INFORMS TutORials in Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Contract Complexity and Performance

Feryal Erhun

TUTORIALS IN OPERATIONS RESEARCH
INFORMS 2011

 | ISBN 978-0-9843378-2-8
doi: http://dx.doi.org/10.1287/educ.1110.0089

Contract Complexity and Performance

Feryal Erhun
Management Science and Engineering, Stanford University, Stanford, California 94305,
feryal.erhun@stanford.edu

Abstract This tutorial discusses theoretical and behavioral aspects of contract design with an emphasis on complexity as a design factor. We first introduce optimal contracts under different complete and asymmetric information settings. We then demonstrate how simple quantity discount contracts can effectively eliminate inefficiencies under some conditions. We present behavioral aspects of contract design and discuss trade-offs between contract design and contract performance. Finally, we introduce the concept of dynamic procurement as a means to implement nonlinear pricing schemes, and we establish a link between dynamic procurement and quantity discounts.

Keywords contract design; complex contracts; contract performance; supply chain efficiency; price-only contracts; quantity discount contracts; asymmetric information; behavioral operations management

1. Introduction

Outsourcing forces manufacturers to increasingly rely on external suppliers to procure goods and services. Today, approximately half of the revenue generated in the U.S. manufacturing industry is spent on procurement (U.S. Department of Commerce [64]). Moreover, because products and their production processes have become more complex, firms are now buying more complex items and services. Hence it is critical for firms to streamline these intricate procurement processes to maintain a competitive edge in the market. Many factors influence the design of procurement contracts such as price, availability, cost, and delivery schedule. In this tutorial, we provide a brief overview of some of the commonly used procurement contracts (price-only contracts and quantity discount contracts) and discuss *contract complexity* as a design factor.

Price-only contracts are common in the supply chain management literature,[1] motivated by the observation that "many supply-chain transactions are governed by simple [price-only] contracts defined only by a per-unit wholesale price" (Lariviere and Porteus [39], p. 293). Price-only contracts have been used for a wide variety of products and services, including paper (Norske Skog [47]); alarm systems (Visual Defence [65]); pharmaceuticals (Cui et al. [19]); software (Robinson [52]); components for airplanes (Boeing [7]); electronics design, assembly, and components including polycrystaline silicon (Space Daily [58]); healthcare services (FD Wire [23]); and security services in the Iraq war (HT Media [30]).

The price-only contract is simple because it requires only the specification of a single parameter: the wholesale price. On the other hand, the literature praises the more complex quantity discount contract because it has been shown theoretically to increase sales, reduce costs, increase channel efficiency, allow for self-selected price discrimination, and eliminate

[1] See, for example, Anupindi and Bassok [3], Bernstein and DeCroix [5], Bernstein and Federgruen [6], Cachon [12], Cachon and Lariviere [13], Chen et al. [17], Cui et al. [19], Gerchak and Wang [26], Guo et al. [28], Lariviere and Porteus [39], Majumder and Srinivasan [41], Taylor and Plambeck [60], Tomlin [62], Wang and Gerchak [66], and references therein.

inefficiencies as a result of information asymmetry (Altintas et al. [2], Burnetas et al. [11], Chu and Sappington [18], Jeuland and Shugan [31]). And indeed, quantity discount contracts are also observed regularly in practice. However, there is evidence that quantity discounts do not perform consistently in practice (Altintas et al. [2]). One underlying reason for this inconsistency was identified by a recent Adesso Solutions [1] study, which found that fewer than 10% of participating suppliers believe their discounts and promotions management are effective. This lack of guidance in designing quantity discounts is also the reason that discounts and promotions have been described by Sellers and Reese [56] as "the dumbest marketing ploy ever."

Given the conflicting theoretical and practical accounts of simple contracts (price-only contracts) and more complex contracts (quantity discount contracts), in this tutorial our goal is to compare the two and evaluate contract complexity as a design factor. The rest of this tutorial is organized as follows. In §2, we analyze a stylized two-stage supply chain with one supplier and one manufacturer under complete information. In §3, we relax the assumption of complete information and characterize the optimal contracts for a setting with asymmetric cost information by restricting ourselves to screening-based contract design. (We characterize the optimal contracts for a setting with asymmetric demand information in the appendix.) Our goals in studying the settings with asymmetric information are (1) to establish that optimal contracts are in the form of quantity discounts and (2) to demonstrate that, especially under asymmetric information, these contracts may be arbitrarily complex (i.e., the optimal quantity discount scheme may require arbitrarily many prices and discount breaks). In §§4 and 5, we limit the complexity of these quantity discount contracts by restricting the number of prices and discount breaks that the contract may have. We theoretically assess the performance of these simpler contracts under asymmetric cost information in §4. In §5, we assess the performance of these simpler contracts under asymmetric demand information by using human subject experiments. Finally, in §6, we show that firms may implement optimal contracts with repeated interactions of dynamic procurement.

Throughout this tutorial, subscripts under expectation operators denote the random variable over which expectation is taken. We use the convention $(x)^+$ for $\max(x, 0)$. Comparison terms such as *greater*, *smaller*, *increasing*, and *decreasing* should be interpreted in the weak sense (allowing for equality). We have omitted the proofs, and we refer the reader to respective papers for detailed analyses.

2. Optimal Contracts Under Complete Information

Consider a simple stylized two-stage supply chain with one supplier and one manufacturer, where the manufacturer (she) buys goods from the supplier (he) and sells them in an end-consumer market.[2] The supplier's unit cost of production is c. The manufacturer faces a market where the price is inversely related to the quantity sold. For simplicity, we assume a linear demand curve $P = a - bq$, where P is the market price, a is the market potential, b is the price sensitivity, and q is the quantity sold by the manufacturer. Assume that all of this information is common knowledge.

Let us analyze the simple price-only contract where the supplier charges the manufacturer w per unit under this setting. The supplier's (S) and the manufacturer's (M) profits are $\Pi_S = (w - c)q$ and $\Pi_M = (a - bq - w)q$, respectively. The supply chain's profits are $\Pi = \Pi_S + \Pi_M = (a - bq - c)q$. Note that the choice of w only indirectly affects the total supply chain profits through the choice of q.

[2] The material in this section follows Erhun and Keskinocak [20] closely.

TABLE 1. Price-only contract vs. centralized supply chain.

	Decentralized supply chain	Centralized supply chain
Supplier's wholesale price (w)	$w = (a+c)/2$	w
Manufacturer's quantity (q)	$q = (a-c)/(4b)$	$q^* = (a-c)/(2b)$
Market price (P)	$P = (3a+c)/4$	$P^* = (a+c)/2$
Supplier's profit (Π_S)	$\Pi_S = (a-c)^2/(8b)$	$\Pi_S^* = (w-c)q^*$
Manufacturer's profit (Π_M)	$\Pi_M = (a-c)^2/(16b)$	$\Pi_M^* = (P^*-w)q^*$
Total supply chain profits (Π)	$\Pi = 3(a-c)^2/(16b)$	$\Pi^* = (a-c)^2/(4b)$

2.1. Decentralized Supply Chain

As in most real-world supply chains, suppose that the supplier and the manufacturer are two independently owned and managed firms, where each party is trying to maximize his or her own profits. The supplier chooses the unit wholesale price w, and after observing w, the manufacturer chooses the order quantity q. Note that this is a dynamic game of complete information[3] with two players, supplier and manufacturer, where the supplier moves first and the manufacturer moves second. Hence, we can solve this game using backward induction.

Given w, first we need to find the manufacturer's best response $q(w)$. The manufacturer will choose q to maximize $\Pi_M = (a - bq - w)q$. This is a concave function of q, and hence from the first-order condition we get

$$\frac{\partial \Pi_M}{\partial q} = a - 2bq - w = 0 \;\Rightarrow\; q(w) = \frac{a-w}{2b}.$$

Next, given the manufacturer's best response $q(w) = (a-w)/(2b)$, the supplier maximizes $\Pi_S = (w-c)q = (w-c)(a-w)/(2b)$. This is a concave function of w, and from the first-order condition we get

$$\frac{\partial \Pi_S}{\partial w} = a - w - w + c = 0 \;\Rightarrow\; w = \frac{a+c}{2}.$$

The equilibrium solution for this decentralized supply chain is given in Table 1. In this contractual setting, the supplier gets two-thirds of the supply chain profits, and the manufacturer gets only one-third. This is partly because of the first-mover advantage of the supplier.

Now, let us consider a centralized (integrated) supply chain where both the manufacturer and the supplier are part of the same organization and are managed by the same entity.

2.2. Centralized Supply Chain

Assume that there is a single decision maker who is concerned with maximizing the entire chain's profits $\Pi = (a - bq - c)q$. The solution for the centralized supply chain is given in Table 1. We see that the quantity sold as well as the profits are higher and the price is lower in the centralized than in the decentralized supply chain. Hence, both the supply chain and the consumers are better off with the centralized scenario. A closer look would reveal that w has no impact on the market price, quantity, or the supply chain profits. However, the choice of w in the centralized supply chain is still important, because it determines how the profits will be allocated between the supplier and the manufacturer. We can interpret w as a form of transfer payment from the manufacturer to the supplier. By choosing w appropriately, the additional gain can be distributed between the manufacturer and the supplier so that both are (weakly) better off. For example, if we set $w = (a+3c)/4$, then

[3] For a compact and rigorous analytical treatment of game theory, we refer the reader to Cachon and Netessine [15]. For an in-depth analysis, any of the following would be a good start: Fudenberg and Tirole [24], Gibbons [27], Mas-Colell et al. [44], Myerson [46], and Osborne and Rubinstein [48].

$w - c = P^* - w = (a - c)/4$, and each party's profits are $(a - c)^2/(8b)$. Note that this is the same as the supplier's profits in the decentralized supply chain. Hence, if the supplier and the manufacturer split the profits equally in the centralized supply chain, the supplier is at least as well off and the manufacturer is strictly better off than in the decentralized supply chain.

In the decentralized supply chain, the outcomes are worse for all the parties involved (supplier, manufacturer, supply chain, and consumer) compared to the centralized supply chain, because in the decentralized supply chain both the manufacturer and the supplier independently try to maximize their own profits; i.e., they each try to secure a margin, $P - w$ and $w - c$, respectively. This effect is called "double marginalization".

In this stylized model, the profit loss in the decentralized supply chain as a result of double marginalization is 25% (also referred to as the double marginalization loss). It is clearly in the firms' interest to eliminate or reduce double marginalization, especially if this can be done while allocating the additional profits to the firms such that both firms benefit. This simple model suggests that vertical integration could be one possible way of eliminating double marginalization. However, vertical integration is usually not desirable, or not practical. Can we change the terms of the trade (i.e., devise contracts) so that independently managed companies act as if they are vertically integrated? This concept is known as "supply chain coordination." In this stylized model, the manufacturer should choose the optimal centralized supply chain quantity $q^* = (a - c)/(2b)$ in any coordinating contract.

So far we have assumed that w is fixed per unit. However, in many applications suppliers offer quantity discounts. In quantity discount contracts, if the manufacturer purchases more units, the unit price decreases. There are two common forms of quantity discount: the all-unit discount and the incremental discount. In the all-unit discount, the price for every unit ordered decreases if the order size passes a certain threshold. In the incremental discount, only the price for the units ordered beyond the threshold is reduced. We assume an all-unit discount; the supplier charges $w(q)$, where w is a decreasing function of q. In practice, $w(q)$ is typically a step function, but for simplicity, we assume it is continuous and differentiable.

Let $R(q)$ be the manufacturer's (expected) revenue when she stocks q units of the final product, where $R(q)$ is a finite, concave function of q, with $R(0) = 0$. We assume $R'(0) > c$ (otherwise, production is never profitable). Note that we do not make assumptions regarding demand in this case; demand may be deterministic or stochastic. For example, with a deterministic, price-dependent linear demand curve $R(q) = (a - bq)q$, in a newsvendor setting with stochastic demand D, a selling price of p per unit, and a salvage value of zero, $R(q) = pE_D[\min(D, q)]$.

The manufacturer's (expected) profit in the decentralized supply chain is $\Pi_M = R(q) - w(q)q$. The decentralized supply chain has the same optimal quantity as the centralized supply chain if Π_M is an affine transformation of Π. Hence, we need

$$\Pi_M = R(q^*) - w(q^*)q^* = \alpha(R(q^*) - cq^*) \;\Rightarrow\; w(q^*) = (1-\alpha)\left(\frac{R(q^*)}{q^*}\right) + \alpha c.$$

If the supplier sets the wholesale price as $w(q) = (1-\alpha)(R(q)/q) + \alpha c$, there is a trade-off for the manufacturer: If she chooses $q < q^*$, she will pay too much per unit, increasing her marginal cost. If she chooses $q > q^*$, then the unit price will decrease but the marginal revenue will decrease more, making the extra unit unprofitable. Hence, the optimal quantity choice for the manufacturer is q^*.

The supplier's profit in this case is $\Pi_S = (w(q^*) - c)q^*$:

$$\Pi_S = \left(\frac{(1-\alpha)R(q^*)}{q^*} + \alpha c - c\right)q^* = (1-\alpha)R(q^*) - (1-\alpha)cq^* = (1-\alpha)\Pi.$$

Under the quantity discount contract, the manufacturer's (and the supplier's) profit is proportional to the centralized supply chain's profit; that is, quantity discount contracts eliminate double marginalization and coordinate the supply chain. There are other contract structures that achieve coordination in such simple settings; e.g., buyback contracts (Pasternack [50]) where the manufacturer can return leftover units to the supplier at the end of the selling season for b per unit, where $b < w$, and revenue-sharing contracts (Cachon and Lariviere [14]) where the manufacturer shares a fraction $\alpha < 1$ of her revenues with the supplier. We refer the reader to Tsay et al. [63] and Cachon [12] for excellent reviews on supply chain contracts and Laffont and Martimort [36] and Salanié [55] for an in-depth analysis on contract theory.

3. Optimal Contracts Under Asymmetric Cost Information

As we discussed in §2, today's global supply chains are mostly decentralized. This decentralization has many advantages, including lower production costs and faster time to market. However, it also has several disadvantages, including loss of control of supply chain functionalities and inefficiencies as a result of incentive issues among different players. A common cause of these incentive issues is the information asymmetry among supply chain partners. In this section, we are interested in a setting with such incentive issues. We study the manufacturer's procurement contract when the supplier is privately informed about his production cost. The manufacturer's prior information on the supplier's cost plays an important role in the design of the contract. The material in this section follows Kayış [35] closely.

A manufacturer procures materials from her supplier. The supplier incurs a cost of c per unit to produce a component and deliver it to the manufacturer. The supplier's cost (which we also refer as the supplier's *type*) is his private information at the time of contracting. The manufacturer does not know c with certainty but knows only that c has probability density function (pdf) $f(c)$ with associated cumulative distribution function (cdf) $F(c)$ and support $\Omega := [\underline{c}, \bar{c}]$, where $0 \le \underline{c} < \bar{c} < \infty$. This prior information is common knowledge. Let $h(c) := F(c)/f(c)$, and define *total cost* as

$$k(c) := c + h(c), \tag{1}$$

which will play an important role in the subsequent analysis.

The manufacturer has (expected) revenue $R(Q)$ when she stocks Q units of the final product, where $R(Q)$ is a finite, increasing, and strictly concave function of Q, with $R(0) = 0$. Note that, once again, we do not make assumptions regarding demand; it may be deterministic or stochastic. For example, in a newsvendor setting with stochastic demand D for the final product, a selling price of p per unit, and a salvage value of zero, $R(Q) = pE_D[\min(D, Q)]$. Finally, we assume $R'(0) > \underline{c}$ (otherwise, production is never profitable). The manufacturer and the supplier are risk neutral: each seeks to maximize his or her own expected profit.

To find the manufacturer's optimal contract in this setting, we use the revelation principle (Myerson [45]). The revelation principle states that any Bayesian Nash equilibrium of any Bayesian game can be represented by an incentive-compatible direct mechanism, i.e., a direct revelation mechanism. In a direct mechanism, the *only action* of the player with the private information (the supplier in our setting) is to submit a claim about his or her type. A direct mechanism is incentive compatible if it is a Bayesian Nash equilibrium for that player to tell the truth. Using the revelation principle, the manufacturer (the mechanism designer) can simply focus on the mechanisms that have truth-telling equilibria, and if she finds the optimal mechanism among them, then it is the optimal one among the entire set of feasible mechanisms.

Taking the manufacturer's perspective and using the revelation principle (Laffont and Martimort [36], Myerson [45]), we adopt the assumption that the manufacturer can implement arbitrarily complex contracts, contingent on the cost reported by the supplier. The manufacturer designs a menu of contracts $\{Q(c), t(c)\}_{c\in\Omega}$ that maximizes her (expected) ex ante payoff from the contract, where $Q(c)$ is the quantity to be procured from the supplier, and $t(c)$ is the payment from the manufacturer to the supplier, contingent on the cost c reported by the supplier. Then the supplier selects a contract from this menu by announcing his cost c. The program that the manufacturer has to solve to find the optimal menu is

$$
\begin{aligned}
\text{P1:}\quad \Pi_M^{\text{DRC}} := \max_{\substack{Q:\,\Omega\longrightarrow\mathbb{R}^+\\ t:\,\Omega\longrightarrow\mathbb{R}}} \quad & E_c[R(Q(c)) - t(c)]\\
\text{s.t.}\quad & t(c) - cQ(c) \geq 0 \quad \forall\, c\in\Omega \quad (\text{IR}_c),\\
& t(c) - cQ(c) \geq t(\hat{c}) - cQ(\hat{c}) \quad \forall\, c, \hat{c}\in\Omega \quad (\text{IC}_{c\hat{c}}).
\end{aligned}
$$

The manufacturer's objective is to maximize her (expected) profit, which is equal to the (expected) total revenue from having $Q(c)$ units, less the expected transfer payment to the supplier. The (IR_c) constraints are individual rationality (IR), or participation constraints: the supplier with cost c should get at least his reservation profit, which is normalized to 0, to accept the contract. The $(\text{IC}_{c\hat{c}})$ constraints are incentive compatibility (IC) constraints: the supplier with cost c has the option to choose $(Q(\hat{c}), t(\hat{c}))$, the contract designed for type $\hat{c} \neq c$, but prefers to choose $(Q(c), t(c))$. The optimal contract under this simple setting is characterized by the following proposition.

Proposition 1. *Assume that $k(c)$ is increasing in c. The manufacturer's optimal menu of contracts $\{Q^c(c), t^c(c)\}$ satisfies*

$$
\{Q^c(c), t^c(c)\} = \left\{ \left((R')^{-1}(k(c))\right)^+, cQ^c(c) + \int_c^{\bar{c}} Q^c(\tau)\, d\tau \right\}.
$$

In the newsvendor setting, the manufacturer's optimal order quantity simplifies to

$$
Q^c(c) = G^{-1}\left(\frac{p - k(c)}{p}\right),
$$

where G denotes the cdf of demand. Because the manufacturer does not know the supplier's cost, the optimal contract does not coordinate the supply chain. The optimal procurement quantity in Proposition 1 implies that the manufacturer has to pay $k(c)$ per unit for each quantity procured instead of the actual cost c. Thus, the coordinating quantity is produced only when $c = \underline{c}$, because $h(\underline{c}) = 0$. Otherwise, the manufacturer procures less than the coordinating quantity. We also observe that the inefficiency decreases as c decreases, because $Q(c)$ is nonincreasing in c. The manufacturer achieves a higher level of coordination when the supplier's cost is lower by leaving the supplier a higher profit, which is simply equal to $t(c) - cQ(c) = \int_c^{\bar{c}} Q(\tau)\, d\tau$. This *underproduction* is an inefficiency arising from asymmetric cost information. The underlying reason for this inefficiency is that, because of the asymmetric cost information, the manufacturer has to pay an "information rent" $t^c(c) - cQ^c(c) \geq 0$ to the supplier; that is, the manufacturer pays the supplier more than his cost.

In practice, the manufacturer may prefer to offer a quantity-contingent payment scheme rather than asking for the cost report to determine payments and quantities. However, the optimal cost-contingent menu of contracts $\{Q^c(c), t(c)\}$ can easily be transformed to an equivalent quantity-contingent payment $\{T^c(Q)\}$ as follows:

$$
T^c(Q) = \begin{cases} t^c(c) & \text{if } Q = Q^c(c) \text{ for some } c\in\Omega,\\ 0 & \text{otherwise.} \end{cases}
$$

This equivalence of quantity-contingent and cost-contingent payments is known as the taxation principle (Martimort and Stole [43]).

To assess the complexity of the contract from a practical point of view, we need to examine the form of the optimal contract. Proposition 2 shows that the optimal quantity-contingent payment is concave in the purchased quantity.

Proposition 2. *The manufacturer's optimal quantity-contingent payment, $T^c(Q)$, is a concave function of Q for $Q \in (Q^c[\bar{c}], Q^c[\underline{c}]]$.*

Concavity of the payment function $T^c(Q)$ means that the manufacturer's optimal contract is a quantity discount contract. As the total quantity increases, the price per unit that the manufacturer pays the supplier decreases. To implement this optimal contract, however, she may have to specify an infinite number of quantity breaks in the contract. In the appendix, we study a setting where the manufacturer is privately informed about her demand, and we demonstrate that the supplier's optimal contract is once again an arbitrarily complex quantity discount contract. However, it has long been recognized in the economics literature that "real world incentive schemes appear to take less extreme forms than the finely tuned rules predicted by . . . theory" (Holmström and Milgrom [29], p. 304). Hence it may not be very practical for the manufacturer or the supplier to enforce such arbitrarily complex quantity discount contracts. Next, we will limit the complexity of these optimal contracts by restricting the number of price breaks a contract may have, and we will assess the performance of these simpler contracts first theoretically under asymmetric cost information in §4 and then with human subject experiments in §5.

4. Efficiency of Simple Contracts Under Asymmetric Cost Information

It has been suggested in the literature that the optimal complex contracts can be implemented via a menu of linear contracts (i.e., the tangents of the optimal complex contracts; see, for example, Laffont and Martimort [36], Laffont and Tirole [37], Reichelstein [51], Rogerson [53]). Each item in the menu is intended for a particular cost realization of the supplier and has two parts: a fixed payment and a fixed per-unit price. The continuity of the uncertainty about the supplier's cost, however, requires the number of contracts in this menu to be infinite, which raises the question of practicality. Bower [10] and Gasmi et al. [25] address this problem by limiting the number of menu items to one, and they identify conditions where this one-item menu, consisting simply of a fixed-price contract, can perform nearly as well as the optimal complex contract. Rogerson [54] generalizes this contract by limiting the number of menu items to two. In Rogerson's [54] model, the supplier can exert costly effort to decrease the production cost from a supplier-specific level only known to himself. The manufacturer can only observe the actual cost after the supplier's cost-reduction effort. The first item in the menu is a cost-reimbursement contract, where the manufacturer reimburses the supplier's actual cost, and the second item is a fixed-price contract. The author finds that this two-item menu always captures at least 75% of the profit that the optimal complex contract achieves when the cost of effort is quadratic and the prior information on supplier's base cost is a uniform distribution. Chu and Sappington [18] argue that Rogerson's [54] simple two-item menu may capture much less than 75% of the profit when the cost distribution is skewed to the right. The authors propose another simple two-item menu with a 73% lower bound on performance when the prior information is a power distribution. In this section, we are interested in the worst-case performance of simple contracts without any distributional assumptions, as well as the relative effect of increasing the number of parameters, and hence the complexity, in the contracts that the manufacturer offers. We revisit the model in §3. Once again, the material in this section follows Kayış [35] closely.

We relax the assumption that the manufacturer can implement arbitrarily complex contracts as in Proposition 2. In particular, we are interested in the increase in the manufacturer's expected profit as a result of an increase in the complexity of the contract. We measure the complexity of the contract by the number of contract terms. Specifically, we restrict our attention to a menu of contracts $\{Q_i, t_i\}_{i=0,1,\ldots,n}$, where n is a fixed finite integer, and for each i, t_i represents the total payment the manufacturer offers to the supplier in exchange for Q_i units. To eliminate trivial cases, we assume $0 = t_0 \le t_1 \le \cdots \le t_n$ and $0 = Q_0 < Q_1 < \cdots < Q_n$. Inclusion of the null contract, $t_0 = Q_0 = 0$, allows for the possibility that the manufacturer does not procure from the supplier.

To find the manufacturer's expected profit under this menu of contracts, we need to study the supplier's contract selection problem first. The supplier must choose one contract from the menu $\{Q_i, t_i\}_{i=0,1,\ldots,n}$. He will prefer the contract i to the contract $i-1$ if and only if his profit is greater with the contract i, or $t_i - cQ_i > t_{i-1} - cQ_{i-1}$. Thus a supplier prefers the contract i to contract $i-1$ if his cost c is strictly greater than $s_i := (t_i - t_{i-1})/(Q_i - Q_{i-1})$. Without loss of generality, we restrict attention to menus with $s_i \le s_{i-1}$ for $i = 2, 3, \ldots, n$. (If $s_i \le s_{i-1}$ did not hold, the supplier would never choose the contract $\{Q_{i-1}, t_{i-1}\}$, and there would be an equivalent menu with $n-1$ contracts $\{Q_i, t_i\}$.) Then, the supplier selects the contract $\{Q_{i-1}, t_{i-1}\}$ when his cost is between $s_i := (t_i - t_{i-1})/(Q_i - Q_{i-1})$ and s_{i-1}, for $i = 2, 3, \ldots, n$. Also without loss of generality, we assume that $\underline{c} \le s_n$ and $s_1 \le \bar{c}$ to reflect the fact that the supplier's cost is distributed between $\underline{c}$ and $\bar{c}$. For brevity of exposition, in all subsequent analysis, we will specify the menu of contracts in terms of the unit prices s_i rather than the total payments t_i. Given the values s_i and Q_i for $i = 0, 1, \ldots, n$, the total payment t_i is

$$t_i = \sum_{j=1}^{j \le i} s_j (Q_j - Q_{j-1}). \tag{2}$$

In the operations management literature, the menu $\{Q_i, s_i\}_{i=1,\ldots,n}$ is called an *incremental quantity discount* contract.[4] Up to Q_1 units, the manufacturer offers to pay s_1 per unit. For an additional $Q < Q_2 - Q_1$ units, the manufacturer pays the respective unit price s_2, which is less than s_1, and so on. Hence, the total payment to a supplier delivering Q_i units is exactly equal to the right-hand side of Equation (2). The parameter Q_i is called a quantity break. A menu with $n = 1$ is equivalent to offering a single price in exchange for Q_1 units; thus we refer to this special contract as a price-only contract.

Given the supplier's contract selection problem, the manufacturer's expected profit given the unit prices s_i and quantity breaks Q_i can be written as

$$\Pi_M^{(n)}(\mathbf{s}, \mathbf{Q}) := \sum_{i=1}^{n} [F(s_i) - F(s_{i+1})] \left[R(Q_i) - \left(\sum_{j=1}^{j \le i} s_j (Q_j - Q_{j-1}) \right) \right],$$

where $s_{n+1} = \underline{c}$, $\mathbf{s} = (s_1, \ldots, s_n)$, and $\mathbf{Q} = (Q_1, \ldots, Q_n)$. We find the manufacturer's expected profit, conditioning on the cost of the supplier. When the supplier has cost c between s_i and s_{i+1}, he will choose the contract i and deliver Q_i units. The manufacturer receives her total expected revenue from the market from having Q_i units and pays t_i to the supplier.

We are interested in the worst-case performance of an optimal incremental quantity discount contract with n quantity breaks. To establish an upper bound on the optimality gap of a quantity discount contract, i.e., $\Pi_M^{\mathrm{DRC}} - \max_{\mathbf{s},\mathbf{Q}} \Pi_M^{(n)}(\mathbf{s}, \mathbf{Q})$, we need to find a bound over the set of all possible prior information sets of the manufacturer on the supplier's cost. To find this bound, we use the following inequality:

$$\Pi_M^{\mathrm{DRC}} - \max_{\mathbf{s},\mathbf{Q}} \Pi_M^{(n)}(\mathbf{s}, \mathbf{Q}) \le \Pi_M^{\mathrm{DRC}} - \Pi_M^{(n)}(\mathbf{s}^{\mathbf{U}}, \mathbf{Q}^{\mathbf{U}}) \le \min_{\mathbf{s},\mathbf{Q}} \left[\max_{F} (\Pi_M^{\mathrm{DRC}} - \Pi_M^{(n)}(\mathbf{s}, \mathbf{Q})) \right] := U(n),$$

where $\{\mathbf{s}^{\mathbf{U}}, \mathbf{Q}^{\mathbf{U}}\} \in \arg\min_{\mathbf{s},\mathbf{Q}} [\max_F (\Pi_M^{\mathrm{DRC}} - \Pi_M^{(n)}(\mathbf{s}, \mathbf{Q}))]$. The first inequality follows from the optimality principle for $\max_{\mathbf{s},\mathbf{Q}} \Pi_M^{(n)}(\mathbf{s}, \mathbf{Q})$, and the second one follows from the optimality of $\{\mathbf{s}^{\mathbf{U}}, \mathbf{Q}^{\mathbf{U}}\}$. To solve for the inner maximization problem in the right-hand side of (3), we use the general-moment bound problem.

[4] Under all-unit discounts, the manufacturer's and the supplier's objective functions are discontinuous at price breaks. Furthermore, the supplier's decisions are functions of not only the prices but also the quantity breaks. The characteristics of all-unit discounts complicate their analysis considerably; hence, we restrict ourselves to incremental quantity discounts in this analysis.

We are interested in the shape of the upper bound function $U(n)$, which is the worst-case performance of a simple contract with n terms. Proposition 3 shows the recursive relationship between consecutive values of this function.

Proposition 3. *The upper bound for the optimality gap, $U(n)$, has the following properties for $n > 1$:*

$$\max\left(\frac{\pi_M(\underline{c})}{n+1}, \pi_M(\bar{c})\right) \le U(n) \le U(n-1), \tag{3}$$

where

$$\pi_M(c) := \max_Q \{R(Q) - cQ\}. \tag{4}$$

For $n = 1$, $U(1) = \max(\pi_M(\underline{c})/2, \pi_M(\bar{c}))$.

There are two observations: the manufacturer can always decrease the upper bound by adding one more contract to the menu of simple contracts she offers. More importantly, though, the manufacturer is guaranteed a worst-case optimality gap of no more than $\max(\pi_M(\underline{c})/(n+1), \pi_M(\bar{c}))$. Clearly, the decreasing rate of this guarantee for a simple contract is diminishing in the number of quantity breaks n. This result implies that a limited number of quantity breaks in the incremental quantity discount contracts is sufficient to ensure that the maximum expected profit loss is kept below a certain level.

To assess the magnitude of the optimality gap of different simple contracts, we undertake a numerical study in a newsvendor setting with the following parameters. Without loss of generality, we normalize the selling price p to 1. Demand is a uniformly distributed random variable with $D \in [M_D(1 - V_D), M_D(1 + V_D)]$, with M_D normalized to 1 and $V_D \in \{0.1, 0.2, \ldots, 1\}$. We consider a uniform distribution for the supplier's cost c and specify the minimum cost $\underline{c} \in \{0, 0.1, \ldots, 0.9\}$ and maximum cost $\bar{c} \in \{\underline{c} + 0.1, \underline{c} + 0.2, \ldots, 2\}$. This study covers most of the possible parametric combinations with 1,550 experiments.

Table 2 compares the manufacturer's expected profit loss from using a price-only contract and a quantity discount contract with $n = 2$. The results show that the maximum optimality gap of 11% with price-only contracts is reduced to 4% when the manufacturer uses a quantity discount contract. Similarly, the flexibility to offer a quantity discount contract leads to a significant decrease in optimality gap for all summary statistics. To further study the relationship between the efficiency of incremental quantity discount contracts and the number of quantity breaks, we calculate the maximum optimality gap as a function of n in Corollary 1 when both the demand of the newsvendor and the supplier's cost are uniformly distributed.

Corollary 1 supports the result presented in Proposition 3 that having only a limited number of quantity breaks may be sufficient to ensure that the maximum expected profit loss is kept below a certain threshold.

Corollary 1. *Assume that $D \sim \mathrm{U}[0, 2M_D]$ and $c \sim \mathrm{U}[\underline{c}, \bar{c}]$. The maximum percentage loss in the manufacturer's expected profit from using an incremental quantity discount contract with n quantity breaks rather than the optimal complex contract is equal to $1/(2n+1)^2$.*

Table 2. The percentage loss in the manufacturer's expected profit from using (i) the price-only contract and (ii) the quantity discount contract with ($n = 2$), rather than the complex optimal contract.

	Mean	Min	10th percentile	Median	90th percentile	Max	No. of cases
Price-only contract	5.618	0.037	0.792	5.106	10.918	11.115	1,550
Quantity discount contract	1.950	0.023	0.386	1.784	3.871	4.065	1,550

For the case of uniformly distributed demand and supplier cost, the maximum optimality gap is less than 1% with just five quantity breaks. These results suggest that even though the manufacturer is always better off with more quantity breaks, a small number of quantity breaks for the incremental quantity discount contract is sufficient to ensure that she receives near-optimal profit.

5. A Behavioral Analysis of the Efficiency of Simple Contracts Under Asymmetric Demand Information

In the previous section, we showed that under asymmetric cost information simple contracts perform remarkably well. In the numerical analyses, we observed a maximum optimality gap of 11% with price-only contracts. With only two price breaks, this gap can be reduced to below 4% with quantity discount contracts. However, under asymmetric demand information (see the appendix), the gap between simple contracts and optimal contracts can be more substantial. Especially when the manufacturer faces a high demand risk, the value of the optimal contract is high because providing flexibility in terms of quantities and payments helps the manufacturer partially alleviate this risk, and hence increases the total quantity and the supplier's profit. Thus, there are conflicting theoretical and practical accounts of the simple price-only contract and the more complex all-unit quantity discount contracts under asymmetric demand information. Given these conflicting accounts, we compare them in an experimental setting to develop contract design guidelines for suppliers in Kalkancı et al. [34]; the material in this section follows that work closely.

Using human subject experiments, we test the theoretical prediction that, under asymmetric information, suppliers' profits increase as the complexity of the contract (i.e., the number of price breaks and thus the number of decisions about contract parameters) increases. In particular, we study a two-tier supply chain with a single (human) supplier and a single (computerized) manufacturer where the manufacturer faces a newsvendor setting and has better information on demand than the supplier. The supplier determines the *pricing scheme*, and the manufacturer selects an *order quantity*. In this setting, the key theoretical predictions are as follows: (1) Suppliers' profits increase as the complexity of the contract increases. (2) Supply chain efficiency increases as the complexity of the contract increases. (3) Manufacturers' profits decrease as the complexity of the contract increases. Table 3 summarizes the supplier's and the manufacturer's decisions as well as the average profits based on theoretical predictions and experimental observations under different contracts between the supplier and the manufacturer. In a price-only contract (one-price contract), the supplier sets a single wholesale price w_1, and the manufacturer procures the quantity she chooses. We also study two all-unit quantity discount contracts. In the quantity discount contract with two prices (two-price contract), the supplier quotes two prices ($w_1 \geq w_2$) and a single price break Q_1. If the manufacturer orders less than Q_1, she pays w_1 per unit. Otherwise, she pays a cheaper unit price w_2. In the quantity discount contract with three prices (three-price contract), the supplier quotes three prices ($w_1 \geq w_2 \geq w_3$) and two price breaks ($Q_1 \leq Q_2$). If the manufacturer orders less than Q_1, she pays w_1 per unit. If she orders more than Q_1 but less than Q_2, she pays a cheaper unit price w_2. If she orders more than Q_2, she pays an even cheaper unit price of w_3.[5]

Table 4 reports the comparison of the supplier, manufacturer, and total supply chain profits with theoretical predictions. From the table, we conclude that suppliers' profits are significantly lower than predicted by theory; that is, suppliers cannot use the contracts as effectively as theory predicts. Consistent with theoretical results, a two-price contract

[5] Theoretically, increasingly complex contracts would benefit the supplier. In the experimental setting of Kalkancı et al. [34], because the mean demand is one of three types (low, medium, and high), up to three different prices with a minimum quantity commitment Q_0 would be enough to separate these types, and more prices would have no additional benefit for the supplier or the supply chain.

TABLE 3. The observed (Obs.) and optimal (Opt.) contract parameters, procurement quantities, and the supplier, manufacturer, and total supply chain profits for our parameter set under different contracts (rounded).

	Supplier's decisions						Manufacturer's decisions			
	Prices			Price breaks		Expected profit	Procurement quantities			Expected Expected profit
	w_1	w_2	w_3	Q_1	Q_2		q_L	q_M	q_H	
One-price										
Obs.	142					6,722	28	70	108	3,270
Opt.	160					7,200	20	60	100	2,000
Two-price										
Obs.	163	142		102		7,065	21	78	112	2,841
Opt.	184	152		160		8,640	8	48	160	704
Three-price										
Obs.	161	143	126	81	129	6,840	24	83	126	3,619
Opt.	200	177	147	84	160	9,576	0	84	160	669

Notes. In all experiments, 19 subjects were assigned the role of supplier using a one-price, two-price, or three-price contract, and they played against a computerized buyer for 40 periods. The observed values reported are the average values for each treatment.

leads to higher profits for the supplier than a one-price contract ($p = 0.048$). However, the observed improvement is less than what is predicted by theory ($p < 0.001$). Moreover, the difference between two-price and three-price contracts is not significant (two-sided $p = 0.26$). In fact, suppliers do not even make significantly higher profits employing a three-price contract than a one-price contract (two-sided $p = 0.587$). These results demonstrate that there is potential improvement to suppliers' expected profit from using quantity discounts. However, simpler contracts such as a price-only contract or a quantity discount contract with a low number of price breaks may be sufficient for a supplier designing contracts under asymmetric demand information. Similarly, total supply chain profits are not significantly different under different types of contracts (two-sided p-values are 0.724, 0.233, and 0.123 for the comparisons of one-price and two-price, two-price and three-price, and one-price and three-price contracts, respectively); that is, statistically, players do not achieve better channel coordination with more complex quantity discount contracts. Furthermore, as can be seen from Table 4, one-price and two-price contracts perform better and lead to significantly higher total supply chain profits compared to theoretical predictions; that is, when simpler contracts such as a price-only contract or a quantity discount contract with a low number of price breaks are employed, human subjects achieve better channel coordination than

TABLE 4. Means and standard deviations (in parentheses) for the supplier's, manufacturer's, and total supply chain profits and comparison of respective profits with theoretical predictions.

	Supplier's profit		Manufacturer's profit		Total profit	
	Experiment	Experiment vs. theory	Experiment	Experiment vs. theory	Experiment	Experiment vs. theory
One-price	6,722.16 (539.88)	Lower ($p < 0.001$)	3,269.60 (1,313.69)	Higher ($p < 0.001$)	9,991.76 (887.96)	Higher ($p < 0.001$)
Two-price	7,065.03 (480.46)	Lower ($p < 0.001$)	2,841.30 (1,382.18)	Higher ($p < 0.001$)	9,906.34 (1,094.1)	Higher ($p = 0.02$)
Three-price	6,840.1 (726.95)	Lower ($p < 0.001$)	3,619.23 (1,829.79)	Higher ($p < 0.001$)	10,459.33 (1,319.75)	Not significant

theory predicts. Finally, in our experiments, manufacturers' profits are not significantly different under different contracts (two-sided p-values are 0.374, 0.120, and 0.506 for the comparisons of one-price and two-price, two-price and three-price, and one-price and three-price contracts, respectively); that is, suppliers cannot use the additional complexity of their contracts to extract manufacturers' profits as effectively as theory predicts. Table 4 also supports this observation: manufacturers (suppliers) earn significantly more (less) than theory predicts under each contract; that is, we observe a more equitable distribution of profits between suppliers and manufacturers compared to what is theoretically predicted.

Human subjects in our experiments deviate from the expected profit-maximizing behavior and act boundedly rational. They exhibit a bias toward decisions that they have already chosen and inertia that slows the movement to new decisions. These subject behaviors can be described with three decision biases: a *probabilistic choice* bias (which captures bounded rationality and describes how a subject deviates probabilistically from the decision that maximizes his score), a *reinforcement* bias (which captures the inertia and describes why a subject may favor the choices he made in the past, even if those choices were not the best), and a *memory* bias (which captures how far a subject looks into the past to calculate his scores). These three decision biases can be modeled effectively using an experience-weighted attraction EWA model (Camerer and Ho [16], Bostian et al. [9]), in which a subject assigns a score to all of his possible decision choices and then updates these scores over time based on his counterfactual profits and previous decisions. Estimations with the EWA model show that our subjects exhibit all three of these biases. Furthermore, the EWA model effectively captures the relative performance of price-only and quantity discount contracts, as well as the directional changes of decisions from theory. These estimations show that as the contract complexity increases, human subjects increasingly rely on simple heuristics by favoring the choices that they made in the recent past (even if those choices were not the best). This finding is consistent with Simon's [57] seminal work on human reliance on simpler decision rules when faced with complexity and is in line with the intuitive notion that "the harder the decision, the more willing we are to use a less precise rule of thumb if it saves us a headache" (Armour and Daly [4], p. 2).

Kalkancı et al. [33] extend this analysis by experimentally studying the trade-off between the contract complexity and economic inefficiencies under a setting with human-to-human interactions. Their human-to-human experiments provide support for the conclusions above for the human-to-computer experiments and extend these results to more general settings. Human subjects, as suppliers, achieve similar profits under simpler contracts (i.e., a price-only contract or a quantity discount contract with a low number of price breaks) compared to the more complex quantity discount contracts. Based on human-to-computer and human-to-human experiments, we conclude that there is a nontrivial trade-off between the complexity and inefficiency of all-unit quantity discount contracts: the notion that complex contracts can optimize the supplier's profit is flawed and requires deeper consideration. Thus, it is crucial to consider contract complexity as a design factor. Simpler contracts, such as a price-only contract or a quantity discount contract with a low number of price breaks, may be sufficient for a supplier designing contracts under asymmetric demand information.

6. Dynamic Procurement

Our discussions so far suggest that designing quantity discount contracts can be challenging. In this section, we propose that dynamic procurement, a sequential negotiation process prevalent in practice, may be used as a means to implement nonlinear pricing such as quantity discounts and may lead to a simple yet efficient contract design. The material in this section follows Erhun et al. [21, 22] closely.

Dynamic procurement, i.e., simple price-only contracts repeated over time (possibly with different prices), is a commonly observed practice in a vertical channel. To manage demand risk, a manufacturer may prefer to procure goods dynamically over time. Other commonly observed reasons for dynamic procurement include spreading payments over a period of time; minimizing potential capacity risks (supplier's or manufacturer's); supplier's decreasing costs over time, which may translate to lower prices (e.g., as in the electronics industry); and forward buying. In this section, we argue that dynamic procurement can also serve as a tool to influence future prices by showing that dynamic procurement benefits the players even if the demand is deterministic, i.e., when the commonly known and intuitive benefits due to risk hedging are not present.

We study a model where there are N possible periods for procurement before the manufacturer's production/selling season begins. The supplier and the manufacturer maximize their profits. The supplier's decisions are the wholesale prices for each period, w_n ($n = 1, \ldots, N$). The manufacturer's decisions are the procurement quantities for each period, q_n ($n = 1, \ldots, N$), and the production quantity, Q_N. (We set $Q_0 = 0$ and $q_0 = 0$.) The market is characterized by a linear inverse demand function $P(Q_N) = a - bQ_N$, where a is the market potential, b is the price sensitivity, and $P(Q_N)$ is the per-unit market price of the product for Q_N. We assume that the manufacturer's unit cost of production is zero. However, the analysis of a positive, constant production cost ($c < a$) case is trivial. Because $(a - bQ)Q - cQ = ((a - c) - bQ)Q$, we can simply modify the demand intercept a such that $\hat{a} = a - c$, and the analysis follows.

The sequence of events in each period n of the N-period game is as follows: (i) Given previous procurement quantities (q_j, $j = 1, \ldots, n-1$), the supplier determines the wholesale price w_n. (ii) Given previous procurement quantities (q_j, $j = 1, \ldots, n-1$) and the current wholesale price (w_n), the manufacturer determines her procurement quantity q_n. (iii) In the last period N, the manufacturer chooses her production quantity Q_N and procures extra quantity, if necessary. The market clears only once at the end of the Nth period; i.e., there is only a single selling opportunity to end consumers.

We use backward induction and obtain the pure-strategy subgame perfect equilibrium (SPE). Proposition 4 characterizes this equilibrium.

Proposition 4. *The unique pure-strategy SPE for the N-period dynamic procurement model is as follows. For $n = 1, \ldots, N-1$, let $\bar{n} = N - n$. Then,*

$$q_1 = \frac{a}{4Nb}, \quad w_1 = K_N^2 N\left(\frac{a}{2}\right), \quad q_{n+1} = \left(\frac{2\bar{n}+1}{2\bar{n}}\right)q_n, \quad w_{n+1} = \left(\frac{2\bar{n}}{2\bar{n}+1}\right)w_n,$$

where $K_1 = 1$, and $K_N = \prod_{k=1}^{N-1}((2k+1)/(2k+2))$. The production quantity is $Q_N = \sum_{n=1}^{N} q_n = a/(2b) - (a/(4b))K_N$. As N tends to infinity, Q_N tends to $a/(2b)$.

From Proposition 4, the production quantity for the N-period model is $Q_N = a/(2b) - (a/(4b))K_N$. As N increases, K_N decreases, and the production quantity increases. Therefore, the double marginalization effect decreases and the efficiency increases and approaches that of the centralized solution. The last period's wholesale price decreases ($a/2 \searrow 3a/8 \searrow 5a/16 \searrow \cdots \searrow 0$), whereas the first-period wholesale price increases ($a/2 \nearrow 9a/16 \nearrow 75a/128 \nearrow \cdots$), as N increases. The total production quantity, Q_N, increases ($a/(4b) \nearrow 5a/(16b) \nearrow 11a/(32b) \nearrow \cdots \nearrow a/(2b)$), and the market price of the product (following the relationship $P(Q_N) = a - bQ_N$) decreases ($3a/4 \searrow 11a/16 \searrow 21a/32 \searrow \cdots \searrow a/2$). Even though the quantity in the first period decreases as N increases ($q_1 = a/(4Nb)$), the manufacturer procures in each period; that is, trading occurs in all N periods. The manufacturer is willing to procure goods at a given period (at a higher price) because she knows that by doing so the best response for the supplier is to lower the price in the subsequent periods.

For a fixed N, the quantity that the manufacturer procures in each period increases and the wholesale price decreases over time. As N increases, the double marginalization effect decreases, and the supplier's and the manufacturer's profits both increase. Dynamic procurement not only increases the supply chain efficiency but also naturally allocates the surplus to the supplier and the manufacturer such that both parties benefit. Independent of the values of a and b, the supplier's profit converges to approximately 64% of the total profits, and the manufacturer's profit converges to approximately 36% of the total profits in the presence of additional procurement periods. Even for small values of N, dynamic procurement decreases the inefficiency considerably. For example, for $N = 3$, the inefficiency is already less than 10% (compared to 25% for $N = 1$).

Next, we consider a two-period extension of our main model where the end-consumer market can be in one of two states (indexed by i): a "high" demand state, $i = h$, or a "low" demand state, $i = l$. The probability that the demand market will be in state i is ϕ_i, $i = l, h$, such that $\phi_h + \phi_l = 1$. The manufacturer and the supplier learn the state of the demand sometime during the planning horizon. Therefore, the information update splits the planning horizon into two periods, possibly of unequal lengths. The manufacturer can procure goods in both periods before the selling season begins. We assume that $b_h = b_l = 1$; i.e., the market is characterized by a linear inverse demand function $P_i(Q_i) = a_i - Q_i$, where $P_i(Q_i)$ is the price of the product when Q_i is the quantity sold by the manufacturer in the end-consumer market when the demand is in state i, and a_i is the market potential in state i. We assume that $0 < a_l \leq a_h$; that is, the market potential is positive in both states and higher under the high demand state.

In this two-period model, the supplier announces the first-period wholesale price at the beginning of the first period. He announces the second-period wholesale price at the beginning of the second period *after* the demand state is revealed. The manufacturer chooses the first-period procurement quantity *before* the demand state is revealed. She chooses the second-period procurement quantity *after* the demand state is revealed. The market clears at the end of the second period. The proposition below shows how dynamic procurement extends to the situations with uncertain demand.

Proposition 5. *When demand is uncertain, the prices and quantities are set as follows:*

$$(w_1, q_1) = \begin{cases} \left(\dfrac{9(\phi_h a_h + \phi_l a_l)}{16}, \dfrac{\phi_h a_h + \phi_l a_l}{8}\right) & \text{if } \phi_h < \left(\dfrac{3a_l}{a_h - a_l}\right)^2, \\ \left(\dfrac{9\phi_h a_h}{16}, \dfrac{a_h}{8}\right) & \text{otherwise;} \end{cases}$$

$$(w_{2,i}, q_{2,i}) = \left(\left(\frac{a_i}{2} - q_1\right)^+, \left(\frac{a_i}{4} - \frac{q_1}{2}\right)^+\right), \quad i = h, l.$$

Depending on the demand state, dynamic prices may increase (advance purchase discount) or decrease (markdown) over time. Dynamic procurement works for the supplier even under uncertain demand; the supplier's profit increases with dynamic procurement. However, for the manufacturer, the value of dynamic procurement depends on when the supplier sets the price in a single procurement situation. If the supplier sets the price *after* the uncertainty is revealed (no-commitment model), then the manufacturer's expected profits are also higher. Because both players benefit, it is a natural consequence that dynamic procurement eliminates supply chain inefficiencies compared to the no-commitment case. If the supplier sets the price *before* the uncertainty is revealed (early commitment model), the manufacturer's expected profits may either increase or decrease under dynamic procurement. Especially when the difference between the market potentials of the high and low demand states is considerable, the manufacturer becomes worse off with dynamic procurement compared to the early commitment model. However, when the supplier chooses his *capacity as well as the prices*, for a wide range of parameters, dynamic procurement is the best alternative for all parties, including the end consumers.

Our results show that firms may implement nonlinear pricing schemes, such as quantity discounts, through repeated interactions of dynamic procurement. Kalkancı et al. [34] argue that such a sequential negotiation process may also lead to better contracts under asymmetric information. The behavioral academic literature supports their conjecture in two ways. First, Özer et al. [49] show that human subjects reveal some of their private information truthfully in laboratory experiments even when there are incentives to hide information, contrary to predictions regarding the theory of self-interested agents. Second, Lim and Ho [40] argue that subjects can use quantity discount mechanisms quite effectively under complete information. Therefore, under the setting we analyze in the appendix and §5, an opportunity to collect more information about the manufacturer's demand type should benefit the supplier. The sequential negotiation process provides him this opportunity. Note that even when the level of information asymmetry between the supplier and the manufacturer increases, each step of the negotiation process will stay the same; that is, when the supplier's initial uncertainty increases, the negotiation process may take longer but will not be more complex.

Acknowledgments

The author thanks Joe Geunes, Başak Kalkancı, Enis Kayış, Tim Kraft, and Jessica McCoy for their helpful suggestions on earlier drafts of this tutorial. This work was supported in part by National Science Foundation Grants CAREER-0547021 and DMI-0400345.

Appendix. Optimal Contracts Under Asymmetric Demand Information

In this section, we are interested in a supplier's contract when the manufacturer is privately informed about her demand. The supplier's prior information on the manufacturer's demand plays an important role in the design of the contract. The material in this section follows Kalkancı and Erhun [32] closely.

We study a setting consisting of a single manufacturer and a single supplier. The supplier's unit production cost is $c > 0$. The manufacturer sells the final product to end consumers at an exogenously specified market price of p. We normalize the manufacturer's cost to zero, and we restrict our attention to the situation where production is profitable; i.e., $0 < c < p$. Any unmet demand is lost without a stockout penalty. There is no salvage value for leftover products.

We focus on a one-time interaction between the manufacturer and the supplier. The manufacturer and the supplier contract on the production quantity and the manufacturer's corresponding payment to the supplier. Then, the supplier produces the contracted quantity, delivers it to the manufacturer, and receives his payment *before* the end-consumer demand is realized. This process is complicated by the fact that by being closer to the end-consumer market, the manufacturer is better informed about the demand, which we model as

$$\text{Demand} = \mu + \alpha,$$

where μ is the demand forecast, and α is a random fluctuation. The demand forecast (which we also refer to as the manufacturer's *type*) is the manufacturer's private information at the time of contracting. The supplier does not know μ with certainty but does know its distribution. We represent the pdf of this common prior distribution as $g(\mu)$ and the cdf as $G(\mu)$, which has finite and positive support $\Lambda := [\underline{\mu}, \bar{\mu}]$. The mean and standard deviation of μ are μ_0 and σ_0, respectively. Neither the manufacturer nor the supplier observe the value of α during contracting, which makes ours a newsvendor setting. We assume that α follows a continuous distribution with cdf $\Phi(\cdot)$ and pdf $\phi(\cdot)$. The mean and standard deviation of α are 0 and σ_1, respectively. We bound α below by $\underline{\alpha}$, which enables us to study markets with a minimum guaranteed size.

We define the hazard rates of α and μ as $h_\phi(\cdot) := \phi(\cdot)/(1 - \Phi(\cdot))$ and $h_g(\cdot) := g(\cdot)/(1 - G(\cdot))$, respectively. For clarity of exposition, we also define $\bar{h}_g(\cdot) := (1 - G(\cdot))/g(\cdot)$. In the analysis of complex contracts and principal–agent models, it is common to assume that $h_\phi(\cdot)$ and $h_g(\cdot)$ are increasing. The assumption that $h_g(\cdot)$ is increasing (e.g., Laffont and Tirole [38], Martimort [42]) ensures that the manufacturer's order quantity is increasing with respect to her demand forecast. The assumption that $h_\phi(\cdot)$ is increasing (e.g., Taylor and Xiao [61]) ensures that the supplier's objective is unimodal for each demand forecast. Many frequently used distributions have increasing hazard rates, such as the uniform, normal, and exponential distributions (Bolton and Dewatripont [8]). Moreover, we assume that $\underline{\alpha}+\underline{\mu}\geq 0$, which ensures that the demand will be nonnegative even for the lowest demand forecast. Finally, $h_\phi(\cdot)$, $\bar{h}_g(\cdot)$, $\Phi(\cdot)$, and $G(\cdot)$ are assumed to be twice differentiable.

Using the revelation principle (Laffont and Martimort [36], Myerson [45]), we can restrict our attention to the class of truth-telling menus of contracts in the form of $\{t(\mu), y(\mu)\}$, where $t(\mu)$ is the payment to the supplier from the manufacturer with type μ, and $y(\mu)$ is a type-μ, manufacturer's procurement quantity for all $\mu\in[\underline{\mu},\bar{\mu}]$. The supplier's problem can then be defined as

$$\textbf{P2}:\quad \Pi_S^{\mathrm{DRC}} := \max_{t(\cdot),\, y(\cdot)} \int_{\underline{\mu}}^{\bar{\mu}} (t(\mu)-cy(\mu))g(\mu)\,d\mu$$
$$\text{s.t.}\quad pE_\alpha[\min(\mu+\alpha, y(\mu))]-t(\mu)\geq 0,\quad \forall\,\mu\in[\underline{\mu},\bar{\mu}],$$
$$pE_\alpha[\min(\mu+\alpha, y(\mu))]-t(\mu)\geq pE_\alpha[\min(\mu+\alpha, y(\hat{\mu}))]-t(\hat{\mu}),\quad \forall\,\mu\in[\underline{\mu},\bar{\mu}].$$

The supplier's objective is to maximize his expected profit, which is equal to the total payment to him from the manufacturer with type μ, less the production cost. The constraints are the IR and IC constraints, respectively. The IR constraints ensure that for any type, the manufacturer will make as much expected profit as she could without participating in any contract. The IC constraints reflect the fact that a manufacturer with type μ prefers to choose $\{t(\mu), y(\mu)\}$ over any other $\{t(\hat{\mu}), y(\hat{\mu})\}$, which ensures truth telling. To analyze the equilibrium under asymmetric information, let $\Upsilon(y)$ be the derivative of the supplier's objective function with respect to the quantity at type μ in the region where $y>\mu+\underline{\alpha}$:

$$\Upsilon(y) = -c + p(1-\Phi(y-\mu))\big(1-\bar{h}_g(\mu)h_\phi(y-\mu)\big).$$

Let $T(y)$ be the total payment to the supplier when the manufacturer orders y units. The optimal contract under this setting is characterized by the following proposition.

Proposition 6. *A manufacturer of type μ buys the quantity $y(\mu)$, where*

$$y(\mu)=\begin{cases}\mu+\underline{\alpha} & \text{if } \Upsilon(\mu+\underline{\alpha})\leq 0,\\ \{Y\colon\ \Upsilon(y)=0\} & \text{otherwise.}\end{cases}$$

The supplier's payment scheme satisfies the following conditions:

(1) $T(y(\mu)) = pE_\alpha[\min(\mu+\alpha, y(\mu))] - \int_{\underline{\mu}}^{\mu} p\Phi(y(\tau)-\tau)\,d\tau.$

(2a) *If $y(\mu)=\mu+\underline{\alpha}$, the payment satisfies*

$$\max\big(c, p-\bar{h}_g(\mu)p\phi(\underline{\alpha})\big)\leq T'(\mu+\underline{\alpha})\leq \min\big(p, \bar{h}_g(\mu)p\phi(\underline{\alpha})+c\big).$$

(2b) *If $y(\mu)>\mu+\underline{\alpha}$, the payment satisfies*

$$T'(y(\mu)) = \bar{h}_g(\mu)p\phi(y(\mu)-\mu)+c.$$

When the supplier has limited information on the manufacturer's demand forecast, asking the manufacturer for her forecast would lead to a bias because the supplier extracts the manufacturer's total profit when he knows her demand forecast; that is, the manufacturer would be better off declaring that her forecast is low and buying a quantity that is intended for a manufacturer with a low demand forecast. Because the supplier is aware of this incentive problem, he cannot base his decisions on any information shared by the manufacturer. Instead, he differentiates among manufacturers with different forecasts and makes sure that manufacturers with higher demand forecasts do not buy lower quantities. To enable the differentiation, the supplier provides discounts to a manufacturer with a higher forecast and increases her safety stock. This procurement structure is fundamentally different than one would expect. Intuitively, the manufacturer's safety stock should increase with the coefficient of variation, favoring manufacturers with low forecasts. We conclude that, under asymmetric information, the supplier's pricing decisions are affected more by incentive concerns than by concerns as a result of the uncertainty in demand, and once again the optimal contract is a quantity discount contract.

References

[1] Adesso Solutions. A double-edged sword. *Business Credit* 107(9):56–59, 2005.

[2] N. Altintas, F. Erhun, and S. Tayur. Quantity discounts under demand uncertainty. *Management Science* 54(4):777–792, 2008.

[3] R. Anupindi and Y. Bassok. Centralization of stocks: Retailers vs. manufacturer. *Management Science* 45(2):178–191, 1999.

[4] P. Armour and M. Daly. Retirement savings and decision errors: Lessons from behavioral economics. *FRBSF Economic Letter* 16(6):1–3, 2008.

[5] F. Bernstein and G. A. DeCroix. Decentralized pricing and capacity decisions in a multitier system with modular assembly. *Management Science* 50(9):1293–1308, 2004.

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

[7] Boeing. Anaheim supplier payment information, Retrieved March 17, 2008. http://www.boeingsuppliers.com/anaheim/SupplierPaymentInstruction.com.

[8] P. Bolton and M. Dewatripont. *Contract Theory*. MIT Press, Cambridge, MA, 2005.

[9] A. J. Bostian, C. A. Holt, and A. M. Smith. Newsvendor "pull-to-center" effect: Adaptive learning in a laboratory experiment. *Manufacturing & Service Operations Management* 10(4):590–608, 2008.

[10] A. G. Bower. Procurement policy and contracting efficiency. *International Economic Review* 34(4):873–901, 1993.

[11] A. Burnetas, S. M. Gilbert, and C. Smith. Quantity discounts in single period supply contracts with asymmetric demand information. *IIE Transactions* 39(5):465–480, 2007.

[12] G. P. Cachon. Supply chain coordination with contracts. A. G. de Kok and S. C. Graves, eds. *Supply Chain Management: Design, Coordination and Operation*, Chapter 6. Handbooks in Operations Research and Management Science, Vol. 11. Elsevier, Amsterdam, 229–340, 2003.

[13] G. P. Cachon and M. A. Lariviere. Contracting to assure supply: How to share demand forecasts in a supply chain. *Management Science* 47(5):629–646, 2001.

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

[15] G. P. Cachon and S. Netessine. Game theory in supply chain analysis. M. P. Johnson, B. Norman, N. Secomandi, eds. *TutORials in Operations Research*, INFORMS, Hanover, MD, 200–233, 2006.

[16] C. Camerer and T.-H. Ho. Experience-weighted attraction learning in normal-form games. *Econometrica* 67(4):837–874, 1999.

[17] F. Chen, A. Federgruen, and Y. Zheng. Coordination mechanism for a distribution system with one supplier and multiple retailers. *Management Science* 47(5):693–708, 2001.

[18] L. Y. Chu and D. E. M. Sappington. A note on optimal procurement contracts with limited direct cost inflation. *Journal of Economic Theory* 137(1):745–753, 2007.

[19] T. H. Cui, J. S. Raju, and Z. J. Zhang. Fairness and channel coordination. *Management Science* 53(8):1303–1314, 2007.

[20] F. Erhun and P. Keskinocak. Game theory in business applications. Working paper, Department of Management Science and Engineering, Stanford University, Stanford, CA, 2003.

[21] F. Erhun, P. Keskinocak, and S. Tayur. Dynamic procurement in a capacitated supply chain facing uncertain demand. *IIE Transactions* 40(8):733–748, 2008.

[22] F. Erhun, P. Keskinocak, and S. Tayur. Dynamic procurement, quantity discounts, and supply chain efficiency. *Production and Operations Management* 17(5):1–8, 2008.

[23] FD Wire. Q3 2006 Matria Healthcare earnings conference call—Final. Press release, October 26, 2006. http://www.accessmylibrary.com/coms2/summary_0286-33356779_ITM, retrieved March 17, 2008.

[24] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, Cambridge, MA, 1991.

[25] F. Gasmi, J. J. Laffont, and W. W. Sharkey. Empirical evaluation of regulatory regimes in local telecommunications markets. *Journal of Economics and Management Strategy* 8(1):61–94, 1999.

[26] Y. Gerchak and Y. Wang. Revenue-sharing vs. wholesale-price contracts in assembly systems with random demand. *Production and Operations Management* 13(1):23–33, 2004.

[27] R. Gibbons. *Game Theory for Applied Economists.* Princeton University Press, Princeton, NJ, 1992.

[28] P. Guo, J. S. Song, and Y. Wang. Outsourcing structures and information flow in a three-tier supply chain. *International Journal of Production Economics* 128(1):175–187, 2010.

[29] B. Holmström and P. Milgrom. Aggregation and linearity in the provision of intertemporal incentives. *Econometrica* 55(2):303–328, 1987.

[30] HT Media. Rep. Davis issues statement on "Iraqi reconstruction: Reliance on private military contractors, status report." Press release, February 7, 2007.

[31] A. P. Jeuland and S. M. Shugan. Managing channel profits. *Marketing Science* 2(3):239–272, 1983.

[32] B. Kalkancı and F. Erhun. Pricing games and impact of private demand information in decentralized assembly systems. *Operations Research*, forthcoming, 2011.

[33] B. Kalkancı, K.-Y. Chen, and F. Erhun. Complexity as a contract design factor: A human-to-human experimental study. Working paper, Department of Management Science and Engineering, Stanford University, Stanford, CA, 2011.

[34] B. Kalkancı, K.-Y. Chen, and F. Erhun. Contract complexity and performance under asymmetric demand information: An experimental evaluation. *Management Science* 57(4):689–704, 2011.

[35] E. Kayış. Essays on procurement contracting. Ph.D. thesis, Department of Management Science and Engineering, Stanford University, Stanford, CA, 2008.

[36] J. J. Laffont and D. Martimort. *The Theory of Incentives: The Principal Agent Theory.* Princeton University Press, Princeton, NJ, 2002.

[37] J. J. Laffont and J. Tirole. Using cost observation to regulate firms. *Journal of Political Economy* 94(3):614–641, 1986.

[38] J. J. Laffont and J. Tirole. Privatization and incentives. *Journal of Law, Economics, and Organization* 7:84–105, 1991.

[39] M. Lariviere and E. L. Porteus. Selling to the newsvendor: An analysis of price-only contracts. *Manufacturing & Service Operations Management* 3(4):293–305, 2001.

[40] N. Lim and T.-H. Ho. Designing price contracts for boundedly rational customers: Does the number of blocks matter? *Marketing Science* 26(3):312–326, 2007.

[41] P. Majumder and A. Srinivasan. Leader location, cooperation, and coordination in serial supply chains. *Production and Operations Management* 15(1):22–39, 2006.

[42] D. Martimort. Multi-contracting mechanism design. R. Blundell, W. K. Newey, and T. Persson, eds. *Advances in Economics and Econometric: Theory and Applications, Ninth World Congress*, Vol 1. Econometric Society Monographs, Cambridge University Press, Cambridge, UK, 57–101, 2006.

[43] D. Martimort and L. Stole. The revelation and delegation principles in common agency games. *Econometrica* 70(4):1659–1673, 2002.

[44] A. Mas-Colell, M. Whinston, and J. Green. *Microeconomic Theory.* Oxford University Press, New York, 1995.

[45] R. B. Myerson. Incentive-compatibility and the bargaining problem. *Econometrica* 47(1):61–73, 1979.

[46] R. B. Myerson. *Game Theory: Analysis of Conflict.* Harvard University Press, Cambridge, MA, 1997.

[47] Norske Skog. Outlook: Norske Skog H1 Pretax 950 mln-1.2 bln NKR vs. 854 mln. Press release, Lysaker, Norway, August 21, 1996.

[48] M. J. Osborne and A. Rubinstein. *A Course in Game Theory.* MIT Press, Cambridge, MA, 1994.

[49] Ö. Özer, Y. Zheng, and K.-Y. Chen. Trust in forecast information sharing. *Management Science* 57(6):1111–1137, 2011.

[50] B. A. Pasternack. Optimal pricing and return policies for perishable commodities. *Marketing Science* 4(2):166–176, 1985.

[51] S. Reichelstein. Constructing incentive schemes for government contracts: An application of agency theory. *Accounting Review* 67(4):712–731, 1992.

[52] P. Robinson. How do you get a good deal on software? It pays to pay attention. *Seattle Times* (November 6), D7, 1994.

[53] W. Rogerson. On the optimality of linear contracts. Mimeo, Northwestern University, Evanston, IL, 1987.

[54] W. Rogerson. Simple menus of contracts in cost-based procurement and regulation. *American Economic Review* 93(3):919–926, 2003.

[55] B. Salanié. *The Economics of Contracts: A Primer*. MIT Press, Cambridge, MA, 1997.

[56] P. Sellers and J. Reese. The dumbest marketing ploy. *Fortune* 126(7):88–92, 1992.

[57] H. A. Simon. Behavioral model of rational choice. *Quarterly Journal of Economics* 69(1):99–118, 1955.

[58] Space Daily. Yingli Green Energy and Xinguang agree to polysilicon pricing for 2008. Press release, Baoding, China, November 1, 2007. Retrieved March 17, http://business.highbeam.com/435147/article-1G1-170556799/yingli-green-energy-and-xinguang-agree-polysilicon, 2008.

[59] J. Spengler. Vertical integration and antitrust policy. *Journal of Political Economy* 58(4):347–352, 1950.

[60] T. A. Taylor and E. L. Plambeck. Simple relational contracts for capacity investment: Price-only vs. quantity-commitment. *Manufacturing & Service Operations Management* 9(1):94–113, 2007.

[61] T. A. Taylor and W. Xiao. Incentives for retailer forecasting: Rebates vs. returns. *Management Science* 55(10):1654–1669, 2009.

[62] B. Tomlin. Capacity investment in supply chains: Sharing the gain rather than sharing the pain. *Manufacturing & Service Operations Management* 5(4):317–333, 2003.

[63] A. A. Tsay, S. Nahmias, and N. Agrawal. Modeling supply chain contracts: A review. S. Tayur, R. Ganeshan, and M. Magazine, eds. *Quantitative Models for Supply Chain Management*, Chapter 10. Kluwer Academic Publishers, Norwell, MA, 299–336, 1998.

[64] U.S. Department of Commerce. Statistics for industry groups and industries: Annual survey of manufacturers. Accessed August 15, 2008, http://www.census.gov/prod/2006pubs/am0531gs1.pdf, 2006.

[65] Visual Defence Inc. Final results. Press release, Richmond Hill, ON, Canada, April 25, 2008, retrieved April 22, 2009. http://www.reuters.com/article/2008/04/25/idUS81019+RNS20080425.

[66] Y. Wang and Y. Gerchak. Capacity games in assembly systems with uncertain demand. *Manufacturing & Service Operations Management* 5(3):252–267, 2003.