OPTIMAL CAPACITY CHOICE AND ALLOCATION IN DECENTRALIZED SUPPLY CHAINS

Vinayak Deshpande

Krannert School of Management, Purdue University
W. Lafayette, IN 47907
e-mail: deshpandev@mgmt.purdue.edu

Leroy B. Schwarz

Krannert School of Management, Purdue University
W. Lafayette, IN 47907
e-mail: lee@mgmt.purdue.edu

September 2002
revised December 2002, July 2005

Abstract

We consider a supply chain in which a single supplier with fixed capacity sells to several independent retailers or manufacturers. The retailers have private information about their individual markets (e.g., mean market demand), which influences the size of their orders to the supplier. If the sum of all retailer orders exceeds the supplier's capacity, the supplier uses a pre-declared rule, which maps retailer's orders to allocations. A broad class of allocation mechanisms are prone to manipulation by retailers, as shown by Cachon and Lariviere (1999).

We first use a mechanism-design approach to obtain the optimal capacity-allocation rule and pricing mechanism for the supplier. We then answer the following questions: What level of capacity should the supplier provide to maximize its profit, given an implementable profit-maximizing allocation policy? What forms do optimal (i.e., supplier profit-maximizing) allocation polices take in typical business scenarios? Is supplier profit sensitive to the type of allocation policy it employs? What are practical ways of implementing the optimal allocation policy to reduce transaction costs? Our analysis examines business scenarios for which the linear and proportional allocation mechanisms are optimal. We also conclude that both supplier and supply-chain profit can increase significantly if a manipulable allocation policy is replaced by the optimal truth-telling allocation policy. Finally, in order to implement the optimal allocation rule, we design an auction mechanism wherein retailers submit purchasing cost bids for supplier capacity.

The authors acknowledge the support of the e-Enterprise center at Discovery Park, Purdue University

1 Introduction

Global competition is motivating supply chains to be as "lean" as possible, particularly with respect to inventory and capacity. This sometimes forces manufacturers or distributors to allocate inventory or capacity among their customers. Capacity shortages also commonly occur in situations in which capacity expansion is costly or when unexpected demand for a product (e.g., "hot" Christmas toys) or product components (e.g., DRAM's, LCD displays) temporarily outstrips supplier capacity or inventory. Under these circumstances, the supplier will often put its customers "on allocation".

Although specific practices vary, being "on allocation" generally means that customers get some or all of the quantity they ordered, the amount depending on metrics (e.g., past sales, days-of-supply availability) instead of their willingness to pay. The auto industry often uses fixed-price allocation schemes based, at least partially, on turn-and-earn. See (Lawrence 1996, Sawyers 1999). Fixed-price turn-and-earn allocation policies are considered "fair" by the automakers, but not by some auto dealers, as is evident from auto-dealer law suits (Freeman, 1997). Price increases are another approach used to handle tight capacity. For example, the price of flat-panel displays and DRAM chips rose by as much as 70% in the first quarter of 2002 due to capacity shortages (CNN report, March 2002). Another (seldom used) approach is to add capacity to meet the short-term increase in demand.

Management Science has provided some guidelines about the allocation of scarce capacity and inventory in centralized supply chains. Cachon and Lariviere (1999b) is the one of the first papers we know in the supply-chain literature to examine capacity allocation in a decentralized supply chain. They consider a system consisting of a single manufacturer and N independent retailers facing demand for a single product over a single time period. Each retailer places an order with the manufacturer based on that retailer's private information about its market. The manufacturer has fixed capacity and only a probability distribution on each retailer's market. Cachon and Lariviere's analysis, which takes the perspective of the retailers, provides several counter-intuitive results. For example, they prove that if the supplier uses a pre-announced system for allocating capacity at fixed prices, then the retailers will over-order even if a Pareto allocation policy, which maximizes total retailer profits under full information, is employed. Cachon and Lariviere also argue that supplier and supply-chain profit can increase if a truth-telling allocation mechanism is replaced by a manipulable one.

Our research examines the same basic model, but from the perspective of the supplier, who often has the ability to dictate policy when total retailer orders exceed its capacity or inventory. Our goal is to provide insights into the decentralized allocation problem from a supply-chain perspective. In the economics literature, Harris and Raviv (1981) and Maskin and Riley (1989) have analyzed supplier profit-maximizing allocation policies in a decentralized setting. Like Maskin and Riley (1989), we first characterize the optimal allocation policy that will maximize supplier profit under general assumptions about retailer revenue functions. This optimal allocation policy is shown to generate non-negative profits for the individual retailers and, hence, considered implementable.

We then examine the following questions:

1. What level of capacity should the supplier provide to maximize its profit, given an implementable profit-maximizing allocation policy?
2. What forms do optimal (i.e., supplier profit-maximizing) allocation polices take in typical business scenarios?
3. Is supplier profit sensitive to the type of allocation policy it employs?
4. Are there practical ways of implementing the optimal allocation policy to reduce transaction costs?

To answer these questions, like Cachon and Lariviere, we consider a one-period business scenario in which a single supplier with fixed capacity or inventory, sells to $N \geq 2$ independent retailers or manufacturers, henceforth labeled retailers, facing independent demands or requirements. Each retailer has private information about its own market (e.g., forecast of market demand, market potential) but not about the markets of any other retailer. The retailers are assumed to place orders based on a price schedule provided by the supplier, whose only information about each retailer's market is a probability distribution on a single information parameter. If the sum of retailer orders exceeds the available capacity or inventory, henceforth labeled available capacity, then the supplier uses a pre-announced policy for mapping retailer orders to their corresponding allocations. In contrast to Cachon and Lariviere, who assume fixed-price allocation, the optimal allocation policy uses variable prices. Indeed, it employs variable prices to induce truth-telling. In further contrast to Cachon and Lariviere, our analysis shows that both supplier and supply-chain profit can increase significantly if a manipulable allocation policy is replaced by the optimal truth-telling allocation policy. Section 2 of the paper provides a review of prior literature. We provide the model framework and analysis of the optimal allocation policy in section 3. We also derive the supplier optimal capacity in section 3. Section 4 provides structural results on the optimal allocation policy for two specific business scenarios. A numerical study with sensitivity analysis of the optimal allocation policy is provided in section 5, followed by an auction implementation procedure in section 6. We conclude with managerial insights in section 7.

2 Literature Review

Models for allocating inventory have a long history in supply-chain management. The earliest article, and, by far, the largest number or articles involve allocation in centrally managed systems. A growing literature examines allocation in decentralized systems, usually with information asymmetry. In what follows, we will cite the most closely related work.

For centrally managed supply chains, it is intuitively obvious that whenever allocations are necessary, they should be made so that the marginal expected consequences (e.g., cost, profit) are the same across all those entities demanding inventory. However, except for the simplest supply chains, it is quite difficult to identify efficient techniques to accomplish this goal. Allen (1958) showed that for a system of N non-identical newsvendors facing normally distributed demand with identical proportional shortage costs, that total expected system shortage cost is minimized by allocating inventory so that the retailers' normalized net inventories are equalized. If the retailers have identical customer-demand distributions, then such allocations equalize the retailers' stock-out probabilities. For identical retailers, of course, a balancing allocation equalizes retailer net inventories.

Models for more complicated systems either illustrate the computational difficulties when "balancing allocations" can't be achieved (e.g., Clark and Scarf, 1960) or assume that a "balancing allocation" - whose definition varies from model to model - can always be made. See, for example, Eppen and Schrage (1981), Federgreun and Zipkin (1984), Jonsson and Silver (1987), Schwarz (1989), Chen and Zheng (1994), Kumar, et al. (1995). Several authors (e.g., Topkis 1969, Ha 1997, Deshpande et al. 2002) have studied systems for allocating inventory with sequentially arriving customers of different priority classes in a centralized setting.

Several authors have also studied inventory-balancing itself. Most recently, McGavin, et al. 1997, examine the optimality of equalizing the retailers' normalized (with respect to the subsequent period's demand distributions) inventories in a one-warehouse N-retailer distribution system facing stochastic demand for a single product over T time periods. They demonstrate that this type of balancing is optimal for some cost functions for not for others.

More recently, several papers have examined the management of decentralized serial supply chains. Although serial systems do not require allocation per se - since there is only one buyer for each supplier - their analysis provides a framework for understanding allocation in more complex decentralized supply chains. For such systems, it is well known that double marginalization necessitates the use of some mechanism either to achieve first-best (i.e., the same performance level as a centralized system) or to improve performance for the buyer, supplier, or both.

Cachon and Lariviere (1999b) examine a capacity-allocation model identical to ours except that the retailers' per-unit purchase price is fixed. They demonstrate that such fixed-price allocation policies can lead the retailers to order more than they desire in order to receive a favorable allocation. In particular, Cachon and Lariviere prove that Pareto allocation mechanisms, which maximize total retailer profit under full information, induce retailer over-ordering if information is hidden. They also demonstrate several non-intuitive results, among them that "supplier and supply-chain profits can increase when a truth-inducing mechanism is replaced by a manipulable mechanism that creates order inflation." They further demonstrate that a truth-telling mechanism that maximizes total retailer profit does not exist if prices are fixed.

Our perspective is that of the supplier, and our focus is on the design of allocation policies that will maximize supplier profit. As expected, these policies involve variable prices. In particular, we develop an incentive-compatible price-quantity schedule for the supplier that induces retailer truth-telling, and, thereby, maximizes supplier profit. Hence, our analysis and insights are complementary to those of Cachon and Lariviere. For example, whereas Cachon and Lariviere demonstrate that fixed-price allocation policies are manipulable, we demonstrate that variable-price policies can induce truth telling. Other differences are described in our numerical study.

In related work, (1999a), Cachon and Lariviere examine "turn-and-earn" allocation, a fixed-price allocation scheme based on past sales. Cachon and Lariviere demonstrate, among other things, that turn-and-earn does not generally coordinate the system.

Lee and Whang (1996) extend Clark and Scarf's model of a centrally managed serial system to show that a mechanism based on consignment, delivery warranties, and an additional backorder-penalty cost at each supplier site can be used to achieve first-best performance using only locally available information. Porteus (2000) demonstrates that such a policy can be implemented using "responsibility tokens". Chen (1999) designs incentive schemes based on payments between echelon division managers and the firm rather than between echelon managers (as in Lee and Whang) for a serial system with information delays. It is important to note that all of these schemes require a central planner to both determine and enforce the coordinating mechanism.

Cachon and Zipkin (1999) examine an independently managed 2-stage serial system in which both the manufacturer and the retailer choose periodic-review base-stock policies that minimize their own individual expected costs; i.e., a 2-stage supply chain without a central planner to determine and enforce coordination as above. They examine two games, one based on tracking echelon inventory, the other based on individual inventories; and show that while these games almost always have a Nash equilibrium, it differs from that of the optimal (centralized) policy.

Cachon and Lariviere (2001) examine contracting between a manufacturer and a supplier that is the sole source of a critical component. They demonstrate that optimal supply-chain performance requires the manufacturer to truthfully share its initial forecast of customer demand with its supplier, but that it has an incentive to inflate its forecast so that the supplier will build more capacity. The supplier is aware of this bias, and, so, may not trust the manufacturer's forecast. See Tsay, Nahmias, and Agrawal (1999) and Cachon (2002) for a more complete review of supply-chain coordination.

Corbett (2001) examines contracting between a single buyer who places orders using a (Q,r) policy on an independently managed supplier under two different types of information asymmetry. Corbett's analysis concludes that in the absence of a central planner with full information, neither party can induce jointly optimal behavior without sacrificing its own profits. An adverse selection inventory model with demand information asymmetry in a single-period context has been considered by Mallik and Harker (1999).

Adverse selection models due to information asymmetry between principal and agents have been well studied in economics (See Fudenberg and Tirole, 2000). The classic references on adverse selection models and signaling include Akerlof 1970, Rothschild and Stiglitz 1976, and Spence 1974. Our model also draws on classical auction theory as described in the seminal papers by Vickrey (1961), Myerson (1981), Riley and Samuelson (1981), and Milgrom and Weber (1981). See Klemperer (1999) for a more recent review on the theory of auctions. The use of auctions for revenue management has been described by Vulcano, van Ryzin and Maglaras (2001).

The use of auctions for allocating resources such as securities is described by Harris and Raviv (1981). They derive the optimal allocation mechanism for a supplier, under the assumption of a unitary demand function, and a uniformly-distributed marginal willingness-to-pay for each retailer. An optimal auction procedure, where retailers submit quantity bids, was derived by Maskin and Riley (1989), under assumptions similar to ours. We attribute Theorem 2 of our paper, which describes the optimal allocation mechanism, to Maskin and Riley (1989).

However, our paper differs from Maskin and Riley (1989) in several ways: First, we also consider the supplier's capacity-choice problem conditional on the optimal allocation mechanism. Second, since our goal is to provide insights into the decentralized allocation problem from a supply-chain perspective, we describe the structure of the optimal allocation mechanism for several business scenarios. We numerically compare the performance of the optimal allocation mechanism to C&L's fixed-price manipulable allocation mechanism, in order to illustrate the benefit of the optimal allocation mechanism. We also compare the supply-chain profits under the optimal centralized and decentralized settings. Finally, Theorem 1 of our paper is a generalization of Maskin and Riley's results because it provides necessary and sufficient conditions for an implementable allocation policy.

Auctions are commonly used in procurement in which multiple suppliers submit bids to a single buyer. See Dasgupta and Spulber (1989/90) and the recent paper by Chen (2001) for an analysis of supply contracts for procurement. Elmaghraby (2000) provides an excellent review of procurement contracts. Our paper can be considered as an inverse of the procurement problem since we have a single seller selling to multiple buyers.

3 Model Framework and Analysis

In this section, we first formulate our model of a supply chain with one supplier with fixed capacity supplying $N \geq 2$ retailers in a single-period setting. We then derive the supplier's optimal pricing-and-allocation policy. The optimal capacity-choice decision is then derived for the supplier based on the optimal pricing-and-allocation policy.

3.1 Model formulation and the optimal allocation policy

Let K be the supplier's fixed capacity. The supplier recognizes that the order from each individual retailer depends on the price it charges and on retailer's private information parameter, which is hidden from the supplier. We model this private retailer information by the scalar parameter θ. Although θ is not known to the supplier, we assume that the supplier has a prior, with a density $f(\cdot)$ on the support $[\underline{\theta}, \overline{\theta}]$. Thus, the revenue function for retailer i, denoted $R_i(q_i, \theta_i)$, is a function of its allocated quantity q_i and its private information parameter θ_i. Retailer i observes its θ_i, but not the values of other retailers, labeled θ_{-i}. The retailer's profit equals its revenue $R_i(q_i, \theta_i)$ minus the purchasing cost (for q_i units) from the supplier.

We make the following assumptions for our analysis.

Assumption 1: Each retailer's revenue function, R_i, is concave in the quantity allocated to the retailer q_i. Hence $R_{QQ} \leq 0$.

Assumption 2: $R_\theta > 0, R_{\theta Q} > 0, \text{and} R_{\theta QQ} \geq 0.$

Assumption 3: Define $H(\theta) = \frac{1-F(\theta)}{f(\theta)}$ (reciprocal of Hazard rate). $H'(\theta) \leq 0$ and $R_{Q\theta\theta \leq 0}$.

Assumption 1 states that retailer's revenue function is marginally decreasing in its allocation. Assumption 2 states that higher θ indicates higher revenues for each retailer, all other things being equal; and that its marginal revenues (with respect to θ) for the retailer are increasing and convex in its allocation. Assumption 3 states that the inverse of the hazard rate is decreasing in θ. This assumption is satisfied by a broad class of probability distributions, such as the uniform, normal, exponential, logistic and chi-squared (see Barlow and Proschan, 1975). Also, the marginal revenues with respect to Q are increasing concave in θ.

We use the mechanism-design approach (Fudenberg and Tirole, 2000) to formulate the supplier's problem. In this direct revelation mechanism, the supplier asks the retailers to reveal their type θ_i, and implements the pricing-and-allocation policy $\{P(\theta_i), Q(\theta_i, \theta_{-i})\}$ based on the types revealed by the retailers. Here $Q(\theta_i, \theta_{-i})$ represents the quantity-allocation function while $P(\theta_i)$ represents the retailer purchasing cost based on the retailer's announcements of their individual type θ. For example, when the private information parameter is the forecast of market demand, each retailer is asked to announce its forecast to the supplier. The pricing-and-allocation policy is then implemented based on the forecast announcement by all retailers. The contracting steps are as follows.

1. The supplier offers a pricing-and-allocation policy $\{P(\theta_i), Q(\theta_i, \theta_{-i})\}$, linking the quantity allocation q_i and purchase cost p_i to the retailers announcement of their individual types.
2. Each retailer chooses a type to announce (which may or may not be its actual type θ_i) that maximizes its individual profit.
3. The quantity-allocation mechanism Q and the purchasing cost function P is implemented based on the set of types announced by all retailers.

Each retailer's decision is to choose a type to announce to the supplier, given the supplier's pricing-and-allocation policy. The objective for each retailer can then be stated as follows:

$$\max_x \pi_i(x, \theta_i) = E_{-i}[R(Q(x, \theta_{-i}), \theta_i) - P(x)] \tag{1}$$

The retailer's decision is to choose a parameter $x \in [\underline{\theta}, \overline{\theta}]$ to maximize its expected profits. Note that revenues are dependent on the amount allocated, $Q(x, \theta_{-i})$, and the retailer's type θ_i. The purchase cost also depends on the type announced $P(x)$. Note that the expectation is taken over the types of other retailers θ_{-i}. The mechanism $\{P(\theta_i), Q(\theta_i, \theta_{-i})\}$ is *incentive compatible* if it is optimal for each retailer to announce its true type to the supplier; i.e. $x^* = \theta_i$.

The supplier has a wide choice of pricing-and-allocation mechanisms to choose from. From the revelation principle (see Fudenberg and Tirole, 1991, p256.), the supplier can, without any loss in profits, restrict its attention to truth-telling mechanisms, where it is optimal for each retailer to announce his true information parameter θ_i. At the Bayes-Nash equilibrium of the direct revelation game, truth-telling by any retailer is a best response to truth-telling by all other retailers. A pricing-and-allocation policy is implementable if it is incentive compatible, and provides non-negative profits to the retailers. The following theorem states the condition under which a mechanism $\{P(\theta_i), Q(\theta_i, \theta_{-i})\}$ is implementable.

Theorem 1 *The pricing-and-allocation mechanism* $\{P(\theta_i), Q(\theta_i, \theta_{-i})\}$ *is incentive compatible* **if and only if** *conditions (2) and (3) hold. In addition, condition (3) is a sufficient* *condition to guarantee non-negative profits for the retailers.*

$$\int_x^{\theta_i} E_{-i}R_\theta(Q(x,\theta_{-i}),\theta)d\theta \le \int_x^{\theta_i} E_{-i}R_\theta(Q(\theta,\theta_{-i}),\theta)d\theta \tag{2}$$

$$P(\theta_i) = E_{-i}[R(Q(\theta_i,\theta_{-i}),\theta_i) - \int_{\underline{\theta}}^{\theta_i} R_\theta(Q(\theta,\theta_{-i}),\theta)d\theta] \tag{3}$$

Proof: Provided in the appendix.

Note that Theorem 1 provides necessary *and* sufficient conditions for an allocation policy to be implementable. The necessary condition for implementability, equation (3), was first established by Maskin and Riley (1989). A sufficient condition for (2) is $R_{\theta Q}Q' \ge 0$. Since $R_{\theta Q} > 0$ (assumption 2), note that condition (2) is always satisfied if $Q(\cdot)$ is non-decreasing in θ. We will assume throughout the paper that $Q(\cdot)$ is non-decreasing in θ. As shown later, assumptions 2 and 3 are sufficient to guarantee that the optimal allocation is non-decreasing in the announcement θ.

Let R_i^s denote the expected payments received by the supplier from retailer i. Then

$$R_i^s = E_i(P(\theta_i)) = E[R(Q(\theta_i,\theta_{-i}),\theta_i) - R_\theta(Q(\theta_i,\theta_{-i}),\theta_i)H(\theta_i)]$$

Thus, the supplier's allocation problem can be stated as:

$$V(K) = \max_Q E[\sum_{i=1}^n R(Q(\theta_i,\theta_{-i}),\theta_i) - R_\theta(Q(\theta_i,\theta_{-i}),\theta_i)H(\theta_i)] \tag{4}$$

subject to

$$\sum_{i=1}^n Q(\theta_i,\theta_{-i}) \le K \tag{5}$$

Under the stated assumptions, the KKT first-order conditions are sufficient for the supplier's problem. The following theorem establishes the solution to the optimal allocation policy for the supplier.

Theorem 2 *(Maskin and Riley, 1989) The optimal allocation policy $Q^*(\theta_i,\theta_{-i})$ is obtained as the solution to the following conditions*

$$R_Q(Q(\theta_i,\theta_{-i}),\theta_i) - R_{Q\theta}(Q(\theta_i,\theta_{-i}),\theta_i)H(\theta_i) = \lambda \tag{6}$$

$$\lambda(\sum_{i=1}^n Q(\theta_i,\theta_{-i}) - K) = 0 \tag{7}$$

$$\lambda \ge 0 \tag{8}$$

Proof: See Maskin and Riley (1989).

Conditions (6)-(8) may be interpreted as follows. Equation (6) states that the marginal revenues for each retailer minus the marginal information costs for the supplier should be equal to the shadow price of the fixed capacity. Equations (7) and (8) state that the shadow price of the supplier's fixed capacity is positive if capacity is tight, and zero otherwise. Notice that a Pareto allocation policy, which maximizes sum of all retailer revenues, would allocate capacity to equalize marginal revenues for all retailers. Hence, the presence of marginal information rents distinguishes the supplier's optimal allocation policy from the Pareto allocation policy.

Note that in the above formulation we have ignored the non-negativity constraints on the allocation $Q(\cdot)$. In many situations it may be desirable to add this constraint to the supplier's problem.

$$Q(\theta_i, \theta_{-i}) \geq 0 \tag{9}$$

Correspondingly, equation (6) should be modified to:

$$R_Q(Q(\theta_i, \theta_{-i}), \theta_i) - R_{Q\theta}(Q(\theta_i, \theta_{-i}), \theta_i)H(\theta_i) = \lambda - \mu_i \tag{10}$$

where

$$\mu_i \geq 0 \quad \text{and} \quad \mu_i Q(\theta_i, \theta_{-i}) = 0 \tag{11}$$

In deriving the optimal allocation policy, we had assumed that $Q^*(\cdot)$ is increasing in the information parameter θ. A sufficient condition for $Q^*(\cdot)$ to be increasing in θ is $V(K)$ be a supermodular function (i.e, $\delta^2 V/\delta Q \delta\theta \geq 0$). Now

$$\frac{\delta^2 V}{\delta Q \delta \theta} = R_{Q\theta}(Q(\theta_i, \theta_{-i}), \theta_i) - H(\theta)R_{Q\theta\theta}(Q(\theta_i, \theta_{-i}), \theta_i) - H'(\theta)R_{Q\theta}(Q(\theta_i, \theta_{-i}), \theta_i) \geq 0$$

Hence, assumptions 2 and 3 guarantee that $V(\cdot)$ is supermodular, implying that the Q^* obtained from conditions in Theorem 2 is increasing in θ.

3.2 Optimal Capacity Choice

In the first stage of the problem, the supplier chooses its capacity, anticipating how the retailers will behave under the optimal pricing-and-allocation policy. We assume that the supplier incurs a unit cost of c for investing in capacity. Given the optimal pricing-and-allocation policy, the supplier's capacity-choice problem can be stated as follows.

$$\max_K V(K) - cK \tag{12}$$

where

$$V(K) = \max_{Q} E[\sum_{i=1}^{n} R(Q(\theta_i, \theta_{-i}), \theta_i) - R_\theta(Q(\theta_i, \theta_{-i}), \theta_i) H(\theta_i)] \quad (13)$$

subject to (5) and (9).

The following theorem establishes the structure of the supplier's profit function $V(K)$.

Theorem 3 *The supplier's revenue function $V(K)$ is concave in its capacity choice K.*

Proof: Under assumptions 1-3, it is easy to show that the supplier's objective is to maximize a concave function subject to a linear constraint on the capacity choice. Hence $V(K)$ is concave in the capacity choice K. □

The implication of Theorem 3 is that the optimal capacity can be easily found by a simple search mechanism such as golden section search. Let K^* be the supplier's optimal capacity choice. The following theorem compares this capacity choice under information asymmetry to the centralized capacity choice under no-information asymmetry.

Theorem 4 *The supplier's optimal capacity choice, K^*, under information asymmetry is less than or equal to the capacity choice K_c^* which maximizes centralized supply-chain profits (with no information asymmetry).*

Proof: Provided in the appendix.
The intuition on Theorem 4 is as follows: Given truth-telling, by under-investing in capacity, the supplier creates increased competition among the retailers, thereby squeezing information rents from them.

4 Structural results for the optimal allocation policy

In this section we explore the structure of the optimal allocation policy. In particular, we are interested in conditions under which linear and proportional allocation mechanims are optimal for the supplier. For each of these allocation mechanisms we then describe stylized business scenarios which satisfy these optimality conditions. We also introduce a generalized version of linear allocation, and describe conditions under which it is optimal.

4.1 The Linear Allocation Mechanism

The Linear allocation mechanism is defined as follows:

Definition 1 Linear Allocation Mechanism *Index the retailers in decreasing order of their order quantities,i.e.* $q_1 \geq q_1 \geq ... \geq q_N$*. Retailer* i *is allocated* $Q_i(q, \underline{n})$*, where*

$$Q_i(q, \underline{n}) = q_i - \frac{1}{\underline{n}} \max\{0, \sum_{i=1}^{n} q_i - K\} \quad \text{if} \quad i \leq \underline{n}$$

$$Q_i(q, \underline{n}) = 0 \quad \text{if} \quad i > \underline{n}$$

where $\underline{n}$ *is the largest integer such that* $Q_i(q, \underline{n}) \geq 0$ *for all* i*.*

The linear allocation mechanism awards each retailer his order minus a common deduction. The common deduction (defined for retailers who get positive allocations) is equal to the difference by which sum of retailer orders exceeds capacity divided by the number of retailers who get positive allocations.

Also, define $G(Q, \theta_i) = R_Q(Q, \theta_i) - R_{Q\theta}(Q, \theta_i)H(\theta_i)$

Theorem 5 *If there exists a function* $\tau(\theta_i)$ *such that* $\tau(\theta^*) = 0$ *for a fixed type* θ^* *and* $G(Q, \theta_i) = G(Q - \tau(\theta_i), \theta^*)$*, then the linear allocation policy is optimal for the supplier; i.e.,*

$$Q_i^*(\theta_i) = Q_i^*(\theta^*) + \tau(\theta_i) = \frac{K}{\underline{n}} + \tau(\theta_i) - \frac{\sum_{i=1}^{n} \tau(\theta_i)}{\underline{n}}$$

Note that the above condition on $G(\cdot, \cdot)$ is satisfied if the marginal revenues satisfy the property $R_Q(Q, \theta_i) = R_Q(Q - \tau(\theta_i), \theta^*)$, and the inverse Hazard rate satisfies $H(\theta_i) = \frac{1}{\tau'(\theta_i)}$.

Theorem 5 can be viewed as a refinement of Theorem 2, where it adds sufficient conditions under which the linear allocation rule is optimal. Theorem 5 provides the intuition that if the information-rent adjusted marginal revenues are linearly shiftable functions of the parameter θ then the linear allocation scheme is optimal.

Example 1: Retailers face downward-sloping linear demand

Consider the problem when "symmetric" (i.e., identical cost structures and demand functions) retailers face a downward sloping linear demand, i.e.,

$$r(q) = \theta - q \tag{14}$$

where $r(q)$ is the market-clearing price charged by retailers to sell q units. The intercept of the demand curve is retailer's private information and not known to the supplier. Thus, for a given allocation q, the retailer's revenue function is given by

$$R(q, \theta) = q(\theta - q) \tag{15}$$

It is easy to verify that this revenue function satisfies assumptions 1 and 2. The following corollary to Theorem 5, as shown by Maskin and Riley (1989), establishes the optimal allocation policy for the supplier.

Corollary 1 *(Maskin and Riley, 1989) If retailers face deterministic downward sloping linear demand, then the linear allocation mechanism is optimal for the supplier.*

It is interesting to note that in this business scenario - but with fixed prices - Cachon and Lariviere prove that linear allocation is manipulable. It is also interesting to note that in their numerical study, under the assumption that the supplier has chosen its capacity based on truth-telling, Cachon and Lariviere observe supplier profit increases as much as 19% if uniform allocation - which induces truth-telling - is replaced by linear allocation, which induces over-ordering. In our numerical study, using an incentive-compatible linear allocation policy, supplier profits increase a minimum of 842% over the non-incentive compatible implementation of the linear-allocation policy considered by Cachon and Lariviere. Thus, the optimal pricing-and allocation policy leads to dramatic increase in profits for the supplier compared to the policy considered by Cachon and Lariviere.

Example 2: Retailers are Linearly Shiftable Newsvendors

Consider the scenario in which symmetric retailers are newsvendors. It is assumed that one of the parameters of the newsvendor demand distribution is each retailer's private information(e.g., mean demand). Let $f(x, \theta)$ denote the demand distribution for the newsvendor, with θ over the range $[\underline{\theta}, \overline{\theta}]$. The revenue function for the retailers based on the quantity allocated Q and the information parameter θ is given by

$$R(Q, \theta) = r[\int_{\infty}^{Q} x f(x, \theta) dx + \int_{Q}^{\infty} Q f(x, \theta) dx] \tag{16}$$

Theorem 5 tells us that if the newsvendor demand distribution belongs to the family of distributions that can be linearly shifted by a function of θ, then the linear allocation policy is optimal for the supplier. The linear shiftability condition is equivalent to the marginal revenues for the retailers differing by a linear parameter shift. In addition, the condition on H guarantees that the information rents are also linearly shiftable by the parameter θ. Examples of such distributions includes the normal distribution.

Corollary 2 *If retailers are newsvendors with a normal demand distribution with mean θ, and an exponential prior on θ, then the conditions of Theorem 5 are satisfied. Hence, the linear allocation mechanism is optimal.*

The linear allocation mechanism can be generalized as follows:

Definition 2 Generalized Linear Allocation Mechanism *Index the retailers in decreasing order of their order quantities,i.e. $q_1 \geq q_1 \geq ... \geq q_N$. If capacity is tight, Retailer i is allocated $Q_i(q, \underline{n})$, where*

$$Y(Q_i(q, \underline{n})) = Y(q_i) - \lambda$$

$$Q_i(q, \underline{n}) = 0 \quad \textit{if} \quad i > \underline{n}$$

where λ is a constant, and $\underline{n}$ is the largest integer such that $Q_i(q, \underline{n}) \geq 0$ for all i.

Generalized linear allocation mechanism first transforms the order quantities using the function $Y(\cdot)$, and then applies the linear allocation policy. Note that the linear allocation mechanism is a special case of the generalized-linear allocation mechanism, where $Y(q) = q$.

Example 3: Retailers face general downward-sloping demand

Consider the problem when symmetric retailers face general downward-sloping demand, i.e.

$$r(q) = \theta - v(q) \tag{17}$$

where r is the market-clearing price charged by retailers to sell q units, and $v(\cdot)$ is an increasing function of q. The intercept of the demand curve is retailer's private information and not known to the supplier. Thus, for a given allocation q, the retailer's revenue function is given by

$$R(q, \theta) = q(\theta - v(q)) \tag{18}$$

The following corollary follows immediately:

Corollary 3 *If retailers face general deterministic downward-sloping demand, then the generalized linear allocation mechanism is optimal for the supplier, with $Y(q) = v(q) + qv'(q)$.*

For example, if the market clearing-price has the form $r(q) = \theta - q^2$, then a generalized linear allocation mechanism with $Y(q) = 2q^2$ is optimal for the supplier.

4.2 The Proportional Allocation Mechanism

Definition 3 Proportional Allocation Mechanism *Retailer i is allocated $Q_i(q)$ where*

$$Q_i(q) = \min\{q_i, \frac{Kq_i}{\sum_{i=1}^{N} q_i}\}$$

The proportional allocation policy gives each retailer a common fraction of his order if capacity binds. The common fraction is equal to the total capacity divided by sum of all retailer orders.

Theorem 6 *If there exists a function $\tau(\theta_i)$ such that $\tau(\theta^*) = 1$ for a fixed type θ^* and $G(Q, \theta_i) = G(\frac{Q}{\tau(\theta_i)}, \theta^*)$, then the proportional allocation policy is optimal for the supplier; i.e,*

$$Q_i^*(\theta_i) = \tau(\theta_i) Q_i^*(\theta^*) = \frac{\tau(\theta_i)K}{\sum \tau(\theta_i)}$$

Note that the above condition on $G(\cdot, \cdot)$ is satisfied if $R_Q(Q, \theta_i) = R_Q(\frac{Q}{\tau(\theta_i)}, \theta^*)$, and $H(\theta_i) = \frac{\tau(\theta_i)}{\tau'(\theta_i)}$.

Example 4: Retailers are Scalable Newsvendors

Theorem 6 states that if the newsvendor demand distribution belongs to the family of distributions that can be scaled by a function of θ, then the proportional allocation policy is optimal for the supplier. The uniform, exponential and gamma distributions, for example, all satisfy this condition. The scalability condition states that the marginal revenues for retailers differ by a scale parameter. In addition, the condition on H guarantees that the information rents are also scalable by the parameter θ.

Corollary 4 *If retailers are newsvendors with a uniform demand distribution on $[0, \theta]$, and a Pareto supplier's prior on θ, then the conditions of Theorem 6 are satisfied. Hence, the proportional allocation mechanism is optimal.*

5 Numerical study

The starting point for our numerical study is a comparison of the performance of the optimal (i.e., supplier profit-maximizing) variable-price linear-allocation policy with the manipulable fixed-price allocation policy examined by Cachon and Lariviere (1999b) in their numerical study (Table 2, page 1102). We then examine the sensitivity of the optimal allocation policy to different model parameters. The initial parameters used are identical to those used by Cachon and Lariviere (1999b). Specifically, we assume N symmetric retailers (i.e., identical cost structures and linear demand function $Q(P) = \theta - P$), where P is the unit price each retailer charges its customers. Each retailer's market potential, θ, is that retailer's private information. $N = 5$ retailers are independently assigned to be one of five types $\theta \in \{4, 5, 6, 7, 8\}$, with probabilities $\{0.05, 0.25, 0.4, 0.25, 0.05\}$, respectively.

For each parameterization, let π_C represent supply-chain profit for the centralized system, for which a Pareto allocation policy is optimal (see Cachon and Lariviere, page 1095); let K_C be the corresponding optimal centralized supplier capacity. Let $\pi_{Di}, i = T, M$, represent the supply-chain profit for the decentralized system under allocation policy i, where $i = T$ is the optimal truth-telling variable-price linear-allocation policy, and $i = M$ is the manipulable fixed-price linear-allocation policy considered by Cachon and Lariviere. Let $K_{Di}, i = T, M$ be the corresponding optimal supplier capacity; and let $\pi^S_{Di}, i = T, M$, represent the decentralized supplier profit.

The performance measures we examine are: (1) the decentralization penalty, $\pi^i_{Penalty} = 100(\pi_C - \pi_{Di})/\pi_C$; that is, the percentage reduction in maximum possible supply-chain profit from using allocation policy $i = T$ or M in a decentralized system; (2) the supplier percentage of supply-chain profit, $\pi^{Di}_{Supplier} = 100\pi^S_{Di}/\pi_{Di}$; and (3) the decentralized supplier capacity ratio, $K_{Di/C} = 100K_{Di}/K_C$.

5.1 Comparing the Optimal and the Manipulable Allocation Mechanisms

Columns 2 and 3 of Table 1 provide the optimal centralized profit, π_C, and supplier capacity, K_C, respectively, for the supplier's per-unit capacity cost, c, specified in Column 1. Columns 4-6 report the performance of the optimal (i.e., supplier profit-maximizing) variable-price linear-allocation policy; columns 7-9, the performance of Cachon and Lariviere's manipulable fixed-price linear-allocation policy. Column 10 reports the percentage increase in supplier profit in changing from the manipulable policy to the decentralized optimal policy.

Most important, observe (Col. 10) that the optimal allocation policy increases the supplier's profit a minimum of 842% over the manipulable policy considered by Cachon and Lariviere. More specifically, Column 5 reports that the optimal allocation policy provides the supplier with 74.1 to 79.1% of total supply-chain profits (an average of 76.9%) versus the manipulable scheme's share of 0.0 to 8.3% (an average of 1.3%) reported in Column 8. In addition, note (in columns 4 and 7) that in moving from the manipulable to the truth-telling policy, the decentralization penalty is reduced in 13 and increased in 22 of the 35 parameterizations. Since in the optimal allocation policy, the supplier maximizes his individual profit, the supply chain could be worse off as compared to the C&L policy. However, in 37% of the cases considered, both the supplier *and* the supply chain are better off by using the optimal allocation policy as compared to the C&L manipulable policy. These results show that manipulable fixed pricing policies can significantly deteriorate supplier profits.

Recall that Theorem 4 establishes that a profit-maximizing supplier will always provide less capacity than under a centralized system; that is, $K_{DT} \leq K_C$ (or, equivalently, $K_{DT/C} \leq 1$). Col. 6 reports that optimal supplier capacity percentage ranges from 64.7 to 89.3% (an average of 77.2%). Cachon and Lariviere observe that, under the manipulable policy, the supplier has an incentive to provide less capacity than a centralized policy. Col. 9 reports that $K_{DM/C}$ ranges from 46.0 to 90.0% (an average of 72.1%). In what follows, we examine the sensitivity of the optimal allocation policy to different model parameters.

5.2 The Impact of Capacity Cost on the Performance of the Optimal Allocation policy

Everything else being equal, one would expect that as per-unit capacity cost increases, optimal capacity should decrease under either a centralized system or decentralized system. Table 1 confirms this expectation. Observe, in particular (Col. 3) that centralized optimal capacity, K_C, decreases from 15.42 to 5.99. More important, since $K_{DT/C} \leq 1$ (Col. 6), K_{DT} decreases more quickly. This is because, in the decentralized system, reduced capacity increases the supplier's information rents.[1] As a consequence, the decentralization penalty

[1] In particular, note that if capacity were free, then the supplier would provide unlimited capacity, always satisfy total retailer orders, and experience no information rents.

	Centralized Optimal		Decentralized Optimal			Cachon Linear allocation			Comparison Decentralized vs C&L
c	π_C	K_C	$\pi^T_{Penalty}$	$\pi^{DT}_{Supplier}$	$K_{DT/C}$	$\pi^M_{Penalty}$	$\pi^{DM}_{Supplier}$	$K_{DM/C}$	Increase in supplier profit $\frac{(\pi^S_{DT}-\pi^S_{DM})}{\pi^S_{DM}}$
0.1	44.53	15.42	8.38%	79.11%	89.28%	2.40%	0.00%	90.00%	Infinity
0.15	43.76	15.09	8.49%	78.90%	88.31%	5.40%	0.00%	85.00%	Infinity
0.2	43.02	14.87	8.62%	78.72%	86.84%	9.80%	0.00%	77.00%	Infinity
0.25	42.28	14.64	8.72%	78.53%	86.11%	16.60%	0.12%	71.00%	71580%
0.3	41.55	14.44	8.88%	78.40%	84.89%	1.60%	0.10%	90.00%	71197%
0.35	40.83	14.27	8.97%	78.22%	84.36%	31.20%	0.00%	53.00%	71207644%
0.4	40.12	14.09	9.13%	78.11%	83.39%	41.40%	0.00%	46.00%	70984092%
0.45	39.42	13.94	9.24%	77.97%	82.69%	3.90%	0.10%	84.00%	70526%
0.5	38.72	13.84	9.36%	77.84%	81.84%	1.00%	0.50%	89.00%	14023%
0.6	37.36	13.53	9.53%	77.54%	81.14%	7.20%	0.11%	78.00%	69914%
0.7	36.02	13.26	9.77%	77.32%	80.10%	1.10%	0.41%	89.00%	17290%
0.75	35.36	13.12	9.91%	77.24%	79.55%	11.70%	0.11%	71.00%	69068%
0.9	33.41	12.77	10.27%	76.95%	78.05%	17.70%	0.12%	62.00%	68129%
1	32.15	12.53	10.49%	76.76%	77.25%	4.10%	0.73%	80.00%	9723%
1.05	31.53	12.35	10.55%	76.62%	77.41%	25.40%	0.27%	54.00%	33763%
1.2	29.70	12.00	10.92%	76.40%	76.22%	35.10%	0.47%	46.00%	22285%
1.25	29.10	11.86	11.25%	76.50%	75.46%	7.00%	0.97%	73.00%	7456%
1.35	27.93	11.62	11.50%	76.37%	74.80%	3.80%	1.46%	85.00%	4722%
1.4	27.35	11.51	11.67%	76.34%	74.30%	3.30%	0.73%	83.00%	9493%
1.5	26.21	11.25	11.86%	76.17%	74.00%	11.10%	1.35%	65.00%	5504%
1.75	23.48	10.62	12.49%	75.92%	72.87%	5.20%	0.95%	79.00%	7239%
1.8	22.95	10.51	12.86%	76.08%	71.94%	5.10%	1.90%	81.00%	3579%
2	20.90	10.00	13.23%	75.79%	71.44%	24.50%	3.30%	50.00%	2541%
2.1	19.91	9.77	13.67%	75.89%	70.48%	8.10%	1.32%	72.00%	5321%
2.25	18.48	9.32	14.07%	75.81%	70.25%	6.90%	2.58%	77.00%	2612%
2.45	16.65	8.85	14.75%	75.86%	68.77%	10.90%	1.93%	58.00%	3667%
2.7	14.51	8.23	15.16%	75.57%	67.81%	9.50%	3.54%	73.00%	1902%
2.8	13.70	8.02	15.33%	75.48%	67.13%	18.60%	3.00%	56.00%	2521%
3.15	11.05	7.12	15.65%	74.94%	66.13%	13.50%	5.21%	67.00%	1302%
3.6	8.10	5.99	15.81%	74.11%	64.71%	20.00%	8.28%	58.00%	842%

Table 1: Impact of capacity acquisition cost c

increases from 8.38 to 15.81% (column 4).

5.3 The Impact of the Number of Retailers, N, on the Optimal Allocation policy

In order to examine the impact of the number of retailers, N, on the performance of the optimal allocation policy we examined parameterizations of $N = 2, 3, 4$, and 5 retailers with uniformly-distributed market potential, θ. Table 2 summarizes the performance/retailer for three different capacity costs; i.e., c = \$0.10, \$1.85, and \$3.60/unit. By assumption, as the number of retailers increases, the total market potential increases. In the centralized system this provides the opportunity to extract higher profit/retailer, π_C/N, from the marketplace, as reported in Col. 3. In the decentralized system with truth-telling, this increase in total market potential increases competition among the retailers, which permits the supplier to reduce the information rent/retailer, thereby increasing its share of supply-chain profits as N increases, as confirmed in Col. 5, and reducing the decentralization penalty, $\pi^T_{Penalty}$, as confirmed in Col. 4. Finally, note (Col. 6) that the optimal supplier capacity ratio, $K_{DT/C}$, decreases as N increases.

c	N	Centralized Profit/ret π_C/N	Decentralization Penalty $\pi^T_{Penalty}$	Suppliers share $\pi^{DT}_{Supplier}$	Capacity ratio $K_{DT/C}$
0.1	2	9.14	16.40%	73.94%	94.02%
0.1	3	9.15	16.36%	74.18%	88.37%
0.1	4	9.16	16.31%	74.31%	85.09%
0.1	5	9.16	16.29%	74.41%	82.74%
1.85	2	4.56	24.90%	63.99%	63.32%
1.85	3	4.64	22.87%	67.64%	61.30%
1.85	4	4.68	21.71%	69.50%	60.71%
1.85	5	4.71	20.80%	70.48%	60.98%
3.6	2	1.69	30.68%	48.71%	48.99%
3.6	3	1.77	25.58%	61.26%	51.70%
3.6	4	1.81	23.37%	66.21%	53.49%
3.6	5	1.84	22.45%	69.30%	54.48%

Table 2: Impact of number of retailers N

5.4 The Impact of Increasing Mean Market Potential

Given the initial parameterization of $N = 5$ retailers with market potential drawn from a uniform distribution, $\theta \in \{4, 5, 6, 7, 8\}$, we examined the effect of increasing mean market potential by increasing the lower bound on θ from 2 to 6 (with corresponding changes in the upper bound) while keeping the number of states and their ordinal probabilities fixed. Hence, in the first experiment $\theta \in \{2, 3, 4, 5, 6\}$, with mean of 4, while in the last $\theta \in \{6, 7, 8, 9, 10\}$, with a mean of 8. Table 3 provides the results for capacity cost c = \$0.10, \$1.85, and \$3.60/unit.

In the centralized system, one would expect increasing mean market potential to provide higher supply-chain profit and increased supplier capacity. This was confirmed (but not reported here). In the decentralized system with truth-telling, although the supplier experiences the same uncertainty about retailer type, the larger market potential/retailer reduces its information rent, thereby reducing the decentralization penalty, $\pi^T_{Penalty}$, as confirmed in Col. 3; and increasing its percentage of total supply-chain profit, $\pi^T_{Supplier}$, as confirmed in Col. 4. Note, further (Col. 5) that $K_{DT/C}$ increases as mean market potential increases for fixed per-unit capacity cost, c.

5.5 The Impact of Increasing the Variance of Market Potential

Given an initial parameterization of $N = 5$ retailers with five uniformly-distributed values of market potential, $\theta \in \{4, 5, 6, 7, 8\}$, we examined the effect of market-potential variance by varying the number of θ's potential states from 3 to 6 while maintaining the same mean market potential. More specifically, the first parameterization had 3 states, $\theta \in \{5, 6, 7\}$, while the last had 9 states, $\theta \in \{2, 3, 4, 5, 6, 7, 8, 9, 10\}$. Table 4 provides the results for capacity cost c = \$0.10, \$1.85, and \$3.60/unit.

Since in the centralized system the θ's of the retailers are known, increasing the range of θ's provides the centralized system access to higher revenue potential for the same capacity level. Hence, centralized capacity, K_C, and centralized supply-chain profit π_C, should increase. This is confirmed in Cols. 3 and 4. However, increasing the range of θ's increases information rents in the truth-telling decentralized system, thereby increasing the decentralization penalty, $\pi^T_{Penalty}$, (Col. 5) and reducing the supplier's share of supply-chain profit, $\pi^T_{Supplier}$, (Col. 6). Note (Col. 7) that the supplier capacity ratio, $K_{DT/C}$ also decreases.

c	mean demand	Decentralization penalty $\pi^T_{Penalty}$	Suppliers share $\pi^{DT}_{Supplier}$	Capacity ratio $K_{DT/C}$
0.1	4	26.18%	77.04%	78.76%
0.1	5	21.43%	74.33%	80.78%
0.1	6	19.46%	74.41%	82.74%
0.1	7	13.79%	72.78%	84.67%
0.1	8	10.28%	72.83%	86.88%
1.85	4	31.58%	72.13%	56.46%
1.85	5	28.38%	70.41%	57.92%
1.85	6	26.26%	70.48%	60.98%
1.85	7	21.11%	69.75%	64.28%
1.85	8	15.91%	69.47%	68.74%
3.6	4	30.05%	66.40%	54.42%
3.6	5	29.34%	69.03%	53.99%
3.6	6	28.94%	69.30%	54.48%
3.6	7	27.00%	69.12%	56.46%
3.6	8	23.50%	69.04%	60.54%

Table 3: Impact of retailer mean demand

c c	number of states	π_C	K_C	Decentralization Penalty $\pi^T_{Penalty}$	Suppliers share $\pi^{DT}_{Supplier}$	Capacity ratio $K_{DT/C}$
0.1	3	44.26	15.28	4.69%	78.51%	90.61%
0.1	5	45.82	16.06	16.29%	74.41%	82.74%
0.1	7	48.21	16.88	20.24%	72.77%	78.33%
0.1	9	51.45	17.64	22.51%	72.83%	76.54%
1.85	3	22.19	10.37	8.67%	74.34%	75.90%
1.85	5	23.53	10.37	20.80%	70.48%	60.98%
1.85	7	25.53	10.42	24.16%	68.32%	57.29%
1.85	9	28.21	10.64	26.33%	68.05%	55.52%
3.6	3	7.87	6.01	16.72%	75.59%	64.18%
3.6	5	9.20	5.98	22.45%	69.30%	54.48%
3.6	7	11.12	6.24	25.26%	65.59%	50.98%
3.6	9	13.51	6.50	27.27%	63.63%	51.38%

Table 4: Impact of retailer demand variance

6 Auction mechanisms for implementing the optimal allocation policy

A common criticism of variable-price allocation mechanisms is the potentially high transaction costs imposed by such policies (Cachon and Lariviere, 1999b). In this section, we first prove that an auction based on purchasing cost bids by retailers can be used to implement the optimal allocation policy. We then describe an example to illustrate its potential implementation on the internet.

The auction is based on retailers valuation (revenue) of quantity allocated by the supplier. Retailers submit purchasing cost bids on the fixed capacity K (in a multi-unit auction), rather than the supplier dictating a price schedule. Retailers recognize the fixed capacity K and "probabilistically" anticipate other retailers valuations of allocated quantity.

The mechanics of the auction are as follows:

1. The supplier announces an output allocation mechanism conditional on purchasing-cost bids submitted by the retailers.
2. Individual retailers submit bids to the supplier which maximize their individual profits. Individual profits take into account the individual retailers private information, the suppliers announced allocation policy, and an expectation over other retailers bids.
3. Based on the bids submitted by the retailers, the announced allocation policy is implemented and the retailers pay their announced bid to the supplier.

Theorem 7 *Let $Q^*(\theta)$ be the optimal allocation mechanism. Then the following auction implements the allocation Q^*.*

1. *The supplier announces an output allocation mechanism $A(V_i, V_{-i})$ based on the purchasing cost bids submitted by retailers.*
2. *The retailers submit bids V_i^* which maximize expected revenues minus purchasing cost, anticipating allocation mechanism A.*
3. *Allocation mechanism A is implemented based on bids submitted, where*

$$A(V_i, V_{-i}) = Q^*(\theta^*(V_i), \theta^*(V_{-i})) \quad \text{and} \quad \theta^*(V) = P^{*-1}(V)$$

We illustrate the auction mechanism with the following example. Consider the 2-retailer scenario, with downward sloping linear demand, as described in section 4.1, and the following parameters. Retailers are independently assigned to be one of five types $\theta \in \{4, 5, 6, 7, 8\}$ from a uniformly distributed prior. Given capacity-acquisition cost is \$1.85 per unit, the optimal capacity for this example is 2.63 units. The supplier posts this capacity and output-allocation mechanism represented in Table 5.

Retailer 2 bids	0.00	0.70	3.05	7.08	12.29
Retailer 1 bids					
0.00	0.00	0.00	0.00	0.00	0.00
0.70	1.00	1.00	0.82	0.32	0.00
3.05	2.00	1.82	1.32	0.82	0.32
7.08	2.63	2.32	1.82	1.32	0.82
12.29	2.63	2.63	2.32	1.82	1.315

Table 5: Allocation Table for Auction Implementation

The interpretation of this table is as follows. If retailer 1 bids \$0, he gets 0 units of capacity, independent of how much player 2 bids. If retailer 1 bids \$12.29, then he gets all of the 2.63 units of capacity, if retailer 2 bids zero. However, if both retailers bid \$12.29, then each receives 1.315 units (half of the available capacity). Note that the allocation table has the intuitive property that the quantity received by retailer 1 is non-decreasing in retailer 1's bids, but non-increasing in retailer 2's bids.

Each retailer, after examining the allocation rule, submits a bid that is optimal for itself. The allocation table has been constructed such that a retailer with $\theta = 4$ will submit a bid of \$0, while a retailer with $\theta = 8$ will submit a bid of \$12.29. Thus, each retailer's bid indirectly reveals its type, thereby enabling the supplier to implement the optimal quantity-allocation policy.

Although, the above example is fairly simple, a similar auction scheme can be easily posted on the web for the N retailer case. This enables the implementation of the optimal quantity allocation policy, while keeping transaction costs low.

7 Conclusion

Our goal in examining capacity allocation in a decentralized system from the perspective of the supplier was to answer the questions posed in the introduction. The answers provided by our analysis are as follows:

1. *What level of capacity should the supplier provide to maximize its profit, given an implementable profit-maximizing allocation policy?*

   Given an optimal allocation policy, Theorem 3 establishes the concavity of the supplier's profit-maximizing capacity objective under Assumptions 1-3, thereby facilitating the efficient numerical search for the supplier's profit-maximizing capacity. Further, Theorem 4 proves that a profit-maximizing supplier's optimal capacity, K^*, in the decentralized system (with information asymmetry) will always be less than or equal to the capacity which maximizes supply-chain profit in a centralized system (with symmetric information). The optimal capacity for specific business scenarios was numerically computed in section 5.

2. *What forms do optimal (i.e., supplier profit-maximizing) allocation polices take in example business scenarios?*

   For example, when retailers face deterministic linear demand, Theorem 5 proves the optimality of the linear allocation mechanism if supplier's prior on the intercept (θ) of the demand curve is a uniform distribution on $[0, a]$. We also show (Theorems 6 and 7) that the linear or proportional allocation rules are optimal for large classes of problems when retailers are "newsvendors".

3. *Is supplier profit sensitive to the allocation policy it employs?*

   In our numerical study, supplier profit increased a minimum of 842% by using the optimal allocation policy compared to the manipulable fixed-price linear allocation policy examined by Cachon and Lariviere. Of course, this improvement is not necessarily representative of the improvement to be expected from other parameterizations of the same business model or from a different business model. However, given any existing allocation policy, the supplier can never be worse off, assuming that an implementable optimal policy exists. And, as our numerical results demonstrate, the resulting improvement in profit can be quite considerable. Hence, we conclude that the choice of the pricing and allocation policy does matter. In particular, both supplier and supply-chain profit can increase significantly if a manipulable fixed-price allocation policy is replaced by the optimal truth-telling allocation policy.

4. *Are there practical ways of implementing the optimal allocation policy to reduce transaction costs?*

   Given the potential for substantially improved supplier profit, two questions remain regarding the implementability of the optimal (i.e., supplier profit-maximizing) allocation policy. First, whether the transactions cost to implement the optimal allocation policy might be higher than the increased profit it provides. Second, whether the history of the specific industry, and the corresponding expectations of the retailers, permit the supplier to implement such a policy, even if it paid for itself. To address the second question first, history and retailer expectations can be strong defenders of the status quo, even though the status quo doesn't serve either the supplier or the retailers very well. Witness the US auto industry's heavy reliance on turn-and-earn oriented policies even though such policies serve neither sector very well (see Cachon and Lariviere (1999a)). With respect to the cost of implementation, although conventional implementations might be both expensive to develop and time-consuming to use, we believe that web-based auction mechanisms, like the one we describe for the deterministic-demand business scenario, would be both fast and inexpensive to use, provided all parties were experienced web users. And, although the initial development cost for a specific business scenario might be high, existing web-auction providers might well be willing to bear this cost and charge suppliers on a fee-for-service basis.

Our model framework can be extended in several ways in future research. The role of transaction costs in capacity-allocation problems needs to be explored further. Our model examined a single period scenario. An obvious extension would be to extend it to a multi-period setting where capacity may have to be allocated in each period, based on each period's demand. In this setting, the optimal allocation policy would depend on whether excess demand is backordered or lost, and whether supplier capacity (i.e., inventory) can be carried between periods. Our model assumed that the N retailers were independent. It would also be interesting to extend this framework to the case where the private information parameters of the N retailers are correlated.

REFERENCES

Akerlof, G. (1970), "The Market for Lemons: Quality Uncertainty and the Market Mechanism", *Quarterly Journal of Economics*, 89, p488-500.

Allen, S. (1958), "Redistribution of Total Stock Over Several User Locations", *Naval Research Logistics*, 5, pp. 337-348.

Barlow, R. E., and Proschan, F., (1975), "Statistical theory of reliability and life testing: probability models", *Holt, Rinehart and Winston*, New York, NY.

Cachon, G. (2002), "Supply Chain Coordination with Contracts", *Working Paper*, University of Pennsylvania, Philadelphia, PA.

Cachon, G. and M. A. Lariviere. (1999a), "Capacity Allocation Using Past Sales: When to Turn-and-Earn", *Management Science*, 45(5), pp. 0685-0703.

Cachon, G. and M. A. Lariviere. (1999b), "Capacity Choice and Allocation: Strategic Behavior and Supply Chain Performance", *Management Science*, 45(8), pp. 1091-1108.

Cachon, G. and P. Zipkin. (1999), "Competitive and Cooperative Inventory Policies in a Two-Stage Supply Chain.", *Management Science*, 45(7), pp. 936-953.

Cachon, G. and M. A. Lariviere. (2001), "Contracting to Assure Supply: How to Share Demand Forecasts in a Supply Chain", *Management Science*, 47(5), pp. 629-646.

Chen, F. and Y. Zheng. (1994), "Lower Bounds for Multi-Echelon Stochastic Inventory Systems", *Management Science*, 40(11), pp. 1426-1443.

Chen, F. (1999), "Decentralized Supply Chains Subject to Information Delays", *Management Science*, 45(8) , pp. 1076-1090.

Chen, F. (2001), "Auctioning Supply Contracts", *Working Paper*, Columbia University, New York, NY.

Clark, A. J. and H. Scarf. (1960), "Optimal Policies for a Multi-Echelon Inventory Problem", *Management Science*, 6, pp. 475-490.

Corbett, C. J. (2001),"Stochastic Inventory Systems In A Supply Chain With Asymmetric Information: Cycle Stocks, Safety Stocks, And Consignment Stock", *Operations Research*, 49(4), pp. 0487-0500.

Dasgupta, S., and Spulber, D. F. (1989/90), "Managing Procurement Auctions", *Information Economics and Policy*, 4, p5-29.

Deshpande, V., Cohen, M. A., Donohue, K. 2002. "A Threshold Inventory Rationing Policy for Service Differentiated Demand Classes", *Working Paper*, Purdue University.

Elmaghraby, W. J. (2000), Supply Contract Competition and Sourcing Policies, *Manufacturing & Service Operations Management*, 2(4), p350-371.

Eppen, G. and L Schrage. (1981), "Centralized Ordering Policies in Multi-Echelon Systems with Leadtimes and Random Demand", *Multi-Level Production/Inventory Control Systems: Theory and Practice*, L. B. Schwarz (ed.), North-Holland, Amsterdam, pp.

51-84.

Federgruen, A. and P. Zipkin. (1984), "Allocation Policies and Cost Approximations for Multilocation Production and Inventory Problems", *Management Science*, 30(1), pp. 69-84.

Freedman, E. (1997), "Ex-Michigan Dealership loses Toyota Suit", *Automotive News*, September 1, 71, p. 25.

Fudenberg, D. and J. Tirole. (2000), "Game Theory", *MIT Press*, Cambridge, MA.

Ha, A. Y., "Inventory Rationing in a Make-to-Stock Production System with Several Demand Classes and Lost Sales", *Management Science*, 1997a, 43(8), 1093-1103.

Harris, M., and Raviv, A. (1981), "A Theory of Monolpoly Pricing Schemes with Demand Uncertainty", *American Economic Review*, 70, 347-365.

Harris, M., and Raviv, A. (1981), "Allocation Mechanisms and the Design of Auctions", *Econometrica*, 49(6), p1477-1499.

Jonsson, H. and E. Silver. (1987), "Analysis of a Two-Echelon Inventory-Control System with Complete Redistribution", *Management Science*, 33(2), pp. 215-227.

Klemperer, P. (1999), "Auction Theory: A Guide to the Literature", *Journal of Economic Reviews*, 3, p227-260.

Kumar, A., L. B. Schwarz, and J. E. Ward. (1995), "Risk-Pooling Along a Fixed Delivery Route Using a Dynamic Inventory-Allocation Policy," *Management Science*, 41(2), pp. 344-362.

Lawrence, D., (1996), "GM Tests Allocation Change", *Automotive News*, February 12, 142.

Lee, H. and S. Whang. (1999), "Decentralized Multi-Echelon Supply Chains: Incentives and Information", *Management Science*, 45(5), pp. 633-640.

Mallik, S., and Harker, P. (1999), "Coordinating Supply Chains with Competition: Capacity Allocation in Semiconductor Manufacturing", *working paper*, The Wharton School, University of Pennsylvania.

Maskin, E. and J. Riley (1989), "Optimal Multi-Unit Auctions", *The Economics of Missing Markets, Information, and Games*, Oxford University Press, New York.

McGavin, E. J., J. E. Ward, and L. B. Schwarz. (1997), "Balancing Retailer Inventories", *Operations Research*, 45(6), pp. 820-830.

Milgrom, P. R., and Weber, R. J. (1982), "A Theory of Auctions and Competitive Bidding", *Econometrica*, 50(5), p1089-1122.

Myerson, R.B. (1981), "Optimal Auction Design", *Mathematics of Operations Research*, 6(1), p58-73.

Porteus, E. (2000), "Responsibility Tokens in Supply Chain Management", *Manufacturing & Service Operations Management*, 2(2), pp. 203-219.

Riley, J. G. and Samuelson, W. F. (1981), "Optimal Auctions", *The American Economic Review*, 71(3), p381-392.

Rothschild, M., and Stiglitz, J. (1976), "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information", *Quarterly Journal of Economics*, 80, p629-649.

Sawyers, A. (1999), "Turn, earn: A Slow Burn", *Automotive News*, October 11, 74, p34.

Schwarz, L. B. (1989), "A Model for Assessing the Value of Warehouse Risk-Pooling: Risk-pooling Over Outside-Supplier Leadtimes", *Management Science*, 35(7), pp. 828-842.

Spence, A. M. (1974), *Market Signaling*, *Harvard University Press*, Cambridge, MA.

Topkis, D. M., "Optimal Ordering and Rationing Policies in a Non-stationary Dynamic Inventory Model with n Demand Classes", *Management Science*, 1968, 15(3), 160-76.

Tsay, A., S. Nahmias, and N. Agrawal, (1998), "Modeling Supply Chain Contracts: A Review", *Quantitative Models for Supply Chain Management*, Kluwer, Boston, MA.

Vickrey, W., (1961), "Counterspeculation, Auctions and Competitive Sealed Tendors", *Journal of Finance* 16(1), p8-37.

Vulcano, G., van Ryzin, G., and Maglaras, C. (2001), "Optimal Dynamic Auctions for Revenue Management", *Working Paper*, Columbia University, New York, NY.

APPENDIX

Proofs

Proof of Theorem 1: Suppose conditions (2) and (3) hold. Then by (3), expected profit to retailer i with parameter θ_i reporting value x is

$$\begin{aligned}
\pi_i(x,\theta_i) &= E_{-i}[R(Q(x,\theta_{-i}),\theta_i) - R(Q(x,\theta_{-i}),x) + \int_{\underline{\theta}}^{x} R_\theta(Q(\theta,\theta_{-i}),\theta)d\theta] \\
&= E_{-i}[\int_{x}^{\theta_i} R_\theta(Q(x,\theta_{-i}),\theta)d\theta + \int_{\underline{\theta}}^{x} R_\theta(Q(\theta,\theta_{-i}),\theta)d\theta] \\
&\leq E_{-i}[\int_{x}^{\theta_i} R_\theta(Q(\theta,\theta_{-i}),\theta)d\theta + \int_{\underline{\theta}}^{x} R_\theta(Q(\theta,\theta_{-i}),\theta)d\theta] \quad \text{by (2)} \\
&= E_{-i}[\int_{\underline{\theta}}^{\theta_i} R_\theta(Q(\theta,\theta_{-i}),\theta)d\theta]
\end{aligned}$$

$$= \pi_i(\theta_i, \theta_i)$$

Thus, if conditions (2) and (3) hold, then $\pi_i(x, \theta_i) \leq \pi_i(\theta_i, \theta_i)$, implying truth-telling is the optimal strategy for retailer i, in a Bayes-Nash equilibrium.

Now let $(P(\theta), Q(\theta))$ be incentive compatible and define $\pi(\theta) = \pi(\theta, \theta)$. Then $\pi(\theta) \geq \pi(\theta', \theta)$ for all $\theta, \theta' \in [\underline{\theta}, \overline{\theta}]$ implies

$$\pi(\theta, \theta) - \pi(\theta, \theta') \geq \pi(\theta, \theta) - \pi(\theta', \theta') \geq \pi(\theta', \theta) - \pi(\theta', \theta')$$

Hence,

$$E_{-i}[R(Q(\theta, \theta_{-i}), \theta) - R(Q(\theta, \theta_{-i}), \theta')] \geq \pi(\theta, \theta) - \pi(\theta', \theta') \geq E_{-i}[R(Q(\theta', \theta_{-i}), \theta) - R(Q(\theta', \theta_{-i}), \theta')]$$

Divide both sides by $\theta - \theta'$ and take limits as $\theta' \to \theta$,

$$\frac{d\pi(\theta)}{d\theta} = E_{-i}[R_\theta(Q(\theta, \theta_{-i}), \theta)]$$

Since $\pi(\underline{\theta}) = 0$, by integration

$$\pi(\theta_i) = E_{-i}[\int_{\underline{\theta}}^{\theta} R_\theta(Q(\theta, \theta_{-i}), \theta) d\theta]$$

Thus from the definition of the retailer's profit function

$$P(\theta_i) = E_{-i}[R(Q(\theta_i, \theta_{-i}), \theta_i) - \int_{\underline{\theta}}^{\theta_i} R_\theta(Q(\theta, \theta_{-i}), \theta) d\theta]$$

Thus, incentive compatibility implies condition (3) holds. Given $\pi_i(x, \theta_i) = E_{-i}[R(Q(x, \theta_{-i}), \theta_i) - P(x)]$, and using $P(\theta_i)$ from above,

$$\pi(\theta_i) - \pi(x, \theta_i) = \int_x^{\theta_i} E_{-i} R_\theta(Q(\theta, \theta_{-i}), \theta) d\theta - \int_x^{\theta_i} E_{-i} R_\theta(Q(x, \theta_{-i}), \theta) d\theta$$

Since $\pi(\theta_i) \geq \pi(x, \theta_i)$, condition (2) holds. Also, by assumption, $R_\theta \geq 0$, implying $\pi(\theta_i) \geq 0$ for all $\theta_i \in [\underline{\theta}, \overline{\theta}]$. □

Proof of Theorem 2: See Maskin and Riley (1989).

Proof of Theorem 4: The first order condition for the centralized problem can be written as

$$R_Q(Q(\theta_i, \theta_{-i}), \theta_i) = \lambda_c(\theta, K)$$

while the first order condition for the decentralized problem is

$$R_Q(Q(\theta_i, \theta_{-i}), \theta_i) - R_{Q\theta}(Q(\theta_i, \theta_{-i}), \theta_i) H(\theta_i) = \lambda(\theta, K)$$

Note that we have written the shadow price of capacity explicitly as function of θ and K. Now since $R_{Q\theta} \geq 0$, and $H(\theta) \geq 0$, it is immediately clear that $\lambda_c(\theta, K) \geq \lambda(\theta, K)$. Hence

$$E_\theta[\lambda_c(\theta, K)] \geq E_\theta[\lambda(\theta, K)]$$

Now, the expected shadow price at optimal capacity is equal to the marginal cost of capacity c. Hence

$$E_\theta[\lambda_c(\theta, K^*)] \geq E_\theta[\lambda(\theta, K^*)] = c$$

Since λ is non-increasing in K, to ensure $E_\theta[\lambda_c(\theta, K_c^*)] = c$, we have $K_c^* \geq K^*$. □

Proof of Theorem 5: The first order condition

$$G(Q, \theta_i) = R_Q(Q, \theta_i) - R_{Q\theta}(Q, \theta_i)H(\theta_i) = \lambda - \mu_i$$

is equivalent to

$$G(Q - \tau(\theta_i), \theta^*) = \lambda - \mu_i$$

Hence $Q_i^*(\theta_i) = Q^*(\theta^*) + \tau(\theta_i)$. Assuming that when capacity binds the base type θ^* receives stock, it must be that

$$\sum_{i=1}^{\underline{n}} Q_i^*(\theta_i) = \underline{n}Q^*(\theta^*) + \sum_{i=1}^{\underline{n}} \tau(\theta_i) = K$$

which yields $Q^*(\theta^*) = \frac{1}{\underline{n}}\{K - \sum_{i=1}^{\underline{n}} \tau(\theta_i)\}$. Hence linear allocation is the optimal mechanism.

□

Proof of Theorem 6: The first order condition

$$G(Q, \theta_i) = R_Q(Q, \theta_i) - R_{Q\theta}(Q, \theta_i)H(\theta_i) = \lambda - \mu_i$$

is equivalent to

$$G(\frac{Q}{\tau(\theta_i)}, \theta^*) = \lambda - \mu_i$$

Hence $Q_i^*(\theta_i) = \tau(\theta_i)Q^*(\theta^*)$. If capacity availability binds, we have $Q^*(\theta^*) = \frac{K}{\sum_{i=1}^{N} \tau(\theta_i)}$. Hence proportional allocation is optimal. □

Proof of Theorem 7: Using equation (3), the optimal pricing rule for the supplier can be written as:

$$P^*(\theta_i) = E_{-i}[R(Q^*(\theta_i, \theta_{-i}), \theta_i) - \int_{\underline{\theta}}^{\theta_i} R_\theta(Q^*(\theta, \theta_{-i}), \theta)d\theta]$$

Hence

$$\frac{\delta P^*(\theta_i)}{\delta \theta_i} = E_{-i}[R_Q(Q^*(\theta_i, \theta_{-i}), \theta_i)\frac{\delta Q^*}{\delta \theta_i}] \geq 0 \quad (19)$$

since Q^* is increasing in θ_i. Hence, P^{*-1} used in Theorem 8, is well defined. To prove theorem 8, we show that the optimal strategy for retailer i is to bid $P^*(\theta_i)$, in a Bayes-Nash equilibrium. The objective function for each retailer in a bidding auction, defined by theorem 8, can be written as

$$\max_{V_i} E_{-i}[R(Q^*(P^{*-1}(V_i), \theta_{-i}), \theta_i) - V_i]$$

The first order necessary condition is

$$E_{-i}[R_Q(Q^*(P^{*-1}(V_i), \theta_{-i}), \theta_i) \frac{\delta Q^*}{\delta P^{*-1}} \frac{\delta P^{*-1}}{\delta V_i}] = 1$$

Using equation (19), it is easy to see that $V_i = P^*(\theta_i)$ satisfies the above condition. Hence, it is optimal for each retailer to bid $V_i = P^*(\theta_i)$ in a Bayes-Nash equilibrium. Also, by definition of the allocation function $A(V_i, V_i)$, the optimal allocation policy is implemented if each retailer bids $P^*(\theta_i)$. □