Use of Game Theory and Stochastic Programming for Supply Chain Optimization

BY

AMY DAVID
BS, University of Illinois at Urbana-Champaign, 2003
MBA, Lake Forest Graduate School of Management, 2008

THESIS

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Industrial Engineering and Operations Research
in the Graduate College of the
University of Illinois at Chicago, 2014

Chicago, Illinois

Defense Committee:

Lin Li, Chair
Urmila Diwekar, Advisor
Houshang Darabi
David He
Elodie Goodman, University of California-Riverside
Bin Jiang, DePaul University

ACKNOWLEDGMENTS

Special thank you to USG for its financial support in completing this dissertation. Thank you also to my advisors, Dr. Elodie Goodman and Dr. Urmila Diwekar, and my industry collaborators Ross Januszyk and David Farr.

PREFACE

Please note that Chapter 3 describes a supply chain optimization project executed for a private corporation under a non-disclosure agreement, and as such, raw data is not included this document. The author has instead provided summary data and believes this to be sufficient in conveying the significance of the work.

TABLE OF CONTENTS

TABLE OF CONTENTS (Continued)

CHAPTER **PAGE**

LIST OF TABLES

LIST OF FIGURES

LIST OF ABBREVIATIONS

PDF	Probability density function
CDF	Cumulative density function
LP	Linear Program

SUMMARY

This thesis investigates supply chain optimization: managerial decisions related to the design of a supply chain with the goal of maximizing profits. The work is divided into two parts, the first focusing on the use of game theory and the second focusing on the use of stochastic programming.

In the first part (Chapter 2), we study competition and coordination in a supply chain in which a single supplier both operates a direct channel and sells its product through multiple differentiated retailers competing in quantities. We construct theoretical models of the supply chain, and use these models to develop managerial insights. Because it is difficult to obtain real-world pricing data to verify the analytical results in this part, the impact of this research is limited to theory that may be applied in future work.

We first study the supply chain with symmetric retailers and find that the supplier generally prefers to have as many retailers as possible in the market, even if the retailers' equilibrium retail price is lower than that of the supplier, and even if the number of retailers and their cost or market advantage prevents sales through the direct channel.

We also find that the two-channel supply chain may be subject to inefficiencies not present in the single-channel supply chain and study coordination. We show that several contracts known to coordinate a single-channel supply chain do not coordinate the two-channel supply chain; thus we propose a linear quantity discount contract and demonstrate its ability to perfectly coordinate the two-channel supply chain with symmetric retailers. We study numerically the supply chain with asymmetric retailers and find that the supplier still benefits from having more retailers in the market and that linear quantity discount contracts can mitigate supply chain inefficiency, though they no longer achieve perfect coordination.

We then extend our investigation of the two-channel supply chain to that in which the supplier has limited capacity, either known or subject to probabilistic constraints. We show that the supplier cannot benefit by selling through the retailer if the maximum possible capacity is small, or if the lower bound on capacity is in a range that would cause the indirect channel's capacity to be zero when the supplier optimizes its retail price to maximize profits. We further show that, under some conditions, the supply chain profits can actually benefit from uncertainty, and that the retailer typically absorbs most of the profit loss when the profits instead decrease.

In the second part (Chapter 3), we present a case study of a large-scale stochastic optimization problem for USG, a building supplies manufacturer with plants and customers throughout North America. USG seeks to minimize total delivered cost (including production and freight costs) of products in its Durock® product line, subject to capacity constraints and uncertainties in both demand and production costs. We first demonstrate that demand uncertainty, rather than production cost uncertainty, is the main cause of month-to-month variations in total cost. We then use the chance constraint method to optimize the network, and propagate uncertainty through the cost models, applying a penalty cost for unfulfilled constraints. We show that we can reduce theoretical costs by approximately 4.8% by optimizing the network for the 50th percentile of demand, and reduce costs by approximately 1.6% as implemented, as compared to the base case using a single month's demand and cost data.

1. INTRODUCTION

All firms that produce physical goods engage in three main areas of supply chain management, broadly defined as source, make, and deliver, as shown in Figure 1. The category of "source" includes any activities related to procuring raw materials to enable production, while "make" involves the actual making of the physical goods. "Deliver" encompasses any activities involved in providing the goods to the customer, such as finished goods stocking, order fulfillment, and transportation management.

Effective management of each area requires not just the day-to-day execution of supply chain-related tasks, but the design of a supply chain appropriate for a firm's strategy. Given the high cost of the supply chain in manufacturing industries (a few representative examples include the automobile industry, which spends 67% of revenue in the supply chain, and the paper industry, which spends 55% of revenue (1)), tools for optimizing the supply chain to reduce costs or increase profits are of considerable value to firms engaged in these activities.

This thesis focuses on two supply chain design problems in the area of delivery. The first is that of channel selection: should a supplier sell directly to its end customers, sell through an

Figure 1: Activities in supply chain management

independent retailer, or both? We address this question in Chapter 2 by looking at a variety of factors that may influence channel strategy, including costs, relative market sizes, price sensitivity, and the number of retailers participating in the market. We also investigate how the decision is affected by capacity constraints, both known and probabilistic. This portion of the thesis is theoretical, as it is difficult to obtain real world data to verify the theoretical results of this part. Therefore, in the second part, we concentrated on a real world supply chain management problem related to network planning.

This second problem is that of designing a distribution network: given a set of plants, warehouses, and customers throughout North America, which plants should deliver to which customers to minimize total costs? Chapter 3 presents a case study in which we optimize the North American distribution network for USG, a major building supplies manufacturer, taking into consideration production, transportation, and warehousing costs, as well as capacity constraints at each production facility. We present an approach to this problem that accounts for the uncertainty in customer demand, and find that a considerable cost reduction can be achieved, as compared to the present-state.

2. USE OF GAME THEORY TO OPTIMIZE THE TWO-CHANNEL SUPPLY CHAIN

2.1 Introduction

2.1.1 Background and Motivation

The increasing consolidation of consumer goods retail into “superstores” has given these large retailers significant power to drive consumer choices. Close to 6% of U.S. retailing happens through a single retailer, Walmart, in both its online and bricks-and-mortar store (2). In addition to Walmart and other general superstores (Target, Meijer, etc.), the last two decades have also seen the growth of specialty superstores, such as Lowe’s and Home Depot in the home improvement sector, Best Buy in consumer electronics, and Barnes & Noble in book-selling.

The combination of these two factors presents a dilemma for suppliers with the resources to operate their own retail channels. Some, like Nike and Apple, have exclusive bricks-and-mortar stores that sell their products directly to end consumers in addition to a presence in superstores. Others take advantage of e-commerce to reach customers in a wide geographical region with comparatively low investment. In either case, these suppliers that extend into the retail space find themselves competing with the retailers to whom they sell products for the dollars of the end consumer. In some cases, the supplier may even find its price being undercut.

If the supplier is operating under capacity constraints, as when they are under strict time constraints for a new product launch, these constraints add another element of complexity to the supplier’s decision. In this case, any amount sold through an independent channel reduces the amount available in the direct channel. This is common in the consumer electronics industry: for example, when Apple released the “new iPad” in 2012, the product was backordered through the direct channel’s website, while available at some of the independent retailers.

We have studied the behavior of these two channel supply chains, those in which a supplier both operates a direct channel and sells its product through multiple independent retailers, with the goal of developing insights into the effects of competition between the channels; the “price of anarchy,” or profits lost to local optimization; opportunities for coordination to improve total supply chain profits; and the effects of capacity constraints.

In this section, our contributions are strictly theoretical: we are unable to verify them through the use of empirical data, yet we hope these insights lay a foundation for future applications:

1. We determine the supply chain structure at equilibrium, and we obtain the equilibrium retail prices, wholesale price, quantities and efficiency in closed form when the retailers are symmetric in cost and demand characteristics.

2. We obtain the supply chain efficiency in closed form when the retailers are symmetric and show that the two-channel supply chain is subject to sources of inefficiency not present in the single-channel supply chain.

3. We demonstrate that many of the contracts that perfectly coordinate the traditional one-channel supply chain fail to coordinate the symmetric two-channel supply chain, but that a linear quantity discount contract does.

4. We show numerically that when the retailers are not symmetric, some of our analytical findings for symmetric retailers continue to hold. The linear quantity discount contract no longer perfectly coordinates the two-channel supply chain, but contract parameters may be found that significantly improve the total supply chain profit.

5. We determine the supply chain structure at equilibrium, and we obtain the equilibrium retail prices, wholesale price, quantities and efficiency in closed form when a supplier competes with a single retailer and capacity is uncertain.

2.1.2 Review of Literature

Since (3) introduced the idea of "double marginalization", many have studied how vertical integration affects a supply chain's quantities, prices, and profits (4). In the last two decades, the number of suppliers vertically integrating their supply chain by creating their own "direct" channels to reach end customers has increased, and with it, the volume of literature on the topic (5), a trend largely due to the ease with which a supplier may compete online (6).

(7) assume that the direct channel is at an inherent disadvantage to the retail channel with respect to customer preference, exclusive of price, i.e given the same price in both channels, no customer would choose to purchase directly from the manufacturer. Using this type of model, they find that it is most profitable for the manufacturer to set its wholesale and retail prices such that nothing is ever sold through the direct channel, but to maintain the direct channel as a means of influencing the independent retailer's retail price. (8) similarly award a customer preference advantage to the retailer, while assuming a cost advantage for the supplier, and develop conditions under which both a retail channel and a direct channel are active in

equilibrium. They show that this occurs when the difference in marginal costs between the two channels falls within a specific range, outside of which either the supplier cost advantage or the retailer's customer preference advantage will allow one firm to set its prices aggressively enough to achieve whole market coverage on its own. (9) focus on the supplier's desire to preserve its relationship with the existing retailer, and therefore analyze a scenario in which a supplier commits to selling through its direct channel at the same price as in the retail channel. They conclude that both firms may benefit from the addition of the direct channel as long as the equal-pricing policy is maintained, the supplier benefiting from additional revenue and the retailer benefiting from a wholesale price reduction. They note, however, that as the direct channel gains acceptance among consumers, the supplier is increasingly motivated to set a retail price lower than that of the retailer, putting the retailer's profits in jeopardy. By contrast to these three papers, our intended model does not assume that any channel possesses an absolute advantage in customer preference, nor in cost, as both the direct and the traditional retail channels may be online, bricks-and-mortar, or both, and rising consumer acceptance of the internet suggests a heterogeneity of consumer preference (6).

Another stream of literature focuses on the challenges and opportunities for a retailer operating more than one channel (e.g. physical store, internet, catalog) in competition with other retailers. (10) present many of the challenges inherent in operating an internet channel alongside a physical store or catalog channel: cannibalization of higher-margin sales, implementation and on-going operational costs, and customer retention issues, but conclude that a consumer-centric view of channel offerings ultimately allows for profitable operation in a multi-channel environment. (11) similarly investigate the decision for retailers to add internet channels alongside bricks-and-mortar stores, showing that, in equilibrium, all retailers will add an internet channel, but typically at a loss of profits. More recently, (12) examine the qualitative impact of channel multiplicity, and suggest further research on how products, customers, channel leadership, and distribution intensity are viewed in an environment with many channels. Other authors exploring the role of several channels operated by the same firm include (13), who find that a retailer-owned online channel may increase that retailer's investment in service components at the physical store and (14), who show that customers who migrate to an online channel from a catalog channel ultimately purchase more over time. While these works investigate each retailer's options in choosing their channel strategy, we will focus, by contrast, on the difference between a supplier-owned channel and a retailer-owned channel, and the resulting channel conflict *independent* of whether the channel is a physical store, internet, catalog, or other type of channel, noting that often the conflict occurs between a supplier's channel and a retailer's channel of the same form.

Other authors similarly focus on the specific use of one type of channel versus another, but also consider direct channels in their analysis. (15) look at the situation where a manufacturer sets up a high-cost direct channel with the expectation that many customers will use the direct channel to learn about the product and then instead purchase from a retailer, i.e. the manufacturer deliberately keeps its retail price high, allowing the retailers to undercut its price and "free ride" on its branding efforts. They find that the existence of the direct channel always benefits the manufacturer, and is most valuable when consumers require a high degree of information about the product being sold or when search costs are either very high or very low. Because our model instead assumes no customer search, there is no such “free riding” effect. Further, our model allows the manufacturer’s direct channel to take a variety of media, and thus cost structures. In combination, these two factors lead us to conclude that, when the retailers have a more favorable cost or market position, the manufacturer’s equilibrium retail price will be lower than that of the retailers, and in the most extreme cases, the manufacturer will leave the retail market entirely. (16) examine the addition of an internet channel, owned by either the manufacturer, a retailer, or a brand new entrant, to either a vertically integrated or decentralized supply chain consisting only of physical stores, comparing it to the entry of an additional physical store, assuming that customer disutility for a purchase from a physical store is proportional to the distance from that store. They conclude that the creation of an internet channel has a different impact on prices and customer utility than the addition of another physical store, and the specific costs or benefits depend on the supply chain’s degree of vertical integration prior to the addition, as well as the distance between physical stores when multiple are present. While this paper also addresses channel conflict, none of the scenarios considered in the existing literature described above allow for an internet presence in both the direct and indirect channels, a common situation as noted in our motivating examples, and one our model will accommodate.

Discussions of channel strategy are also prominent in the literature on franchising. (17) find that a manufacturer is better off selling through a franchisee–an independent retailer carrying only the manufacturer’s products–when competition is high, but prefers vertical integration if he can sufficiently differentiate his product to lower competition. (18) extends these results to the more general case of nonlinear demand and supports his findings with empirical data, while (19) argues that the benefits of selling through an independent retailer rely on not just the intensity of retail competition but also the likelihood of competitors at either tier of the supply chain raising their prices in concert. (20) notes that many franchisors have a combination of company-owned and franchised locations, an example of "dual-channel" distribution, and uses empirical data to show that the franchising vs. integration decision is dependent on a variety of factors including population density, the labor intensity, and the annual growth of an industry, thus concluding that this decision must be made anew for each proposed location. (21) present a review in which they argue that most franchisors have a target percentage for the number of company-owned outlets, driven largely by the value of their brand names. As can be seen in the motivating examples above, in our context, the relationship between the supplier of a single product and its potential independent retailers differs considerably from that of a franchisor and its franchisees. Most notably, the supplier is not able to influence demand at the retailer to the degree that a franchisor may do so at a franchisee and the supplier's decision to sell through an indirect channel is therefore based solely on pre-existing information about the retailers' demand curve. Further, a franchisee's initial investment in a franchise has no analogue for a retailer choosing to add a product to its offerings; a retailer's decision to enter the market is influenced mainly by its expected gross margin.

Also of interest is the large body of recent literature focused on supply chain coordination mechanisms, with the goal of increasing efficiency by providing incentives to align decisions with those of a centralized supply chain (22). (23) was one of the first to show how coordination mechanisms could improve supply chain efficiency in his work on buy-back mechanisms. Other early work on supply chain coordination proposed quantity discounts as a means of improving the efficiency of a channel consisting of a single supplier and a single retailer (24), (25). (26) shows that a supply chain with a single supplier and multiple homogenous retailers may be coordinated using quantity discounts along with a fixed payment or “franchise fee”, and shows the equivalence of an all-units quantity discount and an incremental quantity discount policy in doing so. (27) suggest a “linear quantity discount” contract, in which the wholesale price is a linearly decreasing function of quantity, as a means of coordinating the lot sizes of a supplier and a retailer with fixed retail price and demand, while (28) show that such a contract is sufficient to coordinate a supply chain with one supplier and two independent retailers competing on price. (29) use a model of one supplier and many retailers, introducing a “price-discount sharing” (PDS) scheme, in which the wholesale price to a retailer is a function of that retailer’s retail price. They demonstrate that a linear PDS is sufficient for coordinating a supply chain with non-competing retailers, but that a nonlinear scheme may be required when differentiated retailers compete on price. (30) focus on the inefficiencies caused by externalities among retailers, concluding that complementarity aggravates the double marginalization effect, providing an opportunity for significant efficiency improvements through coordination, while substitutability causes the opposite. (31) analyze revenue sharing contracts, finding that the supply chain can be coordinated, and that profits may be arbitrarily allocated in a supply chain consisting of a single supplier and multiple differentiated retailers engaging in Cournot competition. Our work will extend this literature by testing these contracts in a two channel setting to see if they achieve perfect coordination when both the supplier and the independent retailers operate retail channels.

(32) and (33) bring together literature on the two channel supply chain and coordination mechanisms. (33) model a direct channel in competition with a retail channel in which the allocation of demand between the two channels is constant, total demand is dependent on sales effort in both channels, and retail price is fixed at the same value in both channels. They find that a wholesale price contract must be dependent on both sales effort to coordinate such a supply chain, but note the practical difficulty of implementing any scheme that requires the supplier to monitor the retailer's effort. (32) takes price and demand to be exogenous, though both may differ by channel, and finds that many common contracts (including buy-back, rebate, and revenue sharing) fail to coordinate the supply chain. He instead suggests a "penalty" contract in which the retailer pays the supplier a unit penalty per missed sale to achieve perfect coordination, but, like the effort-dependent contract of Tsay and Agrawal, such a contract may be infeasible because of its requirement that the supplier know the retailer's lost sales. Our work differs from both of these in that we assume price to be endogenous (resulting from the retail quantity decisions and price-demand relationship), and the proportion of the demand realized in the direct channel changes with the decisions made in both channels. We also intend to study the viability of contracts that, unlike those proposed by these authors, require no complex tracking of the retailer's activities.

Many authors have addressed supplier uncertainty and its effect on the supply chain by assuming that suppliers are subject to a disruption that results in complete default (34) [1]. (35), (36), (37), and (38) all investigate the effects of information asymmetry, when a supplier knows its true default risk but the retailers do not. (39) investigate a supply chain featuring a single retailer with uncertain demand and multiple suppliers, each with a known probability of delivery failure. They show that, with two suppliers, a high correlation between their default risks benefits the retailer by driving wholesale prices downward. As more suppliers are added to the supply chain, the retailer may be able to benefit from both low wholesale prices and low default risk. From the suppliers' perspective, they benefit from actions that lower their own default risk's correlation with that of other suppliers, as that leads to higher wholesale prices. (40) studies a single retailer ordering from a single risky supplier that is possibly subsidized by the retailer, and develops conditions under which the retailer's optimal ordering decision is independent of the decision to subsidize the supplier.

More closely related to our work is the literature that assumes a capacity or yield uncertainty, i.e. that a supplier may make a partial delivery to the retailer when production conditions are unfavorable, rather than fully default. (41) assume that such a partial delivery is possible, but, like Babich et al., find that a retailer suffers from a positive correlation among the suppliers' yields. (42) and (43) both address the question of when investments in the suppliers should

---

[1]For consistency, we refer to the supplier as the upstream partner and the retailer as the downstream partner in a two-tier supply chain, though various authors use supplier and manufacturer to indicate the same

be made to reduce uncertainty. The former models multiple suppliers with uncertain yields in competition for sales to a retailer, finding that the retailer typically benefits from providing incentives to its suppliers to reduce uncertainty, either through process improvement or the adoption of minimum standards, in spite of the wholesale price inflation that may follow. Wang et al. investigate whether a retailer should dual-source or directly invest in a single supplier, finding that the investment strategy is typically beneficial when the suppliers have a high level of cost heterogeneity, but dual-sourcing is preferred when there is a high level of reliability heterogeneity. (44) study a single supplier and a single retailer with uncertain demand. The supplier's capacity is fixed, but it may outsource excess demand to an OEM with uncertain capacity. They find that the supplier benefits from coordination with the OEM, but the OEM may not see a benefit unless the supplier also coordinates with the retailer.

While these authors all address supply uncertainty, none have done so in a two channel setting. Our proposed work is therefore unique in this respect: the supplier is a retailer, and thus needs to consider its own uncertainty in planning its retail quantity and that of the other retailers. While (45) address the problem of allocating a fixed amount of inventory among numerous retailers, the addition of both capacity uncertainty and conflict between the direct and indirect channel significantly differentiates the research questions addressed.

2.2 Competition and Coordination in a Two-Channel Supply Chain

2.2.1 The Model

We consider a single supplier that exclusively supplies a single product to N retailers. In addition, the supplier also operates a direct retail channel. These $N+1$ firms competing in quantity form an oligopoly over the end market.

The retail price at a given firm (one of the retailers or the supplier direct channel) is affected by the total quantity released to the market, and therefore is a function of both the firm's quantity and its competitors' quantities. We also assume that there is some degree of differentiation in the customer experience at different firms, and that some customers prefer each firm, exclusive of price. This is a reasonable assumption when one considers the many reasons a customer may choose a retail outlet: some specific to bricks-and-mortar stores (location, staff, hours of operation), some specific to online venues (shipping policies, fulfillment times, level of product detail), and some common to both (return policies, credit card acceptance, availability of other products). We therefore use a linear inverse demand model that allows us to capture this differentiation:

$$p = \alpha - Bq, \tag{2.1}$$

where $p = (p_0, p_1, \cdots, p_N)$ is the vector of retail prices, with p_0 representing the retail price at the supplier direct channel and $p_1, \cdots, p_N$ representing the prices at the N retailers, $q = (q_0, q_1, \cdots, q_N)$ is the vector of quantities, with q_0 representing the quantity at the supplier and

$q_1, \cdots, q_N$ representing the quantities at the N retailers, $\alpha = (\alpha_0, \alpha_1, \cdots, \alpha_N)$, with α_0 representing the maximum selling price at the supplier and $\alpha_1, \cdots, \alpha_N$ representing the maximum selling price at the retailers, and B in $\mathbb{R}^{(N+1)\times(N+1)}$ is the symmetric price sensitivity matrix given by:

$$B = \begin{bmatrix} \beta_0 & \gamma_{01} & \cdots & \gamma_{0N} \\ \gamma_{10} & \beta_1 & \cdots & \gamma_{1N} \\ \vdots & \vdots & \ddots & \vdots \\ \gamma_{N0} & \gamma_{N1} & \cdots & \beta_N \end{bmatrix}.$$

The coefficients $\beta_0, \cdots, \beta_N > 0$ represent the price sensitivity of demand at each firm with respect to its own quantity, and $\gamma_{ij} > 0$ represents the cross-sensitivity of demand at firm i with respect to the quantity at firm j. Because we assume that the change in a competitor's quantity affects a firm's price less than a change in that firm's own quantity, we add the restriction that $\gamma_{ij} < \beta_i \ \forall i, j$ (46). This is broadly applicable when both channels exist as physical locations; customers must physically travel from one to another, at a cost. However, even if both channels are online, a customer may face switching costs in the form of setting up a new account, not being able to bundle shipping with other products, etc. For example, it is easy to imagine that a customer purchasing several items from the Sam's Club website may prefer to add a Keurig machine to his existing order over creating a new order at the Keurig website. This demand model is common throughout the recent operations management and economic literature (46), (47), (48)[1]. Note that we model competition via quantities to reflect the fact that order quantity decisions often have to be made in advance of the selling season, and prior to pricing decisions, especially when independent retailers place an order from an external supplier. (49) state that the Cournot competition model fits best situations where retailers "simultaneously and independently make quantity decisions" and then bring these quantities to the market, letting the price be determined by the quantities on the market. They show that price competition a la Bertrand would require that quantity decisions follow the realization of demand, which is not always realistic when retailers order from an external supplier. In addition, they show that quantity competition is equivalent to quantity precommitment followed by price competition, thus our model can be viewed as similar to a price competition model as long as retailers and supplier first commit in terms of quantities.

Production incurs a fixed per-unit cost, c_A, and retailing incurs a fixed per-unit cost, c_i, for firm i. We assume that the cost of retailing includes all variable costs, such as inbound shipping, and that the retailers may be asymmetric in such costs. For ease of notation, we let $c_i' = c_i + c_A$ represent the total cost of a unit sold through firm i. The quantity $\nu_i = \alpha_i - c_i'$ represents the maximum product margin of a unit sold through firm i. This maximum product margin is the difference between one unit's maximum selling price and its total cost to the supply chain.

---

[1]In practice, there is some uncertainty in price. However, as long as the firms are risk-neutral, neither an additive nor a multiplicative uncertainty has an impact on expected utility maximization and our results continue to hold. For ease of exposition, we therefore present the price as a deterministic function of demand.

Without loss of generality, we assume the maximum product margin to be positive for each channel, as otherwise no item would be produced and sold through that channel.

The supplier and the N retailers engage in a Stackelberg game where the supplier is Stackelberg leader. The supplier chooses the total quantity, Q, to be sold through retailers and the retail quantity q_0 to be sold in the direct channel. The retailers then simultaneously react by choosing their retail quantities q_i, which then determine the retail prices of all firms, while the wholesale price w is set to clear the wholesale market.

For tractability in developing our analytical results, in the next two sections we assume that all N retailers are symmetric. In Section 2.2.4, we numerically investigate the effects of retailer asymmetry.

2.2.2 Equilibrium Analysis With Symmetric Retailers

In this section we consider all retailers to be symmetric in terms of cost characteristics and demand parameters, i.e. $\alpha_1 = \cdots = \alpha_N$, and likewise for $\beta_i, \gamma_{ij}, c_i, c_i'$, and ν_i. We denote this by eliminating the subscripts on $\alpha, \beta, \gamma, c, c'$, and ν to indicate they apply to all retailers, while maintaining the use of 0 as the subscript for the supplier. While we realize that this is a restrictive assumption that is not satisfied in a strict sense in reality, it keeps the model tractable and allows us to obtain analytically some managerial insights that we test numerically in Section 2.2.4 when the symmetry assumption is relaxed. Further, the difference in cost and market positions among retailers is likely to be much less significant than the difference in cost and market positions between a supplier and its retailers. For example, Walmart and Target are much more alike than Walmart and Random House.

The ratio $\rho = \frac{\nu_0}{\nu}$ of maximum margins gives an indication of the relative strength of the supplier's direct channel compared to the indirect channel. Also, we introduce the notations $\delta = 2\beta + \gamma(N-1)$ and $\Delta = \beta + \gamma(N-1)$ to simplify expressions that appear often in our results.

2.2.2.1 The Decentralized Case

We first determine the equilibrium solution of the Stackelberg game described above. The supplier first chooses the total wholesale quantity and the direct channel retail quantity, anticipating the market-clearing wholesale price and the retailers' quantities, then all retailers choose their quantities $q_i, \ i = 1, \ldots, N$. The profit to the supplier, Π_0, and the profit to retailer i, Π_i, are therefore given by[1]:

$$\Pi_0 = q_0(\alpha_0 - \beta_0 q_0 - \gamma Q - c_0') + Q(w - c_A) \tag{2.2}$$

$$\Pi_i = q_i(\alpha - \beta q_i - \gamma \sum_{\substack{j=0 \\ j \neq i}}^{N} q_j - c - w), \quad i = 1, \cdots, N, \tag{2.3}$$

and the supplier's optimization problem is constrained by the wholesale market clearing condition: $\sum_{i=1}^{N} q_i(w) = Q$, where $q_i(w)$ denotes the quantity selected by retailer i when the wholesale price is w. We denote $\Pi_T = \Pi_0 + \sum_i \Pi_i$ the total supply chain profits.

[1]While retail prices cannot, in practice, take negative values, we omit this as an explicit constraint, and instead show that at the equilibrium and at the centralized optimum, prices are positive (see Appendix 4.2, Proof of Prop. 2.2.1).

Proposition 2.2.1. *The equilibrium supply chain structure, profits, quantities, and prices are given in Table Table I.*

Domain	$\rho \leq \rho_{min}$	$\rho_{min} < \rho < \rho_{max}$	$\rho \geq \rho_{max}$
Structure	Wholesale Supplier	Two-Channel	Monopoly Retailer
q_0	N/A	$\frac{\delta\nu_0-\gamma N\nu}{2(\beta_0\delta-\gamma^2 N)}$	$\frac{\nu_0}{2\beta_0}$
q_i	$\frac{\nu}{2\delta}$	$\frac{\beta_0\nu-\gamma\nu_0}{2(\beta_0\delta-\gamma^2 N)}$	N/A
p_0	N/A	$\frac{\alpha_0+c_0'}{2}$	$\frac{\alpha_0+c_0'}{2}$
p_i	$\frac{\Delta(\alpha+c')+2\beta\alpha}{2\delta}$	$(\frac{1}{2})(\alpha+c'+\frac{\beta(\beta_0\nu-\gamma\nu_0)}{\beta_0(2\beta-\gamma)+\gamma N(\beta_0-\gamma)})$	N/A
Π_0	$\frac{\nu^2 N}{4\delta}$	$\frac{\delta\nu_0^2+(\beta_0\nu-2\gamma\nu_0)N\nu}{4(\beta_0\delta-\gamma^2 N)}$	$\frac{\nu_0^2}{4\beta_0}$
Π_i	$\frac{\beta\nu^2}{4\delta^2}$	$\frac{\beta(\beta_0\nu-\gamma\nu_0)^2}{4(\beta_0\delta-\gamma^2 N)^2}$	N/A

TABLE I: EQUILIBRIUM QUANTITIES, PRICES, AND PROFITS FOR THE DECENTRALIZED SUPPLY CHAIN

The equilibrium supply chain structure, as shown in Table Table I, depends solely on parameter ρ. When $\rho \leq \frac{\gamma N}{\delta} = \rho_{min}$, the supplier's direct channel is weak compared with the indirect channel and thus the supplier exits the retail market and acts only as a wholesale supplier to its retailers. When $\rho \geq \frac{\beta_0}{\gamma} = \rho_{max}$, the supplier's direct channel is strong enough to make retailers exit the market, leaving the supplier as a monopoly retailer. Therefore, the two-channel case occurs if $\rho_{min} < \rho < \rho_{max}$. This indicates that a two-channel equilibrium with both channels in operation exists only when the quantity available at either a retailer or the supplier has a limited influence on the price at the others (we note that $\rho_{min} < 1 < \rho_{max}$, so the two-channel structure occurs at equilibrium when the supplier channel and the retailer channel are identical in maximum product margins, i.e. $\rho = 1$).

Increasing the value of N increases the retail market size which offers the supplier the potential to improve its indirect channel profit. However, when its direct channel is in operation as well as the indirect channel, the presence of more retailers also intensifies the competition between the indirect and the direct channel and can hurt the direct channel profit. In addition, the value of N affects the structure of the supply chain by determining the value of ρ_{min}. Changing the number of retailers thus has a non trivial overall effect on the supplier's total profit. We next investigate the effect of N on the equilibrium supply chain profits (see Figure 2a).

In the two-channel structure, it is straightforward to show from Table Table I that the supplier's equilibrium profit is monotonically non-decreasing in N. When there are no retailers ($N = 0$), the supplier acts as a monopoly retailer. As retailers enter the market, the supplier earns more in wholesale revenue than is lost in direct retail revenue due to the intensification of retail competition; thus, to the supplier, the expansion of the indirect retail market is worth the loss of direct retail market share. We note that this holds true even when the supplier's price is undercut (see Figure 2b), a finding that helps explain our motivating example of books sold by Random House.

As N becomes large, the behavior of the supplier depends on its margin position as compared to the other retailers. If $\rho \geq 1$, the supplier remains in the retail market regardless of the number of new entrants, and continues to benefit from increases in wholesale revenue as N grows. If

$\rho \leq 1$, additional entrants eventually force the supplier out of the retail market[1]. However, the supplier's profit continues to increase in N, as enough new wholesale revenue replaces the direct retail revenue no longer earned. Therefore, the supplier benefits from the presence of as many retailers as possible in the market, even if it is itself forced out of the retail market.

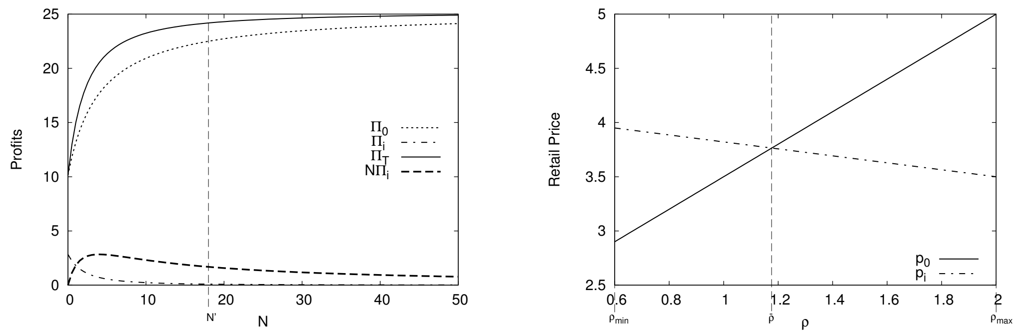

(a) Profits vs. N, where $\alpha_0 = 11$ and $\alpha = 12$ (thus $\nu_0 = 9$, $\nu = 10$, and $\rho = .9$).

(b) Prices vs. ρ, where $N = 3$ and $\alpha = 3$ ($\nu = 1$ and ν_0 varies from .6 to 2, causing ρ to vary from $\rho_{min} = .6$ to $\rho_{max} = 2$).

Figure 2: Profits vs. Number of Retailers and Price vs. ρ in Equilibrium, where $c_0 = c_A = c = 1, \beta_0 = 2$, $\beta = 1.5$, and $\gamma = 1$.

It is easy to show that the growth in total demand as N increases due to additional retailers capturing more consumers has a beneficial effect on the total supply chain profit, though the

[1] ρ_{min} increases with N thus as N increases, ρ becomes lower than ρ_{min} and the structure of the supply chain changes.

profit to each retailer is decreasing because of more intense competition. The combined retailer profit is unimodal, thus, the growth in the total profit beyond the maximum retailer profit accrues entirely as a benefit to the supplier.

2.2.2.2 The Centralized Case and the Price of Anarchy

As a benchmark for the decentralized equilibrium, we consider a centralized supply chain where a single central decision-maker chooses the prices and quantities for all firms, with the goal of maximizing the total system profit. This profit, $\widetilde{\Pi}_T$, is given by:

$$\widetilde{\Pi}_T = q_0(p_0 - c_0') + \sum_{i=1}^{N} q_i(p_i - c'). \tag{2.4}$$

Proposition 2.2.2. *The centralized optimal profits, quantities, and prices are given in Table Table II.*

Domain	$\rho \leq \rho_{min}^C$	$\rho_{min}^C < \rho < \rho_{max}$	$\rho \geq \rho_{max}$
Structure	Wholesale Supplier	Two-Channel	Monopoly Retailer
q_0	N/A	$\frac{\Delta\nu_0 - \gamma N\nu}{2(\beta_0\Delta - \gamma^2 N)}$	$\frac{\nu_0}{2\beta_0}$
q_i	$\frac{\nu}{2\Delta}$	$\frac{\beta_0\nu - \gamma\nu_0}{2(\beta_0\Delta - \gamma^2 N)}$	N/A
p_0	N/A	$\frac{\alpha_0 + c_0'}{2}$	$\frac{\alpha_0 + c_0'}{2}$
p_i	$\frac{\alpha + c'}{2}$	$\frac{\alpha + c'}{2}$	N/A
Π_T	$\frac{\nu^2 N}{4\Delta}$	$\frac{\Delta\nu_0^2 + \beta_0 N\nu^2 - 2\gamma N\nu_0\nu}{4(\beta_0\Delta - \gamma^2 N)}$	$\frac{\nu_0^2}{4\beta_0}$

TABLE II: OPTIMAL QUANTITIES, PRICES, AND PROFITS FOR THE CENTRALIZED SUPPLY CHAIN

As in the decentralized case, the centralized chain structure depends on ρ and the supplier acts as a monopoly retailer when $\rho \geq \rho_{max}$. However, the supplier exits the direct retail market when $\rho < \frac{\gamma N}{\Delta}$, denoted as ρ^C_{min}. The centralized supply chain thus has a two-channel structure if $\rho^C_{min} < \rho < \rho_{max}$. Note that ρ^C_{min} can be interpreted as a measure of the inverse of a retailer's market power: $1/\rho^C_{min} = 1+(1/N)(\beta/\gamma-1)$, and the higher N the more competitive the market is, while the closer γ is to β, the less differentiated retailers are.

We define the efficiency of the system, η, as the ratio of the total profit in the decentralized case and the total profit in the centralized case, a ratio frequently of interest in the supply chain literature (50), (30), (51). The closed form expressions for the efficiency are given in Table Table XI. The efficiency serves as a measure of supply chain performance. Literature on the "price of anarchy," which measures the efficiency lost to selfish behavior (typically, 1 - η: see (52), (53), or (54)), seeks to quantify this inefficiency and uses the price of anarchy as a motivation for coordination mechanisms. In what follows, we study characteristics of the supply chain efficiency and we show that the dual channel efficiency behaves differently than the single-channel efficiency.

Proposition 2.2.3. *The efficiency is not monotone in ρ, N, or γ.*

We summarize the monotonicity properties of efficiency in Table Table XI. We first consider the effect of ρ on the efficiency, illustrated in Figure 3a. In the range $0 \leq \rho \leq \rho_{min}$, η is constant in ρ, and takes a value between $\frac{3}{4}$, at $N = 1$, and 1, at $N = \infty$. This echoes results developed by previous authors for a single-channel supply chain that $\frac{3}{4}$ is the minimum efficiency of a system in which a single supplier interacts with N non-differentiated symmetric retailers in competition ((52), (53), (55), (56)).

When ρ increases beyond ρ_{min}, the decentralized system moves from the wholesale supplier structure to the two-channel structure, and the efficiency is unimodal in this range, reaching a minimum between ρ_{min} and ρ^C_{min}. In particular, we observe that the efficiency of the two-channel supply chain falls below that of a single-channel supply chain when $\rho_{min} < \rho < \rho^C_{min}$. This is surprising, as in a two-channel supply chain, some products are sold through the direct channel, *avoiding* double marginalization. Further, the existence of a direct channel increases the number of retail outlets, and therefore intensifies retail competition, which generally leads to higher efficiency. However, the efficiency in the two-channel supply chain is lower than that of a single-channel supply chain only in the range of ρ for which the optimal centralized supply chain structure differs from the equilibrium decentralized supply chain structure. Therefore, the supplier's sub-optimal presence in the retail market lowers efficiency: products sold through the direct channel yield a lower maximum product margin than the indirect channel.

Figures 3b and 3c illustrate the non-monotonicity of efficiency with respect to both N and γ in the range[1] $\rho^C_{min} < \rho < \rho_{max}$. These results are unique to a dual channel supply chain. For a single-channel supply chain with a single supplier and multiple retailers, (56) show that efficiency is increasing in the number of retailers and decreasing in retailer differentiation (i.e. an

[1]In this numerical example, $\rho = 1$ so it is higher than both $\rho_{min} = \frac{N}{N+9}$ and $\rho^C_{min} = \frac{N}{N+4}$, and lower than $\rho_{max} = 1.5$, so the system has a two-channel structure in equilibrium independently of the number of retailers for both the decentralized and centralized cases.

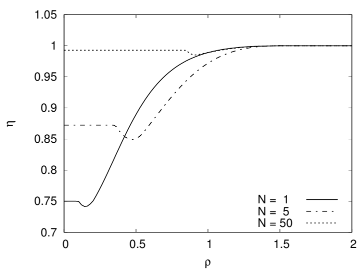

(a) Efficiency vs. ρ, where $\alpha = 12$, $c_0 = c_A = c = 1$, $\beta_0 = 1.5$, $\beta = 5$, and $\gamma = 1$ (α_0 varies from 2 to 22, causing ρ to vary from 0 to 2).

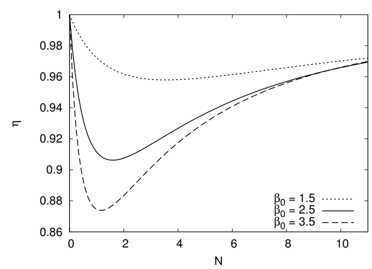

(b) Efficiency vs. N, where $\beta = 2.5$, and $\gamma = 1$, where $\alpha_0 = \alpha = 9$, $c_0 = c_A = c = 1$ (thus $\nu_0 = 7$, $\nu = 7$, and $\rho = 1$).

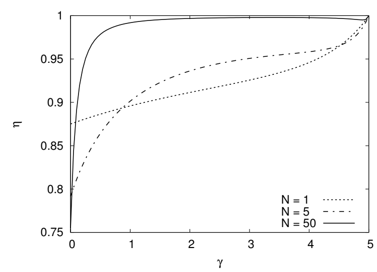

(c) Efficiency vs. γ, where $\beta_0 = \beta = 5$, $\alpha_0 = \alpha = 9$, $c_0 = c_A = c = 1$ (thus $\nu_0 = 7$, $\nu = 7$, and $\rho = 1$).

Figure 3: Efficiency vs. ρ, N, and γ.

increase in γ). In the dual channel supply chain, at $N = 0$ there is a single decision-maker and the efficiency equals 1. As retailers are added to the system, the efficiency initially decreases, as the retail quantity sold through the direct channel (not subject to double marginalization) decreases. When the number of retailers grows large, the double marginalization effect is outweighed by the retail competition effect; the retailers lose all market power and the efficiency approaches 1 (as in single-channel supply chains). The non-monotonicity of efficiency with respect to γ when N is large in a dual channel chain can be explained as follows: the immediate increase in efficiency as soon as differentiation is introduced is very large, reflecting the abrupt change from no competition ($\gamma = 0$) to intense competition ($\gamma > 0$ and large N), and the efficiency becomes very close to 1, indicating that the supplier's dual role gives it almost all of the decision-making power. As γ increases further, the supplier's power is diminished by the decreased differentiation between the two channels, causing a slight drop in efficiency. Finally, as γ approaches β, the retailers' quantity approaches zero, so the supply chain closely mimics that with a single decision-maker, and efficiency approaches 1.

The two-channel supply chain therefore contains several sources of inefficiency not found in the single-channel supply chain, but that can be remedied:

1. When $\rho_{min} < \rho < \rho_{min}^C$, the inefficiency can be reduced by removing the supplier from the retail market.
2. When β_0 is large enough to cause a dip in efficiency for low values of N, the efficiency can be improved by limiting the number of independent retailers.
3. When N is large enough that efficiency is not monotonically increasing in γ, efficiency can be improved by *decreasing* the intensity of competition, through means such as the removal of links to third-party retailers from Random House's website in our introductory example.

Each of these strategies for increasing efficiency is unique to the two-channel supply chain. We caution, however, that these tactics are applicable in a limited number of situations, thus we turn our attention to coordinating contracts as a potentially more robust method of improving efficiency.

2.2.3 Coordination Mechanisms With Symmetric Retailers

Based on the above discussion, we consider ways to improve the efficiency of a system in which the supplier opts to act as both a supplier to other retailers and a direct retailer itself. We exclude from our discussion the range $\rho \leq \rho_{min}$, in which the supplier is acting only as a supplier to other retailers in both the centralized and decentralized cases, because this is equivalent to single-channel supply chains with one supplier and many retailers, on the coordination of which much has been written (22). Likewise, when $\rho \geq \rho_{max}$, we have the trivial case in which both the centralized and decentralized systems are monopolies and efficiency is one. Therefore, we focus on the case when $\rho_{min} < \rho < \rho_{max}$, in which the decentralized supply chain has a two-channel structure in equilibrium, and the centralized supply chain's optimal structure is either two-channel (for $\rho^C_{min} \leq \rho < \rho_{max}$) or single-channel with the supplier acting only as a wholesale supplier (for $\rho_{min} < \rho < \rho^C_{min}$).

For the contracts we discuss herein, we assume that the contract parameters governing compensation from the retailer to the supplier are negotiated before the start of the game (similarly to (31)), then the supplier chooses its retail quantity, and finally the retailers choose their retail quantities. Further, due to symmetry among retailers, we assume that the same contract is offered to all retailers.

2.2.3.1 Applying Common Contracts to the Two-Channel Supply Chain

We find that no fixed per-unit wholesale price coordinates the two-channel supply chain. This result agrees with earlier single-channel literature ((22), (57)) demonstrating that a fixed wholesale price cannot both coordinate a supply chain with a single supplier and a single retailer and preserve a margin for each. However, a notable difference between the one-channel and two-channel supply chains is that the former may be coordinated by eliminating the margin of one of the firms (58). Thus, the trivial cases in which the wholesale price equals the marginal cost or the retail price both coordinate the one-channel supply chain. By contrast, both of these extreme values of the wholesale price ($w = c_A$ and $w = p_i$) fail to coordinate a two-channel supply chain.

Similarly, an all-units quantity discount, in which the supplier offers a wholesale price discount on all units purchased if a retailer orders more than a specified breakpoint (24) also fails to coordinate the two-channel supply chain, though it can achieve perfect coordination in a single-channel supply chain. While the retailers can be incentivized to order their centralized optimal quantity by setting the breakpoint equal to this quantity, the contract does nothing to motivate the supplier to choose *its* quantity at the centralized optimal quantity. Thus, the all-units quantity discount is insufficient in the two-channel case.

Revenue sharing contracts ((31)) and linear price discount sharing (PDS) contracts ((29)) fare only slightly better. In a revenue sharing contract, the retailers get a discount on the wholesale price, in exchange for returning a percentage of their revenue to the supplier. In a linear PDS contract, the wholesale price is discounted by an amount linearly proportional to the discount the retailers offer on their retail prices. These two classes of contracts are equivalent, in that they generate the same prices, quantities, and profits for all firms. When $\rho_{min} < \rho < \rho^C_{min}$, there exists a revenue sharing contract or linear PDS contract that perfectly coordinates the supply chain. Its parameters and resulting profits are given in Table Table XII. However, when

$\rho^C_{min} \le \rho < \rho_{max}$, the supply chain cannot be perfectly coordinated with a revenue sharing contract or linear PDS contract. Therefore, these contracts only coordinate the two-channel supply chain when the supplier is forced out of the retail market, and fail to coordinate the supply chain when both channels are in operation.

2.2.3.2 Linear Quantity Discount Contract

When both channels are active ($\rho^C_{min} \le \rho < \rho_{max}$), the challenge of choosing a contract that induces the supply chain optimal retail quantities in both channels is not easily overcome. A contract that makes the supplier's quantity equal to the centralized quantity by stripping the retailers of a margin cannot penalize the retailers for deviating from the optimal quantity. On the other hand, a contract that gives the retailers an incentive to order their centralized optimal quantities is insufficient because it fails to force the supplier's quantity to match its centralized quantity. We now study a linear quantity discount contract, and show that it successfully coordinates the supply chain for any value of ρ between ρ_{min} and ρ_{max}.

In the linear quantity discount contract, the per-unit discount is a linear function of the number of units purchased by a retailer (27). This contract involves two parameters: w^o, the maximum wholesale price, and s, the discount per unit, resulting in a wholesale price per unit, $w = w^o - sq$. Parameters w^o and s are fixed, and the supplier chooses its retail quantity before the retailers choose theirs. Under this scheme, the profits are given by

$$\widehat{\Pi}_0 = q_0(p_0 - c_0 - c_A) + \sum_{i=1}^{N} q_i((w^o - sq_i) - c_A) \tag{2.5}$$

$$\widehat{\Pi}_i = q_i(p - c - (w^o - sq_i)). \tag{2.6}$$

Theorem 2.2.4. *When $\rho_{min} < \rho < \rho_{min}^{C}$, there exists a linear quantity discount contract that perfectly coordinates the supply chain. Its parameters and resulting profits are given in Table Table XII.*

When $\rho_{min}^{C} \leq \rho < \rho_{max}$, the linear quantity discount contract with $s = \beta - \epsilon$ and $w^o = \frac{\beta_0\delta(\alpha-c+c_A)-\gamma(\beta\nu_0+\gamma N(\alpha-c+c_A))}{2(\beta_0\delta-\gamma^2 N)}$ perfectly coordinates the supply chain when ϵ approaches zero, and this results in limiting profits of $\widehat{\Pi}_0 = \widetilde{\Pi}_T$ and $\widehat{\Pi}_i = 0$.

In the range $\rho_{min} < \rho < \rho_{min}^{C}$, both the maximum wholesale price, w^o, and the discount per unit, s, are increasing in ρ, so the greater ρ, the greater the incentive to buy a large quantity wholesale from the supplier. This is needed to offset the fact that the supplier becomes more competitive in the retail market for larger values of ρ, and the equilibrium quantity of the retailers therefore decreases.

In the range $\rho_{min}^{C} \leq \rho < \rho_{max}$, the wholesale price is uniformly equal to $\frac{\alpha-c+c_A}{2}$. The optimal price for the retailers is $\frac{\alpha+c+c_A}{2}$, and so they are each selling at a price exactly equal to marginal cost. However, unlike in the revenue sharing and linear PDS contracts, this zero margin is not pre-supposed by the contract itself, but is rather a result of the quantity decision. Therefore, the equilibrium quantity for the retailers is equal to the optimal centralized quantity, and a deviation from this quantity would have a negative effect on their profits. Thus, this contract achieves perfect coordination in this range, though at a loss to the retailers, as compared to the decentralized case. This is a significant limitation of the contract, as the retailers have no incentive to participate in a supply chain that offers them no profit. However, this may be remedied with a transfer payment from the supplier, a mechanism previously proposed by a variety of authors ((59), (60), (61)). Because the total supply chain profit is larger with the contract than without, a linear quantity discount contract in combination with a fixed transfer payment can perfectly coordinate the two-channel supply chain at a Pareto improvement to all firms.

With such a two-part contract, it is not guaranteed that the supplier's profit is monotonically increasing in the number of retailers. Though the total profit increases in the number of retailers, the exact profit allocation depends on the supplier's negotiating strength relative to that of the retailers (62). In practice, therefore, a supplier who is able to offer retailers a small transfer payment would be better off with more retailers in the supply chain. By contrast, a supplier with less negotiating power may be able to use the threat of additional retailers to negotiate a lower transfer payment to the existing retailers. In either case, however, the supplier must be more strategic about offering its product to additional retailers than in the non-coordinated decentralized game.

2.2.4 Asymmetry of Retailers

In Sections 2.2.2 and 2.2.3, we assumed that all retailers have similar cost and market positions, i.e. $\alpha_i, \beta_i, \gamma_{ij}, c_i, c_i'$, and ν_i are symmetric among retailers. In this section, we relax this assumption and we examine the effect of asymmetry in these parameters by running a series of numerical experiments to gain insight into the supply chain's behavior when the retailers are asymmetric and differentiated. We find numerically the decentralized equilibrium by formulating the retailers' problem as a linear complementarity problem (LCP) (63) and solving the supplier's decisions as a nonlinear optimization problem subject to the constraint that the retailers' decisions are the solution of this LCP.

2.2.4.1 Effect of the Number of Retailers on Profits and Efficiency

In this section, we confirm that key results from Section 2.2.2 hold when the retailers are asymmetric, namely, that the supplier's profit and the total profit are increasing, while each retailer's profit is decreasing, in the number of retailers. Further, as in the symmetric case, the efficiency is not monotone, but rather, initially declines with the addition of retailers and then recovers as N becomes large. We proceed as follows: we fix the supplier's parameters (α_0, β_0, γ_{0i}, c_0, and c_A), and an upper and lower bound for each of the retailers' parameters. We then assign to each retailer random parameters drawn from a uniform distribution between these bounds, maintaining the assumptions that $\beta_i > \gamma_{ij}$ and $\gamma_{ij} = \gamma_{ji}$ $\forall i, j$ and B is symmetric. The trial is repeated 250 times and we average the relevant output values (e.g. profits, efficiency, etc.). We repeat the process in this manner, varying the supplier's parameters and the bounds on the retailer's parameters, and conclusions are drawn from the complete set of numerical results.

Because the retailers can no longer be assumed to make identical choices about participating or not participating in the retail market (and a central planner may not make the same decision about all retailers), throughout this section we use N to refer to the total number of retailers considered and n to refer to the total number that participate in the retail market, i.e., have a retail quantity greater than zero. We call the latter "active retailers".

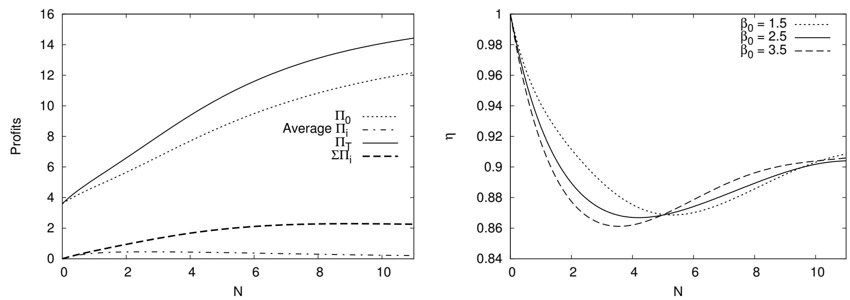

(a) Average Profits vs. N, where $\beta_0 =$ 2.5.

(b) Average Efficiency vs.N.

Figure 4: Average Profits and Efficiency vs. Number of Retailers, where $\alpha_0 = 8, c_0 = c_A = c_i = 1$, and $\alpha_i, \beta_i, i = 1, \cdots, N$, and $\gamma_{ij}, i = 0, \cdots, N, j = 0, \cdots, N$ are drawn from uniform distributions on $[7.5, 10.5], [1, 4]$, and $[.5, 1.5]$, respectively (thus $\nu_0 = 6$ and $\nu_i, i = 1, \cdots, N$, is drawn from a uniform distribution on $[5.5, 8.5]$).

We find that the results obtained numerically in the asymmetric case are consistent with those obtained analytically in the symmetric case: as shown in Figure 4a, the supplier's average profit and the average total profit are increasing in N, while each retailer's average profit is decreasing in N. There may be one or more retailers who leave the retail market, i.e., in general $n < N$, and we find that the larger N, the greater the difference in fraction of active retailers between the centralized and decentralized settings; see Figure 5a. For the numerical example illustrated in this figure, more than 99% of the retailers remain in the retail market in the decentralized case on average for values of N as large as 10, while the centralized case shows a steep drop-off in n, with only 63% of retailers remaining active when $N = 10$. As shown in Figure 5b, the supplier may also choose to leave the retail market (as occurs when the supply chain contains a retailer or retailers with large maximum margins). However, regardless of whether the supplier and/or some retailers are pushed out of the retail market, on average, the supplier still earns more profit when more retailers are present (as shown in Figure 4a), and thus we confirm that it is still in the supplier's best interest to include as many asymmetric retailers as possible in the supply chain.

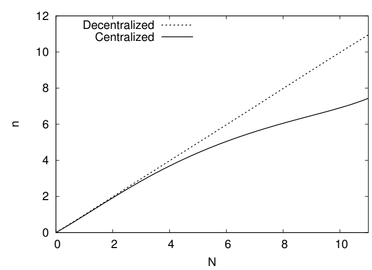

(a) Average Number of Active Retailers vs. N.

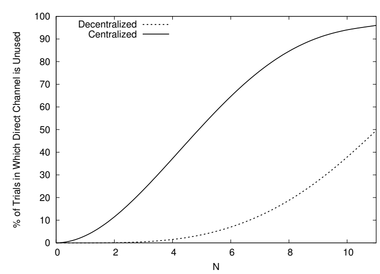

(b) Percentage of Trials in Which Direct Channel is Inactive vs. N.

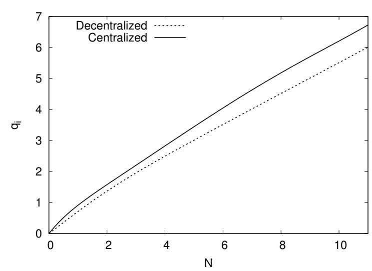

(c) Average Quantity per Active Retailer ($\frac{\sum_{i=1}^{N} q_i}{n}$) vs. N.

Figure 5: Behavior of Firms vs. Number of Retailers, where $\alpha_0 = 8, c_0 = c_A = c_i = 1, \beta_0 = 2.5$, and $\alpha_i, \beta_i, i = 1, \cdots, N$, and $\gamma_{ij}, i = 0, \cdots, N, j = 0, \cdots, N$ are drawn from uniform distributions on $[7.5, 10.5], [1, 4]$, and $[.5, 1.5]$, respectively (thus $\nu_0 = 6$ and $\nu_i, i = 1, \cdots, N$, is drawn from a uniform distribution on $[5.5, 8.5]$).

We also find that, as in the symmetric case, the efficiency is not monotonic in N, but rather drops off upon the initial entry of retailers into the market, and then recovers as additional retailers push the supply chain closer to perfect competition, as shown in Figure 4b. As compared to the symmetric case (Figure 3b), the asymmetric efficiency is lower except for very small values of N. This can be explained by the fact that when retailers are asymmetric a new effect influences efficiency. When retailers are symmetric, either they all participate in the retail market, or none does, both at equilibrium and at the centralized solution. When they are asymmetric, the centralized planner only selects the “strongest” retailers to be active, while, as shown in Figure 5a, unless N is very small, at equilibrium many more retailers decide to be active. This discrepancy between the centralized and decentralized settings lowers efficiency. An implication of this observation is that a contract that would coordinate the asymmetric supply chain must induce only some of the retailers to leave the retail market, which was not the case for a coordinating contract in a symmetric supply chain. We further investigate coordinating contracts under asymmetry of retailers in Section 2.2.4.2.

2.2.4.2 Effect of Asymmetry on the Existence of a Coordinating Contract

In this section, we investigate the effect of a linear quantity discount contract on the supply chain with asymmetric retailers. We find that, contrary to the symmetric case, no such contract achieves perfect coordination. However, we demonstrate that this type of contract significantly improves efficiency and thus may have practical use. We assume that retailer i is offered a linear quantity discount contract characterized by parameters $s_i, w_i^o,\ i = 1, \cdots, N$ and we search numerically for the set of contract parameters that maximizes the total supply chain profit.

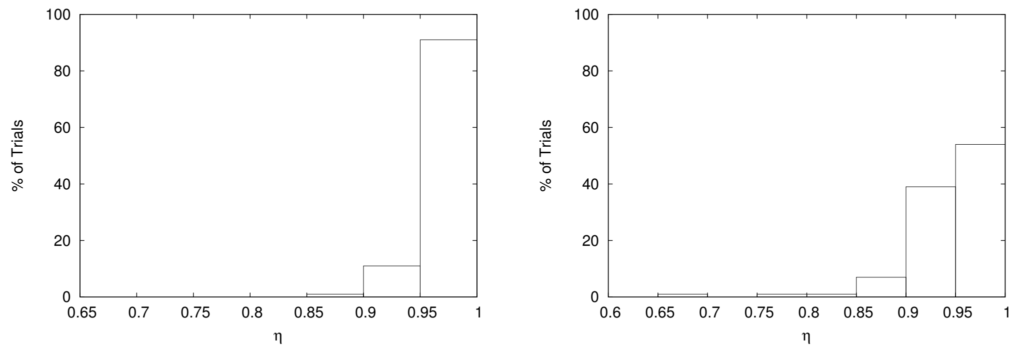

(a) With Fixed Per-unit Wholesale Price.

(b) With Linear Quantity Discount Contracts.

Figure 6: Distribution of Efficiency of the Asymmetric Supply Chain in 100 Randomized Trials, where N is drawn from a discrete uniform distribution from 2 to 6, $c_0 = c_A = 1$, $\beta_0 = 2.5$, and α_0, $\alpha_i, i = 0, \cdots, N$, c_i, $\beta_i, i = 0, \cdots, N$, and $\gamma_{ij}, i = 0, \cdots, N, j = 0, \cdots, N$ drawn from uniform distributions on $[12, 14]$, $[8, 10]$, $[1.5, 2.5]$, $[1.5, 3.5]$, and $[.5, 1.5]$, respectively.

We can find contract parameters for each retailer that induce retailer quantities equal to the centralized solution; however, the asymmetry of cross-sensitivity (the γ_{ij}'s) makes it impossible to induce the supplier to choose a quantity equal to its centralized optimal quantity, and thus no contract perfectly coordinates the asymmetric chain.

We find that, without achieving perfect coordination, linear quantity discount contracts can still significantly improve the efficiency, and therefore the total profit, of the supply chain. We solved 100 instances of the problem of finding the best possible contract terms with N drawn from a discrete uniform distribution from 2 to 6, and $\nu_i, \beta_i, \gamma_{ij}, i = 0, \cdots, N, j = 0, \cdots, N$ drawn from uniform distributions on $[8, 18]$, $[1.5, 3.5]$, and $[.5, 1.5]$, with a restriction that the matrix B be symmetric positive definite. As shown in Figure 6, for 77% of the trials, a contract was found that resulted in efficiency of 95% or better, while for 48% of the trials, the contract resulted in efficiency of 99% or better. By contrast, the efficiency with a fixed per-unit wholesale price was above 95% for only 50% of the trials, and above 99% in only 15% of the trials. In 62% of trials, the efficiency improved by at least 1% when a contract was applied, and in 30% of trials, it improved by at least 5%.

Nevertheless, the fact that the contract proposed here does not perfectly coordinate the decisions with the centralized setting is a clear limitation of this contract. Further research on contracts that may fully coordinate asymmetric dual supply chains is necessary to improve upon this result.

2.3 The Two-Channel Supply Chain with Capacity Constraints

We next turn our attention to the two-channel supply chain in which capacity is constrained.

2.3.1 The Model

We consider a single risk-neutral supplier that exclusively supplies a single product to a risk-neutral retailer. In addition, the supplier also operates a direct retail channel, thus we refer to him as the supplier-retailer throughout. These two firms competing in price form an oligopoly over the end market. We let the subscript s represent the supplier, while subscript r represents the retailer, and by common convention, we use female pronouns for the supplier and male for the retailer. The supplier does not initially know the exact number of units available, K, but knows that it is equal to either L or U, such that $L < U$.

We assume that the demand at each location is convex and linearly decreasing in price, as commonly found in the literature (64). We therefore let $D_s(p) = a_s - bp$ and $D_r(p) = a_r - bp$ represent the retail demand at the supplier and at the retailer, respectively, such that $a_s, a_r, b > 0$. Further, we assume that the supplier's potential market is larger than that of the retailer, i.e. $a_s \geq a_r$, and that $a_r \geq \frac{a_s+bc}{2}$. The latter condition is reasonable in that a monopoly supplier would be unlikely to engage with a retailer with a very small market potential, and we use it here to greatly simplify our results and allow us to better focus on capacity as the parameter of most interest. Demand that is not fulfilled at a retail location is lost; the selling season is too short to allow for customer search. There is no penalty for lost demand, except the unrealized revenue. There is a constant production cost of c per unit.

The game is played in four stages:

1. The supplier chooses a retail price, p.
2. The retailer chooses his quantity, q, and the maximum wholesale price he is willing to pay, w.
3. The supplier chooses whether or not to sell to the retailer. Her decision is either "yes" or "no," i.e. if she chooses to sell, it is the quantity and wholesale price dictated by the retailer.
4. The capacity, K, is revealed.

2.3.2 The Game with Known Capacity

We first assume that the capacity, K, is known, and we solve the game by backwards induction beginning with stage 3. We assume the supplier chooses to sell to the retailer if her profit is greater than or equal to her profit when she does not sell to the retailer. If she sells to the retailer, her profit is

$$\Pi_s = \begin{cases} K(p-c) - q(p-w) & K - q \leq D_s(p), \\ D_s(p)(p-c) + q(w-c) & K - q > D_s(p,) \end{cases},$$

and if she does not sell to the retailer, her profit is

$$\Pi_s = \begin{cases} K(p-c) & K \leq D_s(p), \\ D_s(p)(p-c) & K > D_s(p), \end{cases}.$$

When capacity is known, we assume that the supplier will choose the channel strategy (sell to retailer or don't sell to retailer) that yields the higher profit.

In phase 2, the retailer's profit is

$$\Pi_r = min[q, D_r(p)]p - qw. \tag{2.7}$$

When demand is deterministic, the retailer will order $q = D_r(p)$, so we can write the retailer's profit as

$$\Pi_r = D_r(p)(p - w)$$

Further, the retailer's profit is linearly decreasing in w, so the retailer will choose the smallest w that causes the supplier to sell to him, i.e. the w that makes the supplier's profit equal in the one-channel and two-channel cases, plus some ϵ value that is mathematically insignificant, but gives the supplier a slight incentive to sell to the retailer.

There are three possible cases for the relationship of capacity and demand:

1. $D_s(p) < K - D_r(p)$
2. $K - D_r(p) \leq D_s(p) < K$
3. $D_s(p) \leq K$

Case 1:

$$D_s(p)(p - c) + D_r(p)(w - c) = D_s(p)(p - c)$$

$$w = c.$$

In this case, the wholesale price is c, meaning that the supplier has excess capacity, and should be willing to sell to the retailer at any price that covers the cost of production.

Case 2:

$$\begin{aligned} K(p-c) - D_r(p)(p-w) &= D_s(p)(p-c) \\ w &= \frac{p(D_s + D_r(p) - K) + c(K - D_s(p))}{D_r(p)} \\ w &= \frac{p(a_r + a_s - K - 2bp) + c(K - a_s + bp)}{ar - bp}. \end{aligned}$$

Case 3:

In this case, the only wholesale price that would make $K(p-c) - D_r(p)(p-w) = K(p-c)$ is $w = p$, which would obviously allow the retailer no profit, and thus we conclude that if the capacity is less than demand at the supplier, nothing is sold through the retailer.

We then turn our attention to the supplier's profit in Phase 1. Because, in Cases 1 and 2, the wholesale price has been calculated to make the supplier's profit equivalent to that in the wholesale case, the supplier's profits are equal under both, and thus only dependent on the relationship between $D_s(p)$ and K.

$$\Pi_s = \begin{cases} K(p-c) & K \leq D_s(p) \\ D_s(p)(p-c) & D_s(p) < K \end{cases}.$$

When $K \leq D_s(p)$, the profit is linearly increasing in p, so it is maximized at the endpoint, $p^* = \frac{a_s - K}{b}$ for a profit of $\frac{K(a_s - K)}{b} - Kc$. When $K > D_s(p)$, the profit function is quadratic,

Domain	p	w	D_s	D_r	Π_s	Π_r
$a_r - bc \leq K$	$\frac{a_s+bc}{2b}$	c	$\frac{a_s-bc}{2}$	$a_r - \frac{a_s+bc}{2}$	$\frac{(a_s-bc)^2}{4b}$	$\frac{(a_s-bc)(2ar-as-bc)}{4b}$
$\frac{a_s-bc}{2} \leq K < a_r - bc$	$\frac{a_s+bc}{2b}$	$\frac{\theta}{b(a_s-2a_r+bc)}$	$\frac{a_s-bc}{2}$	$a_r - \frac{a_s+bc}{2}$	$\frac{(a_s-bc)^2}{4b}$	$\frac{(a_s-bc)(2K-a_s-bc)}{4b}$
$K \leq \frac{a_s-bc}{2}$	$\frac{a_s-K}{b}$	N/A	K	N/A	$\frac{K(a_s-K)}{b} - Kc$	N/A

$$\kappa = a_s - a_r + K$$
$$\theta = a_s(K - a_r) + bc(2a_s - a_r - K)$$

TABLE III: PRICES, MARGINS, QUANTITIES, AND PROFITS AT EQUILIBRIUM WHEN CAPACITY IS KNOWN

and maximized at $p^* = \frac{a_s-bc}{2b}$, for a profit of $\frac{(a_s-bc)^2}{4b}$. These two are equal when $K = \frac{a_s-bc}{2}$, thus, when $K \leq \frac{a_s-bc}{2}$, the supplier will choose the former, and when $K > \frac{a_s-bc}{2}$, the latter.

However, we see that when $\frac{a_s-bc}{2} > a_r$, the retailer's quantity is not positive if $p = \frac{a_s+bc}{2}$. We therefore revisit the retailer's decision in Case 2. The smallest possible value of w that motivates the supplier to sell through the retailer is $w = \frac{p(a_r+a_s-K-2bp)+c(K-a_s+bp)}{ar-bp}$. When $K > \frac{a_s-bc}{2}$, $w > \frac{(a_s-bc)(p-c)+2p(a_r+bc-2bp)}{2(a_r-bp)}$. If we add the condition that $a_s > 2a_r - bc$ (required for q_r to be positive), we have $w > 2p - c$. As we know that $p > c$, we can conclude that the smallest possible wholesale price that would guarantee the supplier equal profits in the one- and two-channel cases must be larger than the retail price, and the retailer's profit would therefore be negative. We can therefore conclude that nothing will be sold through the retailer when $a_r \leq \frac{a_s-bc}{2}$.

Proposition 2.3.1. *The equilibrium prices, profits, and margins when capacity is known are given in Table Table III.*

We see that when $K \geq \frac{a_s - bc}{2}$, the supplier's profit is $\frac{(a_s-bc)^2}{4b}$, and when $K < \frac{a_s-bc}{2}$, the supplier's profit is $\frac{(a_s-K)K}{b} - Kc$. The qualitative result is intuitive: when capacity is large compared to size of the supplier's potential market, the capacity is ultimately not a determinant of the supplier's profit, and when capacity is instead relatively small, it limits the supplier's profit.

Less intuitive, however, are the conditions under which the supplier chooses to sell to the retailer. As we see comparing rows 2 and 3 with row 4, the choice to sell to the retailer may depend on the relative demand parameters of the supplier and retailer, so as long as K is large, the supplier sells to the retailer when the retailer's demand position is at least moderately strong compared to that of the supplier (i.e. $a_r > \frac{a_s-bc}{2}$. If the retailer's demand positions is weak compared to that of the supplier, the supplier does not bother selling through the retailer. Even if the capacity is very large, the resulting decline in price from selling through the retailer harms the supplier's profit more than sales to the retailer benefit her.

2.3.3 The Game with Uncertain Capacity

Assume that capacity can be either high or low: K takes value H with probability α and L with probability $1-\alpha$, such that $L < H$. In the one-channel case, we have three possible scenarios:

1. $D_s(p) < L$

   $\Pi_s = D_s(p)(p-c)$

(a) If $L > \frac{a_s - bc}{2}$

$p^* = \frac{a_s + bc}{2b}$

$\Pi_s = \frac{(a_s - bc)^2}{4b}$

(b) If $L < \frac{a_s - bc}{2}$

$p^* = \frac{a_s - L}{b}$

$\Pi_s = \frac{L(a_s - bc - L)}{b}$

2. $L < D_s(p) < H$

$\Pi_s = \alpha[D_s(p)(p - c)] + (1 - \alpha)[L(p - c)]$

(a) If $L > \frac{\alpha(a_s - bc)}{1+\alpha}$

$p^* = \frac{a_s - L}{b}$

$\Pi_s = \frac{L(a_s - bc - L)}{b}$

(b) If $L < \frac{\alpha(a_s - bc)}{1+\alpha}$, $H > \frac{\alpha(a_s - bc + L) - L}{2\alpha}$

$p^* = \frac{\alpha(a_s - bc - L) + L}{2\alpha b}$

$\Pi_s = \frac{(\alpha(a_s - bc - L) + L)^2}{4\alpha b}$

(c) If $H < \frac{\alpha(a_s - bc + L) - L}{2\alpha}$

$p^* = \frac{a_s - H}{b}$

$\Pi_s = \frac{(a_s - bc - H)(\alpha(H - L) + L)}{b}$

3. $H < D_s(p)$

$\Pi_s = \alpha[H(p - c)] + (1 - \alpha)[L(p - c)]$

Domain	Price	Profit
$\frac{a_s-bc}{2} \leq L$	$\frac{a_s+bc}{2b}$	$\frac{(a_s-bc)^2}{4b}$
$\frac{\alpha(a_s-bc)}{1+\alpha} \leq L < \frac{a_s-bc}{2}$	$\frac{a_s-L}{b}$	$\frac{L(a_s-bc-L)}{b}$
$L < \frac{\alpha(a_s-bc)}{1+\alpha}, \frac{\alpha(a_s-bc+L)-L}{2\alpha} \leq H$	$\frac{\alpha(a_s-bc-L)+L}{2\alpha b}$	$\frac{(\alpha(a_s-bc-L)+L)^2}{4\alpha b}$
$H < \frac{\alpha(a_s-bc+L)-L}{2\alpha}$	$\frac{a_s-H}{b}$	$\frac{(a_s-bc-H)(\alpha(H-L)+L)}{b}$

TABLE IV: OPTIMAL PRICES AND PROFITS IN THE ONE CHANNEL CASE WITH UNCERTAIN CAPACITY

$p^* = \frac{a_s-H}{b}$

$\Pi_s = \frac{(a_s-bc-H)(\alpha(H-L)+L)}{b}$

Proposition 2.3.2. *Optimal prices and profits for the one channel supply chain with uncertain capacity are given in Table Table IV.*

When the supplier instead chooses to sell to the retailer as well as operating a direct channel, we must consider not only the relationship among L, $D_s(p)$, and H, but also q $(= D_r(p))$. We therefore see that each of the one-channel cases (indicated by a number) encompasses several two channel cases (indicated by a lowercase letter).

1. $D_s(p) \leq L$

   $\Pi_s = D_s(p)(p - c)$

   (a) $D_s(p) < L - q$

   $\Pi_s = D_s(p)(p - c) + q(w - c)$

   (b) $L - q < D_s(p) < H - q$

   $E[\Pi_s] = (1 - \alpha)[(L - q)(p - c)] + \alpha[D_s(p)(p - c)] + q(w - c)]$

(c) $L - q < H - q < D_s(p) < L$

$$E[\Pi_s] = (1 - \alpha)[(L - q)(p - c)] + \alpha[(H - q)(p - c)] + q(w - c)]$$

The condition $D_s(p) < L$ only results if the demand curve is such that the optimal price for maximizing revenue results in a demand at the supplier less than the lower possible capacity, L, i.e. $\frac{a_s - bc}{2} < L$. If the optimal price is large enough that demand at the supplier is smaller than what is leftover in the low capacity case once the retailer's demand has been fulfilled ($D_s(p) < L - q$), then 1a results, and it is clear that the supplier benefits from selling through the retailer, in addition to selling through the direct channel, at any wholesale price greater than the unit cost. This occurs when $a_r - bc < L$.

If the optimal price results in demand at the supplier such that $L - q < D_s(p) < L$ (1b or 1c), the supplier gives up some quantity of sales (the difference between $D_s(p)$ and $L - q$), in return for an increase in revenue gained by selling wholesale through the retailer. This occurs when $\frac{a_s - bc}{2} < L < a_r - bc$. The supplier is concerned only that the retailer may require a quantity large enough that the supplier is not able to meet her own demand, and needs an increased wholesale price to make up the lost revenue from lower direct sales. Therefore, a retailer instead wishing to purchase at a lower cost could instead propose an agreement in which its quantity is limited, i.e. $q < D_r(p)$ so that the condition $D_s(p) < L - q$ may be maintained. Interestingly, this operates in opposition to the all-units quantity discount contract commonly found in the literature (citation here), in which a retailer receives a lower price for purchasing a larger quantity.

In the case where $D_s(p) < H - q$ (1b), equivalent to $\frac{a_s - bc}{2} < L < a_r - bc < H$, the supplier only gives up retail revenue if capacity turns out to be low, but can fulfill all direct channel demand if capacity is high. Therefore, the wholesale price must be higher than the cost (see value given in Table Table XIII) for the supplier to be better off selling through the retailer than selling only through the direct channel. By contrast, if $H - q < D_s(p)$ (1c), equivalent to $\frac{a_s - bc}{2} < L < H < a_r - bc$, the supplier cannot fulfill direct channel demand if it sells through the retailer, regardless of whether capacity is low or high. An even higher wholesale price is therefore required to provide the supplier with an expected profit equivalent to that in the one-channel case.

2. $L < D_s(p) < H$

   $E[\Pi_s] = \alpha[D_s(p)(p - c)] + (1 - \alpha)[L(p - c)]$

   (a) $L - q < L < D_s(p) < H - q < H$

       $E[\Pi_s] = \alpha[D_s(p)(p - c)] + (1 - \alpha)[(L - q)(p - c)] + q(w - c)$

   (b) $H - q < D_s(p) < H$

       $E[\Pi_s] = \alpha[(H - q)(p - c)] + (1 - \alpha)[(L - q)(p - c)] + q(w - c)$

The one-channel case $L < D_s(p) < H$ can be divided into two two-channel cases, depending on the relationship between $D_s(p)$ and $H - q$. When $D_s(p) < H - q$, case 2a results, and, as in 1b, the supplier can meet all direct channel demand if capacity is high, but not if capacity is low. Therefore, the wholesale price must again account for the chance that the chance that the capacity will be low. This occurs when $L < \frac{\alpha(a_s + bc)}{1 + \alpha}$ and $H > \frac{\alpha(a_r + bc - L) + L}{\alpha}$.

If $H - q < D_s(p)$, case 2b results, and, as in 1c, the supplier cannot meet all direct channel demand for either value of capacity, high or low. The wholesale price therefore accounts for the certainty that direct channel sales will be lost in order to supply the indirect channel. This occurs when $L < \frac{\alpha(a_s+bc)}{1+\alpha}$ and $\frac{\alpha(a_s+bc+L)-L}{2\alpha} < H < \frac{\alpha(a_r+bc-L)+L}{\alpha}$.

However, in both 2a and 2b, a value of L between $\frac{\alpha(2a_r-a_s-bc)}{1-\alpha}$ and $\frac{\alpha(a_s-bc_L)-L}{1-\alpha}$ would result in a negative quantity in the retailer's channel, i.e. $D_r(p) < 0$. Therefore, the supplier must choose a retail price less than $\frac{a_r}{b}$. The closest the retailer can get to an optimal retail price is therefore $\frac{a_r}{b} - \epsilon$, where ϵ is some small quantity that ensures a positive demand in the retail channel, i.e. $D_r(p) > 0$. The wholesale price is then calculated accordingly to leave the supplier with a profit equal to that of the one-channel case. However, this wholesale price is then greater than the retail price, rendering the retailer unwilling to participate in the supply chain. Therefore, the supplier will not sell through the retailer in this case.

3. $H < D_s(p)$

   $E[\Pi_s] = \alpha[H(p-c)] + (1-\alpha)[L(p-c)]$

   (a) $H < D_s(p)$

      $E[\Pi_s] = \alpha[(H-q)(p-c)] + (1-\alpha)[(L-q)(p-c)] + q(w-c)$

In 3a, is clear that the supplier can only benefit from the two-channel case if the wholesale price is greater than the retail price. As it is assumed that a retailer will not agree to sell a product for which its unit net profit is negative, it can be concluded that the supplier does not sell through the retailer when $H < D_s(p)$, equivalent to $H < \frac{\alpha(a_s-bc+L)-L}{2\alpha}$.

Row	Domain	Channel Structure
1	$L < \frac{\alpha(2a_r - a_s - bc)}{1-\alpha}$, $\frac{\alpha(a_s - bc + L) - L}{2\alpha} < H$	Two-channel
2	$\frac{\alpha(2a_r - a_s - bc)}{2\alpha} \leq L < \frac{\alpha(a_s - bc)}{1-\alpha}$	One-channel
3	$H \leq \frac{\alpha(a_s - bc + L) - L}{2\alpha}$	One-channel
4	$\frac{\alpha(a_s - bc)}{1+\alpha} \leq L$	Two-channel

TABLE V: SUPPLY CHAIN STRUCTURE UNDER UNCERTAINTY WHEN CAPACITY IS UNCERTAIN

Theorem 2.3.3. *The equilibrium supply chain structure under capacity uncertainty is given in Table Table V. Full results for equilibrium prices, demands, and profits are given in Tables Table XIII, Table XIV, and Table XV in the appendix.*

As long as this conditions in either the first or third row are met, there exists a wholesale price at which the supplier can both allow the retailer a profit and improve her profit, as compared to only operating the direct channel. While the particular game described finds the wholesale price at which the supplier's profits are the same in the one- and two-channel supply chains, it is assumed that she will charge slightly higher than this minimum, and thus increase her total profit from both wholesale and retail sales.

The first possible scenario that leads to a two-channel equilibrium is $L < \frac{\alpha(2a_r - a_s - bc)}{1-\alpha}$ and $\frac{\alpha(a_s - bc + L) - L}{2\alpha} < H$. The lower boundary on H is dependent on the unit cost (c), the size of the supplier's market (a_s), the price sensitivity parameter (b), the likelihood of high capacity (α), and the value of low capacity. It is somewhat surprising that the size of the retailer's market (a_r) is not a factor, given that the product available in the direct channel ($L - q$ or $H - q$) is dependent upon both the realized capacity and the amount ordered by the retailer. However, because the wholesale price decision is set up to ensure that the supplier is expected to earn at least as much by selling through both channels as by selling only through the direct channel, a large retailer market can't possibly harm the supplier.

Note that this boundary can be written as $\frac{a_s-bc}{2} + \frac{L(\alpha-1)}{2\alpha}$ for clearer comparison to the corresponding boundary, $K > \frac{a_s-bc}{2}$, when capacity is known. As the quantity $\alpha - 1$ is always negative, we see that the minimum H that motivates the supplier to sell through the indirect channel under unknown capacity is lower than the minimum K under known capacity.

We can explain this difference by examining at the supplier's retail pricing decision. Given infinite capacity, the optimal retail price would be $p^{*inf} = \frac{a_s+bc}{2}$. When capacity is known, this is the optimal price as long as the resulting total demand is less than the capacity, K. When capacity is unknown, however, this is only the optimal price as long as the resulting total demand is less than the *lowest possible realization* of capacity, L. If a retail price of p^{*inf} would result in a total demand greater than L, the supplier maximizes expected profit by choosing a higher retail price, and thereby lowering demand, no matter how large the value of H. Therefore, the total demand at the chosen retail price is lower when capacity is unknown, and the value of H at which demand exceeds price is likewise lower than when capacity is known.

It is also of interest to note the boundary on H is decreasing in L. This is somewhat counterintuitive, as a larger value of L results in a retail price closer to p^{inf*}, and thus a larger total demand, and the expectation that H would need to be large to fulfill demand in both channels. However, this can also be explained by the fact that price is increasing in L when a retail price of p^{*inf} would result in a total demand greater than L. Therefore, an increase in L causes a decrease in demand; thus, H does not need to be as large for all demand to be fulfilled in two channels.

The second condition, $L < \frac{\alpha(2a_r - a_s - bc)}{1-\alpha}$, is required for demand to be positive in the indirect channel. This is because the price, when $L < D_s(p) < H$, is increasing in L, causing demand in both channels to decrease. Unsurprisingly, then, this boundary gets larger when the retailer's total market size gets larger, and gets smaller when the supplier's market size, the price sensitivity, or the per-unit cost increase. Further, the supplier's inability to benefit by selling through the indirect channel can only occur when there is a low chance that capacity will be high, specifically, when $\alpha < \frac{a_s - a_r}{a_r - bc}$. When α is smaller than this bound, there is a range of L that will induce a one-channel equilibrium on its own, regardless of the value of H.

The second possible scenario that leads to a two channel equilibrium is $L > \frac{\alpha(a_s - bc)}{1+\alpha}$. This condition indicates that L is larger than demand in the direct channel when the retail price is equal to the price that maximizes the supplier's profit, unconstrained by the relationship between capacity and demand ($p = p^{*inf}$), i.e. the supplier's demand will not exceed capacity, regardless of whether capacity is low or high. Given that $c < p$, the bound on L is therefore less than half of the supplier's total demand. It therefore seems realistic that many real-world scenarios will fall into this case.

Note, however, in the situation described in row 4, the quantity available in the direct channel, once the indirect demand has been fulfilled, may be less than the direct demand.

Therefore, there is still a risk that direct channel sales will be lost if capacity is low, while no such risk would exist if the supplier only sold direct to its market. Hence, the wholesale price is decreasing in α as long as $L - q < D_s(p)$, i.e. the greater the chance that sales will be missed in the direct channel, the larger the wholesale price paid by the retailer.

Returning to our motivating example, we can assume that, given Apple's inability to fulfill orders placed at their own website, $L - q < D_s(p)$, but it remains unknown whether $L < D_s(p) < H$ or $D_s(p) < L$. We can, however, speculate that the production of a brand new type of screen at a third-party supplier has a significant amount of risk, and therefore, α would take a low value. This results in a lower bound on L as given in row 4 that is only a small fraction of Apple's total direct channel demand, and it is therefore likely that the New iPad launch fell into this scenario.

2.3.4 Comparison Between Known and Unknown Capacity

We next investigate the effect that capacity uncertainty has on profits by comparing each domain in the known capacity case to corresponding domains in the unknown capacity case, and find that uncertainty does not always have a detrimental effect on total profit. We assume, for purposes of comparison, that the known capacity is equal to the expected capacity in the unknown case, i.e. $K = \alpha H + (1 - \alpha)L$. This allows us to isolate the effects of uncertainty from the effects of differences in the capacity values themselves. We refer to the various domains when capacity is uncertain by their row number in Table Table III, along with the letter C, and the various domains when capacity is unknown by their row number in Table Table XIV, along with the letter U.

We first compare 1C to 1U. The supplier's, retailer's, and total profits are equal as long as $K \geq L$, indicating, unsurprisingly, that as long as the lowest possible realization of capacity is at least as large as capacity in the known case, neither the highest realization of capacity nor its probability have any effect on profits. It is also possible that L and H could meet the conditions in the second row of the uncertain case, while K remains larger than $a_r - bc$, so we also compare 1C with 2U. Here, we find that the total profit loss due to uncertainty is $\frac{(1-\alpha)(a_s-bc)(L-(a_r-bc))}{2b} > 0$. Unsurprisingly, the profit loss is decreasing in α and L; as the probability of high capacity or the lowest value of capacity increase, so too do the profits come closer to those under known capacity. Similarly, profit loss due to uncertainty is increasing in both a_s and a_r; as the market size in either channel increases, so too does the quantity of missed sales in the direct channel if capacity is low, while all demand is fulfilled when capacity is known and satisfies $a_r - bc < K$.

Comparing the 2C to 2U, the supplier's profit is the same in both, but the retailer's expected profit, and therefore the expected total supply chain profits, is smaller when capacity is uncertain, by a quantity of $\frac{\alpha(H-(a_r-bc))(a_s-bc)}{2b} > 0$. Here, the profit loss is actually *decreasing* in α, a counterintuitive result illustrated in Figure 7a. This is because, as we've defined K in this section to be the expectation of capacity, we find that an increase in α leads to an increase in K, and therefore an increase in the quantity available in the direct channel when $\frac{a_s-bc}{2} < K < a_r - bc$. This increases total supply chain profits at a faster rate than in the case of unknown capacity, where an increase in α merely leads to an increasing probability of a larger quantity available in the direct channel. Also, the profit loss is decreasing in a_r, a

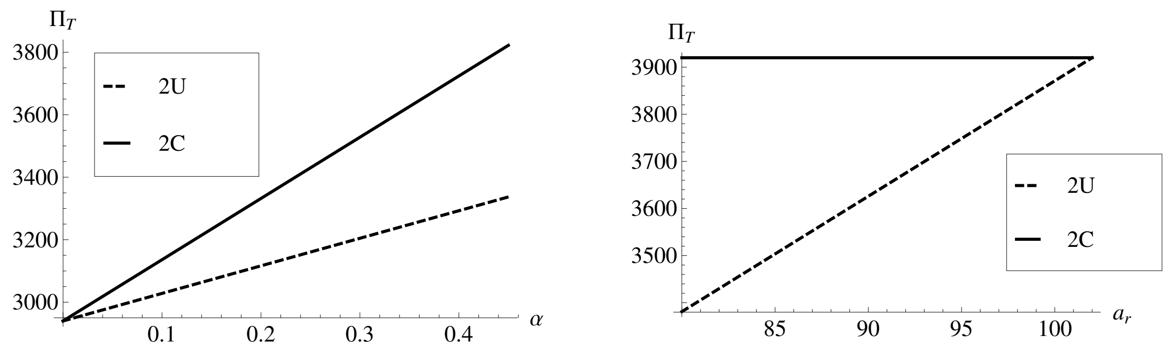

(a) Total profits vs. α for cases 2C and 2U, where $a_r = 80$.

(b) Total profits vs. a_r for cases 2C and 2U, where $\alpha = \frac{1}{2}$.

Figure 7: Profits vs. α and a_r, where $a_s = 100$, $b = 1$, $c = 2$, $L = 60$, $H = 100$, and $K = \alpha H + (1-\alpha)L$.

notable difference from 1C/2U, but easily explained as the total profit increases with a_r when capacity is uncertain, but remains constant when capacity is known, as shown in 7b.

When $\frac{a_s - bc}{2} < L < H < a_r - bc$, the resulting expected capacity must obey the condition $\frac{a_s - bc}{2} < K < a_r - bc$. Therefore, we compare 3U to 2C, and find the profits are exactly equal.

In 4U, H may be greater or less than $\frac{a_s - bc}{2}$; thus K may fall into 2C or 3C. The profit lost to uncertainty between 2C and 4U is $\frac{L(\alpha-1)-H\alpha)(2L-(a_s-bc))}{2b} < 0$, and between 3C and 4U is $\frac{\alpha(H-L)(L(1-\alpha)+H\alpha)}{b} < 0$, indicating that, in both comparisons, the total profit is actually *larger* under uncertainty than when capacity is known. In 4U/2C, the retailer's profit is substantially larger under uncertainty, while the supplier's profit is just slightly smaller, leading to a net effect of higher total profits under uncertainty, as shown in Figure 8. This is driven by an increase in retail price when capacity is uncertain, as, in domain 4U, the supplier chooses the retail price

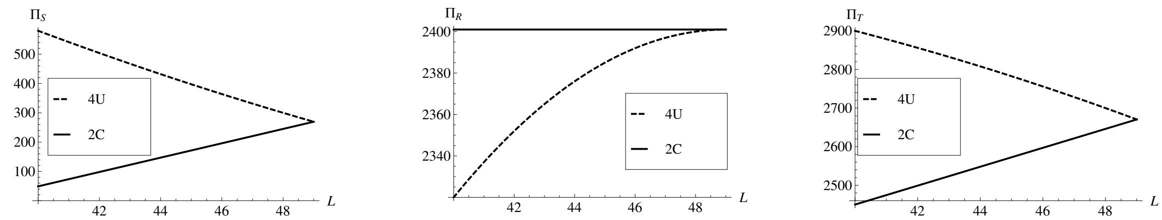

(a) Supplier profits vs. L for cases 2C and 4U.

(b) Retailer profits vs. L for cases 2C and 4U.

(c) Total profits vs. L for cases 2C and 4U.

Figure 8: Profits vs. L, where $a_s = 100$, $a_r = 80$, $b = 1$, $c = 2$, $\alpha = \frac{1}{2}$, $H = 60$, and $K = \alpha H + (1 - \alpha)L$.

so as to induce a demand exactly equal to L. The increased total profit under uncertainty between 4U and 3C has a similar explanation: when capacity is certain, the supplier chooses a retail price such that there is nothing sold through the indirect channel; under uncertainty, a lower retail price is chosen, allowing both channels to be active and thereby increasing total profit.

When capacity is uncertain and $L < D_s(p) < H$ (5U and 6U), profits are lower under uncertainty, and, in contrast to 1U, 2U, and 3U, the loss of total profit due to uncertainty is shared between the supplier and the retailer. Further, the larger α, the larger the percentage of the loss that is borne by the retailer.

The key findings from these comparisons are therefore summarized in the following two theorems:

Theorem 2.3.4. *The profit loss due to uncertainty affects each firm in the supply chain as follows:*

- *When* $\frac{a_s - bc}{2} < L$*, all supply chain profits lost to uncertainty are lost to the retailer.*

- *When* $\frac{\alpha(a_s - bc)}{1+\alpha} < L < \frac{a_s - bc}{2}$*, the supply chain profits are* greater *under uncertainty, and the benefits of uncertainty accrue entirely to the retailer.*

- *When* $L < \frac{\alpha(a_s - bc)}{1+\alpha}$*, the total supply chain profit loss due to uncertainty is shared between the supplier and the retailer.*

Theorem 2.3.5. *The profit loss due to uncertainty is related to the probability of high capacity as follows:*

- *When* $a_r - bc < L$*, the profit loss is decreasing in* α*.*

- *When* $\frac{a_s - bc}{2} < L < a_r - bc$*, the profit loss is increasing in* α*.*

- *When* $\frac{\alpha(a_s - bc)}{1+\alpha} < L < \frac{a_s - bc}{2}$*, there is no profit loss (in fact, there is a profit gain).*

- *When* $L < \frac{\alpha(a_s - bc}{1+\alpha}$*, the profit loss is decreasing in* α*.*

We have therefore demonstrated, by comparing the case with known capacity to the case with unknown capacity, that there is a considerable amount of unexpected behavior: in some cases, the total supply chain profits do not improve with the removal of uncertainty, and in others, the supplier's profits do not improve. Therefore, from the supplier's point of view, efforts to eliminate uncertainty may prove to have no benefit.

3. USE OF STOCHASTIC OPTIMIZATION TO LOWER DISTRIBUTION COSTS AT USG

3.1 Introduction

3.1.1 Background and Motivation

USG is a leading building supplies manufacturer in North America. For the Durock® product line on which this project focuses, several dozen items are produced in three locations in the United States, and may pass through any one of a network of warehouses before being delivered to customers throughout the United States and Canada.

Currently, the production and distribution decisions are made through the use of a large scale linear program (LP). The goal of the LP is to minimize the total delivered cost (production, freight, and handling) for all items in the planning network, while meeting the managerial constraints of capacity at each plant and demand at each customer location. The input parameters (production costs, freight costs, handling costs, and demand) are drawn from a single point in time.

The problem is extremely large, with upwards of 1500 customer locations being used in the current model. In addition, the number of items fluctuates between thirty and forty, and combined with the large number of warehouse locations and the choice between rail and truck modes of freight, over 90,000 decision variables are used in the current state.

Currently, the LP is solved using existing commercial software (INFOR Tactical Planner) to create an input file that is then optimized using IBM's CPLEX solver; the resulting output is then interpreted by the software and manually converted into sourcing rules that link the optimal plan to the order fulfillment system, ensuring that customer orders are placed at the correct plant or warehouse and that warehouse replenishment orders are placed at the correct originating plant.

Therefore, the current planning process does not account for the variability in inputs. Most notably, the customer demand at each location changes from month to month, and at an item/location level, this alone can result in almost 6,000 uncertain input parameters. Further, the cost for each item at each producing plant may vary, adding over 100 more uncertainties to the problem. In the current state, one month of inputs are taken and put into the plan as though they are deterministic parameters, a method that is unlikely to produce the optimal stochastic solution. We therefore seek a method for optimizing the distribution network that accounts for these uncertainties.

3.1.2 Problem Definition

We wish to minimize the total expected value of production and distribution of thirty-seven commodities produced at three plants and distributed to customers throughout the United States and Canada. Each commodity may be shipped directly from a plant to a customer, or shipped through one of fifty-four intermediate warehouses. Outbound shipments to customers are done via truck, but shipments from plants to warehouses may be done via either truck or rail modes.

For a product shipped directly from a plant to a customer, the costs incurred are the production cost, the cost of loading a truck, and the outbound freight. For a product passing through an intermediate warehouse, the costs incurred are the production cost, the rail or truck loading cost, the rail or truck freight from the plant to the warehouse, the cost of unloading the inbound rail or truck at the warehouse, the cost of loading the outbound truck to the customer, and the outbound truck freight to the customer. While the freight and loading/unloading costs incurred are typically larger for a product passing through an intermediate warehouse than one shipping direct from a plant, the absence of freight lanes from every plant to every customer sometimes necessitates the use of this option.

Our objective is therefore to minimize the total cost of production, loading/unloading, and freight by determining which plant or warehouse should ship to each customer, and which plant should ship to each warehouse utilized in meeting customer requirements. The problem is subject to three sets of constraints:

1. Demand constraints: demand is met at each customer location
2. Capacity constraints: capacity is not exceeded at any production location
3. Warehouse balance constraints: the amount of each product leaving each warehouse is less than or equal to the amount entering that warehouse

3.1.3 Approach

To optimize USG's production and distribution network while accounting for demand and cost uncertainty, we first characterize the uncertainty of the historical data, using one year of monthly data for customer demand and three years of monthly data for production costs. This data is fit to probability distributions as described in the section *Statistical Fit of Historical*

Data. In the *Problem Formulation* section, we first solve the base case, to mimic USG's current planning process. We then propagate the demand and cost uncertainties through the base model to determine their relative impacts. Based on our analysis of the historical uncertainty, we determine that the chance constraint method is appropriate for our problem, and use the cumulative probability distributions to convert the deterministic demand constraints into probabilistic constraints. Finally, we sample the cost and demand uncertainties and propagate them through the models obtained using the chance constraint method for several different constraint fulfillment levels, applying a penalty for constraint violations to determine the total cost. We present our findings in the *Results* section.

3.2 Review of Literature

The problem of optimizing distribution through a supply chain has been long studied. In their seminal work, (65) optimize the location of distribution facilities between plants and customers using Bender's decomposition. Other authors who focus on deterministic production and distribution planning are presented in reviews by (66), (67), and (68). More recently, a number of authors have addressed uncertainty in planning a single tier of the supply chain, such as production planning/scheduling or transportation decisions, as reviewed by (69) and (70).

More recently, authors have used stochastic optimization methods to optimize production and distribution decisions in a supply chain under demand uncertainty, similar to the problem under our study. (71) propose a model to handle demand uncertainty by considering the production decisions to be "here-and-now" and the distribution decisions to be "wait-and-see." (72) use a multi-objective two-stage stochastic model to maximize both profit and demand fulfillment, while (73) apply a sample average approximation scheme to optimize the supply chain configuration to minimize total production and distribution cost.

Other literature on supply chain uncertainty is motivated by a desire to formulate optimal responses to natural disasters, i.e. the root cause of uncertainty is a single uncertain event that can be modeled as one of many possible scenarios. In an early paper on the topic, (74) present several heuristic methods for solving a multi-commodity, multi-modal network flow problem under time constraints. (75) develop a two-stage stochastic programming model for distribution of critical goods under uncertainty in both demand and transportation network capacity. (76) assume a discrete number of disaster scenarios to solve a storage location and inventory problem with a vehicle routing subproblem. (77) similarly model both the location and distribution decisions, but assume that the supply chain has several echelons, each containing several entities that must work in cooperation to respond to a disaster.

Another common approach is robust optimization for addressing uncertainty in supply chain planning. (78) use robust optimization to solve a network design problem with uncertainty in both demand and transportation costs, and demonstrate its advantage over worst-case scenario planning. (79) propose a robust optimization approach for solving two-stage network flow problems with uncertain demand, showing that their method is less computationally complex than scenario-based stochastic optimization and allows for parametrization of the conservatism of the solution. (80), investigate a three-tier supply chain with uncertainty in both demand and price preferences among the buyers and sellers and use robustness measures to minimize the degree to which the objective values vary with demand uncertainties.

Our work is most closely related to (81), who model a multi-echelon supply chain network under demand uncertainty, optimizing the number, location, and capacity of warehouses and distribution centers, the transportation links among locations, and the production rates of materials. While we optimize the production and flow of materials through warehouses in already fixed locations, the choice among many warehouses in the USG problem can be thought of as analogous to the decision of choosing warehouse locations from a discrete number of options. Similarly, the choice of freight lanes for moving products in our case is similar to the decision of which transportation links to establish.

3.3 Statistical Fit of Historical Data

The data available to us consists of one year's worth of monthly data for the historical demand and thirty-four month's worth of monthly data for the historical production cost. In this section, we develop probability distributions that fit the data in order to later simulate a large number of scenarios by drawing samples from those distributions.

3.3.1 Historical Demand

Looking at historical demand, we aggregate data into 162 geographical areas (hereafter referred to as "areas") and 37 items. The areas represent states, provinces, or in some cases, portions of states (for example, California is broken down into two areas based on the fact that historical shipments to customers in northern California came from a different warehouse than southern California). This gives us a total of 5994 state/item combinations. However, for the time period studied (2012), only 689 of these combinations have sales in at least one month.

We therefore heuristically eliminate the 5305 combinations that have no historical demand, and work to model the demand of the 689 combinations that do.

For combinations with sales in 10 or more months, we model the demand as a truncated normal distribution in three steps:

1. We calculate the mean and standard deviation (μ and σ) of the historical data.
2. To find the bounds within which 99.9% of samples will fall, we multiply the standard deviation by 3.08 (taken from a standard normal z table), and subtract this from the mean to determine the lower bound. We add the same value to the mean to determine the upper bound.
3. If the lower bound is negative ($L < 0$), we replace the lower bound with zero, and the upper bound with twice the mean (2μ). This ensures that the demand samples are both non-negative and centered around the mean.

Using this method, 41 combinations are modeled with a lower bound greater than zero, and 148 combinations are modeled as a normal distribution truncated at a lower bound of zero.

For combinations with sales in 9 or fewer months, we model the demand as a binned uniform distribution. For this type of distribution, we specify one or more "bins" from which samples should be drawn, each defined by a lower and upper bound, and the percentage of samples that should be drawn from each bin. For each combination, the smallest bin has bounds $[0, \epsilon]$, where ϵ is a mathematically insignificant value (i.e. $1.0E-9$) meant to ensure that a portion of the samples taken for this combination are effectively zero.

For combinations with three or fewer distinct values for demand (201 combinations), we create a separate bin for each value, and then assigned a percentage of samples equal to the percentage of historical data points matching that value, as shown in Table Table VI. For combinations with four or more distinct values (299 combinations), we use a zero bin as described above and a non-zero bin with lower and upper bounds corresponding to the smallest and largest non-zero demand represented in the twelve months of data studied.

TABLE VI: DEMAND COMBINATIONS WITH THREE OR FEWER DISTINCT VALUES

Bins (excluding zero bin)	Number of combinations
1	400
2	49
3	51

3.3.2 Historical Production Cost

Historical production cost data is available for 36 items at three plants, and one additional item at two plants, for a total of 110 item/plant combinations.

As summarized in Table Table VII, we model 108 of these combinations as normal distributions using thirty-four months of monthly historical data (taken in the 2010-2012 timeframe)

TABLE VII: DISTRIBUTIONS USED TO MODEL COST COMBINATIONS

Distribution	Number of data points used	Combinations
Normal	34	91
Normal	33	8
Normal	32	15
Uniform	34	1
Binned Uniform	34	1

to calculate the means and standard deviations. Of these combinations, 76 pass a Pearson's Chi-squared goodness of fit test for a normal distribution at a p-value of 0.05 using all data points available. Of the remaining 32 combinations, 8 pass the test when a single outlier is dropped (and the mean and standard deviations recalculated accordingly), and 15 pass when two outliers are dropped. In the latter case, the two months of extremely high production costs are subsequent, and assumed to be an uncommon production issue that will not recur. The final 9 combinations do not pass at the 5% level, but pass for smaller p-values, and do not more closely resemble any other common distribution, thus the normal approximation is still used.

One combination is modeled as a uniform distribution, and one as a binned uniform distribution as described in the previous section. Both pass the goodness-of-fit test for these distributions at the 5% level.

3.4 Problem Formulation

3.4.1 The Base Case

We first optimized the network using a single month's data (in this case, October 2012, a month in which the total demand was close to the twelve-month average), in order to mimic the current planning process and therefore create a baseline plan against which others can be measured. Our objective is to minimize total cost, which consists of the production costs, the freight costs from plants to warehouses, the freight costs from plants to geographical areas, the freight costs from warehouses to geographical areas, and the handling costs. For the solution to be feasible, we must meet three sets of constraints:

1. Demand constraints: the total amount shipped from all plants to a geographical area plus the total amount shipped from all warehouses to a geographical area must be greater than or equal to demand in that area.

2. Capacity constraints: the total amount produced at each plant must be less than or equal to the capacity at that plant.

3. Warehouse balance constraints: the amount of product shipped from all plants to a warehouse must be greater than or equal to the amount of product shipped from that warehouse to all geographical areas.

The mathematical formulation is given as follows:

$$
\begin{aligned}
Minimize\ &\sum_i \sum_j \sum_l \left(p_{il} + f_{ij} + (\frac{l_i}{n_t}) \right) x_{ijl} + \sum_i \sum_k \sum_l \left(p_{il} + ft_{ik} + (\frac{l_i}{n_t}) \right) xt_{ikl} \\
&+ \sum_i \sum_k \sum_l \left(p_{il} + fr_{ikl} + (\frac{lr_i}{n_r}) \right) xr_{ikl} + \sum_j \sum_k \sum_l \left(fw_{kj} + (\frac{l_i}{n_t}) \right) x_{kjl} \\
subject\ to: &\sum_i x_{ijl} + \sum_k xw_{kjl} \leq d_j \ \forall\ j,l \\
&\sum_j x_{ijl} + \sum_k (xt_{ikl} + xr_{ikl}) \leq c_i \ \forall\ i,l \\
&\sum_i (xt_{ikl} + xr_{ikl}) \leq \sum_j xw_{kjl} \ \forall\ k,l
\end{aligned}
$$

where:

x_{ijl}: amount of item l shipped from plant i to customer j via truck

xt_{ikl}: amount of item l shipped from plant i to warehouse k via truck

xr_{ikl}: amount of item l shipped from plant i to warehouse k via rail

xw_{kjl}: amount of item l shipped from warehouse k to customer j via truck

p_{il}: cost of producing item l at plant i

lt_i: cost of loading a truck at plant i

lr_i cost of loading a rail at plant i

n_t: units per truck

n_r: units per rail

f_{ij}: truck freight per unit from plant i to customer j

ft_{ik}: truck freight per unit from plant i to warehouse k

fr_{ik}: rail freight per unit from plant i to warehouse k

fw_{kj}: truck freight per unit from warehouse k to customer j.

3.4.2 The Average Demand Case

We next optimize the network using average costs and demands, to quantify the cost of the plan when the deterministic optimization techniques are applied to the average of one year's worth of data. In this case, our problem remains the same as in the base case, except that each production cost is replaced with the twelve-month average production cost, and each demand is replaced with the twelve-month average demand.

3.4.3 The Impact of Uncertainty

To determine how much cost and demand uncertainty affect total cost, we create 10,000 scenarios by sampling the demand and cost distributions developed in the previous section and propagate them through the model optimized for average cost and demand, ignoring constraint violations. The were generated using software provided by the Vishwamitra Research Institute to implement a Latin Hypercube sampling technique, and cost propagation was done in MATLAB using standard functions. We first propagate both demand and production cost uncertainty, then repeat the process for only demand, and then for only production cost. As shown in Table Table VIII, the demand uncertainty accounts for most of the variance in total cost: the standard deviation propagating only demand uncertainty through the cost model is more than 97.5% of the standard deviation when both demand and production cost uncertainty are used. We can therefore justify using an optimization method that focuses primarily on finding a best solution with respect to demand uncertainty.

TABLE VIII: TOTAL NETWORK COST UNDER THE PROPAGATION OF DEMAND AND COST UNCERTAINTIES

Uncertainties	Standard Deviation
Demand & Prod. Cost	3.8% of mean
Demand	3.7% of mean
Production Cost	0.8% of mean

3.4.4 The Chance Constraint Method

To account for uncertainty in the demands, we use the chance constraint method, as first proposed by (82). In this method, we seek a solution to our optimization problem that will allow each demand constraint to be met a certain percentage of the time, once implemented. For example, if we choose a 75% chance of meeting demand for a single item in Montana, we would expect that, in any given year, the planned amount of product distributed to Montana would meet or exceed the demand in nine of twelve months, and fall short in the other three. In a practical sense, customers in Montana would still receive the product required in the other three months (as long as total network demand is less than total network capacity, an assumption we will return to later), but from a non-optimal location, and thus at a higher cost.

We therefore divide our calculation of total cost into two phases:

1. We optimize the plan assuming a certain probability (or "chance"), p, that all demand constraints will be met.

2. We sample the demand distributions using a Latin hypercube sampling technique (83) and propagate the uncertainty through the model developed in 1, applying a cost penalty each time a constraint is violated.

The problem under our study lends itself well to this method, due to the zero demands common in our historical data set. Of the 689 area/item combinations with sales in 2012, 201 have no sales in at least nine of twelve months, and 367 have no sales in at least six of twelve months. Therefore, a chance constraint formulation allows us to consider what happens when we focus the optimization on the most likely demands, and decrease the influence of less likely occurrences, i.e. sales in area/item combinations that happen with low frequency.

Further, the chance constraint method enables the conversion of the stochastic program into a standard linear program. Given that USG already has the software to solve a large-scale LP, it is advantageous to use a method such as this one that allows for future planning of the network using existing capabilities.

To optimize the plan under chance constraints (phase 1), we begin with the base case formulation given in the previous section. For each of the demand constraints, we take the cumulative probability distribution developed earlier and locate the p^{th} percentile. We then replace the right side of the demand constraint with this value, thus ensuring that demand would be fulfilled with a probability of p. Because we have monthly data for demand, we begin with $p = \frac{6}{12} = .5$, and then later repeat chance constraint optimization for other fractions of 12: .58, .67, .75, .83, .92, and 1. We formalize this mathematically as

$$P(d_{jl} \geq \sum_i x_{ijl} + \sum_k x_{kjl}) \geq p \; \forall \; j, l$$

$$\sum_i x_{ijl} + \sum_k x_{kjl} \leq F^{-1}_{d_{jl}}(p) \; \forall \; j, l,$$

where $F^{-1}_{d_j}$ is the inverse cumulative probability distribution of demand at location j.

For example, to model the 83rd percentile of demand, we used the 3rd largest demand value out of the 12 available. An example is shown in Figure 9. By using the 83rd percentile of demand, we assume that, were this optimized plan implemented in practice, each month, 83% of the demand constraints would be met.

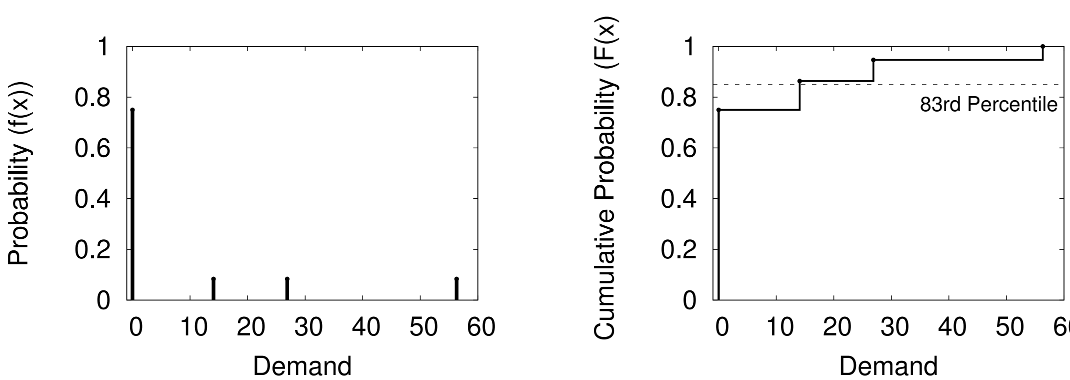

(a) Probability distribution function.

(b) Cumulative distribution function

Figure 9: PDF and CDF of a sample item/area combination.

The chance constraint LP therefore takes the form

$$
\begin{aligned}
Minimize \;\; & E\Big[\sum_i \sum_j \sum_l \left(p_{il} + f_{ij} + (\frac{l_i}{n_t})\right) x_{ijl} + \sum_i \sum_k \sum_l \left(p_{il} + ft_{ik} + (\frac{l_i}{n_t})\right) xt_{ikl} \\
& + \sum_i \sum_k \sum_l \left(p_{il} + fr_{ikl} + (\frac{lr_i}{n_r})\right) xr_{ikl} + \sum_j \sum_k \sum_l \left(fw_{kj} + (\frac{l_i}{n_t})\right) x_{kjl}\Big] \\
subject\ to: & \sum_i x_{ijl} + \sum_k x_{kjl} \leq F_{d_{jl}}^{-1}(p) \;\forall\; j, l, \\
& \sum_j x_{ijl} + \sum_k (xt_{ikl} + xr_{ikl}) \leq c_i \;\forall\; i, l \\
& \sum_i (xt_{ikl} + xr_{ikl}) \leq \sum_j xw_{kjl} \;\forall\; k, l,
\end{aligned}
$$

where p is the probability that each demand constraint is met.

In the second phase, we sampled 10,000 data points from the demand and cost distributions using a Latin Hypercube sampling scheme, and propagated each sample through the optimized model. Because demands in many of the area/item combinations were assumed to be zero in phase one, if they took a non-zero value when sampled, we assumed the product shipped from the location found in the optimization with average demand constraints, provided capacity was available at that location, and applied the appropriate cost. If there was not capacity available at that location, we used the average cost from the other two production locations, and added the average freight/unit of the entire optimized plan to a penalty for shipping from a non-primary location. The penalty value is based on calculations provided by USG for shipments from non-optimized locations.

3.5 Results

All problems described in the previous section were modeled as linear programs in GAMS, and solved using the IBM CPLEX engine. Computational time was under five minutes for each. To protect the confidentiality of USG's cost data, we assume an average total cost in the base case optimization of "$\$BC_1$," and give all other values in terms of a percent increase or decrease from this value. For the case with costs and demand constraints based on average demands, we find an average total cost of $\$BC_1 + 0.3\%$

The total production and distribution cost for the optimized network under chance constraints is given in Table Table IX. Because many of the item/area combinations had sales in fewer than twelve months, the number of non-zero demand combinations are also provided. Note that the 92nd and 100th percentiles could not be optimized in this manner, as for those percentiles, the total network demand exceeds total network capacity. While the lower demand percentiles have lower total costs, they represent a smaller total amount of product, and thus we must proceed to the second phase, in which we apply penalties as described below, to get a realistic cost of applying the chance constraints in the real world.

Table Table X gives the descriptive statistics for monthly cost using a chance constraint-optimized model applied to 10,000 samples of uncertain parameters, and with constraint violation penalties applied. Here we let $\$BC_2$ represent the base case, and present the other numbers relative to this value.

Optimizing at the 50th percentile and applying the penalty for violated constraints gives us a 4.86% improvement over the average, the lowest cost of all the demand schemas used. This

TABLE IX: NON-ZERO COMBINATIONS AND MONTHLY COST OF THE NETWORK OPTIMIZED USING CHANCE CONSTRAINTS WITHOUT VIOLATION PENALTY

Percentile	Non-Zero Combinations	Monthly Optimized Cost
42nd	279	$BC - 29.19%
50th	322	$BC - 21.72%
58th	347	$BC - 12.25%
67th	424	$BC - 1.80%
75th	488	$BC + 14.48%
83rd	557	$BC + 37.44%

indicates that USG should use the median demands when optimizing the network, rather than average demands or demands from a single month, as they do now.

However, as expected, the standard deviation is larger when optimizing at the 50th percentile than at any of the higher percentiles. This is because, when propagate uncertainty through the model optimized at the 50th percentile of demand, more constraints are violated and then penalized by a higher cost. Because the constraints violated differ from sample to sample across the 10,000 used, the results at the 50th percentile show more variance than those at higher percentiles. It is therefore useful to note that there is a trade-off between the expected lowest total cost and the risk of a high total cost in any individual month. For comparison, at the 58th percentile, we get a 4.4% improvement over the base case, but the standard deviation is almost 8% smaller. While optimization at the 50th percentile gives us the lowest total cost, it forces us to accept a greater risk of high network costs in any given month.

TABLE X: DESCRIPTIVE STATISTICS FOR MONTHLY COST WITH CONSTRAINT VIOLATION PENALTY.

Percentile	Mean	Std. Dev.	Avg. # of violations
Base Case	$\$BC_2$	6.05% of mean	262
Average	$\$BC_2$-1.37%	5.70% of mean	256
42nd	$\$BC_2$-4.58%	7.07% of mean	307
50th	$\$BC_2$-4.86%	6.84% of mean	295
58th	$\$BC_2$-4.35%	6.28% of mean	287
67th	$\$BC_2$-1.37%	5.75% of mean	264
75th	$\$BC_2$+9.08%	4.30% of mean	251
83rd	$\$BC_2$+29.07%	3.15% of mean	230

Looking at the capacity constraints, only one is active in the initial optimization at the 50th percentile, indicating that all capacity will be used at the lowest cost plant (LCP). In many cases, it is more cost-effective to distribute product from the LCP through an intermediate plant to a customer than to service that customer directly from a more expensive plant. Qualitatively, this is the main difference between the output from the base case and from that of the 50th percentile optimization: the latter ensures that the capacity at the LCP is better-utilized, in fact, when uncertainty is propagated through the base case model, the full capacity at the LCP is used 87% of the time, while through the 50th percentile model, this increases to 96%.

3.5.1 Value of the Stochastic Solution

The value of the stochastic solution is the difference between the solution obtained by using average values for each of the uncertain inputs and the solution obtained by considering uncertainty in our choice of modeling methods (84). In this case, it is therefore the difference between the base case solution and the solution obtained using the chance constraint method at the 50th percentile. We therefore calculate the value of the stochastic solution to be 4.86%.

In a practical sense, this indicates that the stochastic solution is valuable to USG if it can be executed and implemented for less than 4.86% of the total monthly cost. As previously mentioned, the chance constraint method utilizes USG's existing software capabilities, and thus the execution and implementation costs are very small.

3.5.2 Expected Value of Perfect Information

The expected value of perfect information is the difference between the stochastic solution and the total cost if all costs and demands were known prior to network optimization (85). Optimizing the network individually for each of the months with available demand data gives us an average cost approximately 3.61% less than the expected cost when optimizing the network at the 50th percentile. This indicates that a perfect forecast of both production costs and demands would save the company 3.61% over the stochastic solution.

We can interpret this to mean that USG could save up to 3.61% by reducing the uncertainty in production costs and demands, either through an actual reduction in the month-to-month variability of production costs and demands or through improved forecasting methods that would allow for network optimization based on narrower sampling distributions. For example, production cost variability may be reduced through stricter process control and demand variability may be reduced through the use of safety stock inventories. However, the cost of investing in these improvements must be weighed against the EVPI, which we see offers a less than 3.61% improvement over the stochastic solution.

3.6 Implementation

USG implements its network plan through a set of "sourcing rules" in its order fulfillment system (OFS). When a customer places an order, the OFS places the order at the plant or warehouse location dictated by the rule. Therefore, we converted the output of the CPLEX solver into a table of sourcing rules by creating a sourcing rule where the decision variable had a positive value (i.e. if the decision variable representing the volume of Durock shipped from plant i to the customer area j was positive, we created a sourcing rule in OFS so that Durock orders placed by customers in area j would be shipped from plant i).

Going forward, a member of USG's Network Optimization Planning team will generate the sourcing rules semi-annually as follows:

1. Find the 50th percentile demands:
   (a) Compile historical data on sales for each customer area and item, in monthly buckets, in an Excel spreadsheet.
   (b) Sort the monthly demands from high to low.
   (c) Choose the 7th largest monthly demand, representing the 50th percentile of demand.
2. Find the average production costs:
   (a) Compile historical data on production costs, by plant and item, in monthly buckets, in an Excel spreadsheet.
   (b) Calculate the average (mean) demands.

3. Input the problem parameters to the INFOR Tactical Planner software via Excel spreadsheet:
    - Demands as found in step 1
    - Average production costs as found in step 2
    - Current freight rates
    - Current plant production capacities
4. Run the INFOR Tactical Planner to create a CPLEX input file.
5. Upload the CPLEX input file to the CPLEX engine and initiate the LP solver, producing a CPLEX output file.
6. Upload the CPLEX output file to the INFOR Tactical Planner to view output in a GUI.
7. Make changes to any sourcing rules in the OFS program by keying in a new source where necessary.

Note that steps 4 through 7 remain unchanged from the original process, and the only difference in step 3 is the use of 50th percentile demands and average costs, instead of data taken from the previous month. Therefore, the only additional work is in steps 1 and 2, both of which involve a larger data query and some brief spreadsheet manipulation not done in the original process. The additional workload, for a member of the Network Optimization Planning team, is estimated at less than one full day, annually, the cost of which is therefore negligible.

As discussed in the previous section, the theoretical results suggest a potential 4.9% average savings in each month. To capture these savings, USG began implementing the new sourcing

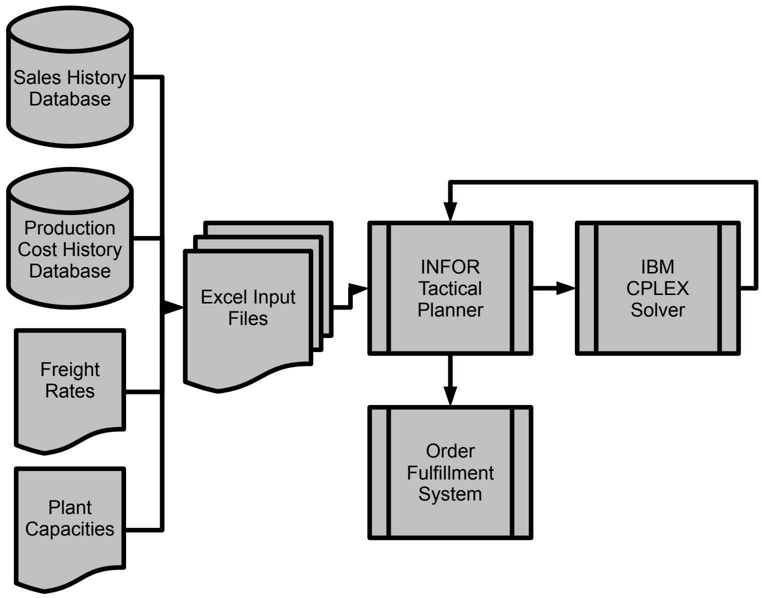

Figure 10: Flowchart of the semi-annual sourcing rule creation process.

rules in January, 2014. However, in actually implementing the sourcing rules, there were found to be several limitations that prevent USG from obtaining the entire cost savings. The largest is the practice of shipping Durock orders in less than full truckload quantities. Only about half of Durock orders ship in full truckload quantities, with the remaining Durock shipping in mixed truckloads with items from other product lines. Typically, the other product line is Sheetrock, and its volume exceeds that of the Durock, so the Sheetrock sourcing rules supersede those of the Durock product line.

Further, other customer service limitations prevent orders from shipping from the optimal plant. For example, for customers who pick up their product at a specific warehouse, the warehouse is chosen by the customer. Taking all of these limiting factors into consideration, the average monthly savings is approximately 1.6% of network costs, with some variation in this number from month to month. While the actual savings are only about one-third of the theoretical savings, they still represents a substantial annual cost reduction to USG.

4. CONCLUSION

In Chapter 2, we showed that a supplier with a direct retail channel achieves higher profit in a two-channel supply chain with symmetric retailers than by acting as a monopoly retailer, even if the equilibrium retail price is lower at the retailers than at the supplier. Further, when the retailers have a stronger cost or price position, the supplier does best by leaving the retail market entirely. We also showed that an increase in the number of retailers has a beneficial effect on the supplier's profit. We therefore conclude that the supplier benefits from additional competition in the retail market provided it is the sole supplier to those competitors and new retailers capture new demand; thus it should sell its product to as many retailers as possible, even if doing so causes nothing to be sold through the direct channel or its own retail price to be undercut.

We also showed that the presence of a direct channel may introduce inefficiencies not present in the single-channel supply chain. Inefficiency is often higher in the two-channel supply chain than in the one-channel supply chain, if the relative maximum product margins of the two channels are such that the supply chain takes different structures under centralization and decentralization. In this case, the efficiency of the supply chain can be increased by eliminating the direct channel. Further, the inefficiency may decrease with differentiation when the number of retailers is high, or increase in the number of retailers when differentiation is low, two phenomena not seen in the single-channel supply chain, and thus creating counterintuitive strategies for mitigating the loss of efficiency.

We demonstrated that the symmetric two-channel supply chain may not be perfectly coordinated using revenue sharing or linear PDS contracts if both channels are to remain in operation. By contrast, we prove that a linear quantity discount contract coordinates this supply chain. However, we also note that these contracts often achieve coordination at a loss to one or more of the participating firms, and in the most extreme cases, cause the retailers to earn no profit at all. Therefore, a fixed transfer payment may be required in concert with one of these three types of contract to perfectly coordinate the supply chain at no loss to any firm. Under such a two-part contract, the total profit is increasing in the number of retailers, but each firm's profit is dependent on its relative negotiating strength for determining the value of the transfer payment. We further showed that our key results for the supplier's profit, the total profit, and supply chain efficiency hold in numerical experiments with asymmetric retailers. Further, the linear quantity discount contract is capable of improving the supply chain efficiency, even though the asymmetry of retailers prevents it from achieving perfect coordination.

These insights have practical implications for manufacturers that operate direct channels through which they sell products that are either exclusive or clear market leaders; specifically, we showed that their exclusivity may better serve them at the wholesale level, selling to many retailers, rather than at the retail level, acting as a monopoly retailer. Additionally, forging relationships with retailers through coordinating contracts further maximizes the leverage of the manufacturer's exclusivity advantage. As these manufacturers adapt to the increasing prominence of superstores in the U.S. marketplace, they are wise to recognize the value of a strategy inclusive of multiple retailers.

In considering the two-channel supply chain with constrained capacity, we found that the supplier can benefit by selling through the retailer, in addition to selling through a direct channel, as long as known capacity is sufficiently high. When capacity is uncertain, however, a one-channel supply chain may result from an upper bound on capacity that is low, relative to the lower bound, the supplier's total market size, and the probability of a high capacity *or* from a lower bound on capacity at which the retail price will induce a zero quantity in the indirect channel.

Further, we showed that, in comparing the profits with known capacity to those under capacity uncertainty, the suppliers does not always benefit from the elimination of uncertainty, and in fact, may be harmed by it. Similarly, the retailer usually benefits from the elimination of uncertainty, but for a range of capacity values, actually sees a decrease in profits if capacity is known. These insights are useful to suppliers both in determining their distribution channel strategy and deciding how much to invest in reduction of uncertainty.

In Chapter 3, we sought to incorporate the uncertainty of production costs and demand into the optimization of a distribution network to find the lowest possible cost of servicing 162 geographical areas from three production sites, either directly or through one of fifty-four intermediate warehouses. To do so, we first fit the historical data for each of 800 uncertainties to either a normal or binned uniform distribution. We then used the chance constraint method to convert probabilistic demand constraints to deterministic demand constraints based on their cumulative distribution functions. Uncertainty was propagated through the chance constrained model and a penalty was applied for each constraint violation. Optimizing the network for the

50th percentile of demand with constraint violation penalties provided the lowest total network cost, at a 4.8% theoretical improvement over the base case, in which a single month's demands were used. As implemented, the savings are an average 1.6% improvement over the base case. This represents a significant monthly cost savings to USG.

In summary, we have investigated several types of decisions that firms make in developing their supply chain strategies, and shown how tools from the fields of game theory and stochastic programming can be used to optimize those decisions. The methodologies and results presented have the potential to improve managerial decisions in the area of supply chain management.

APPENDICES

4.1 Appendix A: Tables

TABLE XI: MONOTONICITY PROPERTIES OF EFFICIENCY

	$0 \leq \rho \leq \rho_{min}$	$\rho_{min} < \rho \leq \rho_{min}^C$	$\rho_{min}^C < \rho < \rho_{max}$	$\rho \geq \rho_{max}$
η	$\frac{(\beta+\delta)(\beta_0\delta-\gamma^2N)N}{\delta^2(\beta_0N-2\gamma N\rho+\Delta\rho^2)}$	$\frac{\Delta R}{N(\beta_0\delta-\gamma^2N)^2}$	$\frac{(\beta_0\Delta-\gamma^2N)R}{2(\beta_0\delta-\gamma^2N)^2(2\beta_0N-(4\gamma N+\delta\rho-\gamma(N-1)\rho)\rho)}$	1
In ρ	Constant	Unimodal	Monotone Increasing	Constant
In N	Monotone Increasing	Not Monotone	Not Monotone	Constant
In γ	Monotone Increasing	Not Monotone	Not Monotone	Constant

$$R = \beta_0N(\beta_0(\beta+\delta) - \gamma^2N) - 2\gamma N(\beta_0(\beta+\delta) - \gamma^2N)\rho + (\beta_0\delta^2 - \gamma^2N\Delta)\rho^2$$

TABLE XII: CONTRACT PARAMETERS AND RESULTING PROFITS FOR $\rho_{min} < \rho < \rho_{min}^C$

	Revenue Sharing	Linear PDS	Linear Quantity Discount
	$\phi = \frac{\delta(\rho_{min}^C-\rho)}{\beta\rho_{min}^C}$	$\zeta = \frac{\delta(\rho-\rho_{min})}{\beta\rho_{min}^C}$	$s = \frac{\delta(\rho-\rho_{min})}{2\rho-\rho_{min}^C}$
	$w^r = \frac{A}{2\beta\gamma N}$	$z = \frac{A}{2\beta\gamma N}$	$w^o = \alpha - c - \frac{\gamma\rho_{min}^C(N-1)\nu}{2\Delta(2\rho-\rho_{min}^C)}$
$\widehat{\Pi}_0$	$\frac{\nu^2(\delta\rho-\beta\rho_{min}^C)}{4\gamma\Delta}$	$\frac{\nu^2(\delta\rho-\beta\rho_{min}^C)}{4\gamma\Delta}$	$\frac{\nu^2N((2\beta+3\gamma(N-1))\rho-(\beta+2\gamma(N-1))\rho_{min}^C)}{4\Delta^2(2\rho-\rho_{min}^C)}$
$\widehat{\Pi}_i$	$\frac{\gamma\nu^2(N-1)(\rho-\rho_{min}^C)}{4\Delta^2(2\rho-\rho_{min}^C)}$	$\frac{\nu^2\delta(\rho_{min}^C-\rho)}{4\gamma\Delta N}$	$\frac{\nu^2\delta(\rho_{min}^C-\rho)}{4\gamma\Delta^2N}$

$$A = (2\beta(\Delta(c' + c_A) + \alpha\gamma(N-1)) + \gamma^2(\alpha + c')(N-1)^2)\rho_{min}^C - \delta\rho(\delta c' + \alpha\gamma(N-1))$$

One-channel Equivalent	Domain	Retail Price	Wholesale Price
$D_s(p) \leq L$	$a_r - bc \leq L$	$\frac{a_s+bc}{2b}$	c
$D_s(p) \leq L$	$\frac{a_s-bc}{2} \leq L < a_r - bc < H$	$\frac{a_s+bc}{2b}$	$\frac{bc(LA+bc\alpha)+(-bc(-2+\alpha)-LA)a_s-a_r(bc(1+\alpha)-Aa_s)}{b(bc-2a_r+a_s)}$
$D_s(p) \leq L$	$\frac{a_s-bc}{2} \leq L < H < a_r - bc$	$\frac{a_s+bc}{2b}$	$\frac{bc(LA-H\alpha)+(2bc+H\alpha-LA)a_s-a_r(bc+a_s)}{b(bc-2a_r+a_s)}$
$D_s(p) \leq L$	$\frac{\alpha(a_s-bc)}{1+\alpha} \leq L < \frac{a_s-bc}{2}$	$\frac{a_s-L}{b}$	$\frac{-L^2+(H-L)(bc+L)\alpha+(-H\alpha+L(2+\alpha)-a_s)a_s+a_r(-L+a_s)}{b(L+a_r-a_s)}$
$L < D_s(p) \leq H$	$L < \frac{\alpha(2a_r-a_s-bc)}{1-\alpha}$, $\frac{\alpha(a_r+bc-L)+L}{\alpha} \leq H$	$\frac{L+\alpha(a_s+bc-L)}{2b\alpha}$	$\frac{LA^2+bc\alpha(1+\alpha)-A\alpha a_s}{2b\alpha}$
$L < D_s(p) \leq H$	$L < \frac{\alpha(2a_r-a_s-bc)}{1-\alpha}$, $\frac{\alpha(a_s-bc+L)-L}{2\alpha} \leq H < \frac{\alpha(a_r+bc-L)+L}{\alpha}$	$\frac{L+\alpha(a_s+bc-L)}{2b\alpha}$	$\frac{L^2A^2(1+\alpha)+\alpha^2 B-2LA\alpha(a_s+bc+H\alpha)-2\alpha(L+(a_s+bc-L)\alpha)a_r}{2b\alpha(L+(a_s+bc)\alpha-L\alpha-2\alpha a_r)}$
$L < D_s(p) \leq H$	$\frac{\alpha(2a_r-a_s-bc)}{1+\alpha} \leq L < \frac{\alpha(a_s-bc)}{1-\alpha}$, $\frac{\alpha(a_s-bc+L)-L}{2\alpha} \leq H$	$\frac{L+\alpha(a_s+bc-L)}{2b\alpha}$	N/A
$H < D_s(p)$	$H < \frac{\alpha(a_s-bc+L)-L}{2\alpha}$	$\frac{a_s-H}{b}$	N/A

$A = \alpha - 1$

$B = (a_s + bc)^2 - (a_s - bc)(a_s - bc - 2H)\alpha$

TABLE XIII: PRICES AT EQUILIBRIUM IN THE TWO CHANNEL SUPPLY CHAIN WITH CAPACITY UNCERTAINTY

One-channel Equivalent	Domain	Supplier Demand	Retailer Demand	Total Demand
$D_s(p) \leq L$	$a_r - bc \leq L$	$\frac{a_s - bc}{2}$	$a_r - \frac{a_s + bc}{2}$	a_r
$D_s(p) \leq L$	$\frac{a_s - bc}{2} \leq L < a_r - bc < H$	$\frac{a_s - bc}{2}$	$a_r - \frac{a_s + bc}{2}$	a_r
$D_s(p) \leq L$	$\frac{a_s - bc}{2} \leq L < H < a_r - bc$	$\frac{a_s - bc}{2}$	$a_r - \frac{a_s + bc}{2}$	a_r
$D_s(p) \leq L$	$\frac{\alpha(a_s - bc)}{1+\alpha} \leq L < \frac{a_s - bc}{2}$	L	$a_r - a_s + L$	$a_r - a_s + 2L$
$L < D_s(p) \leq H$	$L < \frac{\alpha(2a_r - a_s - bc)}{1-\alpha}$, $\frac{\alpha(a_r + bc - L) + L}{\alpha} \leq H$	$\frac{\alpha(a_s - bc + L) - L}{2\alpha}$	$a_r - \frac{\alpha(2a_r - a_s + bc - L) + L}{2\alpha}$	$a_r - bc + L(1 - \frac{1}{\alpha})$
$L < D_s(p) \leq H$	$L < \frac{\alpha(2ar - a_s - bc)}{1-\alpha}$, $\frac{\alpha(a_s - bc + L) - L}{2\alpha} \leq H < \frac{\alpha(a_r + bc - L) + L}{\alpha}$	$\frac{\alpha(a_s - bc + L) - L}{2\alpha}$	$a_r - \frac{\alpha(2a_r - a_s + bc - L) + L}{2\alpha}$	$a_r - bc + L(1 - \frac{1}{\alpha})$
$L < D_s(p) \leq H$	$\frac{\alpha(2a_r - a_s - bc)}{1+\alpha} \leq L < \frac{\alpha(a_s - bc)}{1-\alpha}$, $\frac{\alpha(a_s - bc + L) - L}{2\alpha} \leq H$	$\frac{\alpha(a_s - bc + L) - L}{2\alpha}$	N/A	$\frac{\alpha(a_s - bc + L) - L}{2\alpha}$
$H < D_s(p)$	$H < \frac{\alpha(a_s - bc + L) - L}{2\alpha}$	H	N/A	H

$A = \alpha - 1$

TABLE XIV: DEMANDS AT EQUILIBRIUM IN THE TWO CHANNEL SUPPLY CHAIN WITH CAPACITY UNCERTAINTY

One-channel Equivalent	Domain	Expected Supplier Profit	Expected Retailer Profit	Expected Total Profit
$D_s(p) \leq L$	$a_r - bc \leq L$	$\frac{(a_s-bc)^2}{4b}$	$\frac{(a_s-bc)(2a_r-a_s-bc)}{4b}$	$\frac{(a_r-bc)(a_s-bc)}{2b}$
$D_s(p) \leq L$	$\frac{a_s-bc}{2} \leq L < a_r - bc < H$	$\frac{(a_s-bc)^2}{4b}$	$\frac{(bc-a_s)(2LA+bc(2\alpha-1)+a_s-2\alpha a_r))}{4b}$	$\frac{(LA+\alpha(b_c-a_r))(bc-a_s)}{2b}$
$D_s(p) \leq L$	$\frac{a_s-bc}{2} \leq L < H < a_r - bc$	$\frac{(a_s-bc)^2}{4b}$	$\frac{(a_s-bc)(bc+2(L+H\alpha-L\alpha)-a_s)}{4b}$	$\frac{(LA-H\alpha)(bc-a_s)}{2b}$
$D_s(p) \leq L$	$\frac{\alpha(a_s-bc)}{1+\alpha} \leq L < \frac{a_s-bc}{2}$	$\frac{L(a_s-bc-L)}{b}$	$\frac{(H-L)\alpha(a_s-bc-L)}{b}$	$\frac{(LA-H\alpha)(bc+L-a_s)}{b}$
$L < D_s(p) \leq H$	$L < \frac{\alpha(2a_r-a_s-bc)}{1-\alpha}$, $\frac{\alpha(a_r+bc-L)+L}{\alpha} \leq H$	$\frac{(LA+bc\alpha-\alpha a_s)^2}{4b\alpha}$	$\frac{(LA+bc\alpha-\alpha a_s)(L+bc\alpha-L\alpha-2\alpha a_r+\alpha a_s)}{4b\alpha}$	$\frac{(bc-a_r)(LA+bc\alpha-\alpha a_s)}{2b}$
$L < D_s(p) \leq H$	$L < \frac{\alpha(2a_r-a_s-bc)}{1-\alpha}$, $\frac{\alpha(a_s-bc+L)-L}{2\alpha} \leq H < \frac{\alpha(a_r+bc-L)+L}{\alpha}$	$\frac{(LA+bc\alpha-\alpha a_s)^2}{4b\alpha}$	$\frac{(LA+bc\alpha-\alpha a_s)(LA-(bc+2H)\alpha+\alpha a_s)}{4b\alpha}$	$\frac{(LA-\text{as}\alpha+bc\alpha)(LA-H\alpha)}{2b\alpha}$
$L < D_s(p) \leq H$	$\frac{\alpha(2a_r-a_s-bc)}{1-\alpha} \leq L < \frac{\alpha(a_s-bc)}{1+\alpha}$, $\frac{\alpha(a_s-bc+L)-L}{2\alpha} \leq H$	$\frac{(LA+bc\alpha-\alpha a_s)^2}{4b\alpha}$	N/A	$\frac{(LA+bc\alpha-\alpha a_s)^2}{4b\alpha}$
$H < D_s(p)$	$H < \frac{\alpha(a_s-bc+L)-L}{2\alpha}$	$\frac{(a_s-bc-H)(\alpha(H-L)+L)}{b}$	N/A	$\frac{(a_s-bc-H)(\alpha(H-L)+L)}{b}$

$A = \alpha - 1$

TABLE XV: PROFITS AT EQUILIBRIUM IN THE TWO CHANNEL SUPPLY CHAIN WITH CAPACITY UNCERTAINTY

4.2 Appendix B: Proofs of Propositions and Theorems

Proof of Proposition 2.2.1. We first consider the case that $q_0 > 0$ and $q_i > 0$. Since Π_i is concave in q_i, maximizing Π_i over q_i to find the retailers' reaction function, the first order condition is

$$\frac{d\Pi_i(q_i)}{dq_i} = \alpha - 2\beta q_i - \gamma \sum_{\substack{j=1 \\ j\neq i}}^{N} q_j - w - c = 0.$$

Using the symmetry of the retailers, this becomes

$$\frac{d\Pi_i(q_i)}{dq_i} = \alpha - 2\beta q_i - \gamma((N-1)q_i + q_0) - w - c = 0,$$

giving a reaction function of

$$q_i(w) = \frac{\alpha - c - w - \gamma q_0}{2\beta + \gamma(N-1)}.$$

The wholesale market clearing condition implies $Q = Nq_i(w)$, or $w = \alpha - c - \gamma q_0 - \frac{2\beta+\gamma(N-1)}{N}Q$. We next maximize Π_0 over Q and q_0. It is straightforward to obtain that the Hessian is negative definite and hence, after simplifying the first order conditions, we obtain $q_0^* = \frac{\delta\nu_0 - \gamma N\nu}{2(\beta_0\delta - \gamma^2 N)}$, $q_i^* = Q^*/N = \frac{\beta_0\nu - \gamma\nu_0}{2(\beta_0\delta - \gamma^2 N)}$ and $w^* = \frac{\alpha - c + c_A}{2}$.

If the resulting q_i is such that $q_i \leq 0$, (i.e. $\frac{\beta_0\nu - \gamma\nu_0}{2(\beta_0\delta - \gamma^2 N)} \leq 0$), we then take $q_i = 0$: the supplier acts as a monopoly retailer. In this case, Π_0 is concave in q_0, and therefore, we use the first-order condition to maximize the profit over q_0, and find that $q_0^* = \frac{\nu_0}{2\beta_0}$.

If the resulting q_0 is such that $q_0 \leq 0$, (i.e. $\frac{\delta\nu_0 - \gamma N\nu}{2(\beta_0\delta - \gamma^2 N)} \leq 0$), we then take $q_0 = 0$: the supplier acts as a wholesale supplier, and Π_0 is concave in Q and Π_i is concave in q_i. We proceed as we did in the two-channel case to get a reaction function of $q_i = \frac{\nu - w}{2\beta + \gamma(N-1)}$. Maximizing Π_0 over the wholesale quantity, we get, from the first order conditions, $Q^* = \frac{\nu N}{2\delta}$, $w^* = \frac{\nu}{2} + c_A$ and $q_i^* = \frac{\nu}{2\delta}$. The prices and profits follow from the quantities and wholesale prices in all three cases, and it is easily verified that the assumption of $\nu_0 > 0$ and $\nu > 0$ is sufficient to ensure the non-negativity of all prices and profits. □

Proof of Proposition 2.2.2. The total profit is given in Equation (Equation 2.4). When $q_0 > 0$ and $q_i > 0$, the total profit function is concave in q_0 and q_i, and the first-order conditions are

$$\frac{\partial \Pi_T(q)}{\partial q_0} = \alpha_0 - 2\beta_0 q_0 - 2\gamma \sum_{i=1}^{N} q_i - c_0' = 0$$
$$\frac{\partial \Pi_T(q)}{\partial q_i} = \alpha - 2\beta q_i - 2\gamma \sum_{\substack{j=1 \\ j \neq i}}^{N} q_j - c' = 0.$$

Using the symmetry of retailers, the first-order conditions become

$$\frac{\partial \Pi_T(q)}{\partial q_0} = \alpha_0 - 2\beta_0 q_0 - 2\gamma N q_i - c_0' = 0$$
$$\frac{\partial \Pi_T(q)}{\partial q_i} = \alpha - 2\beta q_i - 2\gamma((N-1)q_i + q_0) - c' = 0.$$

Solving for quantities yields $q_0^* = \frac{\Delta\nu_0 - \gamma N\nu}{2(\beta_0\Delta - \gamma^2 N)}$ and $q_i^* = \frac{\beta_0\nu - \gamma\nu_0}{2(\beta_0\Delta - \gamma^2 N)}$.

As in the decentralized case, if the expression obtained above for q_0 is negative, we let $q_0 = 0$ ($\rho \leq \rho_{min}^C$); the total profit function is then concave in q_i and maximizing over q_i, the first order condition and the symmetry of retailers give $q_i^* = \frac{\nu}{2\Delta}$. If the expression obtained above for $q_i 0$ is negative, we let $q_i = 0$ ($\rho \geq \rho_{max}$); the centralized case is then identical to the monopoly case.

The optimal retail prices and maximum profits can be found from the optimal quantities for all three ranges of ρ. □

Proof of Proposition 2.2.3. In the range $0 \leq \rho \leq \rho_{min}$, the efficiency is given by

$$\eta = \frac{\Delta(3\beta + \gamma(N-1))}{\delta^2},$$

which is constant in ρ and increasing in N and γ ($\frac{\partial \eta}{\partial \gamma} = \frac{2\beta^2}{\delta^3} > 0$).

In the range $\rho_{min} < \rho \leq \rho_{min}^C$, the efficiency is

$$\eta = \frac{\Delta\left(\beta_0 N(\beta_0(3\beta + \gamma(N-1)) - \gamma^2 N) + 2\gamma N(\beta_0(\gamma - 3\beta) + \gamma N(\gamma - \beta_0))\rho + \kappa\rho^2\right)}{N(\beta_0\delta - \gamma^2 N)^2},$$

where $\kappa = \beta_0(2\beta-\gamma)^2 + (2\beta_0(\beta-\gamma) + \beta(2\beta_0 - \gamma + \gamma^2)\gamma N + (\beta_0 - \gamma)\gamma^2 N^2$. Clearly, η is a quadratic function of ρ, and because κ is positive, η is convex in ρ. At $\rho = \frac{\gamma N(\beta_0(\gamma - 3\beta) + \gamma N(\gamma - \beta_0))}{\kappa}$, the partial derivative of η with respect to ρ is zero. We next show that this value is between ρ_{min} and ρ_{min}^C:

$$\frac{\gamma N(\beta_0(\gamma - 3\beta) + \gamma N(\gamma - \beta_0))}{\kappa} - \rho_{min} = \frac{(\beta_0(2\beta - \gamma) + \gamma N(\beta_0 - \gamma))\beta\gamma N}{\delta\kappa} \geq 0,$$

$$\rho_{min}^C - \frac{\gamma N(\beta_0(\gamma - 3\beta) + \gamma N(\gamma - \beta_0))}{\kappa} = \frac{\beta_0\beta^2\gamma N}{\Delta\kappa} \geq 0.$$

We therefore conclude that efficiency is unimodal over ρ in this range, as illustrated in Figure 3a.

In the range $\rho^C_{min} < \rho \leq \rho_{max}$, the efficiency and its partial derivative with respect to ρ are

$$\eta = \frac{(\beta_0\Delta - \gamma^2 N)(\beta_0 N(\beta_0(3\beta + \gamma(N-1)) - \gamma^2 N) + 2\gamma N\rho\eta_1 + (\beta_0\delta^2 - \gamma^2 N\Delta)\rho^2)}{(\beta_0(\gamma - 2\beta) + \gamma N(\gamma - \beta_0))^2(\beta_0 N + \rho(-2\gamma N + \beta\rho + \gamma(N-1)\rho))}$$
$$\frac{\partial\eta}{\partial\rho} = \frac{2\beta_0\beta^2 N(\beta_0(\beta - \gamma) + \gamma N(\beta_0 - \gamma))\rho(\beta_0 - \gamma\rho)}{(\beta_0\delta - \gamma^2 N)^2(\beta_0 N + \rho(\gamma(\rho(N-1) - 2N) + \beta\rho))^2} \geq 0,$$

where $\eta_1 = (\beta_0(\gamma - 3\beta) + \gamma N(\gamma - \beta_0))$ so efficiency is monotonically increasing in ρ.

In the range $\rho > \rho_{max}$, the centralized system acts exactly as the decentralized system, and the efficiency is therefore 1. □

Proof of Theorem 2.2.4. If we assume a linear quantity discount contract where $w = p^o - s.q_i$, where $p^o = \frac{\beta_0(\delta(\alpha - c + 2c_A) + \gamma(N-1)c_A)) - \gamma(\beta\nu_0 + \gamma N(\alpha - c + 3c_A))}{2(\beta_0\Delta - \gamma^2 N)}$, and $s = \beta_i - \epsilon$, retailer i's profit is

$$\Pi_i = q_i(\alpha - \gamma(\sum_{j=0}^{N} q_j) - p^o + \epsilon q_i).$$

Finding the reaction function, plugging it into the supplier's profit function, and solving for both quantities gives us quantities for q_0 and q_i as functions of ϵ. As ϵ goes to zero, the quantities become

$$q_0 = \frac{\Delta\nu_0 + \gamma N\nu}{2(\beta_0\Delta - \gamma^2 N)}$$
$$q_i = \frac{\beta\nu - \gamma\nu_0}{2(\beta_0\Delta - \gamma^2 N)},$$

which are equal to the centralized quantities. □

Proof of Proposition 2.3.1. Given the supplier's profit function as:

$$\Pi_s = \begin{cases} K(p-c) & K \leq D_s(p) \\ D_s(p)(p-c) & D_s(p) < K \end{cases}.$$

it is clear that, when $K \leq D_s(p)$, the profit is linear in p, thus the optimal retail price is found at the endpoint $K = D_s(p)$. When $K > D_s(p)$, the profit function is quadratic in p, thus p^* is found when the first order condition is met:

$$\frac{d\Pi_s}{dp} = a_s - 2bp + bc = 0. \tag{4.1}$$

This gives us an optimal retail price that can be expressed as

$$p^* = \begin{cases} \frac{a_s-K}{b} & K \leq D_s(p) \\ \frac{a_s+bc}{2b} & D_s(p) < K \end{cases} .Demands and profits naturally follow. \tag{4.2}$$

□

Proof of Proposition 2.3.2. The supplier's profit is:

$$\Pi_s = \begin{cases} D_s(p-c) & D_s(p) \leq L \\ \alpha D_s(p)(p-c) + (1-\alpha)L(p-c) & L < D_s(p) \leq H \\ \alpha H(p-c) + (1-\alpha)L(p-c) & H < D_s(p) \end{cases} \cdot$$

We proceed as in the proof of Proposition 2.3.1 to find retail prices; demands and profits naturally follow. □

Proof of Theorem 2.3.3. In this case, we first find the wholesale price that will make the supplier's two-channel profits equal to the supplier's one channel profits found in Proposition 2.3.2 by setting the two profit expressions equal to one another and solving analytically for w. We then proceed as in the proof of Proposition 2.3.1 to find retail prices; demands and profits naturally follow. □

Proof of Theorem 2.3.4. When comparing 2C to 2U, under the assumption $K = \alpha H + (1-\alpha)L$,

$$\Pi_s(2C) - \Pi_s(2U) = 0 \tag{4.3}$$

$$\Pi_r(2C) - \Pi_r(2U) = \frac{\alpha(H - a_r + bc)(a_s - bc)}{2b} \tag{4.4}$$

$$\Pi_T(2C) - \Pi_T(2U) = \frac{\alpha(H - a_r + bc)(a_s - bc)}{2b}. \tag{4.5}$$

Thus, we see that all profit loss due to uncertainty accrues to the retailer. We proceed similarly for all valid comparisons between the one and two channel case to obtain the full table of results. □

Proof of Theorem 2.3.5. When comparing 2C to 2U, under the assumption $K = \alpha H + (1-\alpha)L$,

$$\Pi_T(2C) - \Pi_T(2U) = \frac{\alpha(H - a_r + bc)(a_s - bc)}{2b} \tag{4.6}$$

$$\frac{d[\Pi_T(2C) - \Pi_T(2U)]}{d\alpha} = \frac{(a_s - bc)(H - a_r + bc)}{2b} > 0 \tag{4.7}$$

(4.8)

$$\Pi_s(2C) - \Pi_s(2U) = 0 \tag{4.9}$$

(4.10)

$$\Pi_r(2C) - \Pi_r(2U) = \frac{\alpha(H - a_r + bc)(a_s - bc)}{2b} \tag{4.11}$$

$$\frac{d[\Pi_r(2C) - \Pi_r(2U)]}{d\alpha} = \frac{(a_s - bc)(H - a_r + bc)}{2b} > 0 \tag{4.12}$$

(4.13)

(4.14)

Therefore, the total supply chain profit loss is increasing in α, and the same is true of the retailer's profit loss. We proceed similarly for all valid comparisons between the one and two channel case to obtain the full table of results. □

CITED LITERATURE

1. Heizer, J. H., Render, B., and Weiss, H. J.: Operations management, volume 8. Pearson Prentice Hall, 2004.

2. Walmart: 2011 annual report, March 2012.

3. Spengler, J. J.: Vertical integration and antitrust policy. J. Political Economy, 59(4):347 – 352, 1950.

4. Perry, M. K.: Vertical integration: Determinants and effects. In Handbook of Industrial Organization, eds. M. Armstrong and R. H. Porter, chapter 4. North-Holland, 1989.

5. Tsay, A. A. and Agrawal, N.: Modeling conflict and coordination in multi-channel distribution systems: A review. In Supply Chain Analysis in the eBusiness Era, eds. D. Simchi-Levi, S. D. Wu, and Z. M. Shen, chapter 13. Kluwer Academic Publishers, 2004.

6. Peterson, R., Balasubramanian, S., and Bronnenberg, B. J.: Exploring the implications of the internet for consumer marketing. J. Acad. Marketing Sci., 25(4):329 – 346, 1997.

7. Chiang, W. K., Chhajed, D., and Hess, J. D.: Direct marketing, indirect profits: A strategic analysis of dual-channel supply-chain design. Management Sci., 49(1):1 – 20, 2003.

8. Dumrongsiri, A., Fan, M., Jain, A., and Moinzadeh, K.: A supply chain model with direct and retail channels. Eur. J. Oper. Res., 187(3):691 – 718, 2008.

9. Cattani, K., Gilland, W., Heese, H. S., and Swaminathan, J.: Boiling frogs: Pricing strategies for a manufacturer adding a direct channel that competes with the traditional channel. Production Oper. Management, 15(1):40 – 56, 2006.

10. Schoenbachler, D. D. and Gordon, G. L.: Multi-channel shopping: understanding what drives channel choice. Journal of Consumer Marketing, 19(1):42 – 53, 2002.

11. Bernstein, F., Song, J., and Zheng, X.: “bricks-and-mortar” vs. “clicks-and-mortar”: An equilibrium analysis. Eur. J. Oper. Res., 187(3):671 – 690, 2008.

12. Bruggen, G. H. V., Antia, K. D., Jap, S. D., Reinartz, W. J., and Pallas, F.: Managing marketing channel multiplicity. Journal of Service Research, 13(3):331 – 340, 2010.

13. Ofek, E., Katona, Z., and Sarvary, M.: Bricks and clicks: The impact of product returns on the strategies of multichannel retailers. Marketing Sci., 30(1):42 – 60, 2011.

14. Ansari, A., Mela, C. F., and Neslin, S. A.: Customer channel migration. J. of Marketing Research, XLV(2):60 – 76, 2008.

15. Bernstein, F., Song, J., and Zheng, X.: Free riding in a multi-channel supply chain. Naval Research Logistics, 56(8):745 – 765, 2009.

16. Yoo, W. S. and Lee, E.: Internet channel entry: A strategic analysis of mixed channel structures. Marketing Sci., 30(1):29 – 41, 2011.

17. McGuire, T. and Staelin, R.: An industry equilibrium analysis of downstream vertical integration. Marketing Sci., 2(2):161 – 191, 1983.

18. Coughlan, A.: Competition and cooperation in marketing channel choice: Theory and application. Marketing Sci., 4(2):110 – 129, 1985.

19. Moorthy, K.: Strategic decentralization in channels. Marketing Sci., 7(4):335 – 355, 1988.

20. Norton, S.: An empirical look at franchising as an organizational form. J. of Business, 61(2):197 – 217, 1988.

21. Lafontaine, F. and Shaw, K. L.: Targeting managerial control: evidence from franchising. The RAND J. of Econ., 36(1):131 – 150, 2005.

22. Cachon, G. P.: Supply chain coordination with contracts. In Handbooks in operations research and Management Science: Supply chain management: Design, coordination, and operation, eds. S. Graves and A. D. Kök, chapter 6, pages 229 – 331. North-Holland, 2003.

23. Pasternack, B.: Optimal pricing and return policies for perishable commodities. Marketing Sci., 4(2):166 – 176, 1985.

24. Dolan, R. J.: Quantity discounts: Managerial issues and research opportunities. Marketing Sci., 6:1 – 22, 1987.

25. Munson, C. L. and Rosenblatt, M. J.: Theories and realities of quantity discounts: An exploratory study. Production Oper. Management, 7(4):352 – 369, 1998.

26. Weng, Z. K.: Channel coordination and quantity discounts. Management Sci., 41(9):1509 – 1522, 1995.

27. Rosenblatt, M. J. and Lee, H. L.: Improving profitability with quantity discounts under fixed demand. IIE Transactions, 17(4):388 – 395, 1985.

28. Ingene, C. and Parry, M.: Channel coordination when retailers compete. Marketing Sci., 14(4):360 – 377, 1995.

29. Bernstein, F. and Federgruen, A.: Decentralized supply chains with competing retailers under demand uncertainty. Management Sci., 51(1):18 – 29, 2005.

30. Netessine, S. and Zhang, F.: Positive vs. negative externalities in inventory management: Implications for supply chain design. Manufacturing Service Oper. Management, 7(1):58 – 73, 2005.

31. Cachon, G. P. and Lariviere, M. A.: Supply chain coordination with revenue-sharing contracts: Strengths and limitations. Management Sci., 51(1):30 – 44, 2005.

32. Boyaci, T.: Competitive stocking and coordination in a multiple-channel distribution system. IIE Transactions, 37(5):407 – 427, 2005.

33. Tsay, A. A. and Agrawal, N.: Channel conflict and coordination in the e-commerce age. Production Oper. Management, 13(1):93 – 110, 2004.

34. Tang, C.: Perspectives in supply chain risk management. International Journal of Production Economics, 103(2):451–488, 2006.

35. Yang, Z., Aydın, G., Babich, V., and Beil, D.: Supply disruptions, asymmetric information, and a backup production option. Management Science, 55(2):192–209, 2009.

36. Yang, Z., Aydın, G., Babich, V., and Beil, D.: Using a dual-sourcing option in the presence of asymmetric information about supplier reliability: Competition vs. diversification. Manufacturing & Service Operations Management, 14(2):202–217, 2012.

37. Chaturvedi, A. and Martínez-de Albéniz, V.: Optimal procurement design in the presence of supply risk. Manufacturing & Service Operations Management, 13(2):227–243, 2011.

38. Gümüş, M., Ray, S., and Gurnani, H.: Supply-side story: Risks, guarantees, competition, and information asymmetry. Management Science, 58(9):1694–1714, 2012.

39. Babich, V., Burnetas, A., and Ritchken, P.: Competition and diversification effects in supply chains with supplier default risk. Manufacturing & Service Operations Management, 9(2):123–146, 2007.

40. Babich, V.: Independence of capacity ordering and financial subsidies to risky suppliers. Manufacturing & Service Operations Management, 12(4):583–607, 2010.

41. Tang, S. and Kouvelis, P.: Supplier diversification strategies in the presence of yield uncertainty and buyer competition. Manufacturing & Service Operations Management, 13(4):439–451, 2011.

42. Federgruen, A. and Yang, N.: Competition under generalized attraction models: Applications to quality competition under yield uncertainty. Management Science, 55(12):2028–2043, 2009.

43. Wang, Y., Gilland, W., and Tomlin, B.: Mitigating supply risk: Dual sourcing or process improvement? Manufacturing & Service Operations Management, 12(3):489–510, 2010.

44. Hsieh, C. and Wu, C.: Capacity allocation, ordering, and pricing decisions in a supply chain with demand and supply uncertainties. European Journal of Operational Research, 184(2):667–684, 2008.

45. Cachon, G. and Lariviere, M.: Capacity choice and allocation: strategic behavior and supply chain performance. Management Science, 45(8):1091–1108, 1999.

46. Vives, X.: Oligopoly pricing: Old ideas and new tools. Cambridge, Massachussetts, MIT Press, 2001.

47. Carr, S. M. and Karmakar, U. S.: Competition in multiechelon assembly supply chains. Management Sci., 51(1):45 – 59, 2005.

48. Cachon, G. and Kök, A. D.: Competing manufacturers in a retail supply chain: on contractual form and coordination. Management Sci., 56(3):571 – 589, 2010.

49. Kreps, D. M. and Scheinkman, J. A.: Quantity precommitment and bertrand competition yield cournot outcomes. The Bell Journal of Economics, pages 326–337, 1983.

50. Lariviere, M. A. and Porteus, E. L.: Selling to the newsvendor: An analysis of price-only contracts. Manufacturing & service operations management, 3(4):293–305, 2001.

51. Farahat, A. and Perakis, G.: Profit loss in differentiated oligopolies. Oper. Res. Letters, 37(1):43 – 46, 2009.

52. Roughgarden, T. and Tardos, E.: How bad is selfish routing? J. ACM, 49(2):236 – 239, 2002.

53. Johari, R. and Tsitsiklis, J.: Efficiency loss in a network resouce game. Math. Oper. Res., 29(3):407 – 435, 2002.

54. Perakis, G. and Roels, G.: The price of anarchy in supply chains: Quantifying the efficiency of price-only contracts. Management Sci., 53(8):1249 – 1268, 2007.

55. Martinez de Albeniz, V. and Simchi-Levi, D.: Competition in the supply option market. Oper. Res., 57:1082 – 1097, 2008.

56. Adida, E. and DeMiguel, V.: Supply chain competition with multiple manufacturers and retailers. Oper. Res., 49(1):1 – 31, 2011.

57. Zhang, X., Ou, J., and Gilbert, S. M.: Coordination of stocking decisions in an assemble-to-order environment. Eur. J. Oper. Res., 189(2):540 – 558, 2008.

58. Gallego, G., Hu, M., Beltran, J., Ward, J., and Jain, S.: Channel management for basic products and ancillary services. working paper, 2011.

59. Dada, M. and Srikanth, K.: Pricing policies for quantity discounts. Management Sci., 33(10):1247 – 1252, 1987.

60. Bernstein, F. and Federgruen, A.: Pricing and replenishment strategies in a distribution system with competing retailers. Oper. Res., 51(3):409 – 426, 2003.

61. Ha, A. and Tong, S.: Revenue sharing contracts in a supply chain with uncontractible actions. Naval Research Logistics, 55(5):419 – 431, 2008.

62. Nagarajan, M. and Sošić, G.: Game-theoretic analysis of cooperation among supply chain agents: Review and extensions. Eur. J. Oper. Res., 187:719 – 745, 2008.

63. Cottle, R. W., Pang, J.-S., and Stone, R. E.: The linear complementarity problem, volume 60. Siam, 2009.

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

65. Geoffrion, A. M. and Graves, G. W.: Multicommodity distribution system design by benders decomposition. Management science, 20(5):822–844, 1974.

66. Thomas, D. J. and Griffin, P. M.: Coordinated supply chain management. European journal of operational research, 94(1):1–15, 1996.

67. Vidal, C. J. and Goetschalckx, M.: Strategic production-distribution models: A critical review with emphasis on global supply chain models. European Journal of Operational Research, 98(1):1–18, 1997.

68. Klose, A. and Drexl, A.: Facility location models for distribution system design. European Journal of Operational Research, 162(1):4–29, 2005.

69. Sahinidis, N. V.: Optimization under uncertainty: state-of-the-art and opportunities. Computers & Chemical Engineering, 28(6):971–983, 2004.

70. Mula, J., Poler, R., Garcia-Sabater, J., and Lario, F.: Models for production planning under uncertainty: A review. International journal of production economics, 103(1):271–285, 2006.

71. Gupta, A. and Maranas, C. D.: Managing demand uncertainty in supply chain planning. Computers & Chemical Engineering, 27(8):1219–1227, 2003.

72. Guillén, G., Mele, F., Bagajewicz, M., Espuna, A., and Puigjaner, L.: Multiobjective supply chain design under uncertainty. Chemical Engineering Science, 60(6):1535–1553, 2005.

73. Santoso, T., Ahmed, S., Goetschalckx, M., and Shapiro, A.: A stochastic programming approach for supply chain network design under uncertainty. European Journal of Operational Research, 167(1):96–115, 2005.

74. Haghani, A. and Oh, S.-C.: Formulation and solution of a multi-commodity, multi-modal network flow model for disaster relief operations. Transportation Research Part A: Policy and Practice, 30(3):231–250, 1996.

75. Barbarosoğlu, G. and Arda, Y.: A two-stage stochastic programming framework for transportation planning in disaster response. Journal of the Operational Research Society, 55(1):43–53, 2004.

76. Mete, H. O. and Zabinsky, Z. B.: Stochastic optimization of medical supply location and distribution in disaster management. International Journal of Production Economics, 126(1):76–84, 2010.

77. Chang, M.-S., Tseng, Y.-L., and Chen, J.-W.: A scenario planning approach for the flood emergency logistics preparation problem under uncertainty. Transportation Research Part E: Logistics and Transportation Review, 43(6):737–754, 2007.

78. Mudchanatongsuk, S., Ordóñez, F., and Liu, J.: Robust solutions for network design under transportation cost and demand uncertainty. Journal of the Operational Research Society, 59(5):652–662, 2007.

79. Atamtürk, A. and Zhang, M.: Two-stage robust network flow and design under demand uncertainty. Operations Research, 55(4):662–673, 2007.

80. Chen, C.-L. and Lee, W.-C.: Multi-objective optimization of multi-echelon supply chain networks with uncertain product demands and prices. Computers & Chemical Engineering, 28(6):1131–1144, 2004.

81. Tsiakis, P., Shah, N., and Pantelides, C. C.: Design of multi-echelon supply chain networks under demand uncertainty. Industrial & Engineering Chemistry Research, 40(16):3585–3604, 2001.

82. Charnes, A., Cooper, W. W., and Symonds, G. H.: Cost horizons and certainty equivalents: an approach to stochastic programming of heating oil. Management Science, 4(3):235–263, 1958.

83. McKay, M. D., Beckman, R. J., and Conover, W. J.: Comparison of three methods for selecting values of input variables in the analysis of output from a computer code. Technometrics, 21(2):239–245, 1979.

84. Birge, J. R.: The value of the stochastic solution in stochastic linear programs with fixed recourse. Mathematical programming, 24(1):314–325, 1982.

85. Howard, R. A.: Information value theory. Systems Science and Cybernetics, IEEE Transactions on, 2(1):22–26, 1966.

VITA

Amy David

Email: amydavid@gmail.com

Education

Ph.D. Operations Research, University of Illinois at Chicago, expected August 2014
GPA: 3.75/4.0 Dissertation: Use of Game Theory and Stochastic Programming for Supply Chain Optimization

MBA, Lake Forest Graduate School of Management, 2008
GPA: 3.97/4.0 Graduate with Highest Honors

B.S. General Engineering, University of Illinois at Urbana, 2003
GPA: 2.97/4.0

Awards

USG Supply Chain Research Grant, awarded January, 2013

UIC University Fellowship, awarded August, 2010

Teaching Experience

DePaul University Driehaus College of Business

Adjunct Instructor

Operations Management (Core class for MBA students)
Winter 2014
Sole instructor: developed and delivered all course materials for approximately 25 students

Management and Measurement of Quality (Core class for Management students)
Autumn 2013
Sole instructor: developed and delivered all course materials for approximately 25 students

Principles of Operations Management (Core class for all Business students)
Winter 2012, Spring 2012, Autumn 2012, Winter 2013, Summer I 2013, Summer II 2013
Sole instructor: developed and delivered all course materials for average of 40 students

University of Illinois at Chicago

Teaching Assistant

Introduction to Operations Research (Core class for Industrial Engineering students)
Fall 2012
Assisted primary instructor: developed homework assignments, conducted review sessions, and worked individually with approximately 45 students

Statistics and Probability for Engineers (Core class for Industrial Engineering students)
Fall 2011, Spring 2012
Assisted primary instructor: developed homework assignments, conducted review sessions, and worked individually with average of 50 students

Employment

USG Corporation
2007 - 2011 Logistics Planning and Process Development Manager

Managed Oracle Demantra implementation resulting in 900K annual savings through improved sales and operations planning process

Created technology roadmap to document long-term vision for demand, inventory, supply, and distribution planning teams

Used lean principles to re-engineer distribution planning processes for three distinct product lines, resulting in a 15% reduction in planning cycle times

2006 - 2007 Supply Chain Services Manager

Trained and supported plant personnel on Oracle Advanced Supply Chain Planning software and processes

Designed, tested, documented, and implemented patches, upgrades, and custom RICE items

Maintained MRP input parameters to ensure accurate planning output

Medline Industries
2003 - 2006 Logistics Analyst

Managed staff of analysts to achieve departmental project goals

Developed database programs that maintain MRP parameters in SAP to reflect business conditions for twenty-seven distinct product lines

Improved forecasting methods using SAP Advanced Planning and Optimization to reduce inventory investment and improve customer service

Summer 2002 Engineering Intern

Worked with engineering staff to reduce assembly line changeover time between production runs from 15 to 5 minute average

Implemented new quality control processes to reduce number of missing component defects in sterile procedure trays

Consulting

USG Corporation 2012 Project Arrow Business Process Consultant

Developed functional requirements, test plans, and process documentation for new distribution planning software

Managed software defect process throughout QA and UAT cycles

Academic Publications and Presentations

Book: Diwekar, Urmila and David, Amy. *Better Optimization of Nonlinear Uncertain Systems (BONUS) Algorithm for Large Scale Real World SNLP Problems.* New York, NY: Springer Briefs; 2014.

David, Amy and Adida, Elodie. "Competition and Coordination in a Two-Channel Supply Chain." *Invited for third review at POMS Journal*

David, Amy and Diwekar, Urmila. "Network Optimization Under Uncertainty." *Accepted at Interfaces*

David, Amy and Adida, Elodie. "The Two-Channel Supply Chain with Capacity Uncertainty." *In preparation*

David, Amy and Diwekar, Urmila. "Sampling Methods for Large-Scale Uncertainty." *In preparation*

Presentations

David, Amy and Adida, Elodie. "Competition and Coordination in a Two-Channel Supply Chain." POMS 25th Annual Conference, March 2014, Atlanta, GA.

David, Amy and Diwekar, Urmila. "Network Optimization Under Uncertainty: A Case Study." AIChE Annual Meeting, November 2013, San Francisco, CA.

David, Amy and Diwekar, Urmila. "Network Optimization Under Uncertainty: A Case Study." INFORMS Annual Meeting, October 2013, Minneapolis, MN

David, Amy and Adida, Elodie. "Competition and Coordination in a Two-Channel Supply Chain." INFORMS Annual Meeting, October 2013, Minneapolis, MN

David, Amy and Adida, Elodie. "Competition and Coordination in a Two-Channel Supply Chain." INFORMS Annual Meeting, November 2012, Phoenix, AZ.

Industry Publications and Presentations

Publications

David, Amy. "Planning Process Particulars." APICS Magazine Vol. 17, No. 1 (January, 2007): 24 - 27.

Presentations

David, Amy and Adida, Elodie. "Competition and Coordination in a Two-Channel Supply Chain." POMS Annual Conference, May 2014, Atlanta, GA.

David, Amy. "Maximizing the Value of Your S&OP Process." Oracle Application User Group Demantra SIG Annual Meeting, October 2009, San Francisco, CA.

David, Amy. "Value Chain Planning in Uncertain Times." Oracle OpenWorld, October 2009, San Francisco, CA.

David, Amy and Januszyk, Ross. "Demantra Demand Management Implementation at USG." North Central Oracle Application User Group Training Day, August 2009, Oakbrook, IL.

David, Amy and Geilen, Frank. "Advanced Supply Chain Planning." Oracle OpenWorld, September 2008, San Francisco, CA.

David, Amy and Januszyk, Ross. "Advanced Planning Implementation at USG." APICS International Conference and Expo, October 2007, Denver, CO.

Other

Editorial Advisory Board, Management Decision, April 2014 - Present

Reviewer, Annals of Operations Research, May 2013

Member, AIChE, August 2013 - Present

Member, INFORMS, August 2010 - Present

APICS Certified in Production and Inventory Management, August 2004