Can Coordinating Contracts Improve Channel Efficiency?

Ernan Haruvy
School of Business, University of Texas and Dallas, Richardson, TX 75080
eharuvy@utdallas.edu

Elena Katok
Smeal College of Business, Penn State University, University Park, PA 16802
ekatok@psu.edu

Valery Pavlov
Smeal College of Business, Penn State University, University Park, PA 16802
vpavlov@psu.edu

A growing literature shows that coordinating contracts may not improve efficiency in the laboratory to the extent prescribed by theory. We show that this result is largely due to offer rejection where the bargaining procedure is structured as an ultimatum offer, which is the only structure studied in the lab hitherto in relation to coordinating contracts. We show that a less restrictive procedure does not involve this feature and allows coordinating contracts to coordinate. Specifically, we look at three contract formats—wholesale price, two-part-tariff and minimum order quantity. The wholesale price leads to loss of firm surplus because of double marginalization. The other two contracts—the coordinating contracts-- allow the manufacturer to coordinate the channel, either by pricing at cost and extracting surplus through a lump sum payment (two part tariff) or through announcing a minimum order quantity which is equal to the efficient quantity and extracting the surplus through price (minimum order quantity contract). Even though for fully rational players these two coordinating contracts are equivalent, these two contracts are different under mild bounded rationality assumptions. Proposals in the minimum order quantity treatment are far more efficient than two-part-tariff proposals in terms of the overall surplus they imply. But in the ultimatum context such efficient proposals tend to get rejected, leading to lower ex-post efficiency. With *structured negotiations* bargaining, however, rejection rate drastically falls leading to a more direct relationship between proposal efficiency and ex-post efficiency.

Keywords: Supply Chain Contracts, Experimental Economics, Bargaining

1. Introduction

The most basic channel structure involves a manufacturer and a retailer. The manufacturer sets a wholesale price for the retailer and the retailer sets the price for consumers. This sequence results in efficiency loss, known as double marginalization, relative to the outcome of a single decision maker or an integrated channel. The double marginalization problem has a number of solutions (e.g., Tirole 1988) that typically involve the manufacturer employing a mechanism (e.g., two-part tariff) that induces the retailer to charge the retail price of the integrated channel (e.g., Jeuland and Shugan 1983, Moorthy 1987). There are many alternative mechanisms that are theoretically equivalent in equilibrium. However, these mechanisms may result in different behavioral responses and they are not necessarily equivalent off-equilibrium[1].

The laboratory can be useful in identifying both behavioral implications of different mechanisms and off-equilibrium implications that may be driven by non-pecuniary utility considerations, errors, or the expectation of these elements.

Laboratory tests typically use the simplest possible implementation of bargaining—a take-it-or-leave-it ultimatum offer made by the supplier to the retailer. This is an extreme setting in a sense that it maximizes supplier's bargaining power, and offers a simple prediction that the supplier should be able to not only coordinate the channel, but also to extract the entire channel profit.

Laboratory studies using this implementation report two sets of conclusions. First, coordinating contracts fail to coordinate, and this failure is primarily due to retailer rejections.

[1] For example, we will look at a two-part tariff versus minimum order quantity. In the latter, errors by the retailer are left-truncated.

Suppliers offer contracts that are quite efficient but a large proportion of them are rejected by retailers (rejections range from 11% (Lim and Ho, 2007) to 26% (Ho and Zhang, 2008)). So the resulting efficiency is often not significantly higher than the efficiency of the wholesale price contracts (which are rarely rejected) and is always significantly below 100%. Second, coordinating contracts that are equivalent mathematically should result in similar channel performance, but they do not. For example, Lim and Ho (2007) note that a two-part tariff contract is far less efficient than a three-part tariff though the two are theoretically equivalent. They propose a reference-dependent utility under which each price is compared to the adjacent price. Ho and Zhang (2008) note that two part tariff is less efficient than a quantity discount contract, even though the two contracts differ only in framing. They propose that loss aversion is more salient in the two-part-tariff contract than in the quantity discount contract.

Ours is the first study that investigates the effect of the bargaining implementation on supply chain contract performance. We compare coordinating contracts that use the *ultimatum bargaining* (UB) protocol to another protocol under which the supplier's and the retailer's bargaining power is more equal, that we call *structured bargaining* (SB). In this SB protocol the players have a fixed amount of time in which supplier can propose contracts and retailer can react to these offers. The central feature of the SB protocol is that if a retailer rejects a proposal, the supplier can make a new offer that is either the same or better in terms of the retailer's absolute profit were the retailer to place an profit-maximizing offer. This feature increases the retailer's bargaining power relative to the UB protocol because if a supplier makes an offer, this offer establishes a minimum retailer profit amount—the retailer cannot do any worse than the current offer on the table. If the agreement is not reached within a fixed amount of time, however, the game reverts to the UB setting in which the supplier can make one last and final offer that is at least as good for the retailer as the last offer on the table.

In this study, we investigate the effect of bargaining protocols on contract performance. We investigate two coordinating contracts that have both been examined in the literature: the two-part-tariff (TPT), and the minimum order quantity (MOQ). We compare them to the wholesale-price contract. We chose the TPT and the MOQ contracts because those two contracts induce coordination in very different ways. The TPT contract coordinates by setting the wholesale price to production cost, thus making the retailer's marginal cost the same as the supplier's. It is then up to the retailer to coordinate channel by selecting the order quantity that maximizes his own profit. In contrast, under the MOQ contract, the supplier is responsible for coordination—he coordinates the channel by setting the minimum order quantity to first best. If the retailer accepts supplier's offer, the channel is coordinated. Thus, the TPT and the MOQ contracts represent two extreme coordinating mechanisms in terms of which firm's decisions (the supplier's or the retailer's) are responsible for coordination.

Our main finding is that the structured bargaining protocol significantly improves the efficiency of coordinating contracts relative to the wholesale price contract. In fact, towards the end of the session, channel efficiency of those contracts is nearly 100%. Since the retailer's bargaining power is higher under the structured bargaining protocol than under ultimatum bargaining protocol, the retailer's share under the structured bargaining protocol is higher also. The two coordinating contracts result in different outcomes. Under the ultimatum bargaining protocol, the offers are more efficient under MOQ than under TPT, but rejections are significantly higher under MOQ as well, so the resulting channel efficiency is higher under TPT. Since offers become more efficient and rejections are nearly eliminated under structured bargaining, these differences disappear, although retailers earn a slightly higher share under MOQ.

In the next section we describe the basic analytical model and experimental hypotheses. In section 3 we present the details of our experimental design and protocol. We report our results in section 4, and conclude in section 5 by summarizing our conclusions, discussing limitations, and proposing managerial implications and directions for future research.

2. Theory

2.1. The Setting and Contracts

The setting involves a simple distribution channel that consists of a single supplier and a single retailer. The supplier has a constant marginal cost of production c. The supplier proposes a contract to the retailer. Following the bargaining process, the retailer either rejects the offer, in which case both parties earn zero, or places an order for q units. We assume that the retailer faces a linear demand $q = A - p$ where p is the retail prices and A is a constant. Since the retailer faces deterministic demand and the product has no salvage value, we assume that the retailer's order will match the amount sold, given the retail price.

We say that the channel is integrated if the outcome in terms of units produced is the same as the outcome that would have resulted from a single decision maker maximizing the entire channel profit. The channel profit to be maximized under integrated channel is

$$\pi^C = (p - c)q = ((A - q) - c)q \,. \qquad (1)$$

The order quantity that maximizes this channel profit is $q^* = (A - c)/2$, yielding the optimal (first-best) channel profit of $(A - c)^2 / 4$.

If the channel is not integrated—the two firms optimize separately and independently—and they use a wholesale-price contract in which the retailer pays the supplier w per unit, the retailer optimizes his own profit by ordering $q^*_{WP} = (A - w)/2$ which is lower than the first-best order quantity q^* whenever $w > c$. The supplier must set the wholesale price so as to maximize his profit

$$\pi^S = (w - c)q^*_{WP} = (w - c)(A - w)/4, \tag{2}$$

resulting in the profit-maximizing wholesale price $w_{WP} = (A + c)/2$. This optimal wholesale-price contract results in the supplier profit of $\pi^S_{WP} = (A - c)^2/8$, the retailer's profit $\pi^R_{WP} = (A - c)^2/16$, and the total channel profit $\pi^C_{WP} = 3(A - c)^2/16$, representing the efficiency of only 75% relative to the first-best channel profit π^C. This inefficiency of the wholesale price contract relative to the integrated system is known as *double marginalization.*

A variety of different contracts can solve the double marginalization problem. There are fundamentally two different ways in which a coordinating contract can be constructed. One coordinating contract, known as *two-part-tariff* (TPT), is for the supplier to charge the retailer a lump sum as well as a per unit price. Another coordinating contract we will examine is known as *minimum order quantity* (MOQ), under which the supplier sets a minimum order requirement as well as a per-unit price.

For the two-part tariff to result in an efficient outcome, the retailer's marginal cost needs to be the same as the supplier's marginal cost, w=c, thus inducing the retailer to place the first-best order. The profit is allocated via a lump-sum transfer amount F from the retailer to the supplier.[2] If the supplier wants to extract the entire channel profit, he sets the lump sum transfer

[2] The two-block tariff (Lim and Ho 2007) also coordinate by setting $w = c$ in the last block. The profit is allocated by setting a higher wholesale price for orders below a certain break point x.

amount to $F = (A - w)^2 / 4$. Another mechanism for channel coordination is to let the supplier control the order quantity in a more direct way, by setting the minimum order quantity q_{min} and charging a linear wholesale price w per unit, as long as the retailer places an order of at least q_{min} (see Pavlov and Katok (2009) for the analysis of this minimum-order-quantity (MOQ) contract). The profit is allocated by setting a high w. In equilibrium, the supplier sets $q_{\min} = q^* = (A - c)/2$, and if he wishes to extract the entire channel surplus, he sets $w_{MOQ} = (A + c)/2$ (note that $w_{MOQ} = w_{WP}$). Table 1 summarizes optimal parameters for the integrated channel and the contracting mechanisms we are examining in this paper.

Table 1. Properties of the three contracts and the integrated channel

Contract	Total Channel Profit	Optimal w* set by supplier	Lump sum / Minimum order quantity	Optimal q^* by retailer
Integrated Channel	$(A-c)^2/4$	$w=c$	N/A	$(A-c)/2$
Wholesale Price	$3(A-c)^2/16$	$w_{WP} = (A+c)/2$	N/A	$(A-c)/4$
Two Part Tariff	$(A-c)^2/4$	$w=c$	$F = (A-w)^2/4$	$(A-c)/2$
Minimum Order Quantity	$(A-c)^2/4$	$w_{MOQ} = (A+c)/2$	$q_{\min} = (A-c)/2$	$(A-c)/2$

Before we proceed with articulating our behavioral model we note that the standard model delivers sharp theoretical predictions about the relative channel efficiency of the three contracts we are investigating. We will state this prediction as a hypothesis in it's weak form, since there has already been plenty of evidence that coordinating contracts do not result in 100% channel efficiency.

HYPOTHESIS 1a. (EFFICIENCY: WEAK VERSION). TPT and MOQ contracts result in higher average channel efficiency than does the wholesale price contract.

$$\pi^R_{TPT} + \pi^S_{TPT} > \pi^R_{WP} + \pi^S_{WP}$$
$$\pi^R_{MOQ} + \pi^S_{MOQ} > \pi^R_{WP} + \pi^S_{WP}.$$

HYPOTHESIS 1b. (EFFICIENCY: STRONG VERSION). TPT and MOQ contracts result in 100% channel efficiency.

2.2. Fairness and Errors Considerations

Previous laboratory studies of channel coordination report that rejections by the retailer are a major source of inefficiency for coordinating contracts. For example, Ho and Zhang (2008) report 26% rejections in TPT and 18% rejections in another coordinating contract known as quantity discounts (QD). Lim and Ho (2007) report 11% rejections in two-block tariff and 15% rejections in a three-block tariff.

Rejections result in loss of profit for the retailer so they can be explained either by retailer errors, or by inequity aversion, or most likely by the combination of the two. A model of errors is invariably a part of any econometric model, and it remains to model to inequity aversion. Note that inequity aversion is intended to capture a non-pecuniary preference pertaining especially to the lump sum. There are certainly ways other than inequity aversion to capture non-pecuniary preferences with respect to the lump sum. Ho and Zhang (2008) propose an approach of placing a separate coefficient on the lump-sum transfer that they label loss aversion. This coefficient captures the higher rejections in TPT in their experiment relative to QD. Lim and Ho (2007) frame the two-block tariff as consisting of three decisions by the supplier: An initial quantity, a per unit price for that initial quantity, and a per unit price for subsequent units. Using that framing, they propose a reference-dependent utility loss from the difference in the two prices. Note that in both these proposed models, there is an implied non-pecuniary weight that primarily applies to the lump sum. This weight has a strong parallel in models of inequity aversion since the lump-sum is also the main variable responsible of inequity.

A common approach to modeling the retailer's inequity aversion, which we pursue here, is a variation of a model proposed by Fehr and Schmidt (1999). Assume that the supplier made a proposal that if accepted, will yield π^S to the supplier and π^R to the retailer, if the retailer orders q. If $\pi^S \geq \pi^R$ (the supplier makes the proposal that gives himself at least 50% of the channel profit, as is the case in ours, as well as virtually all other experiments with the ultimatum-like structure), then the retailer's utility from accepting and rejecting is:

$$\begin{aligned} U^R(Accept, q) &= \pi^R - \beta\left(\pi^S - \pi^R\right)^+ \\ U^R(Reject) &= 0 \end{aligned} \tag{3}$$

where β is the retailer's inequality aversion parameter. It is reasonable to assume that the retailer knows his own β, but the supplier only know it's distribution but not it's realization.[3] If retailer does not make any errors, then as long as $U^R(Accept)$ is positive, the retailer would accept the supplier's offer, and rejects it otherwise. But if the retailer makes errors, then a common way to model the retailer's decision is to say that his *probability* of accepting the contract is a function of the relative utility of accepting and rejecting—higher utility offers have a higher probability of being accepted:

$$\Pr(Accept) = \frac{e^{\tau U(Accept, q)}}{1 + e^{\tau U(Accept, q)}}$$

where τ is the retailer's *coefficient of certitute*—a parameter that measures the extent to which the retailer makes errors. This parameter is itself assumed to be a random coefficient distributed normally among bidders with mean of $\bar{\tau}$ and standard deviation σ_τ. A reasonable assumption about information available about τ is that the supplier knows it's distribution but not it's realization.

[3] Cui et al (2007) propose a full information model, in which the supplier does know the realization of β, and Pavlov and Katok (2009) show that the Cui et al. (2007) conclusions also generalize to the incomplete information case.

In addition to the retailer's decision to accept or reject the proposed contract, the utility specification is important in affecting the retailer's decision on quantity to order. With the specification of utility in (3), we get can derive the best reply expression for the order quantity q, conditional on accepting the contract under the ultimatum bargaining protocol (the complete derivation is given in the appendix).

$$q^*_{TPT} = \frac{A+(A+c)\beta-(1+2\beta)w}{2(1+\beta)}B+\varepsilon \tag{4}$$

$$q^*_{MOQ} = Max\left\{\frac{A+(A+c)\beta-(1+2\beta)w}{2(1+\beta)}B+\varepsilon, q_{\min}\right\}$$

From the above derivation, we see that q* is independent of β whenever w=c, decreasing in β for w>c and increasing in β for w<c for any general linear demand as specified by $q = A - p$.

The variable ε allows for the possibility that the retailer makes errors in setting the order quantity; the standard assumption is that $\varepsilon \sim N(0, \sigma_\varepsilon)$. There are two things to note with respect to (4). First, it only applies when the offer gives a greater share to the supplier than the retailer, which happens in nearly all cases (100% of all retailers who ever rejected an offer). An offer that gives more than 50% of the channel profit to the retailer would result in the standard best reply order (100-*w*)/2, assuming that the retailer is not averse to advantageous inequality. Second, under the TPT contract, q^* does not depend on the lump sum *F*. This is because *F* does not interact with quantity ordered so the retailer's first order condition with respect to quantity does not include *F*. Of course the probability of accepting the offer is decreasing in *F*.

The two decisions (q and *Accept*) are not necessarily independent. Let us allow for a covariance term between τ and ε.

$$\begin{pmatrix} \tau \\ \varepsilon \end{pmatrix} \sim N\left\{\begin{pmatrix} 0 \\ 0 \end{pmatrix},\begin{pmatrix} \sigma_\tau & \rho_{\tau\varepsilon}\sigma_\tau\sigma_\varepsilon \\ \rho_{\tau\varepsilon}\sigma_\tau\sigma_\varepsilon & \sigma_\varepsilon \end{pmatrix}\right\} \tag{5}$$

Thus, we can construct a likelihood function as follows:

$$\begin{aligned} L = \prod &\left[\varphi((q - f(\beta)) / \sigma_\tau)\Phi^{BVN}(U^R / \sigma_\varepsilon \,|\, q, \rho_{\tau\varepsilon})\right]^{Accept} \\ &\times\left[1 - \Phi^{BVN}(f(\beta) / \sigma_\tau, U^R / \sigma_\varepsilon, \rho_{\tau\varepsilon})\right]^{(1-Accept)} \end{aligned} \tag{6}$$

Where $f(\beta)$ denotes the non-error term on the right-hand side of equation (4), ϕ and Φ denote the standard normal probability distribution function and cumulative distribution function, respectively, and Φ^{BVN} denotes the bivariate standard normal cumulative distribution function.

Note that when $\rho_{\tau\varepsilon}$ is equal to zero (i.e., the errors between the decision to accept and the quantity decision are uncorrelated), the likelihood is simply the product of a simple tobit on the quantity decision and a simple logit on the accept-reject decision. Both are standard approaches to estimation in the literature and since each forms a valid likelihood function, the product of these two likelihood functions is a valid likelihood function as well.

We now move on the characterizing the supplier's decision on contract parameters, when the supplier accounts for the retailer's potentially being averse to disadvantageous inequality, and the retailer's potential for making errors. Pavlov and Katok (2009) characterized the generic optimal contract when both players have inequality averse preferences and show that the optimal contract is (1) conditionally efficient, and (2) pooling. In other words, the supplier should offer a contract that, if accepted, will coordinate the channel, but will be rejected by some too-fair-minded retailers. Conceptually, the supplier's decision is how to best answer two questions: how to coordinate the channel? And how much channel profit to attempt to extract? Let us focus on the first question, of how to coordinate the channel, using either the TPT or the MOQ contract, assuming that the retailer's decision is characterized by expressions (3), (4) and (5) (the retailer is averse to disadvantageous inequality, and makes errors).

To coordinate the channel under the MOQ contract, the supplier has to set $q_{\min} = (A - c)/2$. Regardless of how the supplier sets w, as long as $w > c$, and independently of β, the retailer should set $q^*_{MOQ} = q_{\min}$. The supplier does not face any strategic uncertainty in regards to the retailer's decision ($q^*_{MOQ} = q_{\min}$ is independent of β) and his decision is also not affected by potential retailer's errors (since $E[\varepsilon] = 0$), so to the extent the supplier might deviate from making a conditionally efficient offer, this deviation has to be entirely due to the supplier's errors. A key observation, however, is that if the supplier sets $q_{\min} = (A - c)/2 + \varepsilon_{moq}$, where $\varepsilon_{moq} \sim N(0,\sigma_{moq})$, for reasonably small σ_{moq}, the retailer's best reply would continue to be independent of β.

To coordinate the channel under the TPT contract, the supplier must set $w = c$. It turns out that when $w = c$, (4) becomes $q^*_{TPT} = \dfrac{A-c}{2}$, but if $w = c + \varepsilon_{TPT}$, q^*_{TPT} is no longer independent of β. Let $\varepsilon_{TPT} \sim N(0,\sigma_{TPT})$. As long as $\sigma_{TPT} \neq 0$ (supplier sometimes makes errors in setting the wholesale price under TPT) the retailer's order quantity is function of β, which is retailer's private information—the supplier faces strategic uncertainty in regards to the retailer's best reply. This observation that the MOQ and the TPT contracts differ in the amount of strategic uncertainty the supplier faces, lead to our next hypothesis that deals with conditional efficiency:

HYPOTHESIS 2. (CONDITIONAL EFFICIENCY). The conditional efficiency of supplier offers is lower under the TPT contract than under the MOQ contract. Conditional on being accepted:

$$\pi^R_{TPT} + \pi^S_{TPT} < \pi^R_{MOQ} + \pi^S_{MOQ} .$$

2.3. The Bargaining Format

The two bargaining formats we investigate are Ultimatum Bargaining (UB) and Structured Bargaining (SB). The UB protocol is an extreme case that gives the retailer the least possible amount of bargaining power, because the supplier makes a take-it-or-leave-it offer, and the retailer has no recourse except to reject it, causing both parties to earn zero. If the retailer is averse to disadvantageous inequality, his utility from accepting unfair offers decreases in β. If the supplier knew β, he would offer the retailer a portion of channel profit that the retailer would just barely accept—the supplier would be able to coordinate the channel and extract the maximum profit, giving the retailer only enough profit to induce him to participate. But since β is the retailer's private information, the supplier cannot do this. He has to offer a pooling contract, which will sometimes be rejected and on average, retailers will earn some information rents.

The SB protocol is equivalent to the UB protocol in equilibrium, because the supplier can simply not make any offers during the fixed bargaining period, after which the game reverts to the UB. However, off equilibrium path, retailer has more bargaining power under SB than under UB because if the supplier does make an offer, the offer is binding—the retailer can now earn at least this amount from any subsequent offer. Note that the SB mechanism increases the retailer's bargaining power in a very weak way, since the supplier can either not make any offers, or make offers that he knows to be below the lower support of the β's distribution, or make a single offer that is the same that he would have made under UB, and never improve it. The only thing the supplier is not allowed to do under SB that works to the retailer's advantage is to withdraw offers (in other words, he has to bargain in good faith).

To better understand the SB mechanism, let us consider two extreme examples. First, imagine that the retailer is non-strategic and simply uses (3) to respond to offers—accepting an offer as soon as his utility from accepting becomes higher than from rejecting. The supplier can start with a small offer, and gradually raise it until the retailer finally accepts—he can essentially get the retailer to reveal his β, and as a result channel efficiency would improve by lowering rejections, and the retailer will not earn his information rent. The retailer can avoid this outcome by strategically rejecting small offers even when they are acceptable, and waiting until the game reverts to the UG stage.

At another extreme, imagine that the supplier is acting as if the retailer is not strategic—he starts with small offers and gradually raises them hoping to find the minimum amount the retailer will accept—but the retailer is, in fact, strategic, and rejects all offers. At some point, the supplier will get to a 50/50 split, which should be acceptable to all retailers with utility function specified by (3). The above discussion makes it clear that under the SB mechanism, the retailer has a dominant strategy to reject all offers prior to the final stage. The supplier, recognizing that this is the case, should not make any offers until the last stage. Formally:

HYPOTHESIS 3a. (DOMINANT STRATEGY UNDER STRUCTURED BARGAINING). No offers should be accepted prior to the final stage. No offers should be made prior to the final stage. SB outcomes should be the same as UB outcomes.

But if retailer sometimes accepts offers prior to the final stage, either because he is not acting strategically, or because of errors, the supplier might benefit from making early offers because he may learn private information about β. Alternatively, if the supplier makes early offers because he is not strategic (or he makes errors), the resulting split of channel profit may be fairer to the retailers. In both cases, rejections should decrease and efficiency should increase.

HYPOTHESIS 3b. (IMPROVED EFFICIENCY UNDER STRUCTURED BARGAINING). If supplier makes offers in the first stage of the SB, channel efficiency should be higher under SB than under UB.

Hypotheses 3a and 3b contradict one another, and a controlled laboratory experiment will distinguish between them.

This likelihood representation permits us to form a conceptual model in Figure 1:

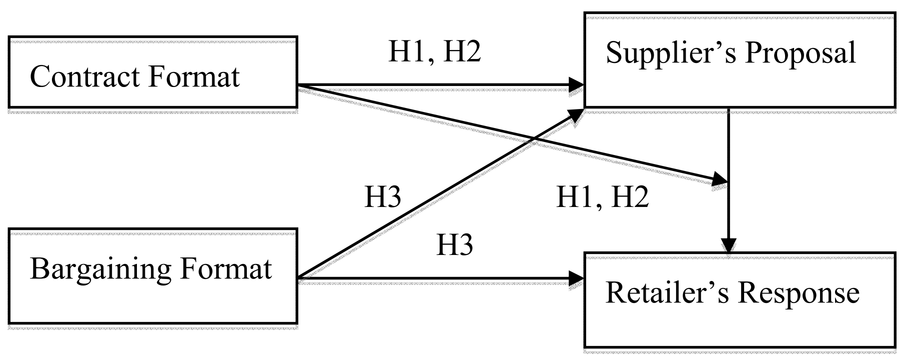

Figure 1. Conceptual model of the research hypotheses. H1, H2 and H3 correspond to hypotheses 1, 2, and 3 respectively.

Hypotheses 1 and 2 deal with the effect of the contract format that primarily affects the supplier's proposal. The retailer's response is affected only through the supplier's proposal. Hypothesis 3 deals with the effect of the bargaining format. It affects both, the supplier's proposal and the retailer's response.

Note that a key difference between this model and previously proposed models (Lim and Ho, 2007; Ho and Zhang, 2008) is that in these other models, the difference between protocols arises from the retailer perceiving the problems differently. In other words, the same proposal under different protocols would receive a treatment response from the retailer. In the present framework, we certainly test the possibility that the retailer perceives the mechanisms and protocols differently, but we also permit for the possibility that the retailer only does this through reacting to the supplier's proposal. In that case, the supplier would be responsible for any differences in the contract and we indeed find this to be the case.

3. Experimental Implementation

3.1 Experimental Design

Contracting arrangements in our experiments are the wholesale price contract (*WP*), the two-part-tariff contract (*TPT*), and the minimum-order-quantity contract (*MOQ*). In all treatments we set the supplier's production cost to be $c = 20$ and the demand function $q = 100 - p$ ($A = 100$). For these parameters, $q^* = 40$, yielding the total channel profit to be allocated between the supplier and the retailer of 1600. Under the wholesale price contract, in equilibrium the supplier charges $w_{WP} = 60$, the retailer orders $q_{WP} = 20$, which results in the supplier profit of 800, the retailer profit of 400, and the channel efficiency of (800+400)/1600 = 75%.

Under the TPT contract, the supplier should set $w_{TPT} = c = 20$ in equilibrium, and charge $F = 1600$. This results in the retailer optimal order of $q_{TPT} = 40$, supplier profit of 1600, retailer profit of 0, and the channel efficiency of 100%. Under the MOQ contract, the supplier should set $q_{\min} = 40$ and $w_{MOQ} = 60$ in equilibrium, and the retailer's best reply is to set $q = q_{\min} = 40$, which results in the supplier profit of 1600, the retailer profit of 0, and the channel efficiency of 100%.

The coordination mechanism is independent of relative bargaining power; in other words, the predictions that $w_{TPT} = 20$ and $q_{\min} = 40$ are independent of the bargaining power of the supplier and the retailer. However, the prediction that suppliers extract the entire channel profit ($F = 1600$ and $w_{MOQ} = 60$) is based on the ultimatum game structure, which gives the supplier all the bargaining power, and also on the assumption that retailers will accept any non-negative profit, no matter how small. We already know that the second assumption does not hold in the laboratory, since all laboratory studies report significant retailer rejections that cannot be explained by errors alone (see Pavlov and Katok, 2009, for an explanation involving fairness preferences).

To investigate the effect of relative bargaining power on contract performance, we tested contracting mechanisms under two bargaining protocols. Under the ultimatum bargaining (UB) protocol, the supplier makes an offer, the retailer either orders q or rejects the supplier's offer; in the latter case, the round ends with both players earning zero. Under the *structured bargaining* protocol (SB), the players have 5 minutes, during which the supplier can make offers to the retailer. If the retailer accepts an offer, the round ends. If the retailer rejects it, the supplier can make another offer that results in the retailer's best profit that is not lower (in absolute term) than the profit in the rejected offer. The rationale for this "good faith" bargaining protocol is to simulate serious negotiations. Note that since the supplier is not obligated to improve on the offer (he can repeat the identical offer), the good faith bargaining restriction does not really force suppliers to make higher offers to retailers. But it does increase retailers' bargaining power, in the sense that a retailer rejection, as long as it is made within the 5-minute period, does not result in zero profit. If the 5-minute period ends without an agreement, the bargaining process enters the second stage, which is identical to the ultimatum protocol—the supplier is allowed to make one last and final offer (which may be identical to the last rejected offer) and the retailer can either order q or declare the final impasse, which now results in zero profits for both players.

In summary, our study manipulates two factors, the contract mechanism (WP, TPT or MOQ) and the bargaining protocol *ultimatum* or *structured bargaining*) for a 3 x 2 full factorial design. We manipulate all factors between subjects. Participants play for 15 rounds, keeping their same role (Supplier or Retailer) for the entire session. Each round a retailer and a supplier are randomly matched with a person of the other role within the same cohort. We conducted all treatments in cohorts of six people (three suppliers and three retailers in each), and in all treatments there were 3 or 4 cohorts in the laboratory at any given time. Participants were not told the cohort size. We summarize all treatments and sample sizes in Table 1. In total, 180 subjects participated in our study.

Table 2. Experimental design and sample sizes

	Bargaining Protocol	
Contract Mechanism	*Ultimatum Bargaining*	*Structured Bargaining*
Wholesale price (WP)	3 cohorts of 6	4 cohorts of 6
Two-Part-Tariff (TPT)	8 cohorts of 6	4 cohorts of 6
Minimum Order Quantity (MOQ)	7 cohorts of 6	4 cohorts of 6

3.2 Experimental Protocol

All experimental sessions followed the same protocol. Participants arrived at the computer lab at a pre-specified time and read experimental instructions that describe the rules of the game, the use of the software, and the payment procedures. After all participants had had a chance to read the instructions, the experimenter read instructions to them aloud to ensure common knowledge. Participants then completed 15 rounds of the game, and were paid their actual earnings accumulated from both sessions, privately and in cash. Participants were not allowed to communicate during the experiment.

All sessions were conducted at the Laboratory for Economic Management and Auctions (LEMA) at the Smeal College of Business, Penn State University. *Structured Bargaining* (SB) protocol session lasted approximately two hours, *Ultimatum Bargaining* (UB) sessions lasted approximately one hour, the average earnings, including a $5 participation fee, were $27. Participants were Penn State students recruited through a web-based recruitment system, with cash being the only incentive offered. The majority of our participants were undergraduates from a variety of majors (77%) and the rest were graduate students.

4. Results

4.1. Do coordinating contracts improve efficiency?

In this section we test Hypothesis 1. Table 3 presents summary statistics.

Table 3. Summary Statistics. *Standard errors* are in parenthesis

Bargaining Protocol:	Ultimatum Bargaining			Structured Bargaining		
Contract type:	WP	TPT	MOQ	WP	TPT	MOQ
Overall Efficiency	66% (2%)	80% (3%)	72% (5%)	81% (2%)	90% (4%)	92% (3%)
Proposal Efficiency	78% (2%)	87% (2%)	94% (2%)	85% (2%)	96% (1%)	97% (1%)
Conditional Efficiency	78% (3%)	88% (2%)	94% (1%)	85% (2%)	96% (1%)	97% (1%)
Share of Retailer	40% (2%)	35% (3%)	26% (3%)	46% (2%)	42% (3%)	48% (7%)
Rejection Rate	6% (3%)	8% (2%)	23% (5%)	2% (1%)	6% (3%)	6% (2%)
w	57.1 (2.1)	43.5 (3.1)	52.2 (1.4)	50.5 (1.7)	32.1 (3.4)	43.0 (0.73)
F	-	387.3 (100.8)	-	-	546.8 (83.4)	-
$q_{\min}$	-	-	36.0 (2.1)	-	-	34.75 (2.0)
q (conditional on acceptance)	19.3 (1.7)	28.2 (1.5)	37.0 (1.6)	23.7 (1.1)	33.6 (1.9)	36.0 (1.1)

Figure 2 plots average channel efficiency over time for UB and SB treatment separately. Under UB there is no statistically significant difference between the efficiency of the WP contract and the MOQ contract (p = 0.6068), but the TPT contract is significantly more efficient that the WP contract (p = 0.0191). Under the SB protocol, both coordinating contracts are more efficient than the WP contract (p = 0.0263 for MOQ and p = 0.0948 for TPT). The data is in line with the week version of hypothesis 1. As to the strong version of hypothesis, MOQ and TPT efficiency is significantly below 100%, strongly so under the UB protocol (p = 0.0007 for MOQ and p = 0.0002 for TPT), and weakly so under the SB protocol (p = 0.0524 for MOQ and p = 0.0789 for TPT), so we can reject the strong version.

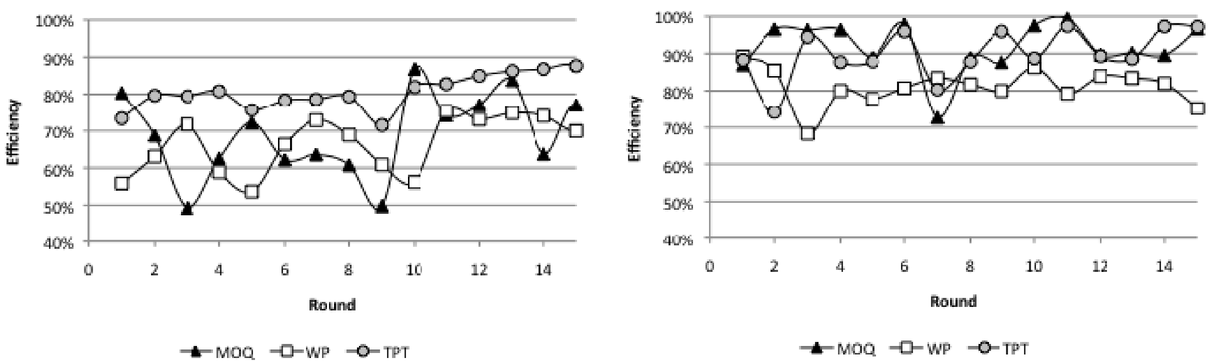

(a) Ultimatum Bargaining (b) Structured Bargaining

Figure 2. Average efficiency over time.

4.2. Do MOQ and TPT contracts differ in their conditional efficiency?

Figure 3 plots conditional efficiency over time for the MOQ and TPT contracts, separately for the UB and SB treatments. In UB treatments MOQ offers are significantly more efficient than TPT offers (p = 0.0414), but in SB treatments there is no significant difference in the efficiency of the offers (p = 0.4902). Our data is consistent with Hypothesis 2 under the UB protocol, but not under the SB protocol.

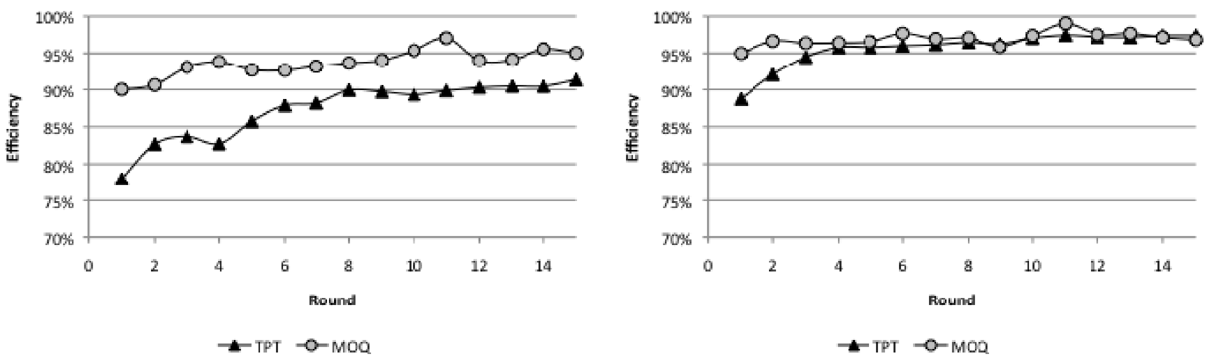

(a) Ultimatum Bargaining (b) Structured Bargaining

Figure 3. Average conditional efficiency over time.

4.3 Does Bargaining Protocol Affect Contract Outcome?

We start by comparing conditional efficiency of suppliers' offers. For MOQ contracts, there is no significant difference in conditional efficiency of the offers ($p = 0.2252$), but for TPT contracts and WP contracts SB offers are significantly more efficient than UB offers ($p = 0.0243$ for TPT and $p = 0.0578$ for WP). When we factor in retailer rejections, however, we find that SB leads to higher efficiency for all three contracts ($p = 0.0106$ for MOQ, $p = 0.0775$ for TPT, and $p = 0.0065$ for WP).

Let us now examine in more detail how contract outcomes differ under the two bargaining protocols. Figure 4 plots the average relative offer the supplier makes to the retailer, over time, separately for MOQ and the TPT contracts.

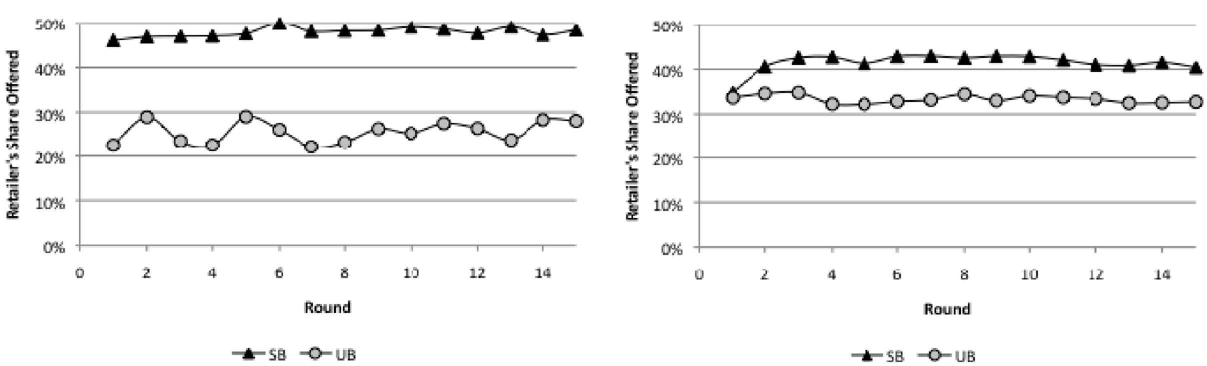

(a) MOQ (b) TPT

Figure 4. Average relative share offered to the retailer over time.

Retailers are offered proportionally more with the SB protocol for both contracts ($p < 0.0001$ for MOQ and $p = 0.0582$ for TPT) but the differences are much larger for MOQ than for TPT. To understand why, we need to examine how suppliers set contract parameters for those two mechanisms. Figure 5 shows how parameters evolve over time under the MOQ contract and under the TPT contract.

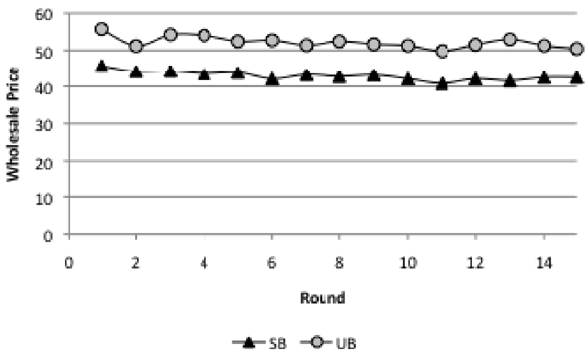

(a) MOQ Contract w over time

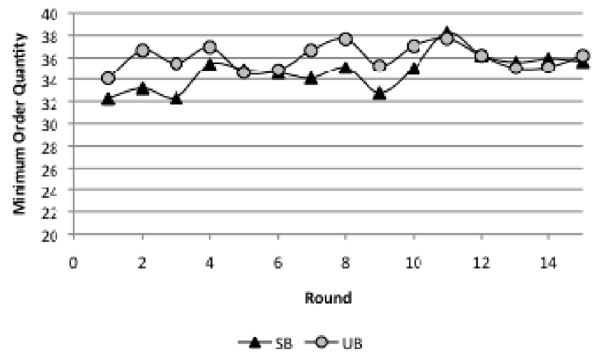

(b) MOQ Contract $q_{\min}$ over time

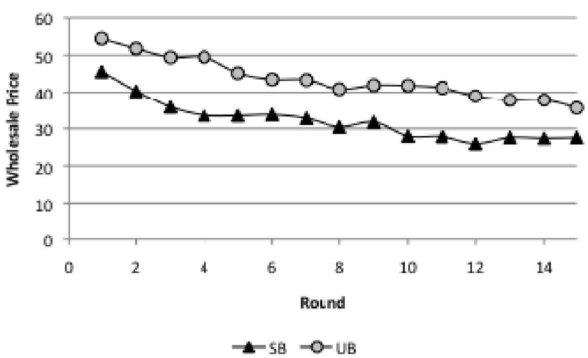

(c) TPT Contract w over time

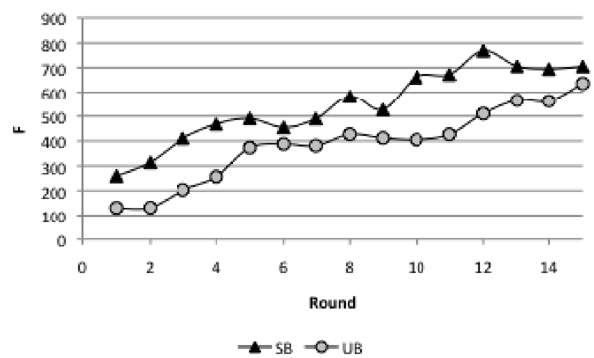

(d) TPT Contract F over time

Figure 5. Contract parameters over time under MOQ and TPT contracts

Recall that under MOQ contract, $q_{\min}$ controls the contract's conditional efficiency, and w determines the retailer's relative share. We can see from Figure 5(a) and 5(b) that those parameters are fairly stable over time. Recall that there is no significant difference in conditional efficiency of the MOQ contract due to the bargaining protocol, which can be verified by noting that the two lines in Figure 5(b) are quite close together. The wholesale prices in Figure 5(a), however, are clearly different, and remain so throughout the session, with only a very small downward trend. Since w determines profit distribution, lower w under SB is consistent with Hypothesis 3(b) that states that the retailer has more bargaining power under SB.

The story is quite different for TPT contract. Recall that w determines conditional efficiency under the TPT contract—the closer w is to c, the closer should the efficiency be to 100%. As evidenced by Figure 5(c), wholesale prices start out much too high under both bargaining protocols, but significantly lower under SB than under UB. Wholesale prices under both bargaining protocols steadily decrease over time, which implies that conditional efficiency increases over time. While suppliers learn that lower w's increase efficiency, they simultaneously learn to extract more of the channel profit, by increasing F, so the retailer share stays remarkably steady over time (Figure 4). TPT contracts that are being offered early in the session (with high w's and low F's) are structurally close to wholesale-price contracts, so the retailer's share, is determined primarily by the contract structure, but by fairness considerations. This explains relatively low rejection rates under UB (only 8%, in contrast to the UB MOQ rejection rates of 26%).

4.4. Is there evidence for deviation due to inequality aversion?

The average deviation from optimal q is small but significant. It is -0.321 (std err 0.121) for TPT UB and 0.566 (std err 0.295) for MOQ UB. These deviations, though small, are significantly different from 0 at p-value= 0.009 for TPT UB and p-value=0.056 for MOQ UB. TPT UB deviations are significantly different from MOQ UB deviations, mean difference is 0.887, std. error 0.289, p-value = 0.001. Thus, deviation from optimal q depends on condition.

Below, in Table 4, we see the estimation for the joint decision model (of quantity and accept/reject decisions) done separately for MOQ and TPT. From Table 4, we see that β, though small is significant for TPT, but not significant for MOQ. This is despite the fact that MOQ displays more rejections than TPT. The result that β is insignificant for MOQ is not merely driven by the insignificant deviations from optimal quantity under MOQ (seen above), but also by an insignificant relationship between Rejection and the profit of the first mover (p-value 0.12 in logistic regression) under MOQ. Thus, the higher rejections under MOQ, while clearly and significantly driven by a profit motive (own profit significantly affects the probability of rejection), do not appear to be driven by a *relative* profit motive.

Coefficient	Estimate TPT	Estimate MOQ
β	0.021 (0.006)	0.003 (0.036)
Sigma1	1.562 (0.055)	4.62 (0.201)
Sigma2	268.4 (18.76)	488.5 (50.74)
Rho12	0	0
LL	695.824	872.114

Table 4. Parameter Estimates for FS Utility estimation using the two-decision joint estimation of eqns 3-6.

5. Conclusions

A growing literature has demonstrated that the format of a bargaining contract matters even when theoretically it should not and that bargaining contracts often fail to reach the 100% efficiency they are intended to achieve. We showed that both results can be primarily attributed to rejections generated by an ultimatum-game-like structure. This finding and its implications allowed us to improve the design of bargaining contracts in two ways.

First, we proposed and tested an alternative format involving minimum order quantity. In that format, the seller announces a minimum order quantity which is equal to the efficient quantity, thus extracting the surplus through price (minimum order quantity contract). Although rejections in the ultimatum version of that format are higher, we find that proposals with the minimum order quantity format are far more conditionally efficient than two-part-tariff proposals in terms of the overall surplus they imply.

Second, we eliminated the ultimatum structure by allowing for structured negotiations. With *structured negotiations*, rejection rate drastically falls leading to a more direct relationship between proposal efficiency and ex-post efficiency. This allows all coordinating contracts to reach near full efficiency and makes the difference between formats less consequential.

Acknowledgments

This research was funded in part by a National Science Foundation grant. Katok gratefully acknowledges support from the Smeal College of Business and the Center for Supply Chain Research (CSCR) at Penn State's Smeal College of Business through the Smeal Summer Grants Program.

References

Azoulay, P. and Shane, S. (2001), Entrepreneurs, contracts, and the failure of young firms, *Management Science*, 47(3), 337-358.

Balachandran, K. R. and Radhakrishnan, S. (2005), Quality implications of warranties in a supply chain, *Management Science*, 51(8), 1266-1277.

Ball, S. B. and P. Cech (1996), Subject pool choice and treatment effects in economic laboratory research, *Experimental Economics*, 6, 239-292.

Benzion, U., Cohen, Y., Peled, R. and Shavit, T. (2007), Decision-making and the newsvendor problem—An experimental study, *Journal of the Operational Research Society*.

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

Bolton, G. and Ockenfels, A. (2000), A theory of equity, reciprocity, and competition, *American Economics Review*, 90(1), 166-193.

Bolton, G., Ockenfels, A. and Thonemann, U. (2008), Managers and Students as Newsvendors: How Out- of-Task Experience Matters, Working paper, University of Cologne, Germany.

Bostian, A. J. A., Holt, C. A. and Smith, A. M. (2007), The newsvendor "pull-to-center effect": Adaptive learning in a laboratory experiment, *Manufacturing and Service Operations Management* 10(4), 590-608.

Burnetas, A. and Ritchken, P. (2005), Option pricing with downward-sloping demand curves: The case of supply chain options, *Management Science*, 51(4), 566-580.

Cachon, G. (2003), Supply chain coordination with contracts, in *Handbooks in Operations Research and Management Science, vol. 11: Supply Chain Management*, A. de Kok and S. Graves, eds., Elsevier, North Holland, 229-340.

Cachon, G. and Lariviere, M. (2005), Supply chain coordination with revenue sharing contracts, *Management Science*, 51(1), 30-44.

Choi, K., Dai, J. G. and Song, J. (2004), On measuring supplier performance under vendor-managed inventory programs in capacitated supply chains, *Manufacturing & Service Operations Management,* 6(1), 53-72.

Corbett, C. J and Fransoo, J. C. (2007), Entrepreneurs and newsvendors: Do small businesses follow the newsvendor logic when making inventory decisions? Working paper, UCLA.

Corbett, C. J., Zhou, D., and Tang, C. S. (2004), Designing supply contracts: Contract type and information asymmetry, *Management Science*, 50(4), 550-559.

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

De Bruyn, A. and Bolton, G. E. (2008), Estimating the influence of fairness on bargaining behavior, *Management Science*, 54(10) 1774-1791.

Donohue, K. (2000), Efficient supply contracts for fashion goods with forecast updating and two production modes, *Management Science*, 46(11), 1397-1411.

Fehr, E. and Schmidt, K. M. (1999), A theory of fairness, competition and cooperation, *Quarterly Journal of Economics,* 114(3), 817-868.

Ferguson, M. E., Guide, V. D. and Souza, G. C. (2006), Supply chain coordination for false failure returns, *Manufacturing & Service Operations Management*, 8(4), 376 - 393.

Gopal, A., Sivaramakrishnan, K., Krishnan, M. S. and Mukhopadhyay, T. (2003), Contracts in offshore software development: An empirical analysis, *Management Science*, 49(12), 1671-1683.

Ho, T. Noah Lim, and Tony Haitao Cui (2010) "Is Inventory Centralization Profitable? An Experimental Investigation," Working paper.

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

Holt, C. A. (1995), Industrial organization: A survey of laboratory results, in *Handbook of Experimental Economics*, J. Kagel and A. Roth, eds., Princeton University Press, Princeton, N.J., 349-443.

Kagel, J. H. and Roth, A. E., eds. (1995), *Handbook of Experimental Economics*, Princeton University Press, Princeton, NJ.

Kahneman, D., Knetsch, J. L. and Thaler, R. H. (1990), Experimental tests of the endowment effect and the Coase theorem, *Journal of Political Economy*, 98, 1325-1348.

Kahneman, D. and Tversky, A. (1979), Prospect theory: An analysis of decision under risk, *Econometrica*, 47, 263-291.

Kamrad, B. and Siddique, A. (2004), Supply contracts, profit sharing, switching, and reaction options, *Management Science*, 50(1), 64-82.

Katok, E., Thomas, D. and Davis, A. (2008), Inventory service level agreements as coordination mechanisms: The effect of review periods, *Manufacturing & Service Operations Management*, 10(4) 609-624.

Katok, E. and Pavlov, V. (2009), Fairness and coordination failures in supply chain contracts, Working paper, Penn State University.

Keser, C. and Paleologo, G. (2004), Experimental investigation of retailer-supplier contracts: The wholesale price contract, Working paper, IBM Research.

Kleindorfer, P. R. and Wu, D. J. (2003), Integrating long- and short-term contracting via business-to-business exchanges for capital-intensive industries, *Management Science*, 49(11), 1597-1615.

Kremer, M., Minner, S. and Van Wassenhove, L. (2007), Regret and framing effects in the newsvendor model—the impact of framing and complexity, Working paper, INSEAD.

Lee, H. and Whang, S. (1999), Decentralized multi-echelon supply chains: Incentives and information, *Management Science*, 45(5), 633-640.

Lim, Noah and Teck-Hua Ho (2007), Designing Price Contracts for Boundedly-Rational Customers: Does the Number of Blocks Matter?, *Marketing Science*, 26 (3), 312-326.

Loch, C. H. and Wu, Y. (2005), Behavioral operations management, in *Foundations and Trends in Technology, Information and Operations Management*, U. Karmarkar, ed., Now Publishers, Los Angeles, CA.

Loch, C. H. and Wu, Y. (2008), Social preferences and supply chain performance: An experimental study, *Management Science*, 54(11), 1835-1849.

Lurie, N. H. and Swaminathan, J. M. (2009), Is timely information always better? The effect of feedback frequency on decision making, *Organizational Behavior and Human Decision Processes*, 108(2), 315-329.

Mortimer, J. H. (2004), Vertical contracts in the video rental industry, Working paper, Harvard University.

Padmanabhan, V. and Png, I. P. L. (1995), Returns policies: make money by making good, *Sloan Management Review*, 37 (1), 65-72.

Plott, C. (1987), Dimensions of parallelism: Some policy applications of experimental methods, in A. Roth, ed., *Experimental Economics: Six Points of View*, Cambridge University Press, New York, NY.

Rudi, N., Kapur, S. and Pyke, D. F. (2001), A two-location inventory model with transshipment and local decision making, *Management Science*, 47(12), 1668-1680.

Schweitzer, M. and Cachon, G. (2000), Decision bias in the newsvendor problem: Experimental evidence, *Management Science*, 46(3), 404-420.

Siegel, S. (1956), *Nonparametric Statistics for Behavioral Sciences*, McGraw-Hill, New York.

Soman, D. (2004), Framing, loss aversion, and mental accounting, in *Blackwell Handbook of Judgment and Decision Making,* D. J. Hoehler and N. Harvey, eds., Blackwell Publishing, Boston, MA.

Spengler, J. J. (1950), Vertical integration and antitrust policy, *The Journal of Political Economy*, 58(4), 347-352.

Wu, D. J. and Kleindorfer, P. R. (2005), Competitive options, supply contracting, and electronic markets, *Management Science*, 51(3), 452-466.

Wu, Y. D. (2009), The impact of long-term relationship on supply chain contracts: A laboratory study, Working paper, University of Kansas.

Appendix. Derivation of the Joint Likelihood for TPTUG

$$Q^* = \frac{100 + 120\beta - (1 + 2\beta)w}{2(1+\beta)} + \varepsilon_1$$

This result only applies when retailer is disadvantaged, which is the case in 100% of all observations used.

Proof:

Let Q = A – Bp. Thus, p = (A – Q)/B

π_S = (w-c)Q

π_R = (p-w)Q

Start with U_R = [(A-Q)/B-W]Q - F - β(π_S-π_R).

We take first order conditions w.r.t Q to get the optimal quantity Q*.

Q* = [A + (A+c) β - (2 β + 1)W]B/[2(1+ β)]

Plugging that back in U_R we get U_R = (1+ β)[(100-Q-w)Q-L]- β [(w-20)+L]

Joint Likelihood: L = $P(Q)^{Accept}*P(Reject)^{1-Accept}$

P(Accept) = P(Q>0 AND U_R>0) = P(Q>0|U_R>0)P(U_R>0)

= exp(Q*)/(1+exp(Q*))exp(U_R)/(1+exp(U_R)

P(Reject) = 1 – P(Accept) = 1- P(-ε_1 < Xb, - ε_2 < U_R)

= 1 – F(Xb,U_R)

=1-cdfbvn(Xb/sigma1,U_R/sigma2,rho1)

=1-cdfbvn(-ε_0/sigma1,U_R/sigma2,rho1)

ε_0 = Q- (β_0+ 100 + 120β_1 - (2β_1 + 1)W)/(2(1+β_1))