{\mathbf{f}} \in \mathbb{R}_+^{2n \times n}} \quad \tilde{\mathbf{C}} \odot \tilde{\mathbf{f}} \tag{A.11}$$

$$\text{s.t.} \quad \sum_{j=1}^{n} \tilde{f}_{i,j} \leq d_{i-n}, i = n+1, n+2, \ldots, 2n, \tag{A.12}$$

$$\sum_{i=n+1}^{2n} \tilde{f}_{i,j} + \sum_{i=(n+1)'}^{(2n)'} \tilde{f}_{i,j} = x_j, j = 1, 2, \ldots, n. \tag{A.13}$$

Let $\hat{\mathbf{f}}^*$ and $\tilde{\mathbf{f}}^*$ denote the optimal solutions to (A.8)–(A.10) and (A.11)–(A.13) respectively. We have $\hat{\mathbf{f}}^* = \left(\tilde{\mathbf{f}}^*\right)^\top$ by the definition of these two optimization problems. In addition, the optimization problem defined in (A.11)–(A.13) is equivalent to the min-cost problem in the network G' defined in Lemma A.0.2, which is illustrated in Figure A.1. $\square$

i) Let $\mathbf{f}^* = \{f_{i,j} : i = 1, 2, \ldots, n, j = (n+1), (n+1)', \ldots, 2n, (2n)'\}$ denote the optimal max-benefit flow on the transportation arcs in the network G and $\mathbf{y}(\mathbf{f}^*) = \{y_{i,j} = f^*_{i,n+j} : 1 \leq i \leq n, 1 \leq j \leq n\}$. It is straightforward to verify that $\mathbf{y}(\mathbf{f}^*)$ is a feasible solution to the optimization problem defined in (2.9) as $\mathbf{f}^*$ is a feasible flow that satisfies capacity and balance constraints.

Based on the definition of the benefits on the transportation arcs in G, we have $\hat{r}_{i,j} = r_{i,j-n} - \Delta_{j-n} = r_{i,j-n} - r_{j-n,j-n} \leq 0$ for $j = (n+1)', (n+2)', \ldots, (2n)'$. Thus, $\sum_{i=1}^{n} \sum_{j=(n+1)'}^{(2n)'} \hat{r}_{i,j} f^*_{i,j} \leq 0$. Therefore, we have,

$$\hat{\mathbf{R}} \odot \mathbf{f}^* = \sum_{i=1}^{n} \sum_{j=n+1}^{2n} \hat{r}_{i,j} f^*_{i,j} + \sum_{i=1}^{n} \sum_{j=(n+1)'}^{(2n)'} \hat{r}_{i,j} f^*_{i,j} \leq \sum_{i=1}^{n} \sum_{j=n+1}^{2n} \hat{r}_{i,j} f^*_{i,j} = \mathbf{R} \odot \mathbf{y}(\mathbf{f}^*) \tag{A.14}$$

The last equality in the above equation holds by the definition of $\hat{\mathbf{R}}$ and $\mathbf{y}(\mathbf{f}^*)$. In addition, let $\mathbf{y}^*$ denote the optimal solution to the optimization problem defined in (2.9). we can construct a feasible solution $\tilde{\mathbf{f}}$ to the max-benefit flow problem in G based on $\mathbf{y}^*$. Specifically, let

$$\tilde{\mathbf{f}} = \begin{cases} \tilde{f}_{i,j} = y^*_{i,j-n} & \text{for } i = 1, 2, \ldots, n,\ j = n+1, n+2, \ldots, 2n, \\ \tilde{f}_{i,j} = x_i - \sum_{k \in N} y^*_{i,k} & \text{for } i = 1, 2, \ldots, n,\ j = (n+i)', \\ \tilde{f}_{j,2n+1} = \sum_{i \in N} y^*_{i,j-n} & \text{for } j = n+1, \ldots, 2n, \\ \tilde{f}_{j,2n+1} = x_i - \sum_{k \in N} y^*_{i,k} & \text{for } i = 1, 2, \ldots, n,\ j = (n+i)', \\ \tilde{f}_{i,j} = 0 & \text{o.w.,} \end{cases} \tag{A.15}$$

and it is straightforward to verify that $\tilde{\mathbf{f}}$ is a feasible flow in the network G by checking the capacity and balance constraints. By the definition of $\hat{\mathbf{R}}$ and the assumption that $\Delta_i = r_{i,i}$, we know that $\hat{r}_{i,(n+i)'} = 0$ for $i = 1, 2, \ldots, n$. Thus, we have $\hat{r}_{i,j} \tilde{f}_{i,j} = 0$ for $1 \leq i \leq n, j = (n+1)', (n+2)', \ldots, (2n)'$. Hence, $\sum_{i=1}^{n} \sum_{j=(n+1)'}^{(2n)'} \hat{r}_{i,j} \tilde{f}_{i,j} = 0$. Therefore, we have

$$\begin{aligned} \mathbf{R} \odot \mathbf{y}^* = \sum_{i \in N} \sum_{j \in N} r_{i,j} y_{i,j} &= \sum_{i=1}^{n} \sum_{j=n+1}^{2n} \hat{r}_{i,j} \tilde{f}_{i,j} + \sum_{i=1}^{n} \sum_{j=(n+1)'}^{(2n)'} \hat{r}_{i,j} \tilde{f}_{i,j} && \text{(A.16)} \\ &\leq \sum_{i=1}^{n} \sum_{j=n+1}^{2n} \hat{r}_{i,j} f^*_{i,j} + \sum_{i=1}^{n} \sum_{j=(n+1)'}^{(2n)'} \hat{r}_{i,j} f^*_{i,j} && \text{(A.17)} \\ &\leq \mathbf{R} \odot \mathbf{y}(\mathbf{f}^*). && \text{(A.18)} \end{aligned}$$

The first inequality holds due to the definition of $\mathbf{f}^*$. Thus, by the above observations, we have,

$$\mathbf{R} \odot \mathbf{y}(\mathbf{f}^*) \geq \mathbf{R} \odot \mathbf{y}^*. \tag{A.19}$$

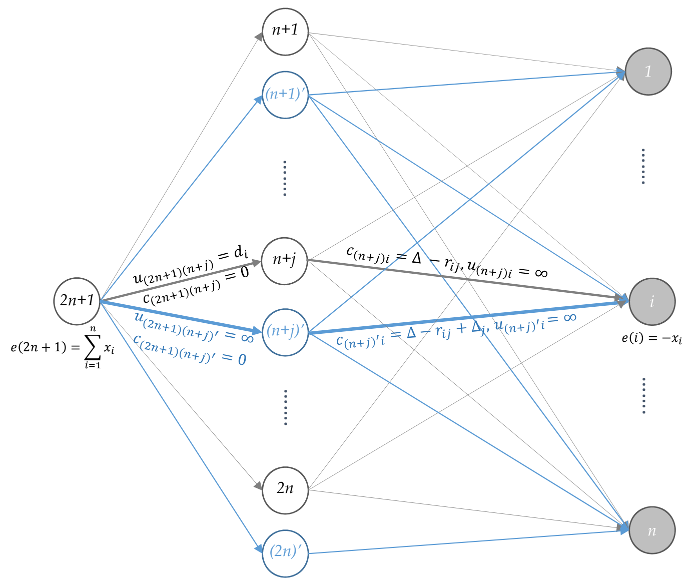

Fig. A.1.: Minimize the total cost in the network G'

By the definition of $\mathbf{y}^*$, which is an optimal solution to (2.9), we have $\mathbf{R} \odot \mathbf{y}^* \geq \mathbf{R} \odot \mathbf{y}(\mathbf{f}^*)$. Thus, we have $\mathbf{R} \odot \mathbf{y}(\mathbf{f}^*) = \mathbf{R} \odot \mathbf{y}^*$. Therefore, $\mathbf{y}(\mathbf{f}^*)$ is the optimal solution to the optimization problem defined in (2.9).

ii) By Lemma 2.4.1 and the definition of $\hat{\mathbf{R}}$, the benefit matrix $\hat{\mathbf{R}}$ also has the reverse Monge property. Thus, the cost matrix $\tilde{\mathbf{C}}$ has the Monge property. Therefore, the following algorithm 2 described in Vaidyanathan (2013) can solve optimally the problem in G'. Finally, by Lemma A.0.2, Algorithm 2 finds the optimal solution in the network G.

By Vaidyanathan (2013), the run time is $O(m\log n_1)$ where $m = n_1 n_2$ is the number of transportation arcs and n_1 is the number of supply nodes in the min-cost flow network G'. The number of transportation arcs in the network G' is $2n^2$ and the number of supply nodes is $2n$, where n is the number of products in the product substitution setting. Therefore, Algorithm 2 can solve the optimization problem defined in G with a run time of $O(n^2\log n)$. $\square$

Algorithm 2 Vaidyanathan (2013)

Initialize the excess flow $\mathbf{e}$

$e(0) = \sum_{i=1}^{n} x_{t,i}$

for $i = 1, 2, \ldots, n$ **do**

  $e(i) = -x_i$

end for

Initialize the flow $\mathbf{f}$ *in network* G'

$\mathbf{f} = \mathbf{0}$

Successive shortest path algorithm

for $k = 1, 2, \ldots, n$ **do**

  while $e(k) < 0$ **do**

    find the shortest non-crossing path P from 0 to k;

    compute $\delta = \min\{u_{i,j} : (i,j) \in P\}$;

    augment $\delta_f = \min\{\delta, -e(k)\}$ along the path P, for $(i,j) \in P$, update $u_{i,j} = u_{i,j} - \delta_f$, $u_{j,i} = u_{j,i} + \delta_f$, and $f_{i,j} = f_{i,j} + \delta_f$;

    update $e(k) = e(k) + \delta_f$, $e(0) = e(0) - \delta_f$.

  end while

end for

Proof of Lemma 2.5.1. Similar to the proof of Lemma 2.4.1, we can show that $s_l^{-1}\mathbf{R}^{(l)}$ has the reverse Monge property. By multiplying with a positive constant s_l, the benefit matrix $\mathbf{R}^{(l)}$ also has the reverse Monge property. □

Proof of Lemma 2.5.2. Let $\delta_{i,0} \in [0, \underline{x}^{(l)}_{t+1,i}]$ and $\delta_{i,k} \in [0,1]$, $k = 0, 1, \ldots, K$, denote the decision variable corresponding to the $(k+1)$-st segment of the piecewise linear function $V^{s_l}_{t+1,i}(x_i)$ and define $\boldsymbol{\delta}_k = (\delta_{1,k}, \delta_{2,k}, \ldots, \delta_{n,k})$ and $\boldsymbol{\delta} = (\boldsymbol{\delta}_0, \boldsymbol{\delta}_1, \ldots, \boldsymbol{\delta}_K)$. We define the following linear program.

$$\begin{aligned}
\max \quad & \mathbf{R} \odot (s_l \mathbf{y}_t) - \mathbf{M} \odot (s_l \tilde{\mathbf{y}}_t) + \sum_{i\in N}\sum_{k=0}^{K} \left(s_l r_{i,i} - \Delta^{(l)}_{i,k}\right)\delta_{i,k} && \text{(A.20)}\\
s.t. \quad & \lfloor \mathbf{z}_t / s_l \rfloor - \mathbf{y}_t \mathbf{1} - \tilde{\mathbf{y}}_t \mathbf{1} = \mathbf{0}, && \text{(A.21)}\\
& \lfloor \mathbf{d}_t / s_l \rfloor - (\mathbf{y}_t)^\top \mathbf{1} \geq \mathbf{0}, && \text{(A.22)}\\
& (\tilde{\mathbf{y}}_t)^\top \mathbf{1} - \underline{\mathbf{x}}^{(l)}_{t+1} - \sum_{k=0}^{K} \boldsymbol{\delta}_k = \mathbf{0}, && \text{(A.23)}\\
& \mathbf{0} \leq \boldsymbol{\delta}_k \leq \mathbf{1}, k = 1, 2, \ldots, K, && \text{(A.24)}\\
& \mathbf{0} \leq \boldsymbol{\delta}_0 \leq \underline{\mathbf{x}}^{(l)}_{t+1}, && \text{(A.25)}\\
& \mathbf{y}_t, \tilde{\mathbf{y}}_t \in \mathbb{R}^{n\times n}_{+}. && \text{(A.26)}
\end{aligned}$$

The optimal solution to (2.14)–(2.18) with the piecewise linear functions $V^{s_l}_{t+1,i}(x_i), i \in N$ are multiples of s_l as the constraint matrix is totally unimodular, and, by Theorem 3.7 in Hochbaum and Shanthikumar (1990), satisfies $s_l \underline{\mathbf{x}}^{(l)}_{t+1} \leq \tilde{\mathbf{x}}_{t+1} \leq s_l \bar{\mathbf{x}}^{(l)}_{t+1}$. Thus, the optimal solution to the above problem satisfies $s_l \sum_{k=0}^{K} \delta_{i,k} \geq s_l \underline{x}^{(l)}_{t+1,i}$ for $i \in N$. In addition, since $\tilde{V}_{t+1,i}(x_i)$ is concave in x_i, $(s_l r_{i,i} - \Delta^{(l)}_{i,k})$ is decreasing in k. We can verify that $\delta_{i,0} \geq \delta_{i,1} \geq \ldots \geq \delta_{i,K}$, and consequently, $\boldsymbol{\delta}_0 = \underline{\mathbf{x}}^{(l)}_{t+1}$. Therefore, the above optimization problem is equivalent to the problem $(\mathrm{IP} - s_l)$.

Next, we show that the optimal solution to the above optimization problem is connected with the max-benefit flow in the network G^{s_l}. Let $\mathbf{f}$ denote the flow in G^{s_l}. We can construct a solution to the above linear program based on $\mathbf{f}$. Specifically,

$$y_t^{i,j}(\mathbf{f}) = f_{i,n+j}, \ i = 1,2,\ldots,n, \ j = 1,2,\ldots,n, \tag{A.27}$$

$$\tilde{y}_t^{i,j}(\mathbf{f}) = \sum_{k=0}^{K} f_{i,(n+j)_k}, \ i = 1,2,\ldots,n, \ j = 1,2,\ldots,n, \tag{A.28}$$

$$\delta_{j,k}(\mathbf{f}) = \sum_{i \in N} f_{i,(n+j)_k}, \ j = 1,2,\ldots,n, \ k = 0,1,2,\ldots,K. \tag{A.29}$$

Since $\mathbf{f}$ satisfies the balance constraints at the excess nodes and the capacity constraints at the demand nodes, it is easy to verify that $\big(\mathbf{y}_t(\mathbf{f}), \tilde{\mathbf{y}}_t(\mathbf{f}), \boldsymbol{\delta}(\mathbf{f})\big)$ satisfies the constraints of the problem defined in (A.20)–(A.26). Moreover, the value of the objective function with $\big(\mathbf{y}_t(\mathbf{f}), \tilde{\mathbf{y}}_t(\mathbf{f}), \boldsymbol{\delta}(\mathbf{f})\big)$ in (A.20) is

$$\mathbf{R} \odot \big(s_l \mathbf{y}_t(\mathbf{f})\big) - \mathbf{M} \odot \big(s_l \tilde{\mathbf{y}}_t(\mathbf{f})\big) + \sum_{i \in N} \sum_{k=0}^{K} \Big(s_l r_{i,i} - \Delta_{i,k}^{(l)}\Big) \delta_{i,k}(\mathbf{f}) \tag{A.30}$$

$$= \sum_{i \in N} \sum_{j \in N} r_{i,j} f_{i,n+j} + \sum_{i \in N} \sum_{j \in N} \sum_{k=0}^{K} \Big(s_l r_{i,j} - \Delta_{j,k}^{(l)}\Big) f_{i,(n+j)_k}, \tag{A.31}$$

where the left side is equal to the total benefit of the flow $\mathbf{f}$ in the network G^{s_l}. Let $\mathbf{f}^*$ denote the optimal benefit-flow in the network G^{s_l}. Then, the solution $\big(\mathbf{y}_t(\mathbf{f}^*), \tilde{\mathbf{y}}_t(\mathbf{f}^*), \boldsymbol{\delta}(\mathbf{f}^*)\big)$ that we construct based on $\mathbf{f}^*$ as described in (A.27)–(A.29) satisfies the constraints and maximizes the objective value of the problem defined in (A.20)–(A.26). Thus, we can conclude that the optimal flow in the network G^{s_l} corresponds to the optimal solution for the problem $(IP - s_l)$. $\square$

Proof of Theorem 2.5.1. By Hochbaum and Shanthikumar (1990), the number of optimization problems solved in Algorithm 1 is $\lceil \log_2\big(\frac{2}{2n+1} \|(\mathbf{z}_t, \mathbf{d}_t)\|_\infty\big)\rceil$. By Lemmas 2.5.1 and 2.5.2, each of these problems can be solved by the network-flow-based algorithm of Vaidyanathan (2013) with a run time of $O(n^4 \log n)$. Thus, the run time of Algorithm 1 is $O\Big(n^4 \log n \log\big(\frac{2}{2n+1} \|(\mathbf{z}_t, \mathbf{d}_t)\|_\infty\big)\Big)$. $\square$

Proof of Proposition 3.4.1. The hindsight profit function $\phi(p,\theta)$ is linear in θ and the maximum of the linear functions of θ is convex in θ. When $\theta_1 \geq \theta_2$, we have $\phi(p, \theta_1) \geq \phi(p, \theta_2)$ for any $p \in [\underline{p}, \bar{p}]$ as $\sigma(p) \geq 0$. Thus, we have $\phi^*(\theta_1) = \max_{\hat{p}} \phi(\hat{p}, \theta_1) \geq \max_{\hat{p}} \phi(\hat{p}, \theta_2) = \phi^*(\theta_2)$. We can conclude that $\phi^*(\theta)$ is increasing and convex in θ.

In addition, $p \min\{D(p, \theta), y\}$ is concave in θ since $D(p, \theta)$ is linear in θ and $h(x) = \min\{x, y\}$ is concave in x. Therefore, the regret function $R(p, y; \theta)$ is convex in θ. To maximize the convex function over the interval $[\underline{\theta}, \bar{\theta}]$, the nature will choose either the upper bound or the lower bound of the interval whichever will lead to a higher regret. Thus, the maximum regret $R(p, y)$ is the maximum of the regret function evaluated at $\underline{\theta}$ and $\bar{\theta}$. □

Proof of Proposition 3.4.2. Let $h_1(p, y) = \phi^*(\bar{\theta}) - p \min\{D(p, \bar{\theta}), y\} + cy$, $h_2(p, y) = \phi^*(\underline{\theta}) - p \min\{D(p, \underline{\theta}), y\} + cy$, and $H(p, y) = h_1(p, y) - h_2(p, y)$. To solve the inner optimization of (7), we restrict the range of y to $[D(p, \underline{\theta}), D(p, \bar{\theta})]$, otherwise it would incur additionally either underage cost or overage cost. We first observe that $h_1(p, y)$ is decreasing in y while $h_2(p, y)$ is increasing in y over the interval $[D(p, \underline{\theta}), D(p, \bar{\theta})]$ when the unit selling price p is fixed. In addition, we can evaluate $H(p, y)$ when $y = D(p, \underline{\theta})$:

$$H\big(p, D(p, \underline{\theta})\big) = \phi^*(\bar{\theta}) - \phi^*(\underline{\theta}) > 0. \tag{A.32}$$

Consequently, we have $h_1(p, y) > h_2(p, y)$ when $y = D(p, \underline{\theta})$. Therefore, the optimal choice of order quantity given a fixed unit selling price p depends on the sign of $H\big(p, D(p, \bar{\theta})\big)$. In other words, if $H\big(p, D(p, \bar{\theta})\big) > 0$, then $h_1(p, y) > h_2(p, y)$ over the interval $[D(p, \underline{\theta}), D(p, \bar{\theta})]$ and thus, in this case, the optimal order quantity would be $y^*(p) = D(p, \bar{\theta})$. If $H\big(p, D(p, \bar{\theta})\big) < 0$, we know that $h_1(p, y)$ and $h_2(p, y)$ will cross once over the interval $[D(p, \underline{\theta}), D(p, \bar{\theta})]$ for a fixed unit selling price p, and the optimal order quantity will satisfy the following:

$$h_1(p, y^*) = h_2(p, y^*). \tag{A.33}$$

Then, we need to check the sign of $H\big(p, D(p, \bar{\theta})\big)$. Specifically, we have,

$$\begin{aligned} H\big(p, D(p, \bar{\theta})\big) &= \phi^*(\bar{\theta}) - \phi^*(\underline{\theta}) - p\big(D(p, \bar{\theta}) - D(p, \underline{\theta})\big) && \text{(A.34)} \\ &= \phi^*(\bar{\theta}) - \phi^*(\underline{\theta}) - p\sigma(p)(\bar{\theta} - \underline{\theta}). && \text{(A.35)} \end{aligned}$$

From the above analysis, the sign of $H\big(p, D(p,\bar{\theta})\big)$ depends on the unit selling price p. Therefore, we can check the sign of $H\big(p, D(p,\bar{\theta})\big)$ over different regions for the unit selling price $p \in [\underline{p}, \bar{p}]$. Evaluating $H\big(p, D(p,\bar{\theta})\big)$ at $\bar{p}^*$, we have:

$$
\begin{aligned}
H\big(\bar{p}^*, D(\bar{p}^*,\bar{\theta})\big) &= \phi^*(\bar{\theta}) - \phi^*(\underline{\theta}) - \bar{p}^*\big(D(\bar{p}^*,\bar{\theta}) - D(\bar{p}^*,\underline{\theta})\big) \\
&= \phi^*(\bar{\theta}) - \phi^*(\underline{\theta}) - (\bar{p}^* - c)D(\bar{p}^*,\bar{\theta}) + \bar{p}^* D(\bar{p}^*,\underline{\theta}) + cD(\bar{p}^*,\bar{\theta}) \\
&= (\bar{p}^* - c)D(\bar{p}^*,\underline{\theta}) - \phi^*(\underline{\theta}) + c\big(D(\bar{p}^*,\underline{\theta}) - D(\bar{p}^*,\bar{\theta})\big) \\
&< 0.
\end{aligned}
$$

Since $H\big(\bar{p}^*, D(\bar{p}^*,\bar{\theta})\big) < 0$ and $p\sigma(p)$ is assumed concave in p, there exists at most two solutions to the equation, $H\big(p, D(p,\bar{\theta})\big) = 0$. Subsequently, we have four cases to analyze:

Case 1: $H\big(p, D(p,\bar{\theta})\big) = 0$ has zero solutions on the interval $[\underline{p}, \bar{p}]$. In this case, we have $H\big(p, D(p,\bar{\theta})\big) < 0, \forall p \in [\underline{p}, \bar{p}]$, and thus, we set $p_l = \underline{p}$ and $p_h = \bar{p}$. Since $H\big(p, D(p,\underline{\theta})\big) > 0$ and $H\big(p, D(p,\bar{\theta})\big) < 0$, there exists a unique $y^* \in \big(D(p,\bar{\theta}), D(p,\underline{\theta})\big)$ such that $H(p, y^*) = 0$. Specifically,

$$\phi^*(\bar{\theta}) - p\min\{D(p,\bar{\theta}), y^*\} + cy^* = \phi^*(\underline{\theta}) - p\min\{D(p,\underline{\theta}), y^*\} + cy^*. \quad \text{(A.36)}$$

By solving the above equation given a fixed unit selling price p, we find the optimal order quantity, $y^*(p) = \frac{1}{p}\big(\phi^*(\bar{\theta}) - \phi^*(\underline{\theta})\big) + D(p,\underline{\theta})$. The corresponding regret is $R\big(p, y^*(p)\big) = \frac{c}{p}\phi^*(\bar{\theta}) + (1 - \frac{c}{p})\phi^*(\underline{\theta}) - (p - c)D(p,\underline{\theta}), \forall p \in [p_l, p_h]$.

Case 2: $H\big(p, D(p,\bar{\theta})\big) = 0$ has one solution $\hat{p}$ on the interval $[\underline{p}, \bar{p}]$ and $\hat{p} < \bar{p}^*$. Then, $H\big(p, D(p,\bar{\theta})\big) > 0$ on $(\underline{p}, \hat{p})$ and $H\big(p, D(p,\bar{\theta})\big) < 0$ on $[\hat{p}, \bar{p}]$. In this case, we set $p_l = \hat{p}$ and $p_h = \bar{p}$. Therefore, when $p \in [p_l(\hat{p}), p_h(\bar{p})]$, by similar reasoning as in *Case 1*, the optimal order quantity, $y^*(p) = \frac{1}{p}\big(\phi^*(\bar{\theta}) - \phi^*(\underline{\theta})\big) + D(p,\underline{\theta})$, and the regret function, $R\big(p, y^*(p)\big) = \frac{c}{p}\phi^*(\bar{\theta}) + (1 - \frac{c}{p})\phi^*(\underline{\theta}) - (p - c)D(p,\underline{\theta})$. When $p \in [\underline{p}, p_l)$, we have $H\big(p, D(p,\bar{\theta})\big) > 0$. Therefore, $h_1(p, y) \geq h_2(p, y)$ over the interval $[D(p,\underline{\theta}), D(p,\bar{\theta})]$. Thus, the optimal order quantity $y^*(p) = D(p,\bar{\theta})$, and the regret function is $R\big(p, y^*(p)\big) = h_1\big(p, D(p,\bar{\theta})\big) = \phi^*(\bar{\theta}) - \phi(p,\bar{\theta})$.

Case 3: $H\big(p, D(p,\bar{\theta})\big) = 0$ has one solution $\hat{p}$ on the interval $[\underline{p}, \bar{p}]$ and $\hat{p} > \bar{p}^*$. In this case, we set $p_l = \underline{p}$ and $p_h = \hat{p}$. The rest of analysis will be the same as in *Case 2*.

Case 4: $H\big(p, D(p,\bar{\theta})\big) = 0$ has two solution $\hat{p}$ and $\hat{\hat{p}}$ on the interval $[\underline{p}, \bar{p}]$ where $\underline{p} \le \hat{p} < \bar{p}^* < \hat{\hat{p}} \le \bar{p}$. In this case, we set $p_l = \hat{p}$ and $p_h = \hat{\hat{p}}$. Then, we have $H\big(p, D(p,\bar{\theta})\big) < 0$ on $(\hat{p}, \hat{\hat{p}})$ and $H\big(p, D(p,\bar{\theta})\big) \ge 0$ elsewhere. Thus, when $p \in [p_l, p_h]$, the optimal order quantity, $y^*(p) = \frac{1}{p}\big(\phi^*(\bar{\theta}) - \phi^*(\underline{\theta})\big) + D(p,\underline{\theta})$. And the corresponding regret is $R\big(p, y^*(p)\big) = \frac{c}{p}\phi^*(\bar{\theta}) + (1-\frac{c}{p})\phi^*(\underline{\theta}) + (p-c)D(p,\underline{\theta})$. Otherwise, when $p \notin [\hat{p}, \hat{\hat{p}}]$, we have $y^* = D(p,\bar{\theta})$ and the regret function is $R\big(p, y^*(p)\big) = h_1\big(p, D(p,\bar{\theta})\big) = \phi^*(\bar{\theta}) - \phi(p,\bar{\theta})$.

By the above analysis, we can summarize the regret function $R\big(p, y^*(p)\big)$ as follows:

$$R\big(p, y^*(p)\big) = \begin{cases} R_1(p) & \text{if } p \in [p_l, p_h], \\ R_2(p) & \text{o.w.}, \end{cases} \tag{A.37}$$

where $R_1(p) = \frac{c}{p}\phi^*(\bar{\theta}) + (1-\frac{c}{p})\phi^*(\underline{\theta}) - (p-c)D(p,\underline{\theta})$ and $R_2(p) = \phi^*(\bar{\theta}) - (p-c)D(p,\bar{\theta})$. Since we assume the concavity of $\phi(p,\theta)$ and $f(p) = \frac{1}{p}\big(\phi^*(\bar{\theta}) - \phi^*(\underline{\theta})\big)$ is convex in p, both $R_1(p)$ and $R_2(p)$ are convex in p. Furthermore, $R_2(p) - R_1(p) = \frac{p-c}{p}\big(\phi^*(\bar{\theta}) - \phi^*(\underline{\theta}) - p\sigma(p)(\bar{\theta} - \underline{\theta})\big)$. By the definition of p_l and p_h, we have $R_2(p) - R_1(p) \ge 0$ when $p \in [\underline{p}, p_l] \cup [p_h, \bar{p}]$, and $R_2(p) - R_1(p) \le 0$ when $p \in [p_l, p_h]$. Therefore, we have $R\big(p, y^*(p)\big) = \max\{R_1(p), R_2(p)\}$. Since the maximum of convex functions is also a convex function, $R\big(p, y^*(p)\big)$ is convex in $p \in [\underline{p}, \bar{p}]$.

In addition, when $p \in [\underline{p}, p_l]$, $R\big(p, y^*(p)\big) = \phi^*(\bar{\theta}) - \phi(p,\bar{\theta})$ is decreasing in p as $p_l \le \bar{p}^*$ by definition. Thus, the unit selling price that minimizes the regret of the firm on $[\underline{p}, p_l]$ is p_l. Similarly, when $p \in [p_h, \bar{p}]$, $R\big(p, y^*(p)\big)$ is minimized at p_h. Thus, we can conclude that the optimal unit selling price p^* lies in the interval $[p_l, p_h]$. $\square$

Proof of Proposition 3.5.1. Before we prove Proposition 3.5.1, we first introduce Lemma A.0.3 that will be used in the proof for Lemma A.0.4, which characterizes the relationship between $\underline{p}^*$ and $\bar{p}^*$.

Lemma A.0.3 *Let* $\phi_1(p) = (p-c)\mu(p)$ *and* $\phi_2(p) = (p-c)\sigma(p)$. *Define* $p_\theta = \arg\max\{p : \phi(p,\theta)\}$ *where* $\theta \neq 0$ *and* $\tilde{p}_i = \arg\max\{p : \phi_i(p)\}$ *where* $i = 1,2$. *If* $\phi_2'(\tilde{p}_1) \geq 0$, *then* $\frac{\phi_1'(p_\theta)}{\theta} \leq 0$. *If* $\phi_2'(\tilde{p}_1) \leq 0$, *then* $\frac{\phi_1'(p_\theta)}{\theta} \geq 0$.

Proof. We first consider the case where $\phi_2'(\tilde{p}_1) \geq 0$ and show the following claim:

Claim 1: $\phi'(\tilde{p}_1,\theta) \geq 0$ when $\theta > 0$ and $\phi'(\tilde{p}_1,\theta) < 0$ when $\theta < 0$.

When $\theta > 0$, we have:

$$\phi'(p,\theta) = \phi_1'(p) + \phi_2'(p)\theta. \tag{A.38}$$

$$\begin{aligned} \phi'(\tilde{p}_1,\theta) &= \phi_1'(\tilde{p}_1) + \phi_2'(\tilde{p}_1)\theta && \text{(A.39)} \\ &= \phi_2'(\tilde{p}_1)\theta && \text{(A.40)} \\ &\geq 0. && \text{(A.41)} \end{aligned}$$

We can conduct similar analysis for the case when $\theta < 0$. Thus, the *Claim 1* is verified.

When $\theta > 0$, we have $\phi'(\tilde{p}_1,\theta) \geq 0$ by Claim 1. Then, we have $p_\theta > \tilde{p}_1$ due to the concavity of ϕ and thus, $\phi'(p_\theta,0) = \phi_1'(p_\theta) < \phi_1'(\tilde{p}_1) = 0$. By similar reasoning, when $\theta < 0$, we have $\phi'(\tilde{p}_1,\theta) \leq 0$ and $p_\theta < \tilde{p}_1$. Consequently, $\phi'(p_\theta,0) > 0$. Therefore we can conclude that if $\phi_2'(\tilde{p}_1) \geq 0$, we have $\frac{\phi_1(p_\theta)}{\theta} = \frac{\phi'(p_\theta,0)}{\theta} \leq 0$.

For the case where $\phi_2'(\tilde{p}_1) \leq 0$ satisfies, we can conduct similar analysis and conclude that $\frac{\phi'(p_\theta,0)}{\theta} \geq 0$. □

Lemma A.0.4 *Let* $\phi_1(p) = (p-c)\mu(p)$ *and* $\phi_2(p) = (p-c)\sigma(p)$. *Let* $\tilde{p}_1 = \arg\max\{p : \phi_1(p)\}$ *and* $p_\theta = \arg\max\{p : \phi(p,\theta)\}$. *If* $\phi_2'(\tilde{p}_1) \geq 0$, *then* p_θ *is increasing in* θ. *Consequently we have* $\underline{p}^* \leq \bar{p}^*$. *If* $\phi_2'(\tilde{p}_1) \leq 0$, *then* p_θ *is decreasing in* θ. *Consequently, we have* $\underline{p}^* \geq \bar{p}^*$.

Proof. When $\phi_2'(\tilde{p}_1) \geq 0$ holds, we first consider the case where $\theta_2 > \theta_1 \neq 0$. Then, let $p_i^* = \arg\max\{p : \phi(p,\theta_i)\}$ where $i = 1,2$. Then, p_1^* satisfies the following:

$$\phi_1'(p_1^*) + \phi_2'(p_1^*)\theta_1 = 0. \tag{A.42}$$

Then, we evaluate $\phi'(p,\theta_2)$ at p_1^*:

$$\begin{aligned}
\phi'(p,\theta_2) &= \phi'(p,\theta_1)+\phi_2'(p)(\theta_2-\theta_1). && \text{(A.43)}\\
\phi'(p_1^*,\theta_2) &= \phi_1'(p_1^*,\theta_1)+\phi_2'(p_1^*)(\theta_2-\theta_1) && \text{(A.44)}\\
&= \frac{\phi_1'(p_1^*)}{-\theta_1}(\theta_2-\theta_1) && \text{(A.45)}\\
&\geq 0. && \text{(A.46)}
\end{aligned}$$

The second equality holds because of (A.42). The last inequality holds because of Lemma A.0.3. Finally, we conclude that if $\phi'(p_1^*,\theta_2) > 0$ then $p_2^* \geq p_1^*$ when $\theta_2 > \theta_1 \neq 0$. When $\theta_2 > \theta_1 = 0$, by the results in Lemma A.0.3, $\frac{\phi'(p_2^*,\theta_1)}{\theta_2} = \frac{\phi_1'(p_2^*)}{\theta_2} \leq 0$ as $\phi(p,\theta_1) = \phi(p,0) = \phi_1(p)$. Thus, we have, $\phi'(p_2^*,\theta_1) \leq 0$ if $\theta_2 > \theta_1 = 0$ and consequently, $p_2^* \geq p_1^*$. Combining the above arguments, we conclude that if $\theta_2 \geq \theta_1$, we have $p_2^* \geq p_1^*$. Consequently, we have $\bar{p}^* \geq \underline{p}^*$ if $\phi_2'(\tilde{p}_1) \geq 0$ holds.

For the case where $\phi_2'(\tilde{p}_1) \leq 0$ holds, we can conduct similar analysis and conclude that p_θ is decreasing in θ and consequently $\bar{p}^* \leq \underline{p}^*$. $\square$

To prove Proposition 3.5.1, WLOG, we assume that $\bar{p}^* \geq \underline{p}^*$. It can be verified that the objective function of the inner optimization (3.9) is convex in θ. The adversarial nature would choose either $\underline{\theta}$ or $\bar{\theta}$ whichever leads to a larger regret of the firm. Thus, the worst-case regret is $\max\{\phi^*(\bar{\theta})-\phi(p,\bar{\theta}), \phi^*(\underline{\theta})-\phi(p,\underline{\theta})\}$. Let $g_1(p) = \phi^*(\bar{\theta})-\phi(p,\bar{\theta})$ and $g_2(p) = \phi^*(\underline{\theta}) - \phi(p,\underline{\theta})$. It can be verified that $g_1(\underline{p}^*) > 0$ and $g_2(\bar{p}^*) > 0$. Additionally, we know that $g_1(\bar{p}^*) = g_2(\underline{p}^*) = 0$. Moreover, $g_1(p)$ is decreasing in p and $g_2(p)$ is increasing in p over the interval $[\underline{p}^*, \bar{p}^*]$. Thus, $g_1(p)$ and $g_2(p)$ will cross once over the interval $[\underline{p}^*, \bar{p}^*]$ and the optimal risk-free price p_0 must satisfy the following:

$$\phi^*(\bar{\theta}) - \phi(p_0,\bar{\theta}) = \phi^*(\underline{\theta}) - \phi(p_0,\underline{\theta}), p_0 \in [\underline{p}^*, \bar{p}^*]. \tag{A.47}$$

Thus, we have,

$$p_0 = \arg\{p : (p-c)\sigma(p) = \frac{\phi^*(\bar{\theta}) - \phi^*(\underline{\theta})}{\bar{\theta}-\underline{\theta}}, p \in [\underline{p}^*, \bar{p}^*]\}. \tag{A.48}$$

When the demand function is of the form $D(p,\theta) = \mu(p) + \theta$ and $\sigma(p) = 1$, from the previous analysis we have $p_0 = c + \frac{\phi^*(\bar{\theta})-\phi^*(\underline{\theta})}{\bar{\theta}-\underline{\theta}}$. When the demand is of the form

$D(p,\theta) = \sigma(p)\theta$, it can be shown that $\bar{p}^* = \underline{p}^*$. Moreover, for the risk-free price, p_0, we have the following:

$$\phi^*(\bar{\theta}) - \phi(p^0, \bar{\theta}) = \phi^*(\underline{\theta}) - \phi(p^0, \underline{\theta}). \tag{A.49}$$

$$(\bar{p}^* - c)\sigma(\bar{p}^*)\bar{\theta} - (\underline{p}^* - c)\sigma(\underline{p}^*)\underline{\theta} = (p_0 - c)\sigma(p^0)(\bar{\theta} - \underline{\theta}). \tag{A.50}$$

Since $\bar{p^*} = \underline{p}^*$ and $p\sigma(p)$ is concave in p, the risk-free price with the multiplicative demand, $p_0 = \bar{p}^* = \underline{p}^*$. □

Proof of Proposition 3.5.2. When the demand is of the form $D(p,\theta) = \mu(p) + \theta$ and the demand function $\mu(p,\theta)$ is assumed to be linearly decreasing in p, we can explicitly express the demand function as $D(p,\theta) = A - Bp + \theta$ where $A, B > 0$. From the proof of Proposition 3.4.2, we have $\sigma(p) = 1$ and thus $p_0 = c + \frac{\phi^*(\bar{\theta}) - \phi^*(\underline{\theta})}{\bar{\theta} - \underline{\theta}}$ with the additive demand. WLOG, we assume that the firm sets the lower bound of the price range $\underline{p} \leq c + \frac{\phi^*(\bar{\theta}) - \phi^*(\underline{\theta})}{\bar{\theta} - \underline{\theta}}$, otherwise it will be trivial to conclude that $p^* \geq \underline{p} \geq p_0$. Therefore, we have $p_l = \max\{\underline{p}, \frac{\phi^*(\bar{\theta}) - \phi^*(\underline{\theta})}{\bar{\theta} - \underline{\theta}}\} < c + \frac{\phi^*(\bar{\theta}) - \phi^*(\underline{\theta})}{\bar{\theta} - \underline{\theta}} = p_0$. We then have two subcases to consider. If $p_0 \geq p_h$, we can conclude that $p^* \leq p_h \leq p_0$. Otherwise, we have $p_0 \in [p_l, p_h]$ and thus, the derivative of the regret function at the risk-free price is calculated as follows. Recall that the regret function on the interval $[p_l, p_h]$ is of the form:

$$R\big(p, y^*(p)\big) = \frac{c}{p}\phi^*(\bar{\theta}) + \big(1 - \frac{c}{p}\big)\phi^*(\underline{\theta}) - (p - c)D(p, \underline{\theta}). \tag{A.51}$$

We take the first-order derivative of $R\big(p, y^*(p)\big)$ with respect to p. WLOG, we assume $\bar{\theta} + \underline{\theta} = 0$.

$$\begin{aligned} \frac{\partial}{\partial p} R\big(p, y^*(p)\big) &= -\frac{c}{p^2}\big(\phi^*(\bar{\theta}) - \phi^*(\underline{\theta})\big) - \phi'(p, \underline{\theta}) && \text{(A.52)} \\ &= -\frac{c}{2Bp^2}(A - Bc)(\bar{\theta} - \underline{\theta}) - (A + Bc - 2Bp + \underline{\theta}). && \text{(A.53)} \end{aligned}$$

By Proposition 3.5.1, the risk-free price p_0 with linear additive demand is as follows:

$$\begin{aligned} p_0 &= c + \frac{\phi^*(\bar{\theta}) - \phi^*(\underline{\theta})}{\bar{\theta} - \underline{\theta}} && \text{(A.54)} \\ &= \frac{1}{2B}(A + Bc). && \text{(A.55)} \end{aligned}$$

Then, we evaluate the derivative of the regret function at the risk-free price p_0.

$$\frac{\partial}{\partial p}R\big(p,y^*(p)\big)\mid_{p=p_0} = \frac{1}{(p_0)^2}\Big((p_0)^2\big(-\phi'(p_0,\underline{\theta})\big)-c\big(\phi^*(\bar{\theta})-\phi^*(\underline{\theta})\big)\Big) \quad (A.56)$$

$$= \frac{\bar{\theta}-\underline{\theta}}{8B^2(p_0)^2}\big((A-Bc)^2+4B^2c^2\big) \quad (A.57)$$

$$> 0. \quad (A.58)$$

By Proposition 3.4.2, we know that the regret function $R\big(p,y^*(p)\big)$ is convex on $[p_l,p_h]$ and we have shown that $\frac{\partial}{\partial p}R\big(p,y^*(p)\big)\mid_{p=p_0}\geq 0$. Thus, we can conclude that $p^*\leq p_0$.

If the demand function is of the form: $D(p,\theta)=\sigma(p)\theta$. By Proposition 3.5.1, we have $p_0=\bar{p}^*=\underline{p}^*$. Moreover, we have $\phi'(p_0,\theta)=0$ due to the optimality condition. By Proposition 3.4.2, we know that $p_l<p_0=\bar{p}^*<p_h$. Then, we can evaluate the derivative of the regret function at the risk-free price p_0:

$$\frac{\partial}{\partial p}R\big(p,y^*(p)\big) = -\frac{c}{p^2}(\bar{p}^*-c)\sigma(\bar{p}^*)(\bar{\theta}-\underline{\theta})-\phi'(p,\underline{\theta}). \quad (A.59)$$

$$\frac{\partial}{\partial p}R\big(p,y^*(p)\big)\mid_{p=p_0} = -\frac{c}{(p_0)^2}(\bar{p}^*-c)\sigma(\bar{p}^*)(\bar{\theta}-\underline{\theta})-\phi'(p_0,\underline{\theta}) \quad (A.60)$$

$$= -\frac{c}{(p_0)^2}(\bar{p}^*-c)\sigma(\bar{p}^*)(\bar{\theta}-\underline{\theta}) \quad (A.61)$$

$$\leq 0. \quad (A.62)$$

By a similar argument as in the additive demand case, the regret function is convex on $[p_l,p_h]$ and $\frac{\partial}{\partial p}R\big(p,y^*(p)\big)\mid_{p=p_0}\leq 0$, thus, we can conclude that the optimal price with the multiplicative demand $p^*\geq p_0$. □

Proof of the Result in Example 3.5.1. With the following demand function,

$$d(p)=\begin{cases}\frac{1}{8}(87-p+\frac{9975}{p-1}), & 1<p\leq 6,\\ 559.5-50p, & p>6,\end{cases} \quad (A.63)$$

and the unit purchase cost $c=1$, the uncertainty set $[\underline{\theta},\bar{\theta}]=[-10.5,10.5]$, the profit function associated with $\underline{\theta}$ is of the form,

$$\phi(p,\underline{\theta})=\begin{cases}\frac{1}{8}\big((p-1)(3-p)+9975\big), & 1<p\leq 6,\\ (p-1)(549-50p), & p>6,\end{cases}$$

and the profit function associated with $\bar{\theta}$ is of the form:

$$\phi(p,\bar{\theta}) = \begin{cases} \frac{1}{8}\big((p-1)(171-p)+9975\big), & 1 < p \le 6, \\ (p-1)(570-50p), & p > 6. \end{cases}$$

It is easy to verify that both $\phi(p,\underline{\theta})$ and $\phi(p,\bar{\theta})$ are concave in p and the optimal profits are $\phi^*(\underline{\theta}) = 1247$ and $\phi^*(\bar{\theta}) = 1352$. By Proposition 3.5.1 and the proof of Proposition 3.4.2, we can find that $p_0 = c + \frac{\phi^*(\bar{\theta})-\phi^*(\underline{\theta})}{\bar{\theta}-\underline{\theta}} = 6$ and $p_l = \frac{\phi^*(\bar{\theta})-\phi^*(\underline{\theta})}{\bar{\theta}-\underline{\theta}} = 5$. Then, by Proposition 3.4.2, we express the derivative of the regret function $R\big(p, y^*(p)\big)$ as follows:

$$\frac{\partial}{\partial p} R\big(p, y^*(p)\big) \;=\; \frac{1}{p^2}\Big(p^2\big(-\phi'(p,\underline{\theta})\big) - c\big(\phi^*(\bar{\theta}) - \phi^*(\underline{\theta})\big)\Big) \tag{A.64}$$

By setting $\frac{\partial}{\partial p}R\big(p, y^*(p)\big) = 0$, we obtain the optimal price $p^* = 6.019$. Therefore, the optimal solution p^* is larger than the risk-free price p_0 with the demand function defined in (3.11). □

Proof of Proposition 3.5.3. When the demand is of multiplicative form, we have $D(p,\theta) = \sigma(p)\theta$ where $\theta \ge 0$. The regret function over the interval $[p_l, p_h]$ is of the form $R_1(p, y^*(p); \delta) = 2c\phi^*\frac{\delta}{p} + (\theta_0 - \delta)\phi^* - (\theta_0 - \delta)\phi(p,1)$. By Proposition 3.5.1, Proposition 3.5.2, and the definitions of $p_l = \max\big\{\arg\{p : p\sigma(p) = \phi^*, p \le \bar{p}^*\}, \underline{p}\big\}$ and $p_h = \min\big\{\arg\{p : p\sigma(p) = \phi^*, p \ge \bar{p}^*\}, \bar{p}\big\}$, we have $p_l \le p_0 \le p^* \le p_h$ and p_l, p_0, p_h are independent of δ. Thus, we have,

$$p^* = \arg\min_{p\in[\underline{p},\bar{p}]} R\big(p, y^*(p); \delta\big) = \arg\min_{p\in[p_0,p_h]} R_1(p, y^*(p); \delta) \tag{A.65}$$

Since $p \in [p_0, p_h]$, where $p_0 = \bar{p}^* = \underline{p}^*$ by Proposition 3.5.1, and $\phi(p,1) = (p-c)\sigma(p)$ is concave by Assumption 1, $\phi(p,1)$ is decreasing in $p \in [p_0, p_h]$. By Corollary 2.6.3 in Topkis (2011), $\delta\frac{1}{p}$ and $\delta\phi(p,1)$ are submodular in (p,δ) on $[p_0,p_h] \times [0,\theta_0]$ respetively. By Lemma 2.6.1 in Topkis (2011), the supermodularity preserves under summation, thus, $R_1\big(p, y^*(p); \delta\big)$ is also submodular in (p,δ) on $[p_0,p_h] \times [0,\theta_0]$. Let $\tilde{R}_1(p,\delta) = -R_1\big(p, y^*(p); \delta\big)$, which is supermodular in (p,δ). By Theorem 2.82 in Topkis (2011), we can claim that the optimal minimax regret price $p^* = \arg\min_{p\in[p_0,p_h]} R_1\big(p, y^*(p); \delta\big) = \arg\max_{p\in[p_0,p_h]} \tilde{R}_1\big(p; \delta\big)$ is increasing in δ.

When the demand is of the linear additive form, we have $D(p,\theta) = \mu(p) + \theta$ and $\mu(p) = A - B \cdot p$ where $A, B > 0$. The regret function over the interval $[p_l, p_h]$ is of the form:

$$R_1\big(p, y^*(p); \delta\big) = \tfrac{c\delta(A-Bc+\theta_0)}{Bp} - (A - Bp + \theta_0 - \delta)(p-c) + \tfrac{(A+\theta_0-\delta-Bc)^2}{4B}. \quad \text{(A.66)}$$

Consider the price $\hat{p} = \sqrt{\frac{c(A-Bc+\theta_0)}{B}}$, we have,

$$\frac{d}{dp} R\big(p, y^*(p); \delta\big) \mid_{p=\hat{p}} = -\big((\sqrt{A - Bc + \theta_0} - \sqrt{Bc})^2 + Bc\big) < 0. \quad \text{(A.67)}$$

By Proposition 2, due to the convexity of the regret function, we have $\hat{p} \leq p^*$. And it is easy to verify that $p_l = \frac{1}{2B}(A + Bc + \theta_0)$ and $p_h = \bar{p}$ with the linear additive demand, therefore, they are independent of δ. Thus,

$$p^* = \arg \min_{p \in [\underline{p}, \bar{p}]} R\big(p, y^*(p); \delta\big) = \arg \min_{p \in [\max\{\hat{p}, p_l\}, \bar{p}]} R_1\big(p, y^*(p); \delta\big). \quad \text{(A.68)}$$

Since $p \geq \hat{p} = \sqrt{\frac{c(A-Bc+\theta_0)}{B}}$, we can verify that $f(p) = p + \frac{c(A-Bc+\theta_0)}{Bp}$ is increasing in $p \in [\max\{\hat{p}, p_l\}, \bar{p}]$. Thus, by Corollary 2.6.3 and Lemma 2.6.1 in Topkis (2011), $R_1\big(p, y^*(p); \delta\big) = \big(p - \frac{c(A-Bc+\theta_0)}{Bp}\big)\delta - (A - Bp + \theta_0)(p - c) - \delta c + \frac{(A+\theta_0-\delta-Bc)^2}{4B}$ is supermodular in (p, δ) over $[\max\{\hat{p}, p_l\}, \bar{p}] \times [0, \theta_0]$. Let $\tilde{R}_1(p, \delta) = -R_1\big(p, y^*(p); \delta\big)$, which is supermodular in $(-p, \delta)$. By Theorem 2.8.2 in Topkis (2011), we can claim that the optimal minimax regret price $p^* = \arg\min_{p \in [\max\{\hat{p}, p_l\}, \bar{p}]} R_1\big(p, y^*(p); \delta\big) = -\arg\max_{-p \in [-\bar{p}, -\max\{\hat{p}, p_l\}]} \tilde{R}_1(p, \delta)$ is decreasing in δ. □

Proof of Proposition 3.5.4. Before we prove Proposition 3.5.4, we first introduce the following two lemmas that will be used in the proof of Proposition 3.5.4.

Lemma A.0.5 *Let $\hat{\phi}(p) = p\sigma(p)$ and p^* denote the optimal price under the minimax regret framework, if the demand function is of multiplicative form and Assumption 1 holds, then $\hat{\phi}'(p^*) \leq 0$.*

Proof. Let $\hat{p}^* = \arg\max\{p : p\sigma(p)\}$, then we have the following optimality condition, $\sigma(\hat{p}^*) + \hat{p}^*\sigma'(\hat{p}^*) = 0$. Then, we can evaluate the derivative of the regret function at $\hat{p}^*$:

$$
\begin{aligned}
R'\big(p, y^*(p)\big) \mid_{p=\hat{p}^*} &= -\frac{1}{(\hat{p}^*)^2}\phi^*(\bar{\theta} - \underline{\theta}) - \phi'(\hat{p}^*)\underline{\theta} && \text{(A.69)} \\
&= -\frac{1}{(\hat{p}^*)^2}\phi^*(\bar{\theta} - \underline{\theta}) + c\sigma'(\hat{p}^*)\underline{\theta} && \text{(A.70)} \\
&< 0. && \text{(A.71)}
\end{aligned}
$$

The second equality follows from the optimality condition for $\hat{p}^*$ and the last inequality is derived by the monotone relationship between the multiplicative demand and the unit selling price. Thus, due to the convexity of regret function, we can conclude that $p^* \geq \hat{p}^*$. Due to the concavity of $\hat{\phi}(p)$ by Assumption 1, $\hat{\phi}(p^*)' \leq 0$. □

Lemma A.0.6 *In the multiplicative demand, let* $\hat{\phi}(p) = p\sigma(p)$, $\phi^* = \max_{\tilde{p}}(\tilde{p}-c)\sigma(\tilde{p})$ *and* p_0 *denote the corresponding risk-free price, if* $p_h = \arg\{p : p\sigma(p) = \phi^*, p \geq p_0\}$ *exists, then* $\hat{\phi}'(p_h) \leq 0$.

Proof. Due to the concavity of $\hat{\phi}(p)$, we only need to prove $p_h \geq \hat{p}^* = \arg\max\{p : \hat{\phi}(p)\}$. We will prove this claim by contradiction by assuming that $p_0 \leq p_h < \hat{p}^*$. We first have $\hat{\phi}(p_0) > \phi(p_0)$ by the definition of these two functions. In addition, by Proposition 3.5.1, we have $p_0 = \bar{p}^* = \underline{p}^*$ and thus, $\phi^* = (p_0 - c)\sigma(p_0)$. Therefore, for $p_h \in (p_0, \hat{p}^*)$, we have $\hat{\phi}(p_h) > \hat{\phi}(p_0) > \phi(p_0) = \phi^*$. In short, $\hat{\phi}(p_h) > \phi^*$, which contradicts the definition of p_h. Thus, we can conclude that if p_h exists, then $p_h \geq \hat{p}^*$ and thus, $\hat{\phi}'(p_h) \leq 0$. □

As shown by Proposition 3.5.2, in the multiplicative demand, the optimal price $p^* \geq p_0 > p_l$. Therefore, we need to consider two cases: (i) the optimal price p^* is an interior point solution on (p_l, p_h). (ii) the optimal solution $p^* = p_h$. Before we discuss these two cases respectively, recall that in the multiplicative case, $\phi(p) = (p - c)\sigma(p)$ and $\phi^* = \max_p(p - c)\sigma(p)$. Given the form of the optimal order quantity $y^*(p)$ with a fixed price p as presented in Proposition 3.4.2, the derivative of the optimal order quantity with respect to the unit ordering cost at the optimal price p^* is of the form:

$$\frac{\partial y^*(p^*)}{\partial c} = \Big(- \frac{1}{(p^*)^2}\phi^*(\bar{\theta} - \underline{\theta}) + \sigma'(p^*)\underline{\theta}\Big)\frac{\partial p^*}{\partial c} + \frac{1}{p^*}(\bar{\theta} - \underline{\theta})\frac{\partial \phi^*}{\partial c}. \tag{A.72}$$

First we consider the case (i) where p^* is an interior solution, then we have the following optimality condition for p^* that always holds in the neighborhood of p^*.

$$R'\big(p^*, y^*(p^*)\big) = \frac{c}{(p^*)^2}\phi^*(\bar{\theta} - \underline{\theta}) - \phi'(p^*)\underline{\theta} = 0. \tag{A.73}$$

Then, we can take the derivative of equation (A.73) with respect to c:

$$\begin{aligned} &\frac{1}{(p^*)^2}\phi^*(\bar{\theta} - \underline{\theta}) + \frac{c}{(p^*)^2}\frac{\partial \phi^*}{\partial c}(\bar{\theta} - \underline{\theta}) - \sigma'(p^*)\underline{\theta} \\ =&\Big(\frac{2c}{(p^*)^3}\phi^*(\bar{\theta} - \underline{\theta}) - \big(2\sigma'(p^*) + (p^* - c)\sigma''(p^*)\big)\underline{\theta}\Big)\frac{\partial p^*}{\partial c}. \end{aligned} \tag{A.74}$$

Let $A = \frac{2c}{(p^*)^3}\phi^*(\bar{\theta} - \underline{\theta}) - \big(2\sigma'(p^*) + (p^* - c)\sigma''(p^*)\big)\underline{\theta}$ and we have $A = \frac{2c}{(p^*)^3}\phi^*(\bar{\theta} - \underline{\theta}) - \big(2\sigma'(p^*) + (p^* - c)\sigma''(p^*)\big)\underline{\theta} > -\big(2\sigma'(p^*) + (p^* - c)\sigma''(p^*)\big)\underline{\theta} = -\phi''(p^*) \geq 0$ due to the concavity of $\phi(p)$. Therefore, we can substitute $\frac{\partial p^*}{\partial c}$ in (A.72) with equation (A.74):

$$\begin{aligned} \frac{\partial y^*(p^*)}{\partial c} =& \frac{1}{A}\bigg(\Big(\frac{1}{(p^*)^2}\phi^*(\bar{\theta} - \underline{\theta}) + \frac{c}{(p^*)^2}\frac{\partial \phi^*}{\partial c}(\bar{\theta} - \underline{\theta}) - \sigma'(p^*)\underline{\theta}\Big)\Big(- \frac{1}{(p^*)^2}\phi^*(\bar{\theta} - \underline{\theta}) + \sigma'(p^*)\underline{\theta}\Big) \\ &+ \frac{1}{p^*}(\bar{\theta} - \underline{\theta})\frac{\partial \phi^*}{\partial c}\Big(\frac{2c}{(p^*)^3}\phi^*(\bar{\theta} - \underline{\theta}) - \big(2\sigma'(p^*) + (p^* - c)\sigma''(p^*)\big)\underline{\theta}\Big)\bigg) \\ =& \frac{1}{A}\bigg((\bar{\theta} - \underline{\theta})\underline{\theta}\Big(\frac{2\phi^*\sigma'(p^*)}{(p^*)^2} + \sigma'(p^*)\frac{c}{(p^*)^2}\frac{\partial \phi^*}{\partial c} - \frac{1}{p^*}\frac{\partial \phi^*}{\partial c}\big(2\sigma'(p^*) + (p^* - c)\sigma''(p^*)\big)\Big) \\ &+ (\bar{\theta} - \underline{\theta})^2\Big(- \frac{(\phi^*)^2}{(p^*)^4} + \frac{c\phi^*}{(p^*)^4}\frac{\partial \phi^*}{\partial c}\Big) - \underline{\theta}^2\big(\sigma'(p^*)\big)^2\bigg). \end{aligned} \tag{A.75}$$

By equation (A.73), we have $\underline{\theta} = -\big(\frac{c}{(p^*)^2}\phi^*(\bar{\theta} - \underline{\theta})\big)/\phi'(p^*)$ since $\phi(p^*)' < 0 \neq 0$. Then, we can substitute $\underline{\theta}$ in equation (A.75):

$$\begin{aligned} \frac{\partial y^*(p^*)}{\partial c} =& \frac{1}{A}\frac{(\bar{\theta} - \underline{\theta})^2}{(p^*)^4\big(\phi'(p^*)\big)^2}\bigg(- \phi'(p^*)^2\phi^{*2} - 2c\phi^{*2}\phi'(p^*)\sigma'(p^*) - \sigma'(p^*)^2c^2\phi^{*2} \\ &+ c\phi^*\phi'(p^*)\frac{\partial \phi^*}{\partial c}\Big(\phi'(p^*) - c\sigma'(p^*) + p^*\big(2\sigma'(p^*) + (p^* - c)\sigma''(p^*)\big)\Big)\bigg) \\ <& \frac{1}{A}\frac{(\bar{\theta} - \underline{\theta})^2}{(p^*)^4\big(\phi'(p^*)\big)^2}\bigg(c\phi^*\phi'(p^*)\frac{\partial \phi^*}{\partial c}\Big(\phi'(p^*) - c\sigma'(p^*) + p^*\big(2\sigma'(p^*) + (p - c)\sigma''(p^*)\big)\Big)\bigg) \\ =& \frac{1}{A}\frac{(\bar{\theta} - \underline{\theta})^2}{(p^*)^4\big(\phi'(p^*)\big)^2}c\phi^*\phi'(p^*)\frac{\partial \phi^*}{\partial c}\Big(\big(p^*\sigma(p^*)\big)' + (p^* - c)\big(p^*\sigma(p^*)\big)''\Big) \\ \leq& 0. \end{aligned} \tag{A.76}$$

By Proposition 3.5.1, we have $p^* \geq p_0$ and consequently, $\phi'(p^*) \leq 0$. Thus, $-\phi'(p)^2(\phi^*)^2 - 2c\phi^*\phi'(p)\sigma'(p) - \sigma'(p)^2c^2(\phi^*)^2 < 0$. We also showed that $A > 0$, therefore, the first inequality is verified. The second equality follows straightforward algebra. For the second inequality, we know that $\big(p\sigma(p)\big)' \mid_{p=p^*} \leq 0$ by Lemma A.0.5 and $\big(p\sigma(p)\big)'' \leq 0$ due to the concavity of $p\sigma(p)$. Finally, it can be shown that ϕ^* is decreasing in c and thus, $\frac{\partial \phi^*}{\partial c} < 0$. Thus, we can conclude that we have $\frac{\partial y^*(p^*)}{\partial c} < 0$ for the case (i) where p^* is an interior solution.

Consider the case (ii) where the optimal price $p^* = p_h$ and p_h is the upper bound of the optimal price. Hence, the optimality condition (A.73) for p^* does not always hold in the neighborhood of p^*. By the definition of $p_h = \arg\{p : p\sigma(p) = \phi^*, p \geq p_0\}$, $p_h\sigma(p_h) = \phi^*$ always holds and we can take derivative of this equation with respect to c:

$$\big(\sigma(p_h) + p_h\sigma'(p_h)\big)\frac{\partial p_h}{\partial c} = \frac{\partial \phi^*}{\partial c}. \tag{A.77}$$

In this case, p_h must exists, otherwise, it reduces to case (i). Then by Lemma A.0.6, $\sigma(p_h) + p_h\sigma'(p_h) = \big(p\sigma(p)\big)' \mid_{p=p_h} < 0$ and $\frac{\partial \phi^*}{\partial c} < 0$. Therefore, we have $\frac{\partial p_h}{\partial c} > 0$, which indicates p_h is increasing in c. By equation (A.72), we have:

$$\frac{\partial y^*(p^*)}{\partial c} = \big(-\frac{1}{p^2}\phi^*(\bar{\theta} - \underline{\theta}) + \sigma'(p)\underline{\theta}\big)\frac{\partial p^*}{\partial c} + \frac{1}{p}(\bar{\theta} - \underline{\theta})\frac{\partial \phi^*}{\partial c} < 0. \tag{A.78}$$

Because $-\frac{1}{p^2}\phi^*(\bar{\theta} - \underline{\theta}) + \sigma'(p)\underline{\theta} < \sigma'(p)\underline{\theta} \leq 0$, $\frac{\partial p^*}{\partial c} = \frac{\partial p_h}{\partial c} > 0$, and $\frac{\partial \phi^*}{\partial c} < 0$. The above inequality is verified consequently. Thus, when p^* is not the interior solution, specifically, $p^* = p_h$, we have $\frac{\partial y^*(p^*)}{\partial c} < 0$. Therefore, we can conclude that $y^*(p^*)$ is decreasing in c. □

Proof of Proposition 3.5.5. Using the max-min robust optimization approach, the optimization problem is of the form:

$$\max_p (p - c)D(p, \underline{\theta}). \tag{A.79}$$

Let $p_{\text{minimax}} = \arg\max\{p : (p-c)D(p,\underline{\theta})\}$, then we can evaluate the derivative of the maximized regret $R\big(p, y^*(p)\big)$ at p_{minimax}.

$$\begin{aligned} \frac{\partial}{\partial p} R\big(p, y^*(p)\big) \mid_{p=p_{\text{minimax}}} &= -\frac{c}{(p_{\text{minimax}})^2}\big(\phi^*(\bar{\theta}) - \phi^*(\underline{\theta})\big) - \phi'\big(p_{\text{minimax}}, \underline{\theta}\big) \\ &= -\frac{c}{(p_{\text{minimax}})^2}\big(\phi^*(\bar{\theta}) - \phi^*(\underline{\theta})\big) < 0. \end{aligned} \tag{A.80}$$

The first equality follows from the definition of the regret function and the second equality follows from the fact that p_{minimax} is the minimizer of the objective function using the minimax robust optimization. And the last inequality is due to the fact that $\phi^*(\bar{\theta}) > \phi^*(\underline{\theta})$ by Proposition 3.4.1. By Proposition 3.5.5 and $\frac{\partial}{\partial p} R\big(p, y^*(p)\big) \mid_{p=p_{\text{minimax}}} < 0$, we can conclude that $p^* \geq p_{\text{minimax}}$. $\square$

Proof of Theorem 4.4.1. Given all firms' $\mathbf{x}$, consider bargaining unit (i,j). Differentiating the Nash product objective in (4.7) with respect to $w_{i,j}$, we have

$$\theta_{i,j}\pi_j(T,\mathbf{x}) - \theta_{j,i}\pi_i(T,\mathbf{x}) = 0,$$

when $\pi_i(T,\mathbf{x}), \pi_j(T,\mathbf{x}) \geq 0$. If $\mathbf{x} = \mathbf{1}$, $\sum_{v\in N} \pi_v(T,\mathbf{1}) = \Pi > 0$. We have that $\pi_i(T,\mathbf{x}), \pi_j(T,\mathbf{x}) > 0$ and the proportional gain allocation condition in (4.8) holds for all bargaining units. Therefore, the net gain of firm v

$$\pi_v(T,\mathbf{1}) = \pi_0(T,\mathbf{1}) \prod_{a\in P_v(T)} \rho_a = \frac{\prod_{a\in P_v(T)} \rho_a}{1 + \sum_{u\in N\setminus\{0\}} \prod_{a\in P_u(T)} \rho_a}\Pi, \forall v \in N \setminus \{0\}.$$

The transfer payment $w_{i,j}(\mathbf{1})$ includes the net gains and costs of firm j and all its successors in the implementation structure T. Therefore,

$$w_{i,j}(\mathbf{1}) = \frac{\sum_{v\in\{j\}\cup\{u:P_{(j,u)}(T)\neq\emptyset\}} \prod_{a\in P_v(T)} \rho_a}{1 + \sum_{v\in N\setminus\{0\}} \prod_{a\in P_v(T)} \rho_a}\Pi + \sum_{v\in\{j\}\cup\{u:P_{(j,u)}(T)\neq\emptyset\}} c_v.$$

If $\mathbf{x} \neq \mathbf{1}$, either $\pi_i(T,\mathbf{x}) = \pi_j(T,\mathbf{x}) = 0$ and $w_{i,j}(\mathbf{x}) = 0$, or the negotiation breaks down.

Given x_l for all $l \neq j$, optimize the Nash product objective with respect to x_j. If $x_l = 0$ for any l, then $x_j = 0$ (otherwise the negotiation breaks down because firm j's disagreement point is violated). If $x_l = 1$ for all $l \neq j$, then $x_j = 1$. Based on the above analysis on the net gains, all firms committing to ESR investment (i.e., $\mathbf{x} = \mathbf{1}$) is the Pareto-dominant equilibrium. $\square$

Proof of Corollary 4.4.1. For any feasible implementation structure T, when $\theta_{i,j} = \theta_{j,i} = 0.5$ for all $i, j \in N$ such that $t_{i,j} = 1$, by (4.8), we have $\pi_i(T, \mathbf{x}) = \pi_j(T, \mathbf{x})$ for all $i, j \in N$. By Theorem 4.4.1, we have the equilibrium $\mathbf{x}^* = \mathbf{1}$ and $\sum_{i \in N} \pi_i(T, \mathbf{x}^*) = R - \sum_{i \in N} c_i$. Therefore, $\pi_i(T, \mathbf{x}^*) = \frac{1}{n+1}(R - \sum_{i \in N} c_i)$, which is equal to the Shapley value given in (4.5). $\square$

Proof of Theorem 4.4.2. Let $\mathcal{P}_v(A^{\text{ESR}})$ denote the set of all directed paths from node 0 to node v in network (N, A^{ESR}). By exchanging the order of minimization and summation, we have

$$\sum_{v \in N \setminus \{0\}} \min_{P \in \mathcal{P}_v(A^{\text{ESR}})} \prod_{a \in P} \rho_a \left(= \sum_{v \in N \setminus \{0\}} \min_{T \in \mathcal{T}(A^{\text{ESR}})} \prod_{a \in P_v(T)} \rho_a \right) \leq \min_{T \in \mathcal{T}(A^{\text{ESR}})} \sum_{v \in N \setminus \{0\}} \prod_{a \in P_v(T)} \rho_a$$

Let $d_a = \log \rho_a$ for all $a \in A^{\text{ESR}}$. Consider each term of the summation in the left-hand side of the above inequality. By minimizing the logarithm of the product, each term is equivalent to finding the shortest path from node 0 to node v in (N, A^{ESR}) with arc weight d_a. When (N, A^{ESR}) is a directed acyclic graph and node 0 is connected with all other nodes, it is well known that there exists a tree T rooted at node 0, such that $P_v(T)$ is a shortest path from 0 to v for all $v \in N \setminus \{0\}$. This shortest path tree is therefore the optimal implementation structure. $\square$

Proof of Theorem 4.4.3. Part i) If the retailer chooses to negotiate with firm j in tier k directly, then the ESR gain ratio of firm j is $\frac{\pi_j(T,\mathbf{1})}{\pi_0(T,\mathbf{1})} = \rho_0$. If there exists tier k' such that $\rho_{k'} < 1$, consider delegating the negotiation responsibility to firm i in tier k' where $P_{(i,j)}(A^{\text{ESR}}) \neq \emptyset$, then the ESR gain ratio of firm j is $\frac{\pi_j(T,\mathbf{1})}{\pi_0(T,\mathbf{1})} = \frac{\pi_i(T,\mathbf{1})}{\pi_0(T,\mathbf{1})} \rho_{k'} < \frac{\pi_i(T,\mathbf{1})}{\pi_0(T,\mathbf{1})} \leq \rho_0$. Therefore, in order to minimize the ESR gain ratio of firm j in tier k, the retailer chooses not to negotiate with firm j in tier k directly.

Part ii) If $\rho_{k'} \geq 1$ for all $k' < k$, then for any directed path that connects the retailer and firm j in tier k in the ESR network, $\{(0, i_{k_1}), (i_{k_1}, i_{k_2}), \ldots, (i_{k_m}, j)\}$, where i_{k_s} is the firm in the intermediate tier k_s, the resulting gain ratio of firm j is $\frac{\pi_j(T,\mathbf{1})}{\pi_0(T,\mathbf{1})} =$

$\rho_0 \prod_{s=1}^{m} \rho_{k_s} \geq \rho_0$. Thus, in the retailer preferred implementation structure, the retailer negotiates with firm i directly to minimize her ESR gain ratio. $\square$

To prove Corollary 4.4.2 using the following stronger version of Theorem 4.4.3.

Theorem 4.4.3A *For the network described in Theorem 4.4.3,*

i) if there exists a $k \in \{1, 2, \ldots, d-1\}$ such that $\rho_k < 1$, then firms in tier $i \in \{0, 1, k-1\}$ do not interact directly with suppliers in tiers $\{k+1, k+2, \ldots, d\}$ in the retailer-preferred ESR implementation.

ii) if there exists a $k \in \{1, 2, \ldots, d-1\}$ such that $\rho_k > 1$, then firms in tier k do not interact with suppliers in tiers $\{k+1, k+2, \ldots, d\}$ in the retailer-preferred ESR implementation.

Proof. The proof is similar to that of Theorem 4.4.3 by replacing the retailer with a supplier in tier $i < k$. $\square$

Proof of Corollary 4.4.2. (i) If $\rho_d > 1$, let $k^D = \arg\min\{s : \rho_s > 1, s = 1, 2, \ldots, d\}$. If $\rho_d \leq 1$, let $k_D = d$. Because ρ_k is increasing, we have $\rho_k \leq 1$, for $k = 0, 1, 2, \ldots, k_D - 1$, and $\rho_k > 1$, for $k = k_D, k_D + 1, \ldots, d-1$. First, consider tier $k \leq k_D$. If $k = 1$, tier k suppliers are dealt with by their immediate downstream, that is, the retailer. If $k > 1$, by the definition of k_D, we have $\rho_{k-1} < 1$. By Theorem 4.4.3A, firms in tiers $0, 1, \ldots, k-2$ do not interact directly with suppliers in tier k. Therefore, negotiations with tier k suppliers are delegated to their immediate downstream, that is, tier $k-1$.

(ii) If $\rho_d \geq 1$, let $k_I = d$. If $\rho_d < 1$, let $k_I = \arg\min\{k : \rho_k < 1,\ k = 1, 2, \ldots, d\}$. Since ρ_k is decreasing, we have $\rho_k \geq 1$, for $k = 0, 1, 2, \ldots, k_I - 1$, and $\rho_k < 1$, for $k = k, k_I + 1, \ldots, d-1$. By Theorem 4.4.3, the retailer negotiates directly with firms in tier $k = 1, 2, \ldots, k_I$. Consider tier $k > k_I$. By the definition of k_I, we have $\rho_{k1} < 1$. Then, by Theorem 4.4.3A, firms in tiers $0, 1, \ldots, k-2$ do not interact directly with firms in tier k. Therefore, negotiations with firms in tier k are delegated to their immediate downstream, that is, tier $k-1$. $\square$

Proof of Theorem 4.5.1. Consider bargaining unit (i, j) for some $t_{i,j} = 1$. For a given $\mathbf{x}$, differentiating the Nash product with respect to $w_{i,j}$, we have

$$\theta_{i,j}\pi_j(T, \mathbf{x}) - \theta_{j,i}\pi_i(T, \mathbf{x}) = 0,$$

when $\pi_i(T, \mathbf{x}), \pi_j(T, \mathbf{x}) \geq 0$. If $R(\mathbf{x}) - \sum_{v \in N} c_v(x_v) > 0$, then the proportional allocation condition in (4.8) holds for all bargaining units. Similar to the proof of Theorem 4.4.1, we have the transfer payment given in Theorem 4.5.1. If $R(\mathbf{x}) - \sum_{v \in N} c_v(x_v) \leq 0$, either $\pi_i(T, \mathbf{x}) = \pi_j(T, \mathbf{x}) = 0$ and $w_{i,j}(\mathbf{x}) = 0$ or the negotiation breaks down.

Given x_l for all $l \neq j$ (denoted by $\mathbf{x}_{-j}$), optimize the Nash product with respect to x_j. If $\max_{x_j}\{R(x_j, \mathbf{x}_{-j}) - \sum_{l \in N} c_l(x_l)\} > 0$, then both π_i and π_j are positive and proportional to $R(\mathbf{x}) - \sum_{v \in N} c_v(x_v)$ when x_j is optimized. Therefore, the optimal $x_j \in [0, 1]$ satisfies the following first-order condition, which is sufficient under our assumption,

$$\frac{\partial R}{\partial x_j} - \frac{\partial c_j}{\partial x_j} + \lambda_j I(x_j = 0) - \mu_j I(x_j = 1) = 0, \quad \lambda_j, \mu_j \geq 0. \tag{A.81}$$

Note that condition (A.81) for all $j \in N$ provides the optimality condition for the centralized optimal investment problem. Therefore, $\mathbf{x}^*$ is an equilibrium investment level of the multi-unit bargaining. Furthermore, because $R(\mathbf{x}) - \sum_{i \in N} c_i(x_i)$ is strictly concave, there does not exist any other equilibrium $\mathbf{x}$ such that $R(\mathbf{x}) - \sum_{v \in N} c_v(x_v) > 0$. If $\max_{x_j}\{R(x_j, \mathbf{x}_{-j}) - \sum_{i \in N} c_i(x_i)\} \leq 0$, then $x_j = 0$. (Otherwise the negotiation breaks down because firm j's disagreement point is violated.) Based on the above analysis on the firms' profits, all firms make the centralized optimal investment $\mathbf{x}^*$ and receive positive gains in the Pareto-dominant equilibrium. $\square$

Proof of Theorem 4.5.2. We first prove the result when in step 2 of the game, all firms in the current partial implementation structure T_k (instead of only the leaf nodes in N_k) can approach their upstream suppliers. We then show that in the equilibrium, the non-leaf nodes at each stage will not approach any upstream suppliers.

Before going into the proof, we introduce some notations and define the *strategy profile corresponding to the retailer-preferred implementation structure*. Let T_k denote the partial implementation structure after round k. Let V_k denote the nodes in T_k, that is $V_k = \cup_{l=1}^{k} N_l$. Let $\bar{V}_k = N \backslash V_k$. Let $\mathcal{T}(T_k)$ denote the set of spanning trees on (N, A^{ESR}) containing subtree T_k. Let $T^*(T_k)$ denote the retailer-preferred implementation structure within $\mathcal{T}(T_k)$, that is,

$$T^*(T_k) = \arg \max_{T \in \mathcal{T}(T_k)} \pi_0(T, \mathbf{1}).$$

Let $I_k(j)$ denote the set of nodes in V_k that offer to negotiate with j for some $j \in \bar{V}_k$. In the strategy profile corresponding to the retailer-preferred implementation structure, given history T_k, a firm $i \in V_k$ offers to negotiate with $j \in \bar{V}_k$, if and only if $(i,j) \in T^*(T_k)$, and a firm $j \in \bar{V}_k$ such that $I_k(j) \neq \emptyset$ chooses to negotiate with node

$$i_k^*(j) = \arg \max_{i \in I_k(j)} \left\{ \rho_{i,j} \prod_{a \in P_i(T_k)} \rho_a \right\}. \tag{A.82}$$

Let $\mathcal{T}_{k+1}(T_k)$ denote the set of all possible partial implementation structures after round $k+1$, that is,

$$\mathcal{T}_{k+1}(T_k) = \left\{ T \in \mathcal{T}(T_k) : T \backslash T_k \subset \left\{ (i,j) \in A^{\text{ESR}} : i \in V_k, j \in \bar{V}_k \right\} \right\}.$$

Let $T^*_{k+1}(T_k)$ denote the partial implementation structure after round $k+1$, under the strategy profile corresponding to the retailer-preferred implementation structure.

The proof is based on the following key observation. Consider round $k+1$ given history T_k. With some abuse of notation, let $\pi_i(T) = \pi_i(T, \mathbf{1})$ for all $i \in N$. For any node i in V_k, and any implementation structure T in $\mathcal{T}(T_k)$, we have

$$\pi_i(T) = \pi_0(T) \prod_{a \in P_i(T_k)} \rho_a. \tag{A.83}$$

In other words, the net gain of any node in V_k is a fixed proportion of the retailer's net gain, regardless of the final implementation structure. This suggests that in each round, all firms in V_k have the same incentive as the retailer does.

Observe that there are at most $|N| - 1$ rounds in the game. Assume there are $m < |N|$ rounds along an equilibrium path. The partial implementation structure before the last round takes place is T_{m-1}. Without loss of generality, assume $T_m^*(T_{m-1}) = T^*(T_{m-1})$.

First, consider the case of simultaneous moves in each round. There are two stages in each round of the game. In Stage I, firms in V_{k-1} approach certain nodes in $\bar{V}_{k-1}$ to initiate ESR negotiation. In Stage II, each firm in $\bar{V}_{k-1}$ who has been contacted by at least one firm in V_{k-1} can choose to partner with one of those firms or refuse to partner with any firm. In the latter case, the game ends and the ESR program fails. Consider Stage II of round m. For any node in $\bar{V}_{m-1}$, the choice specified in (A.82) maximizes its payoff, regardless of others' actions. Therefore, these choices form a Nash equilibrium of the subgame starting from the second stage of round m. Specifically, when firms in V_{m-1} act according to $T_m^*(T_{m-1})$, each firm in $\bar{V}_{m-1}$ will partner with the one and only one firm that has approached it, which is the unique dominant strategy equilibrium. Next consider Stage I. The implementation structure $T_m^*(T_{m-1}) = T^*(T_{m-1})$ maximizes the retailer's net gain among all structures in $\mathcal{T}(T_{m-1})$. By (A.83), it also maximizes the net gain of any node in V_{m-1}. Therefore, it is easy to see that the strategy profile corresponding to the retailer-preferred implementation structure is a Nash equilibrium of the subgame in the last round.

Consider round $k < m$ with history T_{k-1}. Assume that the strategy profile corresponding to the retailer-preferred implementation structure is a Nash equilibrium for any subgame starting from round $k+1$ with history T_k. Furthermore, assume that all players believe that the retailer-preferred implementation structure $T^*(T_k)$ will be the equilibrium outcome of the subgame starting from round $k+1$ with history T_k. First, consider Stage II of round k. Following the argument for round m, it is easy to see that for any firm $i \in \bar{V}_{k-1}$ with $I_{k-1}(i) \neq \emptyset$, the choices specified in (A.82) form a Nash equilibrium of the subgame starting from Stage II of round k. Specifically, if the firms in V_{k-1} act according to $T_k^*(T_{k-1})$, then each firm in $\bar{V}_{k-1}$ will be approached by at most one firm in V_k and will agree to partner with the firm (if any). Next, consider Stage I of round k. Because $T^*\big(T_k^*(T_{k-1})\big) = T^*(T_{k-1})$, we have

$$\pi_0\Big(T^*\big(T_k^*(T_{k-1})\big)\Big) = \pi_0\Big(T^*(T_{k-1})\Big) \geq \pi_0\Big(T^*(T_k)\Big), \;\forall T_k \in \mathcal{T}_k(T_{k-1}).$$

By (A.83), we have,

$$\pi_i\Bigg(T^*\big(T_k^*(T_{k-1})\big)\Bigg) \geq \pi_i\Bigg(T^*(T_k)\Bigg),\ \forall T_k \in \mathcal{T}_k(T_{k-1}),\ \forall i \in V_k.$$

Therefore, the strategy profile corresponding to the retailer-preferred implementation structure is a Nash equilibrium of the subgame starting from round k. Furthermore, note that under the strategy profile corresponding to the retailer-preferred implementation structure, non-leaf nodes in T_k will not approach any supplier. Therefore, we can restrict the game to the original setting that only firms in N_k can approach their suppliers. Using backward induction, this implies that the strategy profile corresponding to the retailer-preferred implementation structure is a subgame perfect Nash equilibrium for the sequential ESR relationship formation game with simultaneous moves, and the equilibrium outcome is the retailer preferred implementation structure, that is, the optimal solution to (4.12).

Next, consider the case of sequential moves in each round. In this case, the k-th round of the game has $2|T_{k-1}|$ stages, including "virtual stages" in which all players have empty action sets. In an odd-number stage, a firm in T_{k-1} approaches certain firms (which could be none) in $\bar{V}_{k-1}$. In the subsequent even-number stage, each approached firm (which could be none) either agrees or refuses to negotiate with the approaching firm. In the latter case, the game stops and the entire ESR program fails, causing all firms to have zero payoff. Therefore, for any approached firm in any even-number stage of any round, agreeing to negotiate with the approaching firm maximizes its net gain regardless of all other actions. Now, consider odd-number stages. Let $\sigma^k = (\sigma_1^k, \ldots, \sigma_{|T_k|}^k)$ denote the random permutation followed by the nodes in T_k in round $k+1$. Let $T_{k,i}$ denote the partial implementation structure after the i-th node in σ^k has established all its negotiation relations, and let $V_{k,i}$ denote the set of nodes in $T_{k,i}$. Start with the last odd-number stage (i.e., stage $2|T_{m-1}|-1$) of the last round (i.e., round m). Without loss of generality, assume $(\sigma_{|T_{m-1}|}^{m-1}, j) \in T^*(T_{m-1,|T_{m-1}|-1})$ for all $j \in \bar{V}_{m-1,|T_{m-1}|-1}$. It is easy to see that acting according to the retailer-preferred implementation structure $T^*(T_{m-1,|T_{m-1}|-1})$ maximizes firm $\sigma_{|T_{m-1}|}^{m-1}$'s payoff. Consider stage $l < T_{m-1}$. Assume acting according to $T^*(T_{m-1,i-1})$ is optimal for σ_i^{m-1}, $i =$

$l+1,\ldots,|T_{m-1}|$. The result in (A.83) implies that acting according to $T^*(T_{m-1,l-1})$ is optimal for σ_l^{m-1}. Therefore, the strategy profile corresponding to the retailer-preferred implementation structure is a Nash equilibrium for the subgame in round m. Using backward induction and following similar arguments, we can show that the retailer-preferred implementation structure is a sub-game perfect Nash equilibrium. Furthermore, a non-leaf node in T_k will not approach any node in $\bar{V}_k$. Therefore, we can restrict the game to the original setting, that is, only the leaf nodes in N_k can approach upstream firms. Consequently, the strategy profile corresponding to the retailer-preferred implementation structure is a subgame perfect Nash equilibrium for the sequential ESR implementation delegation game with sequential moves, and the equilibrium outcome is the retailer preferred implementation structure, that is, the optimal solution to (4.12). □

Proof of Theorem 4.5.3. Suppose that the retailer decides to initiate both ESR programs. Consider a negotiation pair (i,j), with all other bilateral contracts fixed. According to the Nash bargaining solution, the transfer payment allocates the total trade surplus proportionally to firms according to their bargaining power. Thus, in anticipation of a positive share of the total gain, the upstream firm will accept the offer from the downstream firm. Similarly, we can obtain the negotiation outcome of all other bargaining units. By applying the Nash-Nash solution, we derive the equilibrium gains of all firms when both ESR programs are implemented by solving the system of equations in the form of (4.9) and (4.13). Specifically, by (4.14), we derive the following

$$\sum_{i\in N_k}\pi_i(T,\mathbf{1})=\Gamma_k(T)(\pi_0(T,\mathbf{1})-D_0^k(T,\mathbf{1})),\ k\in\{1,2\}. \tag{A.84}$$

In addition, we observe that $\sum_{j\in\{u\in N_{3-k}:t_{0,u}=1\}}w_{0,j}(\mathbf{1})=\sum_{j\in N_{3-k}}\pi_j(T,\mathbf{1})+C_{3-k}$. Then, the disagreement of the retailer $D_0^i, i\in N_k$ can be computed as

$$D_0^i(T,\mathbf{1})=D_0^k(T,\mathbf{1})\equiv R_{3-k}-C_{3-k}-c_{0,3-k}-\sum_{j\in N_{3-k}}\pi_j(T,\mathbf{1})=\Pi_{3-k}-\sum_{j\in N_{3-k}}\pi_j(T,\mathbf{1}) \tag{A.85}$$

Combining the above two equations, we have,

$$\Pi_k-D_0^{3-k}(T,\mathbf{1})=\Gamma_k(T)(\pi_0(T,\mathbf{1})-D_0^k(T,\mathbf{1})),\ k\in\{1,2\}. \tag{A.86}$$

With the above relations between retailer's disagreement points and gain, with equation (4.15), we can derive the retailer's gain, which is given by

$$\pi_0(T,\mathbf{1}) = \frac{1-\Gamma_1(T)\Gamma_2(T)}{(1+\Gamma_1(T))(1+\Gamma_2(T))}\Pi + \frac{\Gamma_2(T)}{1+\Gamma_2(T)}\Pi_1 + \frac{\Gamma_1(T)}{1+\Gamma_1(T)}\Pi_2. \quad (A.87)$$

Moreover, the disagreement point of the retailer is given by,

$$D_0^i(T,\mathbf{1}) = D_0^k(T,\mathbf{1}) \equiv \Pi_{3-k} - \frac{\Gamma_{3-k}(T)}{1+\Gamma_{3-k}(T)}(\Pi-\Pi_k),\ i\in N_k \text{ and } k\in\{1,2\}. \quad (A.88)$$

In addition, we obtain the suppliers' gains.

$$\pi_i(T,\mathbf{1}) = \frac{\Pi-\Pi_{3-k}}{1+\Gamma_k(T)}\prod_{a\in P_i(T)}\rho_a,\ i\in N_k \text{ and } k\in\{1,2\}. \quad (A.89)$$

By the assumption, we have $\Pi-\Pi_{3-k} \geq \Pi-(\Pi_{3-k})^+ \geq (\Pi_k)^+ \geq 0$. As indicated by the above equations (A.89), the equilibrium gains of suppliers are nonnegative. □

Proof of Theorem 4.5.4. By Theorem 4.5.3, the gain of the retailer by implementing both ESR programs is given by

$$\pi_0(T,\mathbf{1}) = \frac{1-\Gamma_1(T)\Gamma_2(T)}{(1+\Gamma_1(T))(1+\Gamma_2(T))}\Pi + \frac{\Gamma_2(T)}{1+\Gamma_2(T)}\Pi_1 + \frac{\Gamma_1(T)}{1+\Gamma_1(T)}\Pi_2. \quad (A.90)$$

Since N_1 and N_2 are disjoint sets, the value of $\Gamma_k(T)$ is independent of that of $\Gamma_{3-k}(T)$. Thus, we have

$$\frac{\partial}{\partial\Gamma_k(T)}\pi_0(T,\mathbf{1}) = \frac{-\Pi+\Pi_{3-k}}{(1+\Gamma_k(T))^2} \leq 0, k\in\{1,2\}. \quad (A.91)$$

The above inequality holds due to the assumption on the gains of ESR programs. Specifically, $\Pi-\Pi_{3-k} \geq \Pi-(\Pi_{3-k})^+ \geq (\Pi_k)^+ \geq 0$. Therefore, the retailer's gain is maximized when both $\Gamma_1(T)$ and $\Gamma_2(T)$ are minimized, that is $T = T_1 \cup T_2$ □

Proof of Corollary 4.5.1. By the assumptions that $\Pi \geq (\Pi_1)^+ + (\Pi_2)^+$, the allocation based on Shapley value can achieve the grand coalition. To show that this allocation is a special case of the multi-unit bargaining approach, we consider the case where the bargaining power of suppliers for the same product family are equal such that the gain is distributed equally within a product family. Thus, $\rho_{i,j} = 1$ if

$i \neq 0$ and $(i,j) \in A^{\text{ESR}}$. Let ρ_k denote the bargaining power of a supplier of product family k relative to the retailer. In other words, $\rho_{0,i} = \rho_k$, $i \in N_k, k \in \{1,2\}$. For the ease of exposition, let $n_k = |N_k|, k = 1,2$. By Theorem 4.5.3, the gains of all firms are given by,

$$\pi_0(T,\mathbf{1}) = \frac{1-\big(n_1\rho_1\big)\big(n_2\rho_2\big)}{\big(1+n_1\rho_1\big)\big(1+n_2\rho_2\big)}\Pi + \frac{n_2\rho_2}{1+n_2\rho_2}\Pi_1 + \frac{n_1\rho_1}{1+n_1\rho_1}\Pi_2, \quad \text{(A.92)}$$

$$\pi_i(T,\mathbf{1}) = \frac{\rho_k}{1+n_k\rho_k}\big(\Pi - \Pi_{3-k}\big), i \in N_k, k \in \{1,2\}. \quad \text{(A.93)}$$

Suppose the gain sharing scheme from multi-unit bargaining coincides with the allocation based on Shapley value. Then, we have,

$$\frac{\rho_k}{1+n_k\rho_k}\big(\Pi - \Pi_{3-k}\big) = \phi_k(N),\ k \in \{1,2\}. \quad \text{(A.94)}$$

Note that $\pi_0(T,\mathbf{1}) = \phi_0(N)$ is redundant as we have $\sum_{i\in N}\phi_i(N) = \sum_{i\in N}\pi_i(T,\mathbf{1})$. Since the grand coalition is achieved according to the allocation based on the Shapley value, we have $\phi_0(N) + n_{3-k}\phi_{3-k}(N) \geq (\Pi_{3-k})^+, k = 1,2$. In addition, $\Pi = \phi_0(N) + n_1\phi_1(N) + n_2\phi_2(N)$. Thus, we have

$$\Pi - \Pi_{3-k} - n_k\phi_k(N) \geq \Pi - (\Pi_{3-k})^+ - n_k\phi_k(N) \geq 0. \quad \text{(A.95)}$$

Thus, we can set

$$\rho_k = \frac{\phi_k(N)}{\Pi - \Pi_{3-k} - n_k\phi_k(N)} \in [0,\infty), k \in \{1,2\} \quad \text{(A.96)}$$

such that the above equations (A.94) hold. This completes the proof. □

A.1 Efficient Algorithms for Finding the Implementation Structure

Theorem 4.4.2 establishes that finding the retailer's preferred implementation structure is equivalent to a shortest path tree problem. Therefore, one can apply known algorithms for the shortest path tree problem to find the optimal implementation structure in order to maximize its gain share from the ESR program. As we have discussed in §§4.4.2, most supply chains are typically acyclic. It is well known that the shortest path tree problem on a directed acyclic graph can be solved in linear time (see, e.g., Lawler, 1976). We describe a linear-time algorithm for finding the optimal implementation structure.

Algorithm 3 The optimal implementation structure when (N, A) is acyclic.

Initialization
$\varrho(0) \leftarrow 1$, $d(i) \leftarrow 0$, $\varrho(i) \leftarrow +\infty$, $p(i) \leftarrow$ NULL, $\forall i \in N \setminus \{0\}$, $\mathcal{D} \leftarrow \{0\}$, $\mathcal{C} \leftarrow \{0\}$, and $\mathcal{U} \leftarrow \emptyset$
for $(i, j) \in A^{\text{ESR}}$ **do**
    $d(j) \leftarrow d(j) + 1$
end for
Construction
while $\mathcal{D} \neq N$ **do**
    for $i \in \mathcal{C}$ **do**
        for $j \in \{k : (i, k) \in A^{\text{ESR}}\}$ **do**
            $d(j) \leftarrow d(j) - 1$
            if $d(j) = 0$ **then**
                $\mathcal{U} \leftarrow \mathcal{U} \cup \{j\}$
            end if
            if $\varrho(j) > \varrho(i) \cdot \rho_{i,j}$ **then**
                $\varrho(j) \leftarrow \varrho(i) \cdot \rho_{i,j}$ and $p(j) \leftarrow i$
            end if
        end for
    end for
    $\mathcal{C} \leftarrow \mathcal{U}$, $\mathcal{D} \leftarrow \mathcal{D} \cup \mathcal{U}$, and $\mathcal{U} \leftarrow \emptyset$
end while

Algorithm 3 iteratively updates the cost $\varrho(i)$ to reach node $i \in N \setminus \{0\}$. Each iteration starts with a subtree described by $\{p(i), i \in \mathcal{D}\}$, where the set $\mathcal{D}$ contains the nodes in the current subtree and $p(i)$ indicates the predecessor of node i in the subtree. The set $\mathcal{C}$ records leaf nodes in the subtree. The set $\mathcal{U}$ consists of the supplier nodes that can be added into the subtree in this iteration. This algorithm iteratively finds the cheapest way to reach the nodes not in the subtree from the leaf nodes.

It is easy to see that the initialization stage of Algorithm 3 takes $\Theta(|N| + |A^{\mathrm{ESR}}|)$ time and the construction stage takes $\Theta(|A^{\mathrm{ESR}}|)$ time. The total running time is therefore $\Theta(|N| + |A^{\mathrm{ESR}}|)$. For general networks with cycles, if there does not exist a cycle in the ESR network with a negative cycle length (in our problem, a negative-length cycle corresponds to a cycle C with $\prod_{a \in C} \rho_a < 1$), then the Bellman-Ford algorithm (Ford, 1956; Bellman, 1958) can be used to find the shortest path tree in $\Theta(|N||A^{\mathrm{ESR}}|)$ time. As we have discussed in §§4.4.2, supply chains with cycles are not common in practice. Therefore, we omit the detailed algorithm for this case.

B. ADDITIONAL NUMERICAL RESULTS

B.1 An Example with Non-unimodal Profit Function

In the classical pricing newsvendor problem, certain conditions are imposed on the random demand to guarantee that the expected profit function is unimodal or quasiconcave. For example, a class of literature assumes that the distribution of the random factor is generalized increasing failure rate (GIFR). Nevertheless, our framework does not require the restrictive assumptions on the distribution of the random factor. We will present a case in which our approach can find more robust solutions compared to the traditional approach when the distribution of the random factor is not necessarily GIFR. Specifically, we consider the additive demand of which the demand function is of the quadratic form, i.e., $d(p) = a - bp^2$. The underlying distribution is the beta distribution with parameters (α, β), which does not have GIFR (i.e., $\alpha, \beta \leq 1$). For the traditional approach, given the unit selling price p, the firm can determine the optimal order quantity y using the quantile information of demand distribution. Then, the firm searches along the dimension of unit selling price to find the optimal one. Since the distribution of the random factor may not be GIFR, the firm may obtain local optimal solutions as the objective function is not necessarily unimodal. On the other hand, the firm can also apply the minimax regret approach with the lower and upper limits of the random factor to determine the unit selling price and the order quantity to minimize the maximum regret. Figure B.1 presents the expected profit function with respect to the unit selling price and marks the optimal minimax price and the corresponding expected profit. We observe that while the traditional approach may choose the decisions that lead to the sub-optimum, the minimax regret approach could find a more robust one.

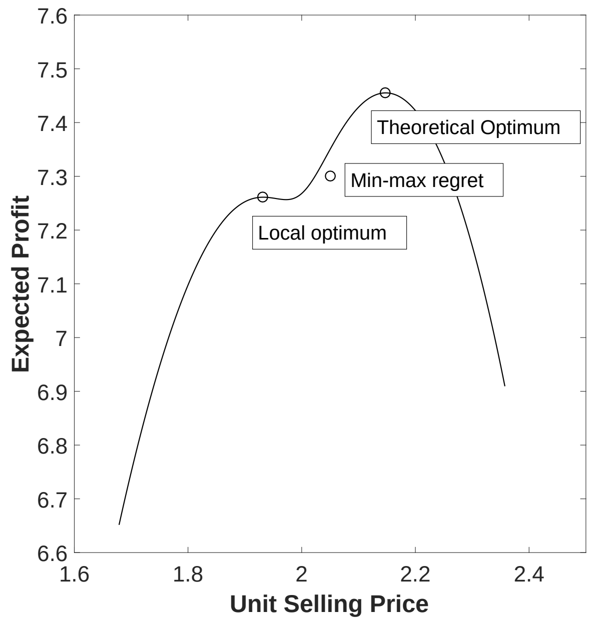

Notes. Consider the additive demand function $D(p, \theta) = a - bp^2 + \theta$, where $a = 20, b = 2.2$, and $[\underline{\theta}, \bar{\theta}] = [-4, 6]$. The unit ordering cost $c = 1$. The underlying random factor follows the beta distribution with $(\alpha, \beta) = (0.01, 0.01)$.

Fig. B.1.: An numerical example when GIFR does not hold

B.2 Numerical Comparison of Minimax Regret and Max-min Robust Frameworks

We conduct a numerical experiment that compares the performance of the minimax regret approach and the max-min robust optimization approach with the benchmark that adopts the traditional newsvendor model with complete demand information. Since the max-min robust optimization approach only focuses on the worst-case scenario, the decisions obtained by the max-min robust optimization approach are conservative and the corresponding expected profit provides a lower bound of the optimal expected profit. Specifically, if we fix the mean of the random factor and vary its variance, we find that the minimax regret approach consistently outperforms the max-min robust optimization approach. Specifically, we set the mean to 0.3, 0.5, and 0.7 where the random factor follows the beta distribution. In Figure B.2, we first observe that expected profits using decisions obtained by both approaches are decreasing in the demand variability with a fixed mean. In addition, when the mean of the random factor is increasing, the expected profit using decisions obtained by the minimax regret approach are improved. On the contrary, the performance of the max-min robust optimization approach cannot benefit from such changes in the underlying distribution.

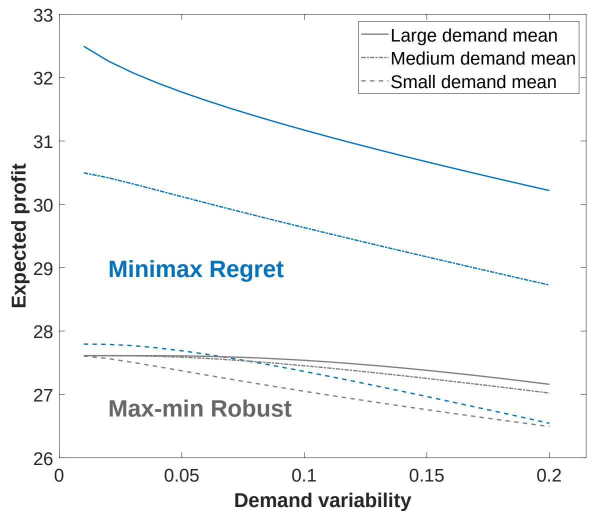

Notes. (i) Consider the additive demand function, $D(p, \theta) = a - bp + \theta$, where $a = 30, b = 5$, and $[\underline{\theta}, \bar{\theta}] = [-2, 2]$. The unit ordering cost $c = 1$. (ii) The underlying random factor follows the beta distribution.

Fig. B.2.: Performance comparison varying variance

In the previous analysis, we use the maximum and minimum of realizations of the random factor as the upper and lower bounds to construct the uncertainty set for the minimax regret framework and the max-min robust optimization approach. In the following numerical study, we investigate the effect of the choice of the uncertainty sets on the performance of both approaches. Specifically, suppose the firm only knows the $\alpha/2$ and $(1 - \alpha/2)$ quantile of the random factor and constructs the uncertainty set with this quantile information. We consider the linear additive demand and α is set to 0.1, 0.2 and 0.3 respectively. As shown in Figure B.3, we compare the performance of the two approaches with different uncertainty sets while varying the variance with a fixed mean. We observe that the performance of the minimax regret framework is robust with different quantile-based uncertainty sets while that of the max-min robust optimization approach highly depends on the choice of uncertainty sets. Because when the lower limits of the uncertainty sets are increasing, the solution derived by the max-min robust optimization approach becomes less conservative as it is forced to order larger quantities.

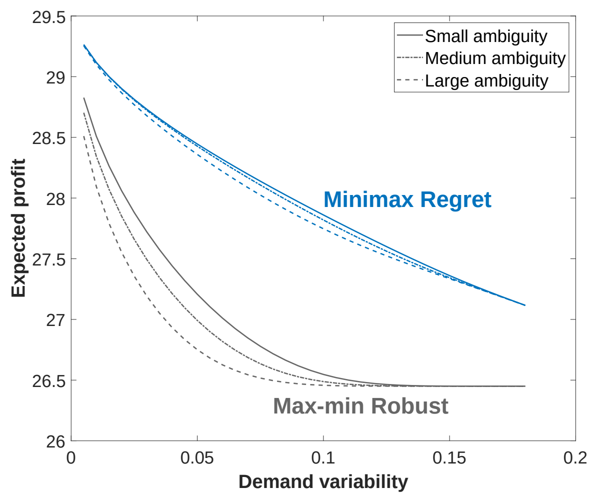

Notes. (i) Consider the additive demand function, $D(p, \theta) = a - bp + \theta$, where $a = 30, b = 5$, and $[\underline{\theta}, \bar{\theta}] = [-2, 2]$. The unit ordering cost $c = 1$. The underlying random factor follows the beta distribution. Confidence level α is set to 0.1, 0.2, and 0.3.

Fig. B.3.: Robust optimization with quantile-based intervals

In summary, we compare the minimax regret approach and the max-min robust optimization approach numerically in this section. We first show that the minimax approach can generate higher expected profits compared to the max-min robust optimization approach when the mean of the random factor is moderate to large. Second, decisions obtained by the minimax regret approach always leads to larger expected consumption. Third, the performance of the minimax regret approach is robust with different choices of the uncertainty sets while that of the max-min robust optimization approach highly depends on such choices. Therefore, the minimax regret approach can provide more robust solutions than the max-min robust optimization approach in the context of pricing newsvendor problem.

B.3 Robustness Check for the Choice of Uncertainty Sets

In this section, we will show that the performance of the minimax regret approach does not highly depend on the choice of uncertainty set in the data-driven setting. We evaluate and compare the performance of the minimax regret approach with different sample sizes, demand variability, and confidence levels α in a data-driven setting described as in Section 3.6 and summarize the results in Figure B.4. We find that the performance of the minimax regret framework is stable when α is ranging from 0.2 to 0.6 with different sample sizes and small-to-moderate demand variability. When the demand variability is large, the performance remains to be stable when α is ranging from 0.3 to 0.6. Specifically, we define the following metric, *relative variation*, to measure the robustness of the minimax regret framework over the interval $[\underline{\alpha}, \bar{\alpha}]$:

$$\frac{\sup_{\alpha \in [\underline{\alpha}, \bar{\alpha}]} \Phi(\alpha) - \inf_{\alpha \in [\underline{\alpha}, \bar{\alpha}]} \Phi(\alpha)}{\max_{p,y} \mathrm{E}[\phi(p, y)]} \times 100\%. \tag{B.1}$$

where $\Phi(\alpha) = \mathrm{E}_\theta \big[\mathrm{E}_{\mathbf{D}}[\phi\big(p^*_\alpha(\mathbf{D}), y^*_\alpha(\mathbf{D})\big)]\big]$ is the expected profit with optimal decisions $\big(p^*_\alpha(\mathbf{D}), y^*_\alpha(\mathbf{D})\big)$ derived from the minimax regret framework using observed data $\mathbf{D}$. The relative variation essentially measures the maximum fluctuation of the performance of the minimax regret framework compared to the hindsight optimal expected profit over a certain range of the confidence level α. When the demand variability is small ($\sigma = 1$), the relative variation is ranging from 0.14% to 0.17%. When the demand variability is moderate ($\sigma = 3$), the relative variation is ranging from 0.28% to 0.37%. When the demand variability is large ($\sigma = 5$), the relative variation lies in between 0.49% and 0.82%. Thus, the minimax regret approach is robust with different choice of uncertainty sets and we choose $\alpha = 0.4$ to construct the uncertainty set for the numerical experiments in the main body.

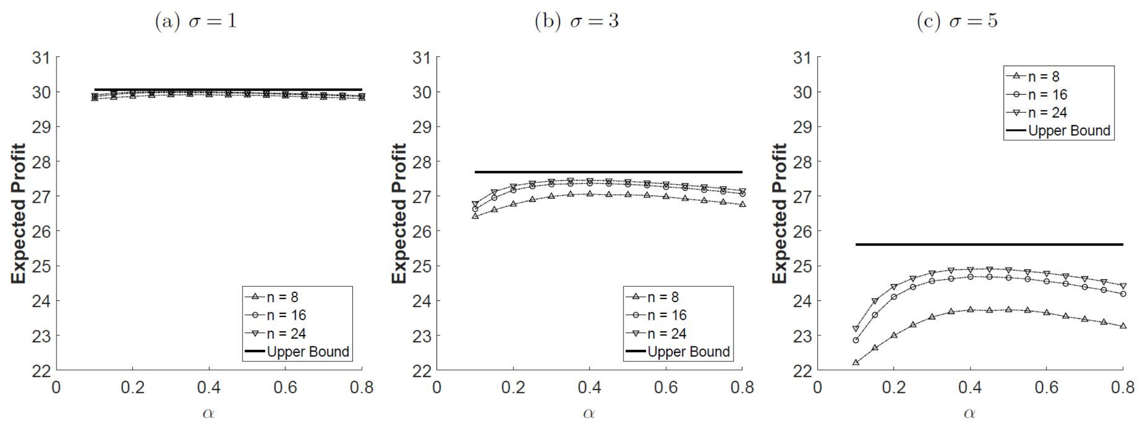

Notes. (i) Consider an additive demand function, $D(p, \theta) = a - bp + \theta$ where $a = 30, b = 5$. The random factor $\theta \sim \mathcal{N}(0, \sigma^2)$, where σ is set to 1, 3, and 5. The unit ordering cost $c = 1$. (ii)The confidence level α ranges from 0.1 to 0.7. The sample size n is set to $8, 16, 24$.

Fig. B.4.: Choice of confidence level

VITA

VITA

Chengzhang Li was born in Wuhan, Hubei province, China, to Mr. Hao Li and Ms. Defen Zhang. He received his bachelor degree in Mechanical Engineering and Automation from Shanghai Jiao Tong University. He received his M.S. degree in Statistics and Computer Science from Purdue University. He is currently a Ph.D. student in Operations Management at Purdue University.