

Auctioning Supply Contracts

Fangruo Chen

Graduate School of Business, Columbia University, New York, New York 10027 and School of Management, Shanghai Jiao Tong University, Shanghai 200052, China, fc26@columbia.edu

This paper studies a procurement problem with one buyer and multiple potential suppliers who hold private information about their own production costs. Both the purchase quantity and the price need to be determined. An optimal procurement strategy for the buyer requires the buyer to first design a supply contract that specifies a payment for each possible purchase quantity and then invites the suppliers to bid for this contract. The auction can be conducted in many formats such as the English auction, the Dutch auction, the first-priced, sealed-bid auction, and the Vickrey auction. The winner is the supplier with the highest bid, and is given the decision right for the quantity produced and delivered. Applying this theory to a newsvendor model with supply-side competition, this paper establishes a connection between the above optimal procurement strategy and a common practice in the retail industry, namely, the use of slotting allowances and vendor-managed inventory. Also discussed in the newsvendor context are the role of well-known supply contracts such as returns contracts and revenue-sharing contracts in procurement auctions, the scenarios where the buyer and suppliers may possess asymmetric information about the demand distribution, and how the cost of supply-demand mismatch is affected by supply-side competition. Finally, this paper compares the optimal procurement strategy with a simpler but suboptimal strategy where the buyer first determines a purchase quantity and then seeks the lowest-cost supplier for the quantity in an auction.

Key words: procurement strategies; decision rights; newsvendor model; auctions; competitive bidding; slotting allowances

History: Accepted by William S. Lovejoy, operations and supply chain management; received March 13, 2003.

This paper was with the author 2 years for 3 revisions. Published online in *Articles in Advance* August 31, 2007.

1. Introduction

Mathematical models in operations management addressing procurement decisions (i.e., when and how much to buy) often make the simplistic assumption that the prices for the items to be purchased are exogenously given. For example, the celebrated newsvendor model typically assumes a given purchase price. The plethora of inventory models in the literature also make a similar assumption. In reality, procurement managers often need to discover the prices for the items they want to buy, and the discovery process typically involves market research, negotiations, auctions/bidding, etc. Procurement models in economics, on the other hand, tend to focus on price discovery, fixing quantity decisions. The purpose of this paper is to understand how price discovery can be integrated with quantity decisions to reach an optimal procurement strategy.

We consider a model with one buyer and a number of potential suppliers. The question facing the buyer is how much of a product should be purchased and from which supplier. The buyer's revenue is a concave, increasing function of the purchased quantity. Each supplier can produce the product at a constant marginal cost and without any capacity limit. The marginal costs of different suppliers all come

from a common probability distribution and each supplier is privately informed about his own marginal cost. The suppliers are risk neutral. We study the procurement problem from the buyer's standpoint, seeking an optimal procurement strategy that maximizes her expected profit, which is her revenue (as a function of the purchased quantity) minus the procurement cost. The structure of such a strategy shows how price discovery should be integrated with the quantity decision.

A key result of this paper is that the following *supply contract auction* is optimal. Here the buyer first commits to a supply contract, which specifies a payment for each quantity the buyer may purchase and delegates the quantity decision to the winning supplier. Each potential supplier views this contract as a business opportunity, taking the payment function offered by the buyer as their revenue function. Based on their own marginal production cost, they each determine an optimal quantity to produce and deliver to the buyer to maximize their own profit. Note that different suppliers derive different values from the contract, with the lowest-marginal-cost supplier deriving the most value. In an auction, the suppliers compete for the contract by submitting an up-front, lump-sum fee they are willing to pay, with the winner being the supplier offering the highest fee. Only the

winning supplier pays an up-front fee.¹ The auction can be conducted in many different forms, including the English auction, the Dutch auction, the first-price, sealed-bid auction, and the Vickrey auction.

The integration of price discovery and the quantity decision therefore requires the design of a supply contract (a quantity-to-payment mapping) followed by an auction that awards the contract to a supplier, who is given the decision right to determine the procurement quantity. This is intuitive. The auction chooses the most efficient supplier. Moreover, the lowest marginal cost among the suppliers together with the supply contract offered by the buyer determine the ultimate quantity decision. In other words, the quantity decision comes from both sides of the trade.

The supply contract auction fits well with a prevalent practice in the retail industry, i.e., the use of slotting allowances and vendor-managed inventory (VMI). A slotting allowance is an up-front, lump-sum fee that a manufacturer pays to a retail store when it introduces a new product to the store. It can also be a fee paid to keep an existing product in the store, although in this case, it is often referred to as a pay-to-stay fee. The total of slotting allowances is estimated at between \$6 and \$9 billion a year (Rao and Mahi 2003). The reasons given for using slotting allowances often focus on the need to signal to the retailer the manufacturer's belief about the product's market potential or the need to compensate the retailer for the risks of carrying a new product. On the other hand, VMI refers to the practice of delegating inventory decisions to the supplier. The arguments for this reallocation of stocking decision rights are multifaceted, including the bargaining power of the retailer (e.g., Wal-Mart), the vendor's expertise in managing inventories, the vendor's ability to coordinate the production and distribution of multiple products, modern communication technologies, etc. While VMI has been hailed as an important strategy for improving supply chain efficiency, the use of slotting allowances is more controversial, attracting antitrust inquiries (see, e.g., Federal Trade Commission 1999). It is therefore interesting that our particular optimal procurement strategy actually requires the use of a slotting allowance (i.e., the lump-sum fee) and the delegation of the production/inventory decision to the chosen supplier. The lump-sum fee results from the suppliers' competition for the retailer's shelf space, while the allocation of decision rights is to take advantage of the suppliers' private cost information.

¹ An alternative way to implement the supply contract auction is to require each supplier to submit a two-dimensional bid consisting of a fee and a quantity, but the winner is chosen solely based on the fee. This implementation eliminates any uncertainty about what a supplier may deliver after being chosen. Thanks to a referee for making this suggestion.

Our results therefore have the potential to contribute to the on-going debate on the purposes and consequences of slotting allowances.

Interpreting the buyer in our procurement model as a newsvendor facing an inventory decision at the beginning of a selling season with uncertain demand, we arrive at a newsvendor model with supply-side competition. When demand is uncertain, we often see supply contracts that are demand dependent. For example, a returns contract calls for a wholesale price paid for each unit delivered to the newsvendor at the beginning of the selling season, and a rebate for each unit of excess inventory at the end of the season. On the other hand, a revenue-sharing contract contains a wholesale price for each unit of inventory at the beginning of the selling season, and an agreement on how the realized revenues are to be split between the newsvendor and the selected supplier. Note that under either a returns contract or a revenue-sharing contract, the (net) transfer payment for any given purchase quantity (or initial inventory) is uncertain as it depends on the realized demand. This is clearly different from the contract considered so far, whereby the transfer payment is a deterministic function of the purchase quantity, but for our risk-neutral suppliers, all that matters is the expected transfer payment as a function of the purchase quantity. It is shown that under fairly general conditions, a returns contract or a revenue-sharing contract can indeed be found that, when auctioned off among the suppliers, constitutes an optimal procurement strategy for the newsvendor. We also discuss the advantages and disadvantages of these demand-dependent contracts if the newsvendor and the suppliers have asymmetric information about the probability distribution of the demand.

An alternative procurement strategy is for the firm to choose a purchase quantity and then run an auction to identify a supplier that can deliver the quantity at the lowest cost. We refer to this strategy as the *fixed-quantity auction*. It is suboptimal because the quantity decision only takes into account the distributional knowledge the buyer has about the suppliers' costs (the actual production costs do not play a role in the quantity decision), but it is simpler than the optimal strategy because it does not involve the specification of a supply contract. There are two decision variables that need to be determined: the purchase quantity and the reserve price for the auction. For the newsvendor setting, we develop an algorithm to determine the optimal values of these variables that maximize the buyer's expected profit. Numerical examples are used to illustrate the difference in the newsvendor's expected profit between the fixed-quantity auction and the optimal procurement strategy.

Closely related to this paper is Dasgupta and Spulber (1990; hereafter DS for brevity), who have provided a different solution to the above procure-

ment problem. We refer to their solution as *the quantity auction*, which works as follows. As in the supply contract auction, the buyer first announces a supply contract (a payment function). The suppliers then submit quantity offers in a sealed, high-bid auction. The supplier with the highest bid (i.e., quantity offer) wins the auction, produces and delivers his bid, and receives a payment according to the preannounced contract. With a properly chosen supply contract, the quantity auction maximizes the buyer's expected profit, but unlike the supply contract auction, the quantity auction must be conducted in the sealed-bid fashion, and would lose its optimality when implemented in other formats. Another distinction between the quantity auction and the supply contract auction is the amount of detail required to determine the optimal payment function: the quantity auction requires the number of potential suppliers, whereas the supply contract auction does not.²

Also closely related to this paper is the literature on multidimensional auctions (see, e.g., Che 1993 and Branco 1997). Here the suppliers compete for a procurement contract on price and quality. For example, the Department of Defense, in procuring a weapons system, cares about the system's performance as well as price. The optimal design of a multidimensional auction typically specifies a scoring rule that combines the multiple dimensions of a bid (e.g., a quality specification and a price) into a single score, which is then used to determine a winner. As mentioned earlier, the supply contract auction discussed in this paper can also be implemented as a multidimensional auction, with an up-front fee and a quantity, and a scoring rule that puts the weight entirely on the former. For further discussions on the relationship between the supply contract auction and Che's scoring rule approach, see §2.2.

This paper is at the intersection of auction theory and operations management. Auction theory has grown tremendously since Vickrey's (1961) seminal work; McAfee and McMillan (1987a) and Klemperer (1999) provide comprehensive surveys of the theory. Supply contracts have recently received much research attention in operations management and the research effort is centered around their role in achieving sup-

ply chain coordination (see Cachon 2003 and Chen 2003b). The contribution of this paper is therefore to demonstrate how supply contracts can be integrated with auction mechanisms to obtain an optimal procurement strategy (for one party of the supply chain). Research efforts to introduce auction theory to the field of operations management have been on the rise. Many papers have started to include an auction mechanism or some other price-discovery mechanism in an operations context (see Mendelson and Tunca 2000, Gallien and Wein 2000, Lee and Whang 2002, etc.). Elmaghraby (2000) provides a survey of this area.

Whereas our paper focuses on integrating a supply contract with an auction mechanism, others have studied the use of incentive contracts in auctions (see, e.g., Laffont and Tirole 1987; McAfee and McMillan 1986, 1987b; and Riordan and Sappington 1987). The setting typically includes one principal and multiple agents. The principal has a project for which the agents compete. Each agent possesses private information and his action is unobservable to the principal. The solution is to auction off an incentive contract among the agents. This part of the literature is also discussed in Laffont and Tirole (1993).

The remainder of this paper is organized as follows: §2 describes and develops different optimal procurement strategies; §3 applies these strategies to a newsvendor model with supply-side competition and compares them with a suboptimal but simpler procurement strategy; §4 contains concluding remarks.

2. Optimal Procurement Strategies

Consider the following procurement problem with one buyer and multiple potential suppliers. The buyer can purchase any quantity of a product from any of the suppliers. Let Q be the total quantity purchased and $R(Q)$ the value the buyer attaches to the total purchase. Assume that $R(\cdot)$ is strictly concave and increasing with $R(0) = 0$. There are n , $n \geq 2$, potential suppliers and each of them is capable of producing the product with a constant marginal cost and an unlimited capacity. Let c_i be the marginal cost for supplier i , $i = 1, \dots, n$. Each supplier is privately informed of his own cost. These costs are independent draws from a common probability distribution F , $c \in [\underline{c}, \bar{c}]$, which is differentiable with $F(\underline{c}) = 0$ and $F(\bar{c}) = 1$. Define

$$\tilde{H}(x) = x + \frac{F(x)}{F'(x)} \quad (1)$$

and assume $\tilde{H}(\cdot)$ is increasing.³ The suppliers are risk neutral. We seek an optimal procurement strategy that maximizes the buyer's expected profit.

³ This is a regularity condition often assumed in auction contexts (see, e.g., Myerson 1981). The assumption is clearly true if $F(x)/F'(x)$ is increasing in x , or equivalently, F is logconcave. Many

² An alternative way to implement the quantity auction is via a wholesale price auction. Here the buyer first announces a purchasing plan, which specifies the quantity the buyer is committed to purchase as a function of the realized wholesale price. The suppliers then engage in a sealed, low-bid auction, where the supplier submitting the lowest wholesale price wins. The transaction is completed at the lowest bid (wholesale price) and the corresponding purchase quantity. The purchasing plan can be shown to be a quantity discount scheme, i.e., lower wholesale prices are associated with larger purchase quantities. Wholesale price auctions have been studied by Hansen (1988) and Jin and Wu (2000), where the purchasing plan is exogenously given.

2.1. Quantity Auction

DS provides an optimal strategy for the above procurement problem. It requires the buyer to first announce a contract $P(\cdot)$, whereby the buyer pays $P(Q)$ to a supplier if Q units are purchased from the supplier for any possible value of Q . Knowing $P(\cdot)$, the suppliers each name a quantity in a sealed bid. The supplier who bids the maximum quantity wins the contract, produces and delivers his bid, and, in return, receives a payment from the buyer according to $P(\cdot)$. (The other suppliers do not produce and do not receive any payment.) With a properly chosen payment function $P(\cdot)$, the above auction maximizes the buyer's expected profit. We shall refer to the above procurement strategy as the quantity auction conducted in the sealed, high-bid format. (For brevity, we may say "the quantity auction" without specifying the format. In such cases, we usually mean the sealed, high-bid format, but other formats are possible. The exact meaning should be clear from the context.)

DS first considered the direct revelation game. The solution to this game is the optimal direct mechanism. They then suggested that the quantity auction implements the optimal direct mechanism and is thus optimal among all possible procurement strategies due to the revelation principle (see, e.g., Kreps 1990). Below we provide a more detailed analysis of this procurement strategy.

Consider the quantity auction (sealed, high-bid). Given the payment function $P(\cdot)$, the suppliers play a game of incomplete information (due to their private cost information) for which the Bayesian-Nash equilibrium is an appropriate solution concept. Assume that there is a symmetric Bayesian-Nash equilibrium strategy; this is plausible because the suppliers are ex ante symmetric (with independent and identically distributed (i.i.d.) costs). Denote this strategy by $Q(\cdot)$: a supplier with marginal cost c bids $Q(c)$, $c \in [\underline{c}, \bar{c}]$. Assume $Q(c)$ is strictly decreasing in c for $c \in [\underline{c}, c_*]$, for some $c_* \in [\underline{c}, \bar{c}]$, and $Q(c) = 0$ for $c > c_*$.

Note that there is a trade if and only if $C_1 \leq c_*$, where C_1 is the lowest marginal cost among the suppliers. Because the lowest-cost supplier always wins the auction, we know that $R(Q(C_1))$ is the system's revenue and $C_1 Q(C_1)$ is the system's production cost. As a result, the expected systemwide profit is

$$\int_{\underline{c}}^{c_*} [R(Q(x)) - xQ(x)] f_{(1)}(x) dx, \quad (2)$$

where $f_{(1)}$ is the probability density function (p.d.f.)

probability distributions are logconcave, including the beta family of which the uniform distribution is a member, and the normal distribution truncated and scaled to a finite interval. See Rosling (2002) for an extensive discussion on logconcave probability distributions and further references on the topic.

of C_1 . To derive the buyer's expected profit, it suffices to determine the suppliers' expected profits.

Take any $i = 1, \dots, n$. Suppose that all the suppliers but supplier i play the strategy $Q(\cdot)$. Consider the problem facing supplier i . For $Q(\cdot)$ to be an equilibrium strategy, it must be the case that supplier i can do no better than bidding $Q(c_i)$ for all $c_i \in [\underline{c}, c_*]$. In particular, supplier i gains nothing by bidding $Q(x)$ for some $x \neq c_i$, which is the same as playing the strategy $Q(\cdot)$, but pretending that his marginal cost is x . Take any $c \in [\underline{c}, c_*]$ and suppose that $c_i = c$. Let $\pi_i(x, c)$ be supplier i 's expected profit if he bids $Q(x)$ while his marginal cost is c , given that all the other suppliers play $Q(\cdot)$. Note that

$$\pi_i(x, c) = [P(Q(x)) - cQ(x)][1 - F(x)]^{n-1}, \quad (3)$$

where $P(Q(x)) - cQ(x)$ is the supplier's profit if he wins the auction (by bidding $Q(x)$) and $[1 - F(x)]^{n-1}$ is the probability of winning, which occurs if and only if every other supplier's marginal cost is greater than supplier i 's "reported" marginal cost x . For $\pi_i(x, c)$ to be maximized at $x = c$, it is necessary that $\partial \pi_i(x, c) / \partial x|_{x=c} = 0$, i.e.,

$$[P'(Q(c))Q'(c) - cQ'(c)][1 - F(c)]^{n-1} - [P(Q(c)) - cQ(c)](n-1)[1 - F(c)]^{n-2}F'(c) = 0.$$

Using this equation in the expression of $\pi_i'(c)$, where $\pi_i(c) \stackrel{\text{def}}{=} \pi_i(c, c)$ for any $c \in [\underline{c}, c_*]$, we have $\pi_i'(c) = -Q(c)[1 - F(c)]^{n-1}$. Setting $\pi_i(c_*) = \pi_i(c_*, c_*) = 0$, we have

$$\pi_i(c) = \int_{c_*}^c \pi_i'(x) dx = \int_c^{c_*} Q(x)(1 - F(x))^{n-1} dx \quad (4)$$

with

$$\begin{aligned} E[\pi_i(c)] &= \int_{\underline{c}}^{c_*} \left\{ \int_c^{c_*} Q(x)(1 - F(x))^{n-1} dx \right\} F'(c) dc \\ &= \int_{\underline{c}}^{c_*} \left\{ \int_{\underline{c}}^x Q(x)(1 - F(x))^{n-1} F'(c) dc \right\} dx \\ &= \int_{\underline{c}}^{c_*} Q(x)(1 - F(x))^{n-1} F(x) dx. \end{aligned}$$

Because the suppliers are symmetric, the sum of the expected profits of the suppliers is simply n times the above expression. Subtracting $nE[\pi_i(c)]$ from (2) gives the buyer's expected profit

$$\int_{\underline{c}}^{c_*} [R(Q(x)) - \tilde{H}(x)Q(x)] f_{(1)}(x) dx. \quad (5)$$

Define

$$Q^*(x) = \arg \max_{Q \geq 0} [R(Q) - \tilde{H}(x)Q], \quad x \in [\underline{c}, \bar{c}]. \quad (6)$$

Because $R(\cdot)$ is concave and $\tilde{H}(\cdot)$ is increasing, $Q^*(x)$ is decreasing in x . Set c_* equal to the minimum $x \in [\underline{c}, \bar{c}]$ with $Q^*(x) = 0$; if no such x exists, set $c_* = \bar{c}$. It is clear from (5) that the buyer's expected profit is maximized if $Q^*(\cdot)$ arises as a Bayesian-Nash equilibrium in the bidding game.

Suppose that $Q^*(\cdot)$ is a Bayesian-Nash equilibrium. Take any $c \in [\underline{c}, c_*]$. From (4),

$$\pi_i(c) = \int_c^{c_*} Q^*(z)(1 - F(z))^{n-1} dz.$$

Because $\pi_i(c, c) = \pi_i(c)$, we have from (3),

$$\begin{aligned} & [P(Q^*(c)) - cQ^*(c)][1 - F(c)]^{n-1} \\ &= \int_c^{c_*} Q^*(z)(1 - F(z))^{n-1} dz. \end{aligned}$$

Therefore,

$$P(Q^*(c)) = cQ^*(c) + \frac{\int_c^{c_*} Q^*(z)(1 - F(z))^{n-1} dz}{(1 - F(c))^{n-1}}. \quad (7)$$

Denote by $P^*(\cdot)$ the payment function that satisfies the above equation for all $c \in [\underline{c}, c_*]$.

The above derivation for P^* is entirely based on necessary conditions for $Q^*(\cdot)$ to be a Bayesian-Nash equilibrium. It remains to verify that under $P^*(\cdot)$, $Q^*(\cdot)$ indeed arises as such. (DS did not address this issue.) To see this, suppose $P^*(\cdot)$ is the payment function and assume that all players but player i follow strategy $Q^*(\cdot)$. Now in (3), replace $P(\cdot)$ with $P^*(\cdot)$ and $Q(\cdot)$ with $Q^*(\cdot)$. We have

$$\begin{aligned} \pi_i(x, c) &= (x - c)Q^*(x)(1 - F(x))^{n-1} \\ &+ \int_x^{c_*} Q^*(z)(1 - F(z))^{n-1} dz, \quad x, c \in [\underline{c}, c_*]. \end{aligned}$$

Note that

$$\frac{\partial \pi_i(x, c)}{\partial x} = (x - c) \frac{\partial}{\partial x} (Q^*(x)(1 - F(x))^{n-1}).$$

Because both $Q^*(x)$ and $(1 - F(x))^{n-1}$ are decreasing in x , the partial derivative on the right side is negative. Therefore, $\partial \pi_i(x, c)/\partial x >$ (respectively, $<$) 0 for $x <$ (respectively, $>$) c . Consequently, $\pi_i(x, c)$ is maximized at $x = c$.

The following theorem, which we attribute to DS, summarizes the above development.

THEOREM 1 (DASGUPTA AND SPULBER 1990). *In the quantity auction defined by the payment function $P^*(\cdot)$ and conducted in the sealed, high-bid format, $Q^*(\cdot)$ arises as a common Bayesian-Nash equilibrium strategy for the suppliers, and the buyer's expected profit is*

$$E[R(Q^*(C_1)) - \tilde{H}(C_1)Q^*(C_1)].$$

This is also the highest expected profit the buyer can achieve among all feasible procurement strategies.

The expression for the buyer's maximum expected profit in Theorem 1 indicates that the buyer's lack of cost information is reflected in the "virtual cost" she pays to the winning supplier, i.e., she pays $\tilde{H}(x)$ to the winning supplier when in fact his marginal production cost is x . The difference $\tilde{H}(x) - x$ is thus the information rent. The outcome is as if the buyer had complete information about the suppliers' costs, but agreed to pay virtual costs.

Several observations are immediate. The quantity traded is $Q^*(C_1)$, a decreasing function of C_1 . Therefore, as competition intensifies, i.e., as n grows, C_1 becomes stochastically smaller, leading to a stochastically larger trade. Moreover, as expected, the buyer's expected profit increases with supply-side competition. To formally show this, define $\Pi(x) = \max_{Q \geq 0} R(Q) - \tilde{H}(x)Q$, a decreasing function of x because $\tilde{H}(x)$ is increasing in x . As n grows, C_1 becomes stochastically smaller, increasing $E[\Pi(C_1)]$, the buyer's expected profit. Finally, note that the efficient trade volume between the lowest-cost supplier and the buyer, i.e., the one that maximizes their joint gains, is $Q^0(x) \stackrel{\text{def}}{=} \arg \max_{Q \geq 0} R(Q) - xQ$. Note that $Q^0(x) > Q^*(x)$ for all x . Hence, asymmetric cost information causes supply chain inefficiencies by reducing trade. This is reminiscent of the well-known double-marginalization phenomenon (see, e.g., Tirole 1988): the marginal cost facing the buyer, i.e., the virtual cost $\tilde{H}(C_1)$, is higher than the system's marginal cost, C_1 .

For independent, private-values auctions, a well-known result is the revenue equivalence theorem, which states that the auctioneer is indifferent among many commonly used auction formats such as the English auction, the Dutch auction, the first-price, sealed-bid auction, and the Vickrey auction. Does a similar result hold for the quantity auction? In other words, if the buyer can freely modify the payment function to suit the auction format used, can she still achieve the optimal expected profit by using an auction format other than sealed, high-bid? For example, the buyer can run the quantity auction in the Vickrey fashion: the suppliers submit quantity offers in sealed bids, the winner is the highest bidder, but the quantity produced by the winning supplier and delivered to the buyer is equal to the second-highest bid. What is the buyer's maximum expected profit in this case? The following theorem shows that revenue equivalence does not hold for the quantity auction. The proof is in Chen (2003a).

THEOREM 2. *The auction format used in the quantity auction matters: while the sealed, high-bid auction and the Dutch auction are optimal (maximizing the buyer's expected profit), the English auction and the Vickrey auction are not. Moreover, the buyer prefers the English auction to the Vickrey auction.*

2.2. Supply Contract Auction

An important feature of the quantity auction is that the buyer first commits to a payment function (supply contract). This payment function represents a potential source of revenue for each supplier. This business opportunity (of trading with the buyer) is likely to be valued differently by different suppliers, with the lowest-cost supplier achieving the highest value. Therefore, an alternative way to select a supplier is to

ask them to bid in terms of an up-front, lump-sum fee. The winner is the supplier offering to pay the highest fee, and only the winning supplier pays an up-front fee. The winning supplier then determines his production quantity (to maximize his profit), delivers it to the buyer, and receives a payment from the buyer according to the supply contract. This procurement strategy will be referred to as the *supply contract auction*. It is also optimal for the buyer, as we show next.

Take any payment function $P(\cdot)$ with $P(0) = 0$. Consider the problem of supplier selection. Define

$$v(c) = \max_{Q \geq 0} P(Q) - cQ. \quad (8)$$

Let $Q(c) = \arg \max_{Q \geq 0} P(Q) - cQ$.⁴ Therefore, supplier i values the business opportunity at $v(c_i) \stackrel{\text{def}}{=} v_i$, $i = 1, 2, \dots, n$. Because the suppliers' marginal costs are independent draws from a common distribution, the values $\{v_i\}_{i=1}^n$ are i.i.d. random variables. Consequently, the problem of choosing a supplier can be thought of as selling an object (i.e., the business opportunity) to the highest bidder, where the bidders have i.i.d. valuations. From the revenue equivalence theorem, the buyer obtains the same expected lump-sum fee (and selects the same supplier) if she uses the English auction, the Dutch auction, the sealed, high-bid auction, or the Vickrey auction.

Suppose the buyer uses the English auction for supplier selection. That is, the suppliers openly bid on the fee they are willing to pay for the privilege to trade, and the supplier with the highest bid wins and pays his bid. Clearly, the supplier with the highest valuation (and the lowest marginal cost) wins the auction and pays (to the buyer) a lump-sum fee equal to the valuation of the second-lowest-cost supplier. Let $V_k = v(C_k)$, $k = 1, \dots, n$, where C_k is the k th lowest cost. Thus, the lump-sum fee the buyer receives is V_2 .⁵

We now proceed to determine the optimal payment function. Consider the buyer's cash flow. First, the buyer collects a lump-sum fee, V_2 , from the lowest-cost supplier. This supplier determines the quantity to maximize his profit, i.e., maximizing $P(Q) - C_1Q$ over Q . The optimal quantity is $Q(C_1)$. The trade gives the buyer revenues in the amount of $R(Q(C_1))$, but costs her $P(Q(C_1))$. Consequently, the buyer's profit is

$$\Pi \stackrel{\text{def}}{=} R(Q(C_1)) - P(Q(C_1)) + V_2.$$

⁴ In the previous subsection, $Q(\cdot)$ was used for a Bayesian-Nash equilibrium strategy in the quantity auction.

⁵ A reader versed in auction theory may think that the buyer is better off with a reserve price (see, e.g., Riley and Samuelson 1981). This is actually not the case, as the subsequent development (summarized in Theorem 3) shows. The reason is that the auction is not regular in the sense of Myerson (1981). Chen (2003a) illustrates this with an example.

Because $V_1 = P(Q(C_1)) - C_1Q(C_1)$,

$$\Pi = R(Q(C_1)) - C_1Q(C_1) - (V_1 - V_2). \quad (9)$$

We next obtain a convenient expression for the expected value of $(V_1 - V_2)$, which is the winning supplier's profit. Note from the optimization problem in (8) that $Q(c)$ is decreasing in c . (This is true for any $P(\cdot)$.) Let c_0 be the minimum c with $Q(c) = 0$. If $Q(c) > 0$ for all $c \in [c, \bar{c}]$, then set $c_0 = \bar{c}$. Take any $c < c_0$. Thus, $Q(c) > 0$. Writing $v(c) = P(Q(c)) - cQ(c)$ and differentiating,

$$v'(c) = P'(Q(c))Q'(c) - Q(c) - cQ'(c).$$

Because $P'(Q(c)) = c$ (the first-order condition for the optimization problem in (8)), we have $v'(c) = -Q(c)$. On the other hand, for any $c > c_0$, $Q(c) = 0$ and thus $v(c) = 0$ because $P(0) = 0$. Consequently, for all $c > c_0$, $v'(c) = 0 = -Q(c)$. Combining the two cases, we have

$$V_1 - V_2 = v(C_1) - v(C_2) = \int_{C_1}^{C_2} Q(x) dx.$$

Using the conditional probability density function of C_2 given $C_1 = c$, we have

$$\begin{aligned} E \left[\int_{C_1}^{C_2} Q(x) dx \right] &= E_{C_1} E \left[\int_{C_1}^{C_2} Q(x) dx \mid C_1 = c \right] \\ &= \int_{\underline{c}}^{\bar{c}} nF'(c)(1 - F(c))^{n-1} dc \\ &\quad \cdot \int_{y=c}^{\bar{c}} \left(\int_c^y Q(x) dx \right) \frac{(n-1)F'(y)(1 - F(y))^{n-2}}{(1 - F(c))^{n-1}} dy \\ &= \int_{\underline{c}}^{\bar{c}} nF'(c) dc \int_{y=c}^{\bar{c}} \left(\int_c^y Q(x) dx \right) (n-1) \\ &\quad \cdot F'(y)(1 - F(y))^{n-2} dy \end{aligned}$$

which, after changing the order of integration twice (first between x and y and then between c and x), becomes

$$\int_{\underline{c}}^{\bar{c}} Q(x)nF(x)(1 - F(x))^{n-1} dx. \quad (10)$$

Substituting (10) for the expected value of $(V_1 - V_2)$ in (9), we have

$$E[\Pi] = \int_{\underline{c}}^{\bar{c}} [R(Q(x)) - \tilde{H}(x)Q(x)]nF'(x)(1 - F(x))^{n-1} dx,$$

which is exactly the same as (5), the buyer's expected profit in the quantity auction. (Note the different meanings of $Q(\cdot)$ in the two places.) Although the above expression is obtained under the assumption that the buyer uses the English auction to select a supplier, we know, from the revenue equivalence theorem, that the same expression holds if she instead uses the Dutch auction, the first-price, sealed-bid auction, or the Vickrey auction.

Note that if $Q(x) = Q^*(x)$, which is defined in (6), for all $x \in [\underline{c}, \bar{c}]$, then $E[\Pi]$ equals the buyer's maximum expected profit (Theorem 1). This is indeed possible. The payment function that achieves this, denoted by $P^{**}(\cdot)$, is the solution to

$$P'(Q^*(c)) = c \quad (11)$$

for all $c \in [\underline{c}, \bar{c}]$ with $Q^*(c) > 0$. Note that this optimal payment function is increasing in Q , concave because $Q^*(c)$ is decreasing in c , and independent of the number of bidders. Also note that it is independent of the auction format. Finally, adding a constant to the payment function does not change its optimality.

THEOREM 3. *The buyer maximizes her expected profit if she uses the supply contract auction with $P^{**}(\cdot)$ as the payment function. Moreover, this payment function is increasing, concave, and independent of the number of potential suppliers.*

We pause here to establish connections between the supply contract auction developed here and the scoring rule auction developed by Che (1993) for multidimensional auctions. It should be evident by now that the procurement problem discussed in this paper has two objectives, i.e., price discovery and quantity decision. The idea of a scoring rule is to ask the suppliers to bid in both price and quantity, and then for each supplier, combine the price bid and the quantity bid into a single number (the supplier's score). The scoring rule is the mechanism that does the second step. The supplier with the highest score wins the auction. Let p and q be the price bid and quantity bid, respectively, submitted by a supplier. Che (1993) gave a scoring rule of the following form: the score for a supplier bidding (p, q) is $s = s(q) - p$ for some function $s(\cdot)$. (Che considered price and quality, but this is just a cosmetic difference.) The supply contract auction with payment function $P(\cdot)$ is closely related to Che's first-score auction, with $s(\cdot) \equiv P(\cdot)$. (Under the first-score auction, the winning supplier delivers his quantity bid to, and receives his price bid from, the buyer.) Note that under Che's scoring rule, a supplier bidding (p, q) is given score $s = P(q) - p$, which corresponds to the up-front, lump-sum fee the supplier offers to pay in the supply contract auction. The two approaches imply different cash flows between the buyer and the winning supplier: under the supply contract auction, the supplier first makes an up-front payment to the buyer (to secure the business), then produces and delivers the quantity, and finally receives a payment from the buyer for the quantity delivered (according to a previously given contract), whereas under Che's design, there is no up-front payment, and the terms of trade (the quantity to be delivered and the corresponding price) are fully specified in the winning bid.

The basic idea of the supply contract auction is to treat a supply contract as an "object" to be auctioned off among the suppliers. The main task is to design this "object" properly, and then invite single-dimensional bids. The supply contract (the "object") serves to convert a quantity decision into a price figure, effectively aggregating price and quantity dimensions into one dimension (price). In contrast, the scoring-rule approach is to first invite multidimensional bids and then aggregate the different dimensions into one. Note that both approaches achieve aggregation, one before bidding and the other after. Interestingly, achieving aggregation before bidding has real benefits: simpler analysis with more general results. For example, with the supply contract auction approach, one could use the revenue equivalence theorem to effortlessly establish the equivalence of different auction formats, whereas Che was confined to sealed bidding. Moreover, the new approach produces an optimal procurement strategy that fits well with a practice in the retail industry (§3.1).

2.3. Comparisons

Although both the quantity and the supply contract auctions are optimal for the buyer, they are different in many ways. Under the supply contract auction, the buyer has the flexibility of using any of the common auction forms mentioned earlier. This flexibility is lost for the quantity auction (Theorem 2). Moreover, the optimal payment function for the supply contract auction is concave, increasing, and independent of the number of bidders, but for the quantity auction, it may actually decrease (a rather unpleasant feature), and it generally depends on the number of bidders. In other words, the optimal payment function in the supply contract auction is more "intuitive" (a supplier is paid more if he delivers more) and more "detail-free" (the buyer does not have to know the exact number of potential suppliers).

3. A Newsvendor Model with Supply-Side Competition

One of the most celebrated models in operations management is the newsvendor model, which succinctly captures the trade-off between buying too much and buying too little. The model assumes complete certainty on the supply side, where, typically, an unlimited quantity can be procured at an exogenously given per-unit wholesale price. In reality, however, most industrial buyers face multiple potential suppliers with private information about their production costs. As a result, the purchase price needs to be discovered, which, of course, influences the purchase quantity. In this section, we introduce a newsvendor model with supply-side competition and characterize the optimal procurement strategies in this setting.

Toward the end of the section, we study a simpler, but suboptimal, procurement strategy and numerically compare it with the optimal strategies.

Consider the following newsvendor model with supply-side competition. A firm (the buyer) must determine how much inventory of a product to stock in advance of a selling season. The total demand for the product over the entire selling season, D , is a non-negative random variable, with p.d.f. $g(\cdot)$ and c.d.f. $G(\cdot)$. The selling price is p per unit, which is exogenously given. (Later, we will discuss the case where the selling price is a decision variable.) If the buyer runs out of stock during the selling season, there are no replenishment opportunities and the excess demand will be lost. On the other hand, if there is excess inventory at the end of the season, it can be salvaged at v per unit. Let e be the processing cost per unit of purchased quantity, such as expenses incurred for space and handling. On the supply side, there are n (≥ 2) potential suppliers for the product. We shall retain all the assumptions made earlier about the suppliers in §2. That is, each supplier is capable of producing an unlimited quantity of the product at a constant but supplier-specific marginal cost. Recall that c_i is supplier i 's marginal cost, $i = 1, \dots, n$, and that these marginal costs are i.i.d. draws from a common probability distribution over the finite interval $[\underline{c}, \bar{c}]$, with p.d.f. $f(\cdot)$ and c.d.f. $F(\cdot)$ satisfying $F(\underline{c}) = 0$ and $F(\bar{c}) = 1$. Each supplier is privately informed of his own marginal cost. Finally, the virtual cost function, $\tilde{H}(\cdot)$, as defined in (1), is strictly increasing. The buyer and the suppliers are all expected-profits maximizers. We seek an optimal procurement strategy for the buyer.

Let Q be the level of inventory at the beginning of the selling season. The buyer's profit, excluding the costs incurred to purchase the inventory, can be expressed as

$$p \min\{Q, D\} + v(Q - D)^+ - eQ.$$

Because $\min\{Q, D\} = Q - (Q - D)^+$, the expected profit is

$$R(Q) \stackrel{\text{def}}{=} (p - e)Q - (p - v) \int_0^Q G(y) dy.$$

For convenience, we will refer to $R(\cdot)$ as the buyer's revenue function. Clearly, this function is concave with $R(0) = 0$. For the rest of the section, we will focus on the case where the revenue function is strictly concave. Note that $\lim_{Q \rightarrow +\infty} R'(Q) = v - e$, which may become negative (unlike the revenue function used in §2).

We assume that

$$p > v, \quad p - e > \underline{c}, \quad \text{and} \quad e + \underline{c} > v, \quad (12)$$

where the first inequality is natural, the second indicates that a profitable supply chain may exist, and the third eliminates the possibility of an arbitrage.

3.1. Optimal Procurement Strategies

The procurement problem facing the newsvendor is identical to the one considered in §2.⁶ Therefore, an optimal procurement strategy is to use either the quantity auction (sealed, high-bid) or the supply contract auction (multiple formats). Recall that this requires the buyer to first specify a payment function and then invite the suppliers to bid in quantity or in fee.

To determine an optimal procurement strategy, first obtain $Q^*(\cdot)$ from (6) (with the newsvendor revenue function). Note that the range of $\tilde{H}(\cdot)$ is $[\underline{c}, \bar{c} + 1/F'(\bar{c})]$ and that $R'(0) = p - e$ (assuming $G(0) = 0$). Because $p - e > \underline{c}$ (see (12)), there are only two possibilities. If $p - e > \bar{c} + 1/F'(\bar{c})$, then $Q^*(x) > 0$ for all $x \in [\underline{c}, \bar{c}]$, and by definition, $c_* = \bar{c}$. Otherwise, c_* is the solution to $p - e = \tilde{H}(c_*)$. Using the specific form of the newsvendor revenue function, we have

$$Q^*(x) = G^{-1} \left(\frac{(p - e) - \tilde{H}(x)}{p - v} \right), \quad x \in [\underline{c}, c_*],$$

where G^{-1} is the inverse of the demand distribution function G . Clearly, $Q^*(x)$ is decreasing in x . Let $[\underline{Q}, \bar{Q}]$ denote the range of $Q^*(\cdot)$.

3.1.1. Quantity Auction. Consider the quantity auction. The optimal payment function is given in (7). Although there is in general not a closed-form solution for $P^*(\cdot)$, it can be easily computed numerically. Note that $P^*(Q)$ is actually a quantity discount scheme because $P^*(Q)/Q$, the average per-unit wholesale price, is decreasing in Q . To see this, simply note from (7) that

$$\frac{P(Q^*(c))}{Q^*(c)} = c + \frac{\int_c^{c_*} Q^*(z)(1 - F(z))^{n-1} dz}{Q^*(c)(1 - F(c))^{n-1}}$$

and that the right side is increasing in c . (The proof, omitted here, does not require the specific form of the newsvendor revenue function.) Hence, under the quantity auction, the buyer demands identical quantity discounts from the suppliers and picks a winner based on the quantities they are willing to supply.⁷

The above observation provides a new rationale for using quantity discounts. Dolan (1987) gives an excel-

⁶ The only difference is that the revenue function here may be unimodal, but this is clearly inconsequential, as there is no loss of optimality if one restricts the procurement quantity to the range where the function is increasing.

⁷ The quantity auction can also be implemented by inviting the suppliers to bid in wholesale price. Define

$$w(Q) = \frac{P^*(Q)}{Q} \quad \text{for all } (Q, \bar{Q}),$$

a decreasing function. Invert this function to obtain a purchase schedule for the buyer. Consider the following sealed-bid auction. Each supplier submits a per-unit wholesale price that they would charge, and the supplier with the lowest wholesale price wins the auction. It is straightforward to show that the Bayesian-Nash equilibrium in this auction, $w^*(c)$ for $c \in [\underline{c}, \bar{c}]$, is such that $w(Q^*(c)) = w^*(c)$ for all c . Therefore, bidding in wholesale price, with the

lent review of the different forms of quantity discounts. He identified logistics system efficiency (e.g., full truck-load economies), price discrimination, and channel coordination as potential reasons for using quantity discounts. Toward the end of the paper, Dolan used competitive bidding among electronic components suppliers (the buyer being IBM) as an example to illustrate a new motivation for quantity discounts. He said “even with constant unit cost, quantity discounts may be optimal if the buyer has a desire to multisource procurement” (p. 14). In Dolan’s example, the buyer’s purchase quantity is fixed, and quantity discounts emerge as a rational bidding strategy when the buyer attaches a large enough penalty to single sourcing. (As the purchase quantity increases, the competition for each additional unit intensifies, triggering discounts.) Note that in our setting, the buyer’s purchase quantity is not fixed a priori and there is no penalty for single sourcing. Consequently, the need to integrate quantity decision with price discovery in variable-quantity procurement is a new reason for quantity discounts. The intuition is clear: the more efficient the winning supplier, the larger the purchase quantity, indicating that the per-unit price the buyer pays should decrease as she buys more. In short, quantity contains cost information.

3.1.2. Supply Contract Auction. Another optimal procurement strategy, which is in many ways more attractive than the quantity auction (see §2.3), is the supply contract auction. Here, the buyer offers a payment function (a supply contract) and invites the suppliers to bid for the right to trade. The supplier willing to pay the most for the right to trade wins and is given the decision right to choose the quantity to deliver to the buyer. The optimal supply contract, $P^{**}(\cdot)$, is characterized by (11), which has a closed-form solution for the newsvendor model. Take any $Q \in [\underline{Q}, \bar{Q}]$. Thus, there exists an $x \in [c, c_*]$ with $Q = Q^*(x)$. Note that $P^{**'}(Q) = x$ and $R'(Q) = \tilde{H}(x)$. Because $R'(Q) = (p - e) - (p - v)G(Q)$, we have

$$P^{**'}(Q) = \tilde{H}^{-1}((p - e) - (p - v)G(Q)),$$

where the inverse function $\tilde{H}^{-1}(\cdot)$ is well defined because $\tilde{H}(\cdot)$ is increasing. Therefore,

$$P^{**}(Q) = \int_{\underline{Q}}^Q \tilde{H}^{-1}((p - e) - (p - v)G(z)) dz, \quad Q \in [\underline{Q}, \bar{Q}]. \quad (13)$$

Note that adding a constant to the above payment function does not alter the marginal calculations facing the suppliers and only shifts their bids by the

same constant. Therefore, auctioning off $P^{**}(\cdot) + K$ for any constant K constitutes an optimal procurement strategy. (Adding a constant allows us to establish connections with some well-known supply contracts, as we will see later.)

The supply contract auction fits well with the use of slotting allowances and VMI, both of which are prevalent in the retail industry. Slotting allowances are lump-sum, up-front payments from a supplier to a retailer when the supplier introduces a new product to the retailer’s stores. Sometimes, such payments are made to keep an existing product on the retailer’s shelves (also called pay-to-stay fees). Therefore, a slotting allowance corresponds to the right-to-trade fee collected by the newsvendor in our model. On the other hand, VMI is a practice where the inventory replenishment decision is delegated to the supplier. In our model, this means that the supplier that has been selected is given the decision right to determine the level of inventory to stock at the beginning of the selling season. Of course, the supplier should not be given complete freedom in making the inventory decision. Instead, the decision is guided by the payment function specified by the buyer. It is interesting that in our framework, slotting allowances and VMI arise as two distinct features of an optimal procurement strategy. The existing literature on the reasons for using slotting allowances tends to focus on the signaling effect (when the supplier has more information about the product’s potential in the marketplace) or the risk-sharing effect (the retailer often incurs real, out-of-pocket costs in new-product introductions and many new products fail) (see Federal Trade Commission 2001 and the many academic papers on slotting allowances such as Shaffer 1991, Chu 1992, Lariviere and Padmanabhan 1997, Sullivan 1997, Bloom et al. 2000, and Desai 2000). Our interpretation is simply that slotting allowances are fees resulting from suppliers competing for the scarce shelf space and that they can be part of an optimal procurement strategy for the retailer. This observation adds a new dimension to the ongoing debate on the purposes and consequences of slotting allowances.

The above interpretation of the supply contract auction relies on a particular method of “bookkeeping,” i.e., the winning supplier first pays a lump-sum fee to the buyer to gain the right to trade, followed by the supplier’s production/delivery decision, and finally ending with a payment from the buyer to the supplier for the delivered quantity. As mentioned earlier, one could also ask each of the suppliers to submit a two-dimensional bid, consisting of a price and a quantity. The suppliers are then ranked according to a given scoring rule and the supplier having the highest score wins the auction. Under this implementation, we do not require the winning supplier to pay the buyer

purchase schedule properly chosen, is also optimal. Hansen (1988) studied a variable-quantity auction where the purchase quantity is an exogenously given decreasing function of the winning wholesale price.

an up-front, lump-sum fee. Consequently, whether or not we observe a slotting allowance depends on the bookkeeping method we use to implement the optimal procurement strategy. The fact that slotting allowances are prevalent in practice suggests that other factors may be at play here. For example, the buyer, given her lack of information about the new product's market potential, may want some insurance before committing her shelf space (a very important resource) to a particular supplier. An up-front fee serves just that purpose. Incorporating the buyer's aversion to risks may therefore tilt the balance toward the use of slotting allowances. This is an interesting topic for future research.

3.1.3. Returns Contracts and Revenue-Sharing Contracts. In both the quantity and the supply contract auctions, the payment function contains no uncertainty, i.e., the payment a supplier receives is a deterministic function of the quantity he delivers to the buyer. In practice, we often see supply contracts where the transfer payment also depends on the realized demand (see Cachon 2003 for a comprehensive review of the supply contract literature). Next, we discuss the implications of using some of these supply contracts in our procurement setting.

Let $P(Q, D)$ be the transfer payment (from the buyer to a supplier) if Q units are delivered to the buyer and the realized demand is D . Consider, e.g., the returns contract, whereby the buyer pays a per-unit wholesale price w for every unit of inventory stocked at the beginning of the selling season, and if there is leftover inventory at the end of the selling season, the buyer obtains a rebate of b per unit of excess inventory. (The buyer does not physically return the leftover inventory to the supplier, but salvages the excess inventory for v per unit. Other arrangements are possible, but are not considered here.) Under such a contract,

$$P(Q, D) = wQ - b(Q - D)^+.$$

Another example is the revenue-sharing contract. Here the buyer pays per-unit wholesale price w for the initial inventory and promises to transfer $\alpha \in (0, 1)$ fraction of her revenues (regular sales and salvage sales) to the supplier. In this case,

$$P(Q, D) = wQ + \alpha(p \min\{Q, D\} + v(Q - D)^+).$$

Suppose our newsvendor auctions off the supply contract $P(Q, D)$. Because the suppliers are risk neutral, they value the contract at $E_D P(Q, D) \stackrel{\text{def}}{=} \tilde{P}(Q)$. If $\tilde{P}(\cdot) = P^*(\cdot)$ (respectively, $P^{**}(\cdot)$), then auctioning off the contract by inviting the suppliers to bid in quantity (respectively, up-front fee) maximizes the newsvendor's expected profit. Below we provide an

example where an optimal procurement strategy can be implemented by auctioning off either a returns contract or a revenue-sharing contract.

Assume $F(x) = (x - \underline{c})/(\bar{c} - \underline{c})$ for all $x \in [\underline{c}, \bar{c}]$, i.e., the marginal production cost is uniformly distributed. Hence, $\bar{H}(x) = 2x - \underline{c}$, $x \in [\underline{c}, \bar{c}]$. From (13),

$$P^{**}(Q) = \int_0^Q \frac{1}{2} [(p - e) - (p - v)G(z) + \underline{c}] dz.$$

(Here we have added a constant, \int_0^Q , to the payment function. As noted earlier, the resulting payment function is still optimal.) It can be easily verified that a returns contract with

$$w = \frac{p - e + \underline{c}}{2} \quad \text{and} \quad b = \frac{p - v}{2}$$

satisfies $\tilde{P} = P^{**}$ for any G .⁸ On the other hand, an alternative way to express the above payment function is

$$P^{**}(Q) = \frac{\underline{c} - e}{2} Q + \frac{p}{2} E_D \min\{Q, D\} + \frac{v}{2} E_D (Q - D)^+.$$

Therefore, a revenue-sharing contract with

$$w = \frac{\underline{c} - e}{2} \quad \text{and} \quad \alpha = \frac{1}{2}$$

also achieves $\tilde{P} = P^{**}$ for any G .⁹ An optimal procurement strategy for the newsvendor is, therefore, to auction off a returns contract or a revenue-sharing contract with the above parameters in the supply contract auction. Note that the contract specification does not require the knowledge of the demand distribution.

Now we have seen an example where specifying the payment function either as $P(Q)$ or $P(Q, D)$ makes no difference for the parties of the trade. This is no longer true if there is information asymmetry regarding the demand distribution or there are differences in risk attitudes.

Suppose that only the buyer knows the demand distribution $G(\cdot)$. In this case, the buyer can compute the optimal payment function $P(\cdot)$ for either the quantity auction or the supply contract auction. Because the payment is independent of demand, there is no need for the suppliers to know the demand distribution. The buyer's expected profits are maximized (as if the suppliers had full knowledge about the demand distribution). In contrast, if the buyer chooses to specify the payment function as $P(Q, D)$ (through, e.g., a returns or a revenue-sharing contract), then the suppliers are

⁸ Note that b may be higher than w . But this is fine because the supplier chooses the quantity to deliver to the buyer in the supply contract auction and because $b + v < p$, which means the buyer does not want to turn away demand.

⁹ If $w < 0$, then the supplier subsidizes the initial inventory before taking a share of the buyer's revenues.

forced to evaluate this contract based on their imperfect information about demand. As a result, their bidding strategies will be different from their strategies under full knowledge about the demand distribution. The buyer’s expected profits depend on the nature of the suppliers’ imperfect knowledge about demand. If the suppliers happen to believe that demand is higher than it really is, then their evaluation of the supply contract may be inflated and the buyer may even be better off (relative to her expected profits under the optimal procurement strategy assuming symmetric demand information).

On the other hand, suppose that the buyer does not know the demand distribution, but the suppliers do. In this case, if the buyer chooses to offer a payment function in the form of $P(Q)$, then she must first calculate the revenue function $R(Q)$ based on her faulty knowledge of the demand distribution. The resulting payment function is likely to be different from what she would offer given true demand distribution. The outcome is that the payment function is suboptimal and the buyer’s expected profit is not maximized. (In this case, the suppliers’ knowledge of the demand distribution goes wasted because, as mentioned above, they have no use for this knowledge when given a deterministic payment function.) Now consider the alternative of specifying the payment function in the form of $P(Q, D)$. As our earlier example demonstrates, when the marginal production costs are uniformly distributed, the design of an optimal procurement strategy does not require the knowledge of the demand distribution if the buyer chooses to write the payment function as a returns contract or a revenue-sharing contract.¹⁰ Therefore, the buyer suffers no loss of optimality even without knowledge of the demand distribution. (When the suppliers evaluate the contract, i.e., $E_D P(Q, D)$, they use the true demand distribution.) This rather surprising result is due to the specific form of the optimal payment function for the supply contract auction, $P^{**}(Q)$, and its close connection to the returns or revenue-sharing contract. (The form of the optimal payment function for the quantity auction, however, does not seem to lend itself to this connection.)

Note that specifying the payment function as $P(Q, D)$ instead of $P(Q)$ makes the suppliers bear more risk. Of course, in the current model setup with risk neutrality, the suppliers are indifferent. An interesting future research topic is to investigate the impact of risk aversion on the form of the optimal procurement strategy.

3.1.4. Impact of Supply and Demand Characteristics. We next consider how the demand- and supply-

¹⁰This is reminiscent of Pasternack (1985), where a coordinating contract is found to be independent of the demand distribution.

side characteristics affect the newsvendor’s profit. To this end, suppose D is normally distributed with mean μ and standard deviation σ . As this assumption implies the possibility of negative demand,¹¹ we modify the revenue function as follows:

$$R(Q) = (p - e)Q - (p - v) \int_{-\infty}^Q G(y) dy. \quad (14)$$

In this case, we have

$$Q^*(x) = \mu + \sigma \Phi^{-1} \left(\frac{p - e - \tilde{H}(x)}{p - v} \right), \quad x \in [\underline{c}, c_*], \quad (15)$$

where $\Phi(\cdot)$ is the c.d.f. of the standard normal and Φ^{-1} its inverse. Recall that c_* is determined as follows. If $p - e > \tilde{H}(x)$ for all $x \in [\underline{c}, \bar{c}]$, then $c_* = \bar{c}$. Otherwise, c_* is the solution to $p - e = \tilde{H}(c_*)$.

With the normal demand distribution, the buyer’s maximum expected profit, given in Theorem 1, can be expressed as

$$\Pi^* = a\mu - b\sigma, \quad (16)$$

where

$$a = \int_{\underline{c}}^{c_*} [p - e - \tilde{H}(x)] f_{(1)}(x) dx \quad \text{and}$$

$$b = (p - v) \int_{\underline{c}}^{c_*} \phi \left(\Phi^{-1} \left(\frac{p - e - \tilde{H}(x)}{p - v} \right) \right) f_{(1)}(x) dx, \quad (17)$$

with $\phi(\cdot)$ being the density function of the standard normal. Note from (16) that the impacts of the demand- and supply-side characteristics on Π^* are “decoupled”: a and b depend on the number of suppliers and their marginal cost distribution, but are independent of μ and σ , which are of course the demand-side characteristics. Note also that $a\mu$ would be the buyer’s expected profit if there were no demand uncertainty (the expectation is with respect to the supply-side uncertainty), where a is the buyer’s expected profit margin.¹² On the other hand, $b\sigma$ is the cost of demand uncertainty, which increases linearly in the demand standard deviation. The coefficient b thus represents the buyer’s marginal benefit from reduction in demand uncertainty.

It is easy to see that a is increasing in supply-side competition (i.e., n), which also means that there is synergy between increasing the number of potential suppliers and increasing the mean demand, i.e., increasing n increases the benefit of increasing μ

¹¹The likelihood of negative demand should be negligible for the normal assumption to be plausible.

¹²Note that $a\mu$ is also the newsvendor’s expected profit under the make-to-order regime, whereby she waits until the demand is revealed and then procures the exact quantity needed from the lowest-cost supplier via an auction. Of course, this mode of operation is not always feasible: it depends on how patient the customers are and how fast the suppliers can deliver the product.

and vice versa. The relationship between b and n is, however, not as simple. Define $\theta(x) = \phi(\Phi^{-1}((p - e - \tilde{H}(x))/(p - v)))$, $x \in [\underline{c}, c_*]$. Note that $\theta(\cdot)$ is independent of n . (Hereafter, all proofs can be found in the online appendix, which is provided in the e-companion.)¹³

LEMMA 1. *If $\tilde{H}(x) \leq (p + v)/2 - e$ for all $x \in [\underline{c}, c_*]$, then $\theta'(x) > 0$ for all $x \in [\underline{c}, c_*]$. Conversely, if $\tilde{H}(x) \geq (p + v)/2 - e$ for all $x \in [\underline{c}, c_*]$, then $\theta'(x) < 0$ for all $x \in [\underline{c}, c_*]$.*

THEOREM 4. *As n increases, the coefficient b decreases (respectively, increases) if $\tilde{H}(x) \leq$ (respectively, \geq) $(p + v)/2 - e$ for all $x \in [\underline{c}, c_*]$.*

For the case with $\tilde{H}(\underline{c}) < (p + v)/2 - e < \tilde{H}(\bar{c})$, the coefficient b is no longer monotone in n (see Chen 2001).

One may interpret $(p + v)/2$ as the average revenue per unit of production (averaging between regular sales and salvage sales). Thus, $(p + v)/2 - e$ is the average profit margin, excluding procurement costs. Theorem 4 indicates that when the margin is high relative to the procurement costs (as measured by the virtual cost $\tilde{H}(\cdot)$), b decreases as n increases. In this case, an increase in supply-side competition decreases the cost of demand uncertainty and reduces the incentive for demand variance reduction. The reverse is true in the low-margin case.

The intuition behind the above results can be better understood by considering the traditional newsvendor model where the per-unit procurement cost is exogenously specified. Let c be the per-unit procurement cost. Then, $\Pi^* = a_0\mu - b_0\sigma$, where $a_0 = p - e - c$ and $b_0 = (p - v)\phi(\Phi^{-1}((p - e - c)/(p - v)))$. It is easy to verify that for $c > (p + v)/2 - e$, b_0 increases as c decreases, and the reverse is true for the range $c < (p + v)/2 - e$. (These are parallel to the observations made in Theorem 4 by replacing c with the virtual cost.) For the traditional newsvendor model, first note that the expected profit is essentially independent of demand uncertainty for either very large Q or very small Q . For example, if the order quantity is very large, the newsvendor never runs out of stock and her expected profit is only a function of the mean demand. Similarly, if the order quantity is very small, the newsvendor always runs out of stock, in which case her expected profit is again only a function of the mean demand. This suggests that b_0 is small for extreme values of c . It may also suggest that b_0 is maximal for “intermediate” values of c . This is indeed true. To see this, note that b_0 is maximized when $c = (p + v)/2 - e$, which corresponds to $b_0 = (p - v)\phi(0)$ and an

order quantity of $Q = \mu$. Therefore, the newsvendor’s profit is most sensitive to demand uncertainty when the order quantity is of “intermediate” values. Finally, note that as c decreases, the order quantity increases. In the range $c > (p + v)/2 - e$ (the low-margin case), a decrease in c moves the order quantity closer to the “middle ground,” making the profit more sensitive to demand uncertainty (i.e., b_0 increases). On the other hand, for the range $c < (p + v)/2 - e$, a decrease in c moves the order quantity away from the “middle ground,” making the profit less sensitive to demand uncertainty (i.e., b_0 decreases).

3.1.5. Newsvendor with Pricing. So far, the selling price p is fixed. Now suppose it is a decision variable. Let $D(p)$ be the total demand during the selling season, a random variable whose distribution depends on p . Suppose the pricing decision is made at the beginning of the season, after the initial inventory level is known. Define

$$\tilde{R}(Q, p) = E[p \min\{Q, D(p)\} + v(Q - D(p))^+ - eQ].$$

Let $R(Q) = \max_p \tilde{R}(Q, p)$. To the extent that $R(\cdot)$ is concave, our previous results can be applied to obtain an optimal procurement strategy by simply using this newly defined $R(\cdot)$ in place of the old revenue function. Limited numerical tests suggest that this $R(\cdot)$ is indeed concave if $D(p) = d(p) + \epsilon$, where $d(p)$ is a downward-sloping, linear function of p and ϵ is a normal random shock. It is interesting to note that the concavity of the above-defined $R(\cdot)$ is seldom discussed in the voluminous literature on the (traditional) newsvendor model with pricing (see, e.g., Porteus 1990, Petruzzi and Dada 1999, and Dana and Petruzzi 2001). Analytical characterization of the behavior of $R(\cdot)$ is a good topic for future research.

3.2. Fixed-Quantity Auctions

An alternative procurement strategy for the buyer is to first determine a purchase quantity, Q , and then in an auction, find the lowest-cost supplier to produce and deliver the Q units of the product. This procurement strategy, where the input quantity is fixed before the auction, is suboptimal because the quantity decision only reflects the distributional knowledge about the suppliers’ costs (as opposed to the actual costs), but it is clearly simpler than auctioning off a supply contract, as required in an optimal procurement strategy. Later, we will compare these two strategies in terms of the newsvendor’s expected profit.¹⁴

Given Q , optimal auction design is easy: simply choose any of the commonly used auction formats

¹³ An electronic companion of this paper is available as part of the online version that can be found at <http://manscijournal.informs.org/>.

¹⁴ There is a body of research investigating the trade-off between control losses and informational gains when a decision is delegated (see Melumad et al. 1997 and the references therein). When a principal transfers a decision right to an agent with private information, the principal loses control, but gains informational advantage

such as the English auction, the Dutch auction, the first-price, sealed-bid auction, or the Vickrey auction and set a (per-unit) reserve price c_* . (This is essentially the well-known single-unit procurement auction, with the “single unit” being the lot Q .) With the optimal design, the buyer’s expected profit as a function of Q and c_* is

$$\begin{aligned} \Pi(Q, c_*) &= \int_{\underline{c}}^{c_*} [R(Q) - Q\tilde{H}(x)]f_{(1)}(x) dx \\ &= R(Q)F_{(1)}(c_*) - Q \int_{\underline{c}}^{c_*} \tilde{H}(x)f_{(1)}(x) dx, \end{aligned} \quad (18)$$

where, as before, $F_{(1)}$ and $f_{(1)}$ are, respectively, the c.d.f. and the p.d.f. of C_1 , the lowest per-unit production cost among the suppliers. The expression is intuitive as $R(Q)$ is the value of the lot (of Q units), acquired at per-unit virtual cost $\tilde{H}(x)$ if $C_1 = x$. The optimal values of the two decision variables, Q and c_* , can be found via an iterative procedure. Due to space limit, the detailed algorithmic development is presented in the online appendix.

3.3. Numerical Examples

We use numerical examples to compare the fixed-quantity auction with the optimal strategy that auctions off a supply contract. The comparison illustrates the gap between the two strategies in terms of the newsvendor’s expected profit and how the gap depends on some of the model parameters.

For the numerical examples, we assumed that the suppliers’ marginal production costs are drawn from a uniform distribution and the demand is normally distributed with mean μ and standard deviation σ . The following parameters were used: $n = 2, 4$; $(\underline{c}, \bar{c}) = (3, 8), (1, 10)$; $p = 10, 15$; $e = v = 0$; $\mu = 10, 15$; $\sigma = 3, 5$. Note that the two ranges for the uniform distribution are mean preserving. There are a total of 32 examples. For each example, we computed the newsvendor’s maximum expected profit (achieved with an optimal procurement strategy), and for the fixed-quantity auction, the optimal quantity, the corresponding reserve cost, and expected profit for the newsvendor. (The iterative procedure described in §3.2 converged to a point in a few iterations. Theorem 5 guarantees that the point is the optimal solution.) The percentage increase

because the decision, now made by the agent, reflects information that was not accessible to the principal. In some cases, the informational gain outweighs the control loss, suggesting that the principal is better off delegating. This is similar to the comparison between the fixed-quantity auction and either the quantity auction or the supply contract auction in our procurement context. With the fixed-quantity auction, the buyer (principal) takes back the quantity decision right, but loses the informational advantage because the buyer’s quantity decision is based on her distributional knowledge about the suppliers’ costs. In our setting, as the paper has shown, the buyer prefers delegating.

Table 1 Comparison between Optimal Strategy and Fixed Quantity Auctions

No.	n	(\underline{c}, \bar{c})	p	μ	σ	Opt. profit	Fixed qty.	Reserve price	Fixed-qty. profit	% diff.
1	2	(3, 8)	10	10	3	28.42	9.58	5.98	27.25	4
2	2	(3, 8)	10	10	5	22.24	9.68	5.55	20.90	6
3	2	(3, 8)	10	15	3	47.35	14.49	6.17	45.91	3
4	2	(3, 8)	10	15	5	41.43	14.36	5.91	39.23	6
5	2	(3, 8)	15	10	3	70.21	10.59	7.93	69.06	2
6	2	(3, 8)	15	10	5	59.47	11.09	7.25	57.43	4
7	2	(3, 8)	15	15	3	114.07	15.59	8.00	112.39	1
8	2	(3, 8)	15	15	5	103.35	15.99	7.81	100.65	3
9	2	(1, 10)	10	10	3	30.64	10.39	4.82	29.00	6
10	2	(1, 10)	10	10	5	25.70	11.18	4.32	24.01	7
11	2	(1, 10)	10	15	3	48.91	15.25	5.07	47.57	3
12	2	(1, 10)	10	15	5	44.58	15.74	4.74	42.19	6
13	2	(1, 10)	15	10	3	68.28	10.80	6.86	64.56	6
14	2	(1, 10)	15	10	5	59.21	11.78	6.08	54.92	8
15	2	(1, 10)	15	15	3	108.35	15.69	7.25	104.61	4
16	2	(1, 10)	15	15	5	98.28	16.41	6.73	94.32	4
17	4	(3, 8)	10	10	3	39.06	10.12	5.88	38.30	2
18	4	(3, 8)	10	10	5	31.95	10.36	5.45	30.72	4
19	4	(3, 8)	10	15	3	64.46	15.08	6.09	63.27	2
20	4	(3, 8)	10	15	5	56.74	15.22	5.81	55.53	2
21	4	(3, 8)	15	10	3	84.27	11.29	7.70	83.64	1
22	4	(3, 8)	15	10	5	73.94	12.16	6.99	72.74	2
23	4	(3, 8)	15	15	3	134.55	16.29	8.00	133.64	1
24	4	(3, 8)	15	15	5	123.85	17.15	7.58	122.73	1
25	4	(1, 10)	10	10	3	46.13	10.92	4.72	44.58	3
26	4	(1, 10)	10	10	5	40.29	11.88	4.21	38.14	6
27	4	(1, 10)	10	15	3	73.50	15.83	4.98	72.00	2
28	4	(1, 10)	10	15	5	68.01	16.58	4.63	65.21	4
29	4	(1, 10)	15	10	3	90.62	11.62	6.59	88.51	2
30	4	(1, 10)	15	10	5	81.36	12.89	5.81	78.61	3
31	4	(1, 10)	15	15	3	142.46	16.58	7.03	140.43	1
32	4	(1, 10)	15	15	5	133.72	17.73	6.46	130.25	3

in profit from a fixed-quantity auction to an optimal strategy is also calculated. These results are summarized in Table 1. The gap ranges from 1% to 8%, with the average at 3%.

Table 2 shows how the gap between the two strategies depends on the model parameters. A larger gap is observed when (1) there are fewer potential suppliers, (2) the marginal-cost distribution is more spread out, (3) the retail price is lower, and (4) the coefficient of variation for the demand distribution is higher. These trends are intuitive. Recall that the fixed-quantity auction only uses the marginal-cost distribution and as a result, the quantity decision is not influenced by the minimum realized cost. To the contrary, the quantity decision under an optimal strat-

Table 2 Sensitivity Results

n	% diff.	(\underline{c}, \bar{c})	% diff.	p	% diff.	% diff. μ	σ	
							3 (%)	5 (%)
2	4.46	(3, 8)	2.68	10	4.13	10	3.25	4.95
4	2.47	(1, 10)	4.25	15	2.80	15	2.14	3.53

egy depends on the minimum realized cost via the $Q^*(\cdot)$ function. This informational loss is a key reason for the profit gap between the two strategies. When there are many potential suppliers, the distribution of C_1 becomes very narrow, implying that the informational loss is quite small. Put differently, when there are fewer potential suppliers, the informational loss becomes larger. This explains (1). The same kind of reasoning explains (2). When the retail price is lower, the cost side plays a relatively larger role in determining the expected profit. Similarly, an increase in the demand's coefficient of variation depresses the newsvendor's revenues, increasing the relative importance of the procurement costs. Both cases magnify the informational loss. This explains (3) and (4).¹⁵

4. Concluding Remarks

A main observation of this paper is that when the purchase quantity is a decision variable and the price needs to be discovered, optimal procurement can be achieved by embedding a supply contract within an auction mechanism. This is done by inviting the suppliers to compete for a supply contract, and the competition can take place in the space of quantity offers, wholesale price quotes, or up-front fees for the right to trade. Under an optimal strategy, the ultimate quantity and price decisions come from both sides of the trade: the buyer specifies the contract and the auction rules, and the suppliers bid with their private cost information in mind.

A special case of the procurement model is a newsvendor model with supply-side competition. Here, an optimal procurement strategy, namely, the supply contract auction, offers a new explanation for the use of slotting allowances and VMI, both of which are prevalent in the retail industry. Under this strategy, the winning supplier is the one willing to pay the highest up-front fee (or slotting allowance) and is given the decision right to make the quantity decision. In some cases, the optimal supply contract to be auctioned off can be written as a returns contract or a revenue-sharing contract. These contracts are well known for their abilities to coordinate supply chains, but as we have shown here, they are also useful in procurement settings. Moreover, when the buyer has inferior information about demand vis-a-vis the suppliers, writing the supply contract as a returns contract or a revenue-sharing contract can mitigate her informational disadvantage.

This paper has also studied an alternative procurement strategy, namely, the fixed-quantity auction,

where the buyer first determines a purchase quantity and then seeks the lowest-cost supplier to deliver the quantity. Although suboptimal, this strategy is simpler than the optimal strategy because announcing a quantity is easier than specifying a contract. In a set of numerical examples (newsvendor problems with supply-side competition), the gap (relative difference in the buyer's expected profit) between the fixed-quantity auction and the optimal strategy ranges from 1% to 8%, with an average of 3%. This gap increases when the number of potential suppliers is fewer, the distribution for the suppliers' marginal costs has a larger variance, the retail price is lower, or the coefficient of variation for demand is higher.

Interesting research questions abound. When the marginal production costs are increasing, sole sourcing is no longer optimal. DS provide an optimal mechanism that uses the information received in the competitive bidding process to allocate production among the bidders. It would be interesting to see if there exist alternative, simpler mechanisms. On the other hand, if production exhibits economies of scale, then sole sourcing should remain optimal, but the optimal procurement strategy is still unknown. One possibility is to ask the suppliers to make supply-curve bids (i.e., quantity-discount schedules) (see Hohner et al. 2003). It is unclear if this design is optimal. Procurement models with asymmetric or risk-averse suppliers are interesting topics as well. For optimal auction designs when the bidders are either asymmetric or risk averse, see Myerson (1981) and Maskin and Riley (1984). These papers deal with single-object auctions, and it would be interesting to see how the optimal designs change when the quantity being sold (or procured) is variable. Finally, the optimal procurement strategy identified for our model represents the best the buyer can do given her information about the suppliers' costs. It is conceivable that if the buyer can extract more information from the suppliers about their costs, she can improve her profit. Seshadri and Zemel (2003) have explored this idea by considering audits as a means to extract information.

5. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at <http://mansci.journal.informs.org/>.

Acknowledgments

The author expresses his gratitude to Bill Lovejoy, the department editor for this paper, and the anonymous reviewers for their constructive comments and suggestions, which have led to significant improvements in both the presentation and the content of this paper. This paper has also benefited from discussions with many seminar participants. The first version of this paper was written during the

¹⁵ For the normal distribution to be a good description of demand, the coefficient of variation cannot be too large. This puts a limit on the gap between the two strategies. With a different distribution, a larger gap may appear.

author's sabbatical leave at the Stanford University Graduate School of Business. Thanks to the operations, information, and technology group of the Stanford Business School for their hospitality. Financial support from the National Science Foundation, the Columbia Business School, and the Stanford Global Supply Chain Management Forum is gratefully acknowledged.

References

- Bloom, P. N., G. T. Gundlach, J. P. Cannon. 2000. Slotting allowances and fees: Schools of thought and the views of practicing managers. *J. Marketing* 64(2) 92–108.
- Branco, F. 1997. The design of multidimensional auctions. *RAND J. Econom.* 28(1) 63–81.
- Cachon, G. P. 2003. Supply chain coordination with contracts. A. G. de Kok, S. C. Graves, eds. *Handbooks in OR and MS*, Vol. 11. Elsevier Science Publishers, North-Holland, Amsterdam, The Netherlands, 229–339.
- Che, Y.-K. 1993. Design competition through multidimensional auctions. *RAND J. Econom.* 24(4) 668–680.
- Chen, F. 2001. Auctioning supply contracts (version 1). Working paper, Columbia Business School, New York.
- Chen, F. 2003a. Auctioning supply contracts (version 2). Working paper, Columbia Business School, New York.
- Chen, F. 2003b. Information sharing and supply chain coordination. A. G. de Kok, S. C. Graves, eds. *Handbooks in OR and MS*, Vol. 11. Elsevier Science Publishers, North-Holland, Amsterdam, The Netherlands, 341–421.
- Chu, W. 1992. Demand signaling and screening in channels of distribution. *Marketing Sci.* 11(4) 327–347.
- Dana, J. D., Jr., N. C. Petruzzi. 2001. The newsvendor model with endogenous demand. *Management Sci.* 47(11) 1488–1497.
- Dasgupta, S., D. F. Spulber. 1990. Managing procurement auctions. *Inform. Econom. Policy* 4(1) 5–29.
- Desai, P. 2000. Multiple messages to retain retailers: Signaling new product demand. *Marketing Sci.* 19(4) 381–389.
- Dolan, R. J. 1987. Quantity discounts: Managerial issues and research opportunities. *Marketing Sci.* 6(1) 1–22.
- Elmaghraby, W. J. 2000. Supply contract competition and sourcing policies. *Manufacturing Service Oper. Management* 2(4) 350–371.
- Federal Trade Commission. 1999. FTC official testifies before House Committee on slotting allowances, <http://www.ftc.gov/opa/1999/10/slotting.htm>.
- Federal Trade Commission. 2001. Report on the Federal Trade Commission workshop on slotting allowances and other marketing practices in the grocery industry, <http://www.ftc.gov/os/2001/02/slottingallowancesreportfinal.pdf>.
- Gallien, J., L. M. Wein. 2000. Design and analysis of a smart market for industrial procurement. Working paper, Massachusetts Institute of Technology, Cambridge, MA.
- Hansen, R. G. 1988. Auctions with endogenous quantity. *Rand J. Econom.* 19(1) 44–58.
- Hohner, G., J. Rich, E. Ng, G. Reid, A. J. Davenport, J. R. Kalagnanam, H. S. Lee, C. An. 2003. Combinatorial and quantity-discount procurement auctions benefit Mars, Incorporated and its suppliers. *Interfaces* 33(1) 23–35.
- Jin, M., D. Wu. 2000. Supply chain contracting in electronic markets: Incentives and coordination mechanisms. Working paper, Lehigh University, Bethlehem, PA.
- Klemperer, P. 1999. Auction theory: A guide to the literature. *J. Econom. Surveys* 13(3) 227–286.
- Kreps, D. 1990. *A Course in Microeconomic Theory*. Princeton University Press, Princeton, NJ.
- Laffont, J.-J., J. Tirole. 1987. Auctioning incentive contracts. *J. Political Econom.* 95(5) 921–937.
- Laffont, J.-J., J. Tirole. 1993. *A Theory of Incentives in Procurement and Regulation*. MIT Press, Cambridge, MA.
- Lariviere, M., V. Padmanabhan. 1997. Slotting allowances and new product introductions. *Marketing Sci.* 16(2) 112–128.
- Lee, H., S. Whang. 2002. The impact of the secondary market on the supply chain. *Management Sci.* 48(6) 719–731.
- Maskin, E., J. Riley. 1984. Optimal auctions with risk averse buyers. *Econometrica* 52(6) 1473–1518.
- McAfee, R. P., J. McMillan. 1986. Bidding for contracts: A principal-agent analysis. *Rand J. Econom.* 17(3) 326–338.
- McAfee, R. P., J. McMillan. 1987a. Auctions and bidding. *J. Econom. Literature* 25(2) 699–738.
- McAfee, R. P., J. McMillan. 1987b. Competition for agency contracts. *Rand J. Econom.* 18(2) 296–307.
- Melumad, N., D. Mookherjee, S. Reichelstein. 1997. Contract complexity, incentives, and the value of delegation. *J. Econom. Management Strategy* 6(2) 257–289.
- Mendelson, H., T. Tunca. 2000. Business to business exchanges and supply chain contracting. Extended abstract, Stanford Business School, Stanford, CA.
- Myerson, R. B. 1981. Optimal auction design. *Math. Oper. Res.* 6(1) 58–73.
- Pasternack, B. 1985. Optimal pricing and returns policies for perishable commodities. *Marketing Sci.* 4(2) 166–176.
- Petruzzi, N. C., M. Dada. 1999. Pricing and the newsvendor problem: A review with extensions. *Oper. Res.* 47(2) 183–194.
- Porteus, E. L. 1990. Stochastic inventory theory. D. P. Heyman, M. J. Sobel, eds. *Handbooks in OR and MS*, Vol. 2. Elsevier Science Publishers, North-Holland, Amsterdam, The Netherlands, 605–652.
- Rao, A. R., H. Mahi. 2003. The price of launching a new product: Empirical evidence on factors affecting the relative magnitude of slotting allowances. *Marketing Sci.* 22(2) 246–268.
- Riley, J. G., W. F. Samuelson. 1981. Optimal auctions. *Amer. Econom. Rev.* 71(3) 381–392.
- Riordan, M. H., D. E. M. Sappington. 1987. Awarding monopoly franchises. *Amer. Econom. Rev.* 77(3) 375–387.
- Rosling, K. 2002. Inventory cost rate functions with nonlinear shortage costs. *Oper. Res.* 50(6) 1007–1017.
- Seshadri, S., E. Zemel. 2003. Supply intelligence. Working paper, New York University, New York.
- Shaffer, G. 1991. Slotting allowances and resale price maintenance: A comparison of facilitating practices. *RAND J. Econom.* 22(1) 120–135.
- Sullivan, M. 1997. Slotting allowances and the market for new products. *J. Law Econom.* 40(2) 461–493.
- Tirole, J. 1988. *The Theory of Industrial Organizations*. MIT Press, Cambridge, MA.
- Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *J. Finance* 16(1) 8–37.