MANAGEMENT SCIENCE
Vol. 57, No. 4, April 2011, pp. 689–704
ISSN 0025-1909 | EISSN 1526-5501 | 11 | 5704 | 0689
DOI 10.1287/mnsc.1110.1318

Contract Complexity and Performance Under Asymmetric Demand Information: An Experimental Evaluation

Basak Kalkanci
Department of Management Science and Engineering, Stanford University, Stanford, California 94305,
kalkanci@stanford.edu

Kay-Yut Chen
HP Laboratories, Hewlett-Packard Company, Palo Alto, California 94304, kychen@hpl.hp.com

Feryal Erhun
Department of Management Science and Engineering, Stanford University, Stanford, California 94305,
ferhun@stanford.edu

Exploring the tension between theory and practice regarding complexity and performance in contract design is especially relevant. The goal of this paper is to understand why simpler contracts may commonly be preferred in practice despite being theoretically suboptimal. We study a two-tier supply chain with a single supplier and a single buyer to characterize the impact of contract complexity and asymmetric information on performance and to compare theoretical predictions to actual behavior in human subject experiments. In the experiments, the computerized buyer faces a newsvendor setting and has better information on end-consumer demand than the human supplier. The supplier offers either a quantity discount contract (with two or three price blocks) or a price-only contract, contracts that are commonplace in practice, yet different in complexity. Results show that, contrary to theoretical predictions, quantity discounts do not necessarily increase the supplier's profits. We also observe a more equitable distribution of profits between the supplier and the buyer than what theory predicts. These observations can be described with three decision biases (the probabilistic choice bias, the reinforcement bias, and the memory bias) and can be modeled using the experience-weighted attraction learning model. Our results demonstrate that simpler contracts, such as a price-only contract or a quantity discount contract with a low number of price blocks, are sufficient for a supplier designing contracts under asymmetric demand information.

Key words: behavioral operations management; all-unit quantity discount contracts; price-only contracts; complex contracts; contract performance; supply chain efficiency; asymmetric demand information; experience-weighted attraction learning model
History: Received December 18, 2008; accepted January 4, 2011, by Christian Terwiesch, operations and supply chain management.

1. Introduction

A common cause of the incentive issues in many decentralized supply chains today is the information asymmetry that exists between different firms. A typical example of such information asymmetry is with regard to the end-consumer demand. Due to being closer to the consumer market, buyers (who buy goods from their suppliers and sell those goods to end consumers) usually have better demand information than their suppliers. Instead of sharing it with their suppliers, buyers may use this superior demand information strategically: for example, personal computer and electronics manufacturers often submit "phantom orders" to guarantee a higher component capacity from their suppliers (Lee et al. 1997). Many times, suppliers cannot obtain this information through market research firms. For example, Wal-Mart does not share its data on DVD sales with research firms such as SoundScan or ACNielsen (Jiang et al. 2011). Hence, suppliers face a challenge: how should they design contracts when they have limited information about end-consumer demand relative to their buyers?

Theory suggests contracts with menus of different quantities and payments corresponding to different demand types of buyers. Although these contracts are theoretically optimal, they are also arbitrarily complex (Laffont and Martimort 2002, Bolton and Dewatripont 2005). This calls into question the feasibility of these optimal contracts in practice, because it has long been argued that decision makers may not be able to cope effectively with such complexity. According to

Zadeh (1973, p. 28), "as the complexity of a system increases, our ability to make precise and yet significant statements about its behavior diminishes." As a result, decision makers start to use simple decision rules instead of considering all of the intricacies of the ordering situation. In light of these discussions, questioning the equivalence of complexity and performance in contract design is especially relevant for practitioners and theorists alike.

Holmström and Milgrom (1987, p. 304) noted that "real world incentive schemes appear to take less extreme forms than the finely tuned rules predicted by . . . theory." Despite its potentially poor theoretical performance, one such simple contract, often the choice of practitioners, is a price-only contract. The price-only contract is simple because it requires only the specification of a single parameter: the wholesale price. On the other hand, scientific literature praises the more complex quantity discount contract because it has been shown theoretically to increase sales, reduce costs, increase the channel efficiency, allow for self-selected price discrimination, and eliminate the inefficiencies due to information asymmetry (Jeuland and Shugan 1983, Burnetas et al. 2007, Chu and Sappington 2007, Altintas et al. 2008). And indeed, quantity discount contracts are also observed regularly in practice. However, there is evidence that quantity discounts do not perform consistently (Altintas et al. 2008). One underlying reason for this inconsistency has been identified by a recent Adesso Solutions (2005) 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 (1992) as "the dumbest marketing ploy."

Given the conflicting theoretical and practical accounts of the simple price-only contract and the more complex all-unit quantity discount contracts,[1] our goal is to compare them in an experimental setting to develop contract design guidelines for suppliers. We test the theoretical prediction that, under asymmetric information, suppliers' profits increase as the complexity of the contract (i.e., the number of price blocks and thus the number of decisions about contract parameters) increases. Research (e.g., Camerer 2003) has shown that human behavior does not conform to strict predictions of theory when behavioral influences are considered. Our goal is to identify these behavioral influences and integrate them into theory. Our results will guide managers in designing contracts by incorporating trade-offs between the level of complexity and inefficiency.

To provide direct evidence regarding the potential trade-off between contract complexity and economic inefficiencies, we turn to experimentation. Because it is hard to conduct field experiments to compare different contracts under the same environment, experimentation in a laboratory setting is the natural choice. Hence, we conducted a series of experiments to compare the price-only contract with quantity discount contracts under asymmetric demand information. In particular, we study a two-tier supply chain with a single (human) supplier and a single (computerized) buyer where the buyer faces a newsvendor setting and has better information on demand than the supplier. The supplier determines the *pricing scheme*, and the buyer 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) Buyers' profits decrease as the complexity of the contract increases. We test each of these predictions.

Experimental results suggest a nontrivial trade-off between complexity and inefficiency of all-unit quantity discount contracts. 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 blocks) compared to the more complex quantity discount contracts. This is strong evidence that the notion that "complex contracts optimize the supplier's profit" is flawed. We also show that the price-only contract achieves comparable total supply chain profits with respect to the more complex quantity discount contracts, which leads us to conclude that when simple contracts are employed, human subjects achieve much better channel coordination than theory predicts. Finally, we observe a more equitable distribution of profits between suppliers and buyers.

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

[1] 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. Because the all-unit discount provides a higher incentive to self-select into the appropriate order range, we focus on the all-unit discount in this paper. Furthermore, Burnetas et al. (2007, theorem 6, p. 473) show that "the supplier can earn at least as much under an all-unit discount than under an incremental discount."

he made in the past, even if those choices may not have been 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 learning (EWA) model (Camerer and Ho 1999, Bostian et al. 2008), 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.

The rest of this paper is organized as follows. In §2, we summarize the related theoretical and behavioral literature. Section 3 provides the details of our model: the experimental setting and theoretical predictions under the experimental setting (§§3.1 and 3.2) and our hypotheses (§3.3). Section 4 details our methodology (§4.1) and displays the consequent results (§4.2). We provide the underlying behavioral observations in §5, and the behavioral model in §6. We conclude this paper with discussion and managerial insights in §7.

2. Literature Review

There are three streams of literature relevant to this research, which we briefly review below. We then highlight our contributions.

2.1. Theoretical Literature on All-Unit Quantity Discounts

All-unit quantity discounts are commonly used in practice because of their well-accepted role in increasing sales, reducing costs, and allowing self-selected price discrimination (Munson and Rosenblatt 1998, Chu and Sappington 2007, Altintas et al. 2008). Because of these benefits and the practical popularity of these contracts, there is an established theoretical literature on quantity discounts. (We refer readers to Benton and Park (1996) and Munson and Rosenblatt (1998) for extensive reviews.) The majority of the studies in this literature analyze quantity discounts with complete information. Jeuland and Shugan (1983), Weng (1995), and Chen et al. (2001) show that centralized channelwide profits can be achieved in a decentralized system by different quantity discount schemes. Corbett and de Groote (2000) and, more recently, Burnetas et al. (2007) consider coordinating the supply chain with quantity discounts when the buyer has private information.

2.2. Behavioral Literature on Contracts

Keser and Paleologo (2004) study the price-only contract in the interaction of a newsvendor buyer with a supplier. The authors conclude that even though the inefficiency of the price-only contract is not significantly different from theoretical predictions, the behaviors of the players are. Katok and Wu (2009) further extend this line of research by comparing the performance of a simple price-only contract with more complex buyback and revenue-sharing contracts. The authors automated one of the decision makers to eliminate the effect of social preferences. Their experiments show that complex contracts improve the supply chain performance significantly compared to a price-only contract; however, the improvement is less than what is theoretically predicted.

The behavioral literature on quantity discount contracts is still in its infancy. Lim and Ho (2007) study the optimal number of price blocks under complete information where demand is deterministic and characterized by a linear-inverse function. In this case, theory predicts that all benefits should accrue when the number of price blocks increases from one to two and that no additional benefits should be observed from increasing the number of blocks further. However, the authors observe that benefits continue to increase in the number of blocks because subjects cannot coordinate the supply chain with only two price blocks. Ho and Zhang (2008) question the role of framing by comparing a two-part tariff with an all-unit quantity discount in a setting similar to that of Lim and Ho (2007). The authors show that the fixed fee fails to increase the efficiency of the channel when framed as a two-part tariff, but achieves a higher efficiency when framed as a quantity discount. This is contrary to theory, which predicts that these two mechanisms should be equal. Thus, the authors conclude that framing matters when designing supply contracts.

2.3. Behavioral Modeling and EWA

To understand the behavioral perspective of human subjects in our experiments, we employ an EWA model. The EWA model, developed by Camerer and Ho (1999), combines belief-based models with reinforcement models to reflect both 6"the history of how others played" and "past successes (reinforcements) of chosen strategies" (p. 828). As such, the complex and dynamic EWA model captures three decision biases (i.e., probabilistic choice/bounded rationality, the reinforcement bias, and the memory effect) and effectively describes deviations from theory. We refer readers to Camerer (2003) for a detailed review of EWA, and to Bostian et al. (2008) for an application in the newsvendor context.

2.4. Our Contributions

Although the existing literature has theoretically studied the role of complexity in contracts (e.g., Wilson 1993, Laffont and Martimort 2002), to the best of our knowledge, this paper is the first to study contract complexity in the number of decisions with behavioral experiments. In addition, we identify, model mathematically, and quantify three major decision biases in this environment, two of which (i.e., the reinforcement bias and the memory effect) have not been studied before in a supply chain contracting setting.

In the behavioral literature, the paper by Lim and Ho (2007) is closest to this paper. With behavioral experiments, the authors study quantity discount contracts but they do not focus on the issue of complexity. In particular, their subjects make exactly the same number of decisions under different numbers of price blocks because the authors fix some of the decisions to their optimal values. In addition, Lim and Ho (2007) study a situation with complete information even though almost all contracting situations in practice exhibit asymmetric information. Indeed our results show that commonly observed results in the behavioral literature under complete information may no longer hold under asymmetric information, thereby validating our focus (see §4.2 for details).

We adapt the EWA model to contracting decisions in a supply chain setting to quantify the decision biases and to examine the effects of contract complexity. Using the EWA model, we verify that as contract complexity increases, human subjects increasingly rely on simple heuristics. The work of Bostian et al. (2008), which applies the EWA model to a newsvendor setting, is the only other EWA research in the supply chain area that we are aware of. Our work is significantly different from that study because we model the supplier's price and price break decisions, whereas Bostian et al. (2008) focus on a single quantity decision made by the buyer. Our model is designed to capture substantial complexity in the supplier's decision space. In particular, we may have up to five decision variables per individual, which complicates the estimation significantly, requiring about a billion calculations just to characterize the probability distribution of decisions.

3. Experimental Design and Expected Profit Maximizing Model

We study a setting with a single supplier (he) and a single buyer (she). The buyer procures a product from the supplier and sells it to the end consumer at p dollars per unit. The supplier has ample capacity and produces the buyer's orders at a cost of k dollars per unit, where $k < p$. Products are ordered before demand is realized. Any unmet demand is lost without a stockout penalty. There is no salvage value or disposal cost for leftover products. The supplier determines the *pricing scheme* of the component and the buyer selects an *order quantity* that maximizes her expected profit.

Our buyer is a newsvendor who faces a random demand, D, which is uniformly distributed between $(\mu - v)$ and $(\mu + v)$, where μ is the mean demand, and v defines the range. The mean demand is the buyer's private information. The supplier only knows that the mean demand is one of three types: high (μ_H), medium (μ_M), or low (μ_L), each with equal probability. We assume that types are equally spaced; i.e., $\mu_M = \mu_L + \delta$ and $\mu_H = \mu_M + \delta$, where δ is the degree of separation between types. We also assume that the lowest possible demand is zero; i.e., $v = \mu_L$.[2]

We study different contracts between the supplier and the buyer. In a price-only contract (one-price contract), the supplier sets a single wholesale price w_1, and the buyer 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 buyer 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 buyer 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.[3] The number of the supplier's decision variables depends on the contract employed, whereas the buyer's only decision is the order quantity. Note that because we assume free disposal, under some conditions for the all-unit quantity discount contract it is optimal for the buyer to procure more units than she needs to achieve a lower wholesale price for all units.

3.1. Expected Profit Maximizing Theory of Price-Only and Quantity Discount Contracts

To derive the expected profit maximizing theory of price-only and quantity discount contracts, we first

[2] Note that our theoretical development coincides with our experimental setting. In our experimental setting, we assume that the lowest possible demand is zero because this assumption increases the incentive to use more complex contracts by sharpening the comparison between simple and more complex contracts. Theoretically, this assumption simplifies the comparisons by removing the guaranteed payoff of the buyer.

[3] Theoretically, increasingly complex contracts would benefit the supplier. In our experimental setting, because the mean demand is one of three types, up to three different price blocks with a minimum quantity commitment Q_0 would be enough to separate these types, and more price blocks would have no additional benefit for the supplier or the supply chain. In reality, there are potentially many more demand types than three, and arbitrary complex contracts may be needed to separate them. Therefore, our design provides a much higher chance for subjects to succeed because the setting (fewer types) is simpler than in the field; that is, any results showing decision biases in reaction to contract complexity would be even more exaggerated in the real world.

present the computerized buyer's optimal behavior. Let

$$S_j(\mu) = F^{-1}\left(\frac{p-w_j}{p}\right)^+ = \begin{cases} (\mu - v) + 2v\left(\dfrac{p-w_j}{p}\right) & p \geq w_j, \\ 0 & \text{otherwise} \end{cases} \tag{1}$$

be the order-up-to level for the wholesale price w_j ($j = 1, 2, 3$), where $F(\cdot)$ is the cumulative distribution function of demand distribution, and $(x)^+ = \max(0, x)$. For a price-only contract, the buyer's optimal order quantity is simply $S_1(\mu)$. Theorem 1 details her optimal order quantity under all-unit quantity discount contracts:

Theorem 1 (Jucker and Rosenblatt 1985). *Under all-unit quantity discounts, the buyer's optimal order quantity is either at one of the order-up-to levels or at one of the price breaks.*

Given the buyer's optimal response to a contract (Equation (1)) when types are equally spaced (i.e., $v = \mu_L$, $\mu_M = \mu_L + \delta$, and $\mu_H = \mu_M + \delta$, where δ is the degree of separation between types) and equally likely, the supplier's optimal wholesale price under a price-only contract is

$$w_1 = \min\left\{\frac{p}{4v}\left(\frac{\mu_L + \mu_M + \mu_H}{3}\right) + \frac{p}{4} + \frac{k}{2}, p\right\} = \frac{p}{2}\min\left\{1 + \frac{1}{2}\left(\frac{\delta}{\mu_L}\right) + \frac{k}{p}, 2\right\}. \tag{2}$$

Therefore, the buyer's and the supplier's optimal decisions under the price-only contract are fully characterized by Equations (1) and (2), respectively.

The supplier's pricing decisions under quantity discounts with asymmetric information are more involved. When the supplier has complete information, he can simply design a two-price contract (which essentially acts as a two-part tariff) and can extract the entire supply chain profit. Thus, the supplier does not have any reason to offer a contract more complex than a two-price contract under complete information. However, in the presence of information asymmetry, this is no longer the case. The supplier can effectively utilize a three-price contract, and even then he has to pay an information rent to the buyer. Under asymmetric information, prices and price breaks serve an additional role: they are utilized to separate the different types of the buyer; that is, the supplier must consider the buyer's incentive compatibility while designing contracts, which increases the complexity of his decision.

Under quantity discounts, contract design requires comparisons among several alternatives. For example, when the supplier offers a two-price contract, he needs to consider cases (such as "all types buy at or above the price break," "high and medium types buy at or above the price break," "the high type buys at or above the price break," and "all types buy below the price break") before he sets his contract parameters. However, when types are equally spaced (i.e., $v = \mu_L$, $\mu_M = \mu_L + \delta$, and $\mu_H = \mu_M + \delta$, where δ is the degree of separation between types), the pricing scheme has a structure as presented in the following proposition (the proof is in Online Appendix A, provided in the e-companion).[4]

Figure 1 **Equilibrium Cases in Two-Price and Three-Price Contracts**

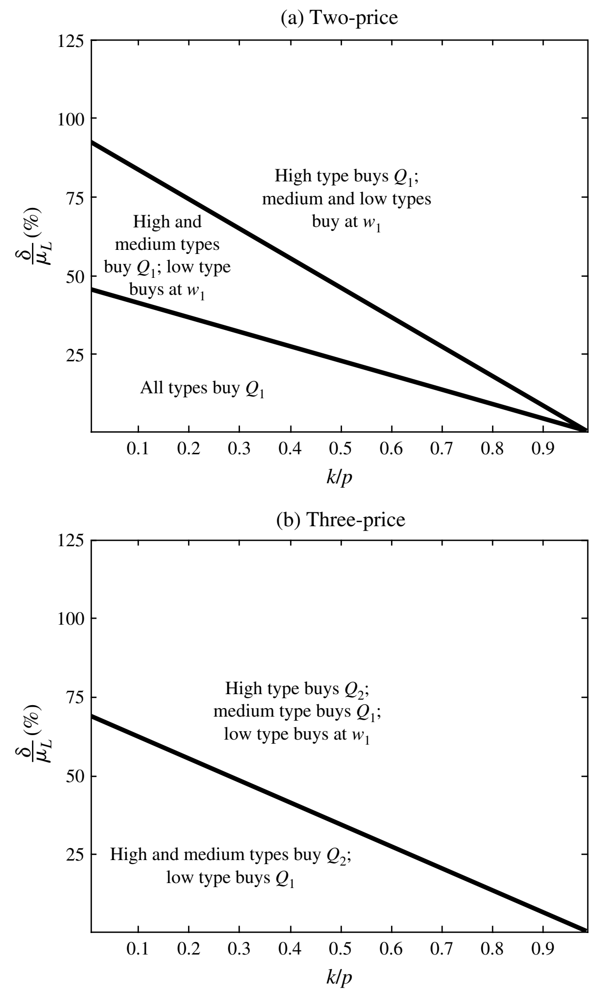

Proposition 1. *In quantity discount contracts, the supplier sets price schemes leading to the procurement structures displayed in Figure* 1.

Our assumptions guarantee that the high type always procures at the price break. How medium and low types act, on the other hand, depends on the degree of separation. If types are close to each other

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

(i.e., δ/μ_L is close to zero), the degree of information asymmetry significantly decreases, and the prices are set without separating types. However, as δ/μ_L increases, the value of separation also increases, and the supplier first separates the low type, followed by the medium type. Another factor in pricing is the relative cost of the supplier. As the cost increases, the profit margin of the supplier decreases, and the supplier is more inclined to separate types to extract the highest profit possible. Note that because of its additional parameters, a three-price contract is more flexible than a two-price contract. Thus, the supplier needs to consider fewer cases when setting the pricing scheme under a three-price contract, and he is able to separate the high type for lower values of δ/μ_L (Figure 1). However, finding the optimal pricing scheme is more complicated under a three-price contract because the supplier must consider many more interactions.

Our goal is to study the role of contract complexity. The parameter values where contract complexity pays off the most are those where the separation between types is high and the profit margin is relatively low, i.e., when the supplier completely separates the high type from the others. In this region, we can characterize the components of the pricing scheme. The next proposition details the pricing scheme for a two-price contract. A similar analysis follows for a three-price contract.

Proposition 2. *If only the high type buys at the price break, the optimal two-price contract is defined as follows*:

$$w_1 = \min\left\{\frac{3p}{10v}\left(\frac{\mu_L+\mu_M+\mu_H}{3}\right)+\frac{3p}{10}+\frac{2k}{5}, p\right\}$$
$$= \frac{p}{5}\min\left\{3+\frac{3}{2}\left(\frac{\delta}{\mu_L}\right)+2\left(\frac{k}{p}\right), 5\right\},$$
$$Q_1 = (\mu_H - v) + 2v\left(\frac{p-k}{p}\right) = 2\mu_L\left(1+\frac{\delta}{\mu_L}-\frac{k}{p}\right),$$

and w_2 is set to make the high-type buyer indifferent between buying at w_1 and w_2.

Now that we have characterized the optimal pricing schemes under price-only and quantity discount contracts, we next discuss our experimental design.

3.2. Experimental Design

Our primary goal is to test the impact of contract complexity on the profits of the supplier, the buyer, and the total supply chain. We used a between-subject design where each subject was faced with one of three different treatments: acting as the supplier and choosing the parameters of the one-price, two-price, or three-price contract defined above. Our human subjects were recruited from the Stanford student body and were provided monetary compensation according to their performance. Subjects were given Web-based instructions and a quiz before the experiment. (The detailed instructions and the quiz are provided in Online Appendix B.) We implemented our experiments in the HP Experimental Economics software platform and conducted them at the Stanford Graduate School of Business Behavioral Lab and the Stanford School of Medicine computer labs. To facilitate subjects' decision process, we provided them with a decision support tool. Before a subject submitted a decision, he could use this "what-if" decision support tool to evaluate his decision. In this tool, he could enter his potential decision as well as his buyer's potential average demand and order quantity.[5] Then, the tool showed the subject his potential profit and his buyer's profit under typical demand conditions. Subjects were able to test several different decisions before submitting their pricing scheme. Subjects could also see a summary of results from past periods in a history table. The table included the subject's wholesale prices/price breaks, the buyer's order quantity, the buyer's realized average demand, and the subject's revenue and profit from that period. The inclusion of the decision support tool was intended to give theory its best shot in determining subjects' behavior in the experiments.

We performed three sets of experiments in the fall of 2009. 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 computerized buyer was programmed to maximize her expected profits. During each period, each computerized buyer was randomly assigned one of three types (high, medium, or low). In all experiments, we set the parameter values as follows: the supplier's unit cost was $k = 40$, the unit selling price was $p = 200$, the mean demand for the low type was $\mu_L = 50$, and the degree of separation of types was $\delta = 40$ (i.e., the mean demand for the medium type was $\mu_M = 90$, and for the high type it was $\mu_H = 130$). The range of the demand distribution was $2v = 100$.

Based on our parameter setting, the supplier should choose his pricing scheme such that for the two-price contract only the high type buys at the price break. Table 1 displays the optimal values of prices and price breaks along with the buyer's procurement quantities under one-price, two-price, and three-price contracts and the expected profits for the supplier, the buyer, and the total supply chain. In our setting, theory predicts a 20% (65%) increase (decrease) in the supplier's (buyer's) expected profits when the supplier moves

[5] Note that the decision support tool does not provide the buyer's optimal procurement quantity for each type.

Table 1 Observed (Obs.) and Optimal (Opt.) Contract Parameters, Procurement Quantities, and Supplier, Buyer, and Total Supply Chain Profits for Our Parameter Set Under Different Contracts (Rounded)

		Supplier's decisions						Buyer's decisions				
		Prices			Price breaks			Procurement quantities				
		w_1	w_2	w_3	Q_1	Q_2	Expected profit	q_L	q_M	q_H	Expected profit	Total expected profit
One-price	Obs.	142					6,722	28	70	108	3,270	9,992
	Opt.	160					7,200	20	60	100	2,000	9,200
Two-price	Obs.	163	142		102		7,065	21	78	112	2,841	9,906
	Opt.	184	152		160		8,640	8	48	160	704	9,344
Three-price	Obs.	161	143	126	81	129	6,840	24	83	126	3,619	10,459
	Opt.	200	177	147	84	160	9,576	0	84	160	669	10,245

from a one-price contract to a two-price contract, and an 11% (5%) increase (decrease) when the supplier moves from a two-price contract to a three-price contract. Hence, as the contract complexity increases, the supplier's expected profits increase, and the buyer's expected profits decrease. Overall, the supply chain efficiency increases 11% when a simple price-only contract is replaced with more complex and efficient quantity discount contracts.

3.3. Hypotheses

Under asymmetric information, theory predicts that as contract complexity increases, the supplier increases his profits; that is, the supplier prefers quantity discount contracts with additional price blocks to simple price-only contracts. Furthermore, as the number of price blocks increases, the inefficiencies in the system are partially eliminated, thereby increasing the expected total supply chain profits. However, this increase in total supply chain profits is at the expense of the buyer. In other words, as the number of price blocks increases, the buyer's advantage due to her superior information decreases; that is, theory predicts that the buyer is worse off under quantity discount contracts compared to a simple price-only contract. Our first hypothesis tests these theoretical assertions.

Hypothesis 1 (Complexity vs. Expected Profits).
(a) *The more complex the contract, the higher the supplier's expected profit.*
(b) *The more complex the contract, the higher the expected total supply chain profit.*
(c) *The more complex the contract, the lower the buyer's expected profit.*

Hypothesis 1 compares contracts to each other and investigates whether or not the supplier can use contract complexity to his benefit to extract more profit from the buyer while decreasing the inefficiencies in the supply chain. Our second hypothesis compares contracts with their theoretical benchmarks. Previous literature has studied the coordination of supply chains both theoretically and empirically under complete information and has concluded that contracts that are known to coordinate supply chains theoretically may not be as effective when tested empirically (Katok and Wu 2009). The following hypothesis tests this finding in a setting with asymmetric information.

Hypothesis 2 (Coordination).
(a) *Under each contract, the supplier's expected profit is equal to what theory predicts* (*Table* 1).
(b) *Under each contract, the expected total supply chain profit is equal to what theory predicts* (*Table* 1).

Note that our hypotheses provide comparisons with theoretical predictions, without taking any underlying behavioral factors into consideration. We revisit those factors in §6.

4. Analysis

In this section, we examine the experimental performance of the three contracts and compare them to theoretical predictions. First, we outline our methodology in §4.1. Then, in §4.2, we test our hypotheses. Before we detail our analysis, Table 1 summarizes the supplier's and the buyer's decisions as well as the average profits based on theoretical predictions and experimental observations.

4.1. Methodology

We use a regression model to characterize the impact of contract types, buyer's demand types, time, and variation among suppliers on supplier, buyer, and total supply chain profits.[6] All three regression equations are similar; we provide the one for the supplier's profit below:

$$\begin{aligned}\Pi_{i,t}^{j} = {} & Intercept + \beta_t \times t + \beta_{QD2} \times QD2 + \beta_{QD3} \times QD3 \\ & + \beta_M \times M + \beta_H \times H + \beta_{QD2\times M} \times (QD2 \times M)\end{aligned}$$

[6] We used regression analysis to be able to address all of the factors that could affect profits. However, average profit analysis using simpler statistics (such as the t-test or Wilcoxon test) leads to the same results as in §4.2.

$$+\beta_{QD2\times H}\times(QD2\times H)+\beta_{QD3\times M}\times(QD3\times M)$$
$$+\beta_{QD3\times H}\times(QD3\times H)+\beta_{QD2\times t}\times(QD2\times t)$$
$$+\beta_{QD3\times t}\times(QD3\times t)+\nu_i+\epsilon_{i,t}.$$

The dependent variable $\Pi^j_{i,t}$ is the supplier's profit under a j-price contract ($j=1,2,3$), considering the effects of each individual supplier (i) and period (t). The independent variables $QD2$ and $QD3$ are dummy variables for two-price and three-price contracts, followed by two dummy variables M and H for medium and high demand types, respectively. The contract–demand-type terms capture possible interactions in the relations of contract and demand types. We represent the effect of period in two ways: an independent variable t captures the direct effect, and interaction terms with contract type capture the incremental effect. We use a random-effects model (ν_i) to control for individual heterogeneity in the subject pool. Both ν_i and the error term $\epsilon_{i,t}$ are assumed to be normally distributed with mean zero and a positive standard deviation. Note that there is no random-effects term for buyers because they are computerized and thus do not contribute to variability. For this reason, even in the regression equations for buyer and total supply chain profits, we only capture suppliers' individual heterogeneity. We provide the regression estimates in Table 2.

Although the regression analysis helps us devise some basic insights, a direct comparison of the average profits under different contracts does not follow immediately from this analysis. The incremental effect of quantity discounts on the profits obtained from different demand types as well as the time trend are represented by separate terms and are not averaged. Therefore, we use linear hypothesis testing (Freund et al. 2006) to analyze the predictions in §3.3. We first calculate estimated profits averaged over period and demand types.[7] The average profits of the supplier under one-price, two-price, and three-price treatments estimated from the regression are defined, respectively, as

$$\pi^1_s := Intercept+\beta_t\times 20+1/3\times\beta_M+1/3\times\beta_H,$$
$$\pi^2_s := Intercept+\beta_t\times 20+\beta_{QD2}+1/3\times\beta_M+1/3\times\beta_H +1/3\times\beta_{QD2\times M}+1/3\times\beta_{QD2\times H}+\beta_{QD2\times t}\times 20,$$
$$\pi^3_s := Intercept+\beta_t\times 20+\beta_{QD3}+1/3\times\beta_M+1/3\times\beta_H +1/3\times\beta_{QD3\times M}+1/3\times\beta_{QD3\times H}+\beta_{QD3\times t}\times 20.$$

We then use these average profits under different contract types to test our hypotheses. Table 3 reports the comparison of the supplier's, buyer's, and total supply chain profits with theoretical predictions using linear hypothesis testing.

[7] Averaging profits over period corresponds to the profits at time period $t=20$. Averaging profits over demand types corresponds to a probability of 1/3 for each demand type.

Table 2 **Regression Coefficients for Supplier, Buyer, and Total Supply Chain Profits**

	Regression coefficient estimates and standard errors (in parentheses)		
Variable	Supplier's profit	Buyer's profit	Total profit
Intercept	2,290.58*** (176.11)	1,735.91*** (370.81)	4,026.7*** (272.94)
t	15.26*** (4.48)	−39.39*** (4.64)	−24.14*** (3.6)
*QD*2	−620.19* (249.46)	−768.77 (524.62)	−1,390.18*** (386.16)
*QD*3	3.64 (249.46)	−781.23 (524.62)	−778.41* (386.16)
M	4,060.82*** (126.52)	2,482.55*** (131.24)	6,543.88*** (101.92)
H	8,282.34*** (127.17)	4,525.1*** (131.89)	12,806.78*** (102.43)
*QD*2 × *M*	1,496.23*** (179.19)	−271.43 (185.88)	1,227.08*** (144.35)
*QD*2 × *H*	856.14*** (179.93)	399.22* (186.61)	1,255.96*** (144.92)
*QD*3 × *M*	969.29*** (179.18)	295.54 (185.86)	1,267.75*** (144.34)
*QD*3 × *H*	128.7 (179.93)	1,847.7*** (186.61)	1,977.5*** (144.92)
*QD*2 × *t*	7.48 (6.34)	14.24* (6.57)	21.72*** (5.10)
*QD*3 × *t*	−13.39* (6.34)	19.84** (6.58)	6.44 (5.11)
Standard deviation of ν_i	532.91***	1,511.36***	1,103.34***

$^*p=0.05$; $^{**}p=0.01$; $^{***}p=0.001$.

4.2. Results

We test the theoretical prediction that suppliers' profits increase as the complexity of the contract increases (Hypotheses 1(a) and 2(a)) in §4.2.1, that the supply chain efficiency increases as the complexity of the contract increases (Hypotheses 1(b) and 2(b)) in §4.2.2, and that buyers' profits decrease as the complexity of the contract increases (Hypothesis 1(c)) in §4.2.3.

4.2.1. Suppliers Do Not Necessarily Benefit from More Complex Contracts. From Table 3, we conclude that suppliers' profits are significantly lower than predicted by theory; that is, contrary to Hypothesis 2(a), suppliers cannot use the contracts as effectively as theory predicts. We find partial support for Hypothesis 1(a): consistent with theoretical results, a two-price contract leads to higher profits for the supplier than a one-price contract ($p=0.048$). However, the observed improvement is less than what is

Table 3 Means and Standard Deviations (in Parentheses) for Supplier, Buyer, and Total Supply Chain Profits and Comparison of Respective Profits with Theoretical Predictions

	Supplier's profit		Buyer's profit		Total profit	
Treatment	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

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 blocks may be sufficient for a supplier designing contracts under asymmetric demand information.

Worth noting is that the demand type of the buyer is a strong predictor of the supplier's profits. This confirms the intuition that higher demand types result in higher profits (Table 2). The interaction terms reveal that suppliers extract significant gains from the medium type (in two-price and three-price contracts) and the high type (in two-price contract), potentially by foregoing some of their profits from the low type. This behavior is in line with how suppliers would design the quantity discount contracts at optimality. Therefore, these significant interaction terms possibly indicate that subjects understand the mechanism of quantity discount contracts but cannot effectively set the optimal parameters, especially when contract complexity increases. We revisit this observation in §6. Time trends of the supplier's profits also show interesting dynamics: the coefficient of the period under the one-price treatment is positive and significant, implying that the supplier's profits tend to increase over time under a price-only contract. Because the incremental effect of the period under the two-price treatment (i.e., $QD2 \times t$) is not significant, a two-price contract follows a similar dynamic; i.e., the suppliers are able to improve their profits as they become more familiar with the contract. However, this effect is no longer significant for a three-price contract because the effect of period is not significantly different from zero (two-sided $p = 0.677$). Thus, suppliers fail to benefit from highly complex contracts even after they use such contracts for many periods.

4.2.2. Simple Contracts Lead to Significantly Higher Supply Chain Efficiencies Than Predicted by Theory and Are Comparable to Complex Contracts. Theory predicts an 11% increase in total supply chain profits when the supplier moves from a one-price contract to a three-price contract. We do not observe such an increase in our experiments. On the contrary, 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 3, 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 blocks are employed, human subjects achieve better channel coordination than theory predicts. Hence, we do not find any support for Hypotheses 1(b) and 2(b) in our experiments.

4.2.3. We Observe a More Equitable Distribution of Profits Between Suppliers and Buyers Than Theory Predicts. Theoretically, buyers' expected profits should decrease 66% as suppliers move from a one-price contract to a three-price contract. In our experiments, we do not observe support for this assertion; buyers' 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 buyers' profits as effectively as theory predicts. Table 3 also supports this observation: buyers (suppliers) earn significantly more (less) than what theory predicts under each contract; that is, we observe a more equitable distribution of profits between suppliers and buyers compared to what is theoretically predicted. Hence, we do not find any support for Hypothesis 1(c) in our experiments. Note that this result shows that equitable distribution of profits may not always be attributed to fairness considerations because buyers are computerized in our experiments.

Until now, we have focused on the performance of price-only and quantity discount contracts. In the next section, we delve into how subjects make their decisions in our experiments.

5. Behavioral Observations

In §4 we observe that the relative performances of the contracts deviate significantly from theoretical predictions. In the rest of this paper, we investigate the underlying behavioral reasons for these deviations. In this section, we provide observations on the decisions of our subjects. In the next section, we develop a behavioral model based on these observations.

Figures 2(a)–(f) show that suppliers chose significantly lower wholesale prices compared to theoretical predictions under all contracts; Wilcoxon tests comparing each wholesale price with its theoretical prediction also support this conclusion at a significance level under 0.05 for all tests. Suppliers also had difficulty setting the correct price breaks (Figure 2, (g) and (h)). At optimality, the supplier should set the discount scheme such that for high and medium types of the buyer, the optimal procurement quantity is at the price breaks in the three-price treatment (Table 1). However, that is not the case in the experiments. Of a total of 252 (256) high (medium) demand type occurrences, the buyer's best-response quantity is equal to the correct quantity break in only 156 (85) cases. We also observe that, except for the initial periods, subjects do not change their decisions significantly. These observations support our finding in §4; i.e., suppliers' low prices and inability to separate types by setting correct quantity breaks explain the more equitable division of profits between the buyer and the supplier relative to theoretical predictions.

In fact, the determination of the price breaks is the component of the supplier's decision making that would lead to the largest increase in profit if improved. We arrive at this conclusion by optimizing the value of a single parameter while keeping the remaining parameters at the values chosen by human subjects. We then compare the supplier's expected profit improvements among all alternatives. Under both quantity discount contracts, a correction in one of the price breaks would lead to the highest improvement (up to 68%) in the supplier's expected profits. We conclude that suppliers understand the dynamics of pricing decisions better than the dynamics of price break decisions and can set the prices more successfully compared to price breaks under quantity discounts, despite a similar level of complexity involved. Therefore, a decision support tool that would help suppliers set their discount schemes effectively would be especially beneficial.

6. Behavioral Model

To understand our subjects' deviations from the expected profit maximizing behavior, we first consider explanations such as social preferences like altruism and fairness, loss aversion, and risk aversion. Social preferences (see Niederhoff and Kouvelis (2008) and Cui et al. (2007) for behavioral and theoretical treatments of such concerns, respectively) cannot be the underlying reason in our setting because by designing an experimental setting with a human supplier facing a computerized buyer, we suppress these preferences to the extent possible (similar to Katok and Wu 2009). Because the supplier does not face any loss in our setting, we also eliminate loss aversion as a possible explanation. We rule out risk aversion after observing that the predictions under risk aversion are inconsistent with our observations. In particular, with a log utility function, risk aversion would push w_1 to its maximum value (200 in our case) in two-price and three-price contracts, but we observe that subjects quote prices even below the risk-neutral optimal values. Given that these alternatives do not fully explain behavior of the subjects and are less robust, we turn to an alternative explanation: bounded rationality.

With his seminal work, Simon suggested that when "faced with complexity beyond his ken" a human subject finds "ways of action that are sufficient unto the day" (Simon 1978, p. 368); i.e., he relies on simpler decision rules (Simon 1955). For example, unless there is a compelling reason to do so, humans tend not to change their established behavior (Simon 1978, Samuelson and Zeckhauser 1988). Similarly, humans have a tendency to weigh recent decisions more heavily than earlier decisions, especially when the complexity of decisions is high (Hogarth and Einhorn 1992). Such behavioral biases clearly indicate that when faced with complex decisions, humans act boundedly rational. A boundedly rational decision maker lacks the propensity to optimize; although he may choose better alternatives more often, his decision making is subject to decision noise and random decision errors (McKelvey and Palfrey 1995, Su 2008).

Qualitatively, the behavior of our subjects shows similarities with these biases; i.e., our subjects deviate from the expected profit maximizing behavior by exhibiting a bias toward the decisions that they have already chosen, and are subject to inertia that

Figure 2 Suppliers' Price and Price Break Decisions Under All Contracts

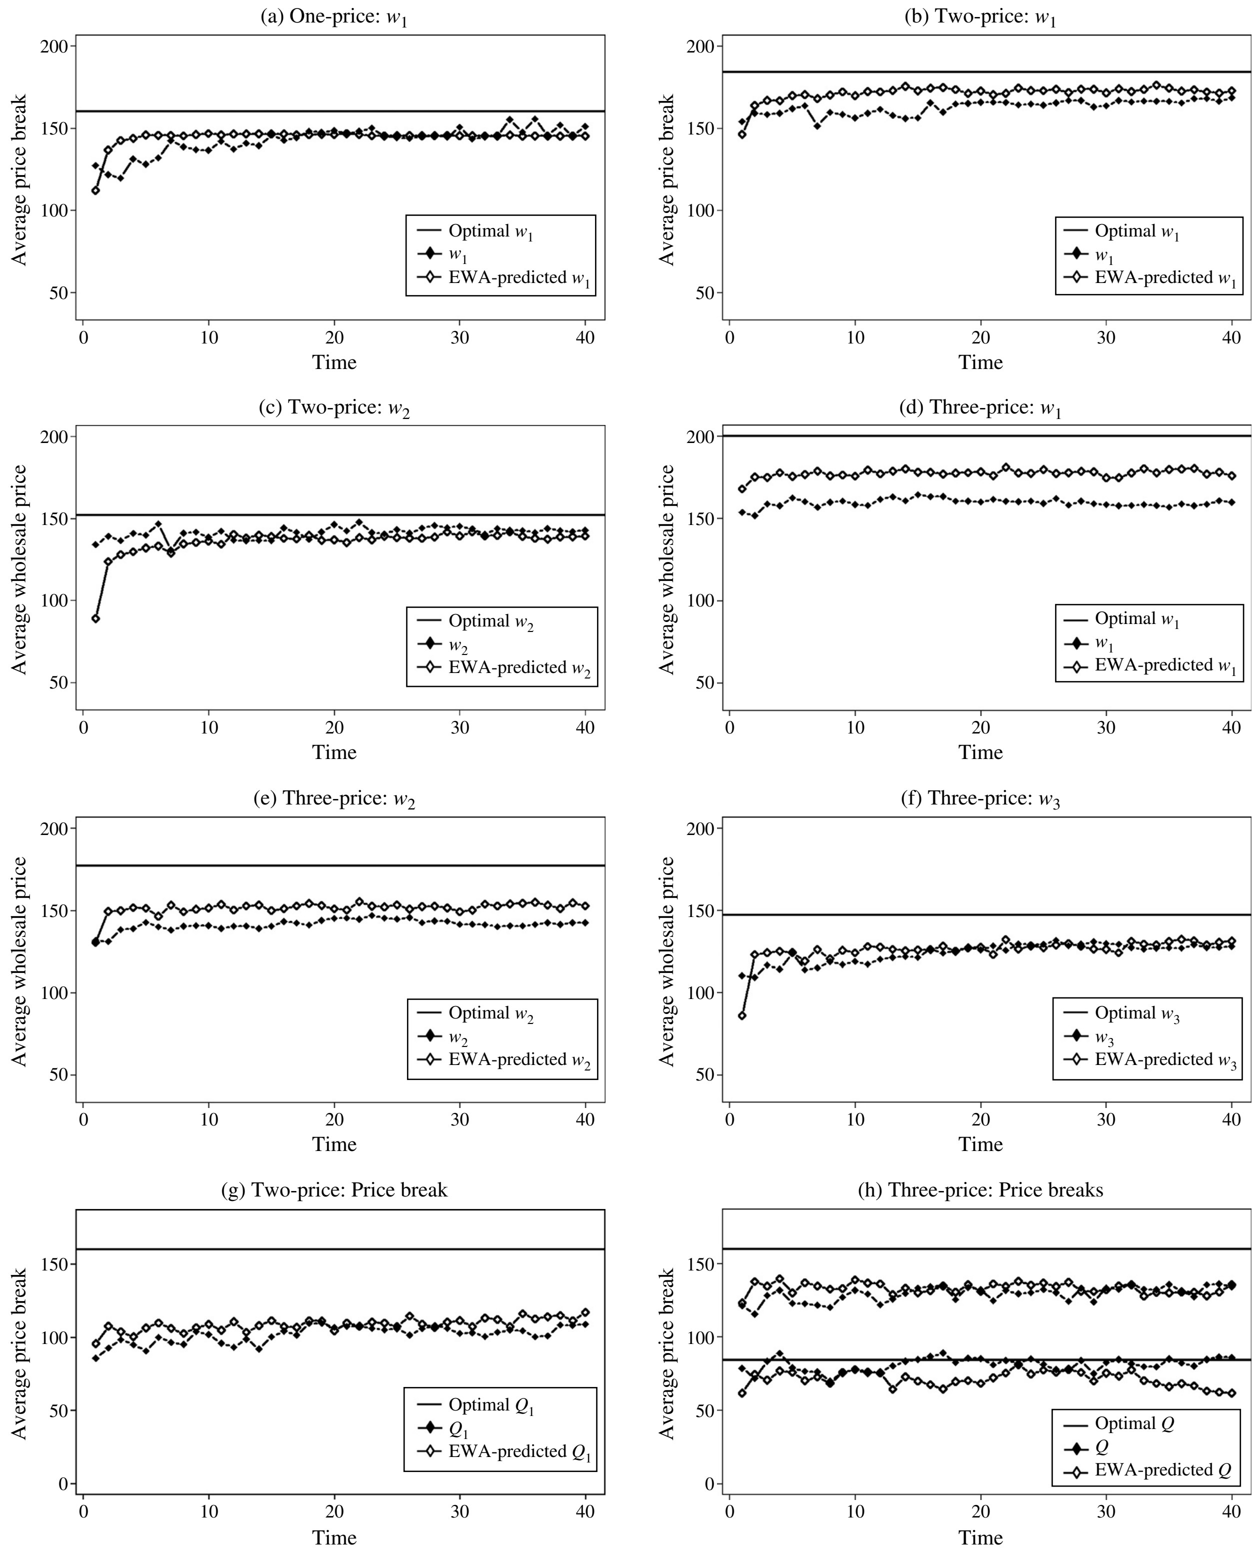

slows the moving to new decisions (see §5). Therefore, we conjecture that the subjects make their decisions probabilistically and rely on a simple heuristic by reinforcing their decisions. Thus, three decision biases govern the subjects' behavior: a *reinforcement* bias that captures the inertia, a *probabilistic choice* bias that captures the bounded rationality (variability/randomness) of the subjects, and a *memory* bias that captures how far a subject looks into the past when making decisions.

With the conjectured heuristic and three decision biases in mind, we ask the following questions: (1) Can the reinforcement, memory, and bounded rationality biases describe the trade-off between complexity and inefficiency and the decision trends observed in our experiments? (2) If so, how does contract complexity shape the decision biases? For consumer choice models, DeShazo and Fermo (2002, p. 127) argue the following:

> [There is] an inverse relationship between complexity and the use of simplifying decision rules. Thus, as complexity increased, so too would the use of the simplifying decision rules employed by the consumer. ... Heiner (1983) argued that increasing choice complexity would increase the gap between a respondent's cognitive ability and the cognitive demands of the decision. Heiner predicts that as this gap grows, consumers will find it welfare-enhancing to restrict the range of decision rules they consider. Although this behavior produces increasingly predictable outcomes, it is not welfare maximizing. Thus, for both Simon (1955) and Heiner (1983), an increase in choice set complexity will add noise to the error term of a random utility function. Finally, de Palma et al. (1994) consider the choice processes used by individuals with an imperfect ability to choose. De Palma et al. predict that as choice complexity increases, so too will the magnitude of sub-optimal mistakes.

We hypothesize that, in our model, the "noise to the error term of a random utility function" and "sub-optimal mistakes" manifest themselves through the increase in the reinforcement and the memory biases.

Next, we describe a framework known as the EWA model (Camerer and Ho 1999, Camerer 2003), which captures the three biases together effectively. We then use this framework to test the hypothesis above in §6.2.

6.1. Experience-Weighted Attraction Learning Model

Under the EWA model, each possible value of a decision of the supplier has a score that determines its probability of being chosen. The score of a decision $\mathbf{x}_i$ at period t, $s_t(\mathbf{x}_i)$, is a function of its previous score at $t-1$ and the profit that the supplier would receive if he chose x_i given the observed average demand of the buyer at the end of t (i.e., his counterfactual profit $\Pi^*(\mathbf{x}_i, \mu_t)$):

$$s_t(\mathbf{x}_i) = \frac{1-\rho}{1-\rho^{t+1}}\Pi^*(\mathbf{x}_i, \mu_t) + \rho\frac{1-\rho^t}{1-\rho^{t+1}}s_{t-1}(\mathbf{x}_i), \qquad (3)$$

where ρ is the memory parameter, $0 \le \rho \le 1$. The individual remembers the scores given to each decision in the past and can use them for the present decision. When ρ is zero, the individual only uses the counterfactual profit and does not use previous scores at all. When ρ approaches one, the weight given to each past score approaches $1/t$. Thus, the individual treats all past calculations equally. A lower ρ favors recency effect, i.e., places more weight on recent calculations.

The counterfactual profit $\Pi^*(\mathbf{x}_i, \mu_t)$ represents a bias for counterfactual alternatives, i.e., a reinforcement bias:

$$\Pi^*(\mathbf{x}_i, \mu_t) = \begin{cases} \Pi(\mathbf{x}_i, \mu_t) & \text{if } \mathbf{x}_i = \mathbf{x}_t, \\ (1-r)\Pi(\mathbf{x}_i, \mu_t) & \text{if } \mathbf{x}_i \neq \mathbf{x}_t, \end{cases} \qquad (4)$$

where $\Pi^*(\mathbf{x}_i, \mu_t)$ is the supplier's profit under decision $\mathbf{x}_i$ given that the buyer's demand type is μ_t, and r is the reinforcement bias parameter, $0 \le r \le 1$. That is, the supplier will place less emphasis on the decisions he has *not* chosen in the past. When r is zero, there is no bias, and each period score, for each decision, is simply the expected profit of that decision. When r reaches 1, the counterfactual profits associated with decisions that have not been chosen for the period will be ignored.

Finally, the probabilistic choice/bounded rationality parameter determines the probability that the supplier chooses $\mathbf{x}_i$ at $t+1$ given the score $s_t(\mathbf{x}_i)$. A commonly used form of this probability function is the logit form (Camerer and Ho 1999, Bostian et al. 2008):

$$\Pr_{t+1}(\mathbf{x}_i) = \frac{\exp(\lambda s_t(\mathbf{x}_i))}{\sum_{x\in X}\exp(\lambda s_t(\mathbf{x}))}, \qquad (5)$$

where λ can be interpreted as the degree of bounded rationality, which is defined as the lack of propensity to maximize. An agent is more boundedly rational at lower λ values. At $\lambda = 0$, the supplier randomly chooses his decision and disregards the scores; he places equal probability on all possible decisions. As $\lambda \to \infty$, he picks the decision with the highest score with probability 1.

6.2. Comparing the EWA Model with Experimental Results

We now compare the EWA model with our experimental results. As such, we first estimate the EWA parameters for our experiments. We then simulate the EWA model with these estimated parameters and compare it with our experiments.

For each subject, we estimate ρ, λ, and r using maximum likelihood estimation. The likelihood function is the product of the probabilities for each realized decision of a subject over all periods. We minimize log-likelihoods for each subject numerically. We note that two of our subjects, Subject 17 in the two-price treatment and Subject 7 in the three-price treatment, consistently set their highest price break to a value higher than the maximum demand,

Table 4 EWA Coefficients for the Suppliers in One-Price, Two-Price, and Three-Price Treatments and Significance Levels of These Coefficients Compared to Their Rational Values (1,000 for λ, 0 for r, and 1 for ρ)

	One-price			Two-price			Three-price		
Player	λ	r	ρ	λ	r	ρ	λ	r	ρ
1	3.06	0.83	0.91	5.29	0.82	0.87	4.39	0.62	**1**
2	1.76	0.57	0.51	4.48	0.75	0.83	3.66	0.60	**0.90**
3	2.09	0.56	0.64	4.64	0.82	0.78	7.79	0.81	**1**
4	5.57	0.25	0.86	4.93	0.75	0.83	5.36	0.88	0.60
5	1.73	0.56	0.63	4.75	0.78	0.84	4.16	0.96	0.21
6	2.87	0.40	**1**	6.24	0.84	0.90	4.34	0.77	0.53
7	1.90	0.98	0.86	5.68	0.82	0.90	N/A	N/A	N/A[a]
8	3.76	0.40	0.93*	4.66	0.76	0.81	5.02	0.83	0.57
9	8.60	0.11	0.78**	5.94	0.80	0.89	7.09	0.90	0.89
10	1.84	0.72	0.82	4.48	0.77	0.78	2.80	0.81	0.30
11	4.18	0.25	0.71	4.12	0.72	0.78	5.98	0.93	0.55
12	23.68	0.05	0.75*	5.45	0.81	0.84	5.90	0.88	0.65
13	2.40	0.63	0.81	4.77	0.78	0.81	3.53	0.91	0.11
14	2.76	0.56	**0.94**	4.82	0.78	0.84	4.98	0.86	0.62
15	2.60	0.39	**0.67**	4.82	0.81	0.78	4.97	0.93	0.40
16	2.14	0.81	0.93**	4.41	0.77	0.81	3.8	0.75	0.67
17	2.15	0.74	0.85	N/A	N/A	N/A[a]	5.15	0.82	0.75
18	4.01	0.97	**0.97**	4.91	0.80	0.85	6.96	0.87	0.80
19	4.11	0.59	0.95*	4.99	0.78	0.84	6.19	0.75	0.76
Average	4.27	0.55	0.82	4.97	0.79	0.83	5.12	0.83	0.63

[a]Subject 17 in the two-price treatment and Subject 7 in the three-price treatment set the highest price break to a value higher than the maximum demand, reducing the contracts that they offered to one-price and two-price contracts, respectively. Therefore, we dropped these subjects from our estimations.

*$p = 0.05$; **$p = 0.01$. All nonsignificant values are bold. All other values are significant at $p = 0.001$.

reducing the contracts that they offered to one-price and two-price, respectively. We eliminate these 2 subjects (of 57) from our estimations. In addition, there are a few instances in the three-price treatment where a subject set only a few of his price breaks higher than the maximum demand in earlier periods. In those instances, we truncate the price breaks to the maximum demand. Using the average values of estimations, we then simulate the EWA model using 100 simulation runs to describe (1) the relationship between contract complexity and performance and (2) deviations from theory under the EWA model. For these simulations we rely on discretization of the decision space under a three-price contract. Without discretization, our simulations run beyond the maximum memory that can be utilized at once in MATLAB. Thus, we sample pricing decisions with a step size of 3 (i.e., we only consider wholesale prices such that $w_i = 40 + 3s$, where s is an integer, and $w_i \leq 200$).

We observe that subjects exhibit the decision biases of the EWA model. Table 4 provides the individual estimates of the bias parameters ρ, λ, and r. As can be seen, for all of our subjects, λ and r are highly significantly different from their rational values (0 for r and estimated as 1,000 for λ). The memory parameter ρ is significantly lower than its rational value (which is 1) for 48 of 55 subjects. A few observations are worth noting. First, the reinforcement parameters are very high. Even in the one-price treatment, subjects are reinforcing their earlier decisions and discounting the profits of untried options by more than 50%. As the complexity of the contract increases, this behavior grows even stronger. In the three-price treatment, discounting reaches 83% on average. Second, our subjects exhibit bounded rationality because the parameter λ is significantly less than its value for a theoretical rational player. We should note that our subjects are not completely irrational, because the parameter λ is also significantly greater than 0. Finally, our subjects also exhibit significant memory bias. As noted earlier, interpreting the memory parameter requires more care. An average value of 82% does not suggest that there is no memory effect; with $\rho = 0.82$, the current period's information is discounted by about 50% in fewer than four periods.

When we compare the parameters across treatments, we find support for our main thesis regarding the nontrivial trade-off between complexity and inefficiency of all-unit quantity discount contracts; i.e., EWA describes the relationship between contract complexity and performance. Table 5 shows that as the complexity of the contract increases, subjects' reinforcement parameters increase and their memory parameters decrease; that is, *subjects rely more on simple heuristics by reinforcing more*. Probabilistic choice/bounded rationality parameters increase as subjects move from a one-price contract to quantity

Table 5 Difference Between Bias Parameters Under Different Contracts

	Probabilistic choice/ bounded rationality		Reinforcement		Memory	
	Alternative	p-value	Alternative	p-value	Alternative	p-value
One-price vs. two-price	Lower	<0.001	Lower	<0.001	Not significant	
Two-price vs. three-price	Not significant		Lower	0.02	Greater	0.003
One-price vs. three-price	Lower	0.001	Lower	<0.001	Greater	0.008

Notes. We applied the Wilcoxon test for comparison across treatments. We report the results for significant parameters only.

discount contracts, without any significant difference between two-price and three-price contracts. This is likely a response to a much higher level of reinforcement in quantity discount contracts.

Finally, we observe that the EWA model is also powerful in describing deviations from theoretical predictions (Figure 2). When we fit the EWA model and simulate suppliers' decisions, we observe that the EWA model captures the deviations of *all* decisions in the right direction. More importantly, the model describes the decision that has the highest impact on supplier profits precisely. In a two-price (three-price) contract, correctly setting Q_1 (Q_2) leads to the highest improvements in the supplier's profits and the EWA model fits Q_1 (Q_2) effectively (Figure 2, (g) and (h)). Having said that, we note a mismatch between experiments and the EWA model in Figure 2(d). This mismatch hints that there may be additional behavioral drivers affecting subjects' decisions.

Given EWA's success of describing our subjects' deviations from theory, one expects that the fitted EWA model would perform similarly to our subjects. That is indeed the case. The average profits estimated under the EWA model are very close to those in our experiments; a Wilcoxon test comparison of these average profits shows that they are not significantly different from each other under all contract types (two-sided p-values are 0.62, 0.96, and 0.39 for the comparisons under the EWA model and the experiments for one-price, two-price, and three-price contracts, respectively). We also observe that the relative performance of contracts observed in experiments is preserved with the EWA model: a Wilcoxon test comparison of the average profits of each simulation under different contracts shows that a two-price contract leads to significantly higher profits than a one-price contract ($p < 0.001$); however, a three-price contract does not have any additional benefit over a two-price contract (two-sided $p = 0.69$).

7. Conclusions and Managerial Insights

This paper experimentally studies the trade-off between contract complexity and economic inefficiencies under a setting with asymmetric demand information. We test the key theoretical predictions that increasing the complexity of contracts employed increases suppliers' profits and supply chain efficiency and decreases buyers' profits. Experimental results contradict these predictions. Increasing the complexity of the contract does not always increase suppliers' profits. Suppliers benefit from quantity discounts, but even a single price break (two-price contract) is enough to capture these benefits. More complex contracts do not lead to better channel coordination either; we do not observe any statistical difference between total supply chain profits under the price-only contract and more complex quantity discount contracts. We therefore conclude that there is a nontrivial trade-off between complexity and inefficiency of all-unit quantity discount contracts: the notion that complex contracts can optimize the supplier's profit in a supply chain is flawed and requires deeper theoretical consideration.

It is crucial to consider contract complexity as a design factor. Our estimations using the EWA model 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 may not be the best). This finding is consistent with theory established by Simon's (1955) 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 2008, p. 2).

Our experiments suggest that subjects find quantity break decisions particularly challenging. In a setting with asymmetric information, quantity breaks help the supplier extract the buyer's private information. Given subjects' inability to utilize quantity breaks effectively, one might conjecture that a screening-based contract design is not ideal for the supplier, especially when the separation of types is important (i.e., when the margins are low and type dispersion is high). Under such conditions, a sequential negotiation process, which is prevalent in practice, may lead to better contracts for the supplier. The behavioral academic literature supports this conjecture in two ways. First, Özer et al. (2011) 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 (2007) argue that subjects can use quantity discount mechanisms quite effectively under complete information. Therefore, an opportunity to collect more information about the buyer's demand type should benefit the supplier. The sequential negotiation process provides him this opportunity. We also note that even when the level of information asymmetry between the supplier and the buyer 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.

This study is based on a between-subject experimental design. We chose the between-subject design to eliminate order effects. However, because of individual heterogeneity, this design results in higher treatment variability. Parallel experiments (available from the authors upon request) with an alternate within-subject design where each subject was faced with three different treatments, provide the same general conclusions with respect to profit comparisons. Another potential experimental limitation is our reliance on student subjects. However, the decision biases we observe with student subjects have similarities with the flaws in managerial decision making identified in the literature, such as the status quo bias and the recency bias (Roxburgh 2003, McKinsey Global Survey 2009). Therefore, our conclusions are robust with respect to these experimental design choices.

There are several directions of potential future research. This paper provides strong empirical evidence refuting the predictions of standard expected-payoff-maximization-based contract theory. In particular, with an EWA model, we show that reinforcement and bounded rationality are key biases that impact human subjects' decisions. It would be interesting to incorporate the issue of complexity-dependent bounded rationality into maximization-based theory. Second, we note that the rich environment of our experiments hints that there may be other behavioral drivers that affect subjects' decisions. Although we have explored and ruled out loss aversion and risk aversion with log utility function as possible explanations, it would be beneficial to test some additional alternative behavioral explanations, as well as the mechanics behind the EWA model for future research. For example, Bolton and Katok (2008) show that in a newsvendor setting, subjects do not take foregone payoffs into account. The high reinforcement parameters that we observe in our experiments imply that the same may be true in our setting. Thus, extending such observations to our setting in a systematic way would be valuable. However, it is highly likely that any behavioral explanations in our setting will be a combination of several behavioral factors. We have observed significant probabilistic choice/bounded rationality parameters in the EWA model, which implies that a modification of the utility function cannot fully explain the observations by itself. Third, in our experiments, we eliminated the effect of reputation by studying human–computer interaction. However, understanding the role of reputation in contract design would also be an interesting future research direction. Especially for products with a limited number of suppliers, buyers and suppliers interact repeatedly over time and may develop a better understanding of each other's strategies. Finally, we would like to acknowledge that because it is difficult to conduct field experiments to compare different contracts under the same environment, we turned to experimentation in a laboratory setting in this paper. Our experiments are a necessary first step, but they of course suffer from lack of external validation. Therefore, a field study on contract complexity versus inefficiency would be a welcome addition to the literature.

Our study shows that the price-only contract is very effective under experimental settings. Even though there is strong theoretical support for quantity discounts, managers should be cautious when employing these contracts because the additional complexity may be an unwelcome burden during decision making. However, by limiting the number of price blocks they offer, managers may achieve the perfect balance between efficiency and complexity. A decision support tool that helps them especially in determining how to segment their buyers by using quantity breaks effectively would be useful.

8. 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

This research was funded in part by National Science Foundation Grants CAREER-0547021 and DMI-0400345. The authors gratefully acknowledge partial support from HP Laboratories.

References

Adesso Solutions. 2005. A double-edged sword. *Bus. Credit* **107**(9) 56–59.

Altintas, N., F. Erhun, S. Tayur. 2008. Quantity discounts under demand uncertainty. *Management Sci.* **54**(4) 777–792.

Armour, P., M. Daly. 2008. Retirement savings and decision errors: Lessons from behavioral economics. *FRBSF Econom. Lett.* (June 6) 1–3.

Benton, W. C., S. Park. 1996. A classification of literature on determining the lot size under quantity discounts. *Eur. J. Oper. Res.* **92**(2) 219–238.

Bolton, G. E., E. Katok. 2008. Learning-by-doing in the newsvendor problem: A laboratory investigation of the role of experience and feedback. *Manufacturing Service Oper. Management* **10**(3) 519–538.

Bolton, P., M. Dewatripont. 2005. *Contract Theory*. MIT Press, Cambridge, MA.

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

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

Camerer, C. 2003. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press, Princeton, NJ.

Camerer, C., T.-H. Ho. 1999. Experience-weighted attraction learning in normal-form games. *Econometrica* **67**(4) 827–874.

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

Chu, L. Y., D. E. M. Sappington. 2007. A note on optimal procurement contracts with limited direct cost inflation. *J. Econom. Theory* **137**(1) 745–753.

Corbett, C. J., X. de Groote. 2000. A supplier's optimal quantity discount policy under asymmetric information. *Management Sci.* **46**(3) 444–500.

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

de Palma, A., G. M. Myers, Y. Y. Papageorgiou. 1994. Rational choice under an imperfect ability to choose. *Amer. Econom. Rev.* **84**(3) 419–440.

DeShazo, J. R., G. Fermo. 2002. Designing choice sets for stated preference methods: The effects of complexity on choice consistency. *J. Environ. Econom. Management* **44**(1) 123–143.

Freund, R. J., W. J. Wilson, P. Sa. 2006. *Regression Analysis*, 2nd ed. Academic Press, Oxford, UK.

Heiner, R. A. 1983. The origin of predictable behavior. *Amer. Econom. Rev.* **73**(4) 560–595.

Ho, T.-H., J.-J. Zhang. 2008. Designing price contracts for boundedly rational customers: Does the framing of the fixed fee matter? *Management Sci.* **54**(4) 686–700.

Hogarth, R. M., H. J. Einhorn. 1992. Order effects in belief updating: The belief-adjustment model. *Cognitive Psych.* **24**(1) 1–55.

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

Jeuland, A. P., S. M. Shugan. 1983. Managing channel profits. *Marketing Sci.* **2**(3) 239–272.

Jiang, H., S. Netessine, S. Savin. 2011. Technical note: Robust newsvendor competition under asymmetric information. *Oper. Res.* **59**(1) 254–261.

Jucker, J. V., M. J. Rosenblatt. 1985. Single-period inventory models with demand uncertainty and quantity discounts: Behavioral implications and a new solution procedure. *Naval Res. Logistics* **32**(4) 537–550.

Katok, E., D. Y. Wu. 2009. Contracting in supply chains: A laboratory investigation. *Management Sci.* **55**(12) 1953–1968.

Keser, C., G. Paleologo. 2004. Experimental investigation of retailer-supplier contracts: The whole-sale price contract. Working paper, IBM T. J. Watson Research Center, Yorktown Heights, NY.

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

Lee, H. L., V. Padmanabhan, S. Whang. 1997. Information distortion in a supply chain: The bullwhip effect. *Management Sci.* **43**(4) 546–558.

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

McKelvey, R. D., T. R. Palfrey. 1995. Quantal response equilibria for normal form games. *Games Econom. Behav.* **10**(1) 6–38.

McKinsey Global Survey. 2009. Flaws in strategic decision making. *McKinsey Quart.* (January) 1–6.

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

Niederhoff, J., P. Kouvelis. 2008. Supplier and retailer behavior and the effectiveness of coordinating contracts. Working paper, Syracuse University, Syracuse, NY.

Özer, Ö., Y. Zheng, K.-Y. Chen. 2011. Trust in forecast information sharing. *Management Sci.* Forthcoming.

Roxburgh, C. 2003. Hidden flaws in strategy. *McKinsey Quart.* (May) 26–39.

Samuelson, W., R. Zeckhauser. 1988. Status quo bias in decision making. *J. Risk Uncertainty* **1**(1) 7–59.

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

Simon, H. A. 1955. Behavioral model of rational choice. *Quart. J. Econom.* **69**(1) 99–118.

Simon, H. A. 1978. Rational decision-making in business organizations. A. Lindbeck, ed. *Nobel Lectures, Economics 1969-1980.* World Scientific Publishing, Singapore, 343–371.

Su, X. 2008. Bounded rationality in newsvendor models. *Manufacturing Service Oper. Management* **10**(4) 566–589.

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

Wilson, R. B. 1993. *Nonlinear Pricing*. Oxford University Press, New York.

Zadeh, L. 1973. Outline of a new approach to the analysis of complex systems and decision processes. *IEEE Trans. Systems, Man and Cybernetics* **3**(1) 28–44.