

Bargaining Process and Channel Efficiency

Ernan Haruvy

*Naveen Jindal School of Management, University of Texas at Dallas, Richardson, Texas 75080,
eharuvy@utdallas.edu*

Elena Katok

*Naveen Jindal School of Management, University of Texas at Dallas, Richardson, Texas 75080,
ekatok@utdallas.edu*

Valery Pavlov

*Business School, University of Auckland, Auckland, New Zealand,
v.pavlov@auckland.ac.nz*

The behavioral literature has demonstrated that the format of supply chain contracts matters even when theoretically it should not, and that contracts that in theory coordinate channels fail to do so in laboratory experiments. The existing body of experimental evidence uses an ultimatum bargaining protocol to test analytical models, but there is no reason to think that bargaining in supply chains is in the form of ultimatum offers. We investigate the effect of bargaining on contract performance by extending the bargaining protocol to allow the manufacturer to make concessions. We test coordinating contract with bargaining in the laboratory by comparing wholesale price and the two-part-tariff contracts using two different bargaining protocols. We then develop and estimate a statistical model of behavior with bargaining and find that this model organizes our data well. Our main finding is that the contracts we study are more efficient when participants are allowed to make concessions. The additional channel efficiency is due to more efficient offers made by manufacturers. The higher channel efficiency primarily benefits the retailer-- the weaker party. Our main contribution is the observation that when testing analytical models of contracts in the laboratory, the way the bargaining process is implemented, such as the ability to make concessions, has a critical effect on conclusions.

Keywords: Behavioral Operations, Supply Chain Contracts, Experimental Economics, Bargaining

Draft Date: February 2019

1. Introduction

Channel coordination is an important and widely studied topic in the supply chain literature (see Cachon 2003 for a review). A key message that comes out of the analytical modeling literature on channel coordination is that the structure of supply contracts has a large effect on channel efficiency. But conclusions based on laboratory experiments, such as Ho and Zhang (2008), Lim and Ho (2007), Kalkanci et al (2011), Katok and Pavlov (2013) among others) find that contract structure does not affect channel efficiency very much. Many of these studies find that coordinating contracts do not even improve efficiency relative to the wholesale price contract, and others find the improvement observed is substantially smaller than what the theory predicts. The two streams of research result in different managerial implications: the modeling literature suggests that managers should devote effort to carefully structuring contracts. In contrast, behavioral literature suggests that managerial effort is better spent on thinking how to frame contracts (Ho and Zhang (2008), Lim and Ho (2007)), or managing some other aspects of the relationship, such as private information Kalkanci et al (2011).

Our work provides a link between analytical and behavioral research on channel coordination. Particularly, we are interested in the efficiency implications of structure, with careful attention to behavior (e.g., Gode and Sunder, 1993). Most published behavioral studies on channel coordination, that report little if any improvement due to contract structure, implement bargaining as a take-it-or-leave-it offer from the manufacturer to the retailer, that the retailer can only either reject completely, or accept as is. In contrast, we find that in our experiments contract structure matters and improves coordination, but only when the experiment implements bargaining in a way that captures some amount of give-and-take, specifically, the ability to make *concessions*.

The standard take-it-or-leave-it implementation of the manufacturer's proposal makes the contracting setting similar to the Ultimatum Game that has been thoroughly studied in behavioral economics since Güth et al. (1982, and see a review of early literature in Roth 1995). In the Ultimatum Game, one player-- the proposer-- makes an offer to divide a fixed sum of money between himself and a recipient. The recipient can either accept the offer, in which case both players receive the amounts specified in the proposer's offer, or reject it, in which case both players earn zero.

The advantage of the take-it-or-leave-it implementation is that it offers clear theoretical predictions. Specifically, under coordinating contracts, such as two-part-tariff, the manufacturer should be able to not only coordinate the channel, but also extract the entire channel profit. In contrast, ultimatum bargaining laboratory studies report both, high levels of rejections and relatively equal profit divisions -- features that are also evident in laboratory studies of contracting. The findings from our experiments show that the bargaining protocol has a large effect on channel efficiency (defined as actual channel profit divided by the first-best channel profit) and on profit distribution.

In this study we directly compare two bargaining protocols. Under the *Ultimatum Bargaining* (UB) protocol, the manufacturer proposes contract terms to the retailer. These contract terms result in some profit for the retailer; we call the maximum profit the retailer can obtain given the manufacturer's proposed contract terms, the manufacturer's *offer*. Contract terms depend on the contract type. For example, under the wholesale price (WP) contract, the manufacturer offers the wholesale price w , so we can write manufacturer's offer as $(\pi_R^*(w))$. To streamline notation, we will omit contract parameters and refer to retailer's profit conditional on placing the optimal order as simply π_R^* , and the corresponding manufacturer profit as π_M^* . The retailer observes manufacturer's offer and either accepts it by placing some order quantity q , or rejects it. If the retailer rejects the contract, both parties earn their outside option, which is zero in our experiments. If the retailer orders quantity q , the retailer earns π_R and the manufacturer earns π_M . The difference between π^* and π is in that the former is based on the optimal order quantity, while the latter is based on the actual order quantity.

An important realistic feature missing from Ultimatum Bargaining is the ability of the parties to make *concessions*. There is evidence in the literature (Cialdini et al. 1975) that concessions increase the likelihood of reaching an agreement. The explanation Cialdini et al. 1975 provides, is called *reciprocal-concessions*. The idea is that during the bargaining process, when one side makes a concession, it activates a social norm to reciprocate.

We study the effect of the bargaining procedure by introducing a *Structured Bargaining* (SB) protocol that augments the UB protocol by adding a stage in which the manufacturer can make concessions. We do this by adding a bargaining stage prior to the final ultimatum stage, that we call Stage 1. Stage 1 lasts a pre-specified amount of time during which the manufacturers can

make offers to the retailer.¹ If the retailer accepts an offer, the game ends with the players earning π_R and π_M , but if the retailer rejects an offer, then rather than both players earning zero, the manufacturer can make another offer to the retailer.

If the players fail to reach an agreement during Stage 1's allotted time, the game reverts to the Ultimatum Stage, in which the manufacturer can make one last and final offer. If the retailer rejects this final ultimatum offer, the round ends in an *impasse*. The feature that the SB protocol can revert to the ultimatum bargaining is similar to some of the treatments in Bolton and Karagözoğlu (2014). The SB mechanism is intentionally simple (simpler than what a real bargaining process is likely to be). It's only substantive different from UB is that a retailer rejection does not automatically end the negotiation, but instead, it gives the manufacturer an opportunity to make a *concession*. It is the effect of this feature of bargaining on contract outcomes that we are most interested in exploring.

We begin to see some of work in behavioral operations that consider the effect of the bargaining process. Leider and Lovejoy (2016) test the balanced-principal bargaining model of Lovejoy (2010) and use free form bargaining, as well as allow communication, in a supply chain involving three tiers and multiple players on each tier. Davis and Hyndman (2017) use free bargaining with structured communication in the wholesale price contract in a setting with uncertain demand, and show that efficiency improves when players can negotiate wholesale price and order quantity at the same time. Becker-Peth et al. (2017) find a similar result for the buyback contract. Davis and Leider (2016) use free bargaining in their investigation of the capacity investment problem. Katok and Tan (2017) compare behavior in a setting with supply disruptions under different bargaining and communication protocols and find that communication by itself is sufficient to restore most of the efficiency that is lost in a setting when neither bargaining nor communication is allowed. Villa and Katok (2017) report on an experiment that uses free bargaining to negotiate transfer prices in their test of the Rudi et al. (2001) model. Our paper contributes to the literature by identifying concessions as a key feature of the bargaining process

¹ The fixed duration of Stage 1 is reminiscent of the deadline effect in the behavioral economics literature (for example see Ochs and Roth 1989, Roth and Xing (1994), Gneezy et al. (2003) among others). In our experiment Stage 1 is followed by the ultimatum stage, so the deadline effect is not pronounced.

that leads to improved efficiency, and by providing a direct and clean test of the effect of bargaining on contract performance in a fundamental simple setting.

Two streams of literature are relevant for our study: behavioral economics literature on ultimatum games, and the literature on channel coordination and contracting. There are two critical findings from the ultimatum game literature (see overview in Roth, 1995): (1) low offers are often rejected, and (2) high offers (in the context of the ultimatum game this is usually 50% of the total pie) are rarely rejected.

Laboratory experiments have been instrumental in recent behavioral operations literature in identifying behavioral implications of using different contracting arrangements. The two central findings from the experimental work are that (1) contracts that in theory should coordinate channels do not generally do so in the laboratory, and (2) the profits tend to be divided more equally than theory predicts. One source of channel inefficiency is retailer rejections that range from 26% in Ho and Zhang's (2008) two-part-tariff contracts to 11% in Lim and Ho's (2007) two-block tariff contract. Another source of channel inefficiency is offer inefficiency. *Offer efficiency* is defined as the channel efficiency if the retailer accepts the offer and places the optimal order. Offer efficiency ranges from 93.62% in Ho and Zhang's (2008) TPT contract to 80.8% in Lim and Ho's (2007) two-block tariff contracts.

We find that SB improves channel efficiency primarily by improving offer efficiency while maintaining essentially the same rejection rates as UB. Interestingly, SB starts with offers that resemble UB offers, but offer efficiency increases as Stage 1 progresses, and the main beneficiary from the added offer efficiency turns out to be the retailer. This happens in both TPT and WP contracts, and the effect persists in late rounds.

In the next section we formally describe our setting and the theoretical benchmarks when players are fully rational expected-profit maximizers. We also summarize behavioral predictions for our settings based on what we know from the literature. In section 3 we describe the details of the design of our experiments. In Section 4 we present data analysis. In section 5 we develop and estimate a statistical behavioral model that includes concessions and loss aversion, and show that this model predicts actual bargaining dynamics quite well. We conclude in section 6 by summarizing our results, mentioning some limitations, and discussing managerial implications.

2. Analytical Background with Full Rationality

We examine a one-period bilateral channel commonly used in modeling papers, in which a single manufacturer sells its product to a single retailer that then sells it to consumers (e.g., Biyalogorsky and Koenigsberg, 2010; Iyer and Bergen 1997; Iyer and Villas-Boas, 2003; Padmanabhan and Png 1997; and see Desai et al. (2004) and Iyer et al. (2007) for settings with competing retailers).

In this section we begin by summarizing the contract structure and analytical results about the two-part-tariff (TPT) and wholesale price (WP) contracts when the players are fully rational expected profit maximizers. We consider a channel with a single manufacturer and a single retailer. The manufacturer has a constant marginal production cost c . The retailer faces a linear demand $q = A - p$ where p is the retail price and A is a constant. The product has no salvage value; therefore the retailer sells the entire quantity ordered.

The channel is coordinated if the outcome in terms of units produced (and therefore the total channel profit) is the same as the *first best* amounts – the outcome that would have resulted from an integrated channel with a single decision maker maximizing the channel profit. We investigate two contracts within this framework: wholesale price (WP) and two-part-tariff (TPT). We chose the WP as one of the contracts we examine because it is the simplest contract, having only a single parameter, and it provides a good baseline in that a substantial amount of laboratory data on WP has already been published and its performance is quite consistent across studies. We chose the TPT contract to compare with WP because TPT is a coordinating contract and is most like the Ultimatum Game, in that if the manufacturer sets the wholesale price to coordinate the channel, the fixed fee corresponds to the proposer's demand in the ultimatum game. TPT is one of many coordinating contracts, some, such as the quantity discount, are mathematically equivalent to it, while others, such the block tariff, or the minimum order quantity contract, are not. We leave extending this research program to other contracts to future research.

Under WP, the retailer pays to the manufacturer a wholesale price w per unit. Under standard rationality and profit maximization assumptions, it is well known that theoretical channel efficiency under the WP contract is below that of the integrated contract. This channel inefficiency of the wholesale price contract relative to the integrated channel is known as *double marginalization*.

While coordinating contracts are numerous, there is a large class of these contracts that, in one way or another, induce the retailer to place the first best order by making the retailer's marginal cost w equal to the manufacturer's marginal cost c . The manufacturer then extracts some of the profit from the channel. For example, the two-block tariff (Lim and Ho 2007) coordinates the channel 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. In the TPT which we investigate here, the manufacturer coordinates the channel by setting $w = c$, and extracts channel profit by charging the retailer a fixed fee F . The equilibrium properties of WP and TPT contracts are established, and we summarize in Table 1 the expressions for optimal decisions in the integrated channel, the wholesale-price contract, and the TPT contract.

Table 1. Properties of the WP and TPT contracts and the integrated channel.

Contract	Channel Profit	Optimal wholesale price w^*	Optimal Fixed Fee F^*	Optimal order quantity q^*
Integrated Channel	$\frac{(A-c)^2}{4}$	c	N/A	$\frac{A-c}{2}$
Wholesale Price	$\frac{3(A-c)^2}{16}$	$\frac{A+c}{2}$	N/A	$\frac{A-w}{2}$
Two Part Tariff	$\frac{(A-c)^2}{4}$	c	$\frac{(A-w)^2}{4}$	$\frac{A-c}{2}$

Under the UB protocol, the manufacturer moves first and makes an offer (w, F) in the TPT setting or (w) in the WP setting. The retailer then determines the optimal order quantity given the offer, computes the corresponding profits, and decides whether to accept or reject the offer. If both players are profit maximizers, the manufacturer's proposal and the retailer's order follow the expressions in Table 1, and the retailer accepts any offer that results in a non-negative profit.

SB involves two stages. The last stage, that we call the Ultimatum Stage, is identical to the UB setting. The first stage, Stage 1, involves sequential offers by the manufacturer.² The retailer can either accept a standing offer, ending the round, or reject it, in which case the manufacturer

² We conducted some sessions with "good faith bargaining" restriction, which limits manufacturers to offers that are non-decreasing in terms of retailer's profit conditional on retailer's profit-maximizing order. We also conducted sessions without this restriction. Qualitatively, neither behavior nor any of our conclusions are affected by this restriction, so we pooled the data for the purpose of analysis.

can make another offer, as long as there is time remaining in Stage 1. A retailer motivated by profit (or one with reference-dependent preferences such as in Ho and Zhang 2008) may accept a Stage 1 offer if she believes that subsequent offers are unlikely to result in the higher utility, or reject it, hoping for better offers.

The total channel profit under the TPT contract depends on the wholesale price w —the closer is w to the production cost c , the larger is the channel profit. Ho and Zhang (2008) already demonstrated that manufacturers under the TPT indeed set wholesale prices above production costs ($w > c$), and attribute this to retailers' reference dependence (or loss aversion) with respect to the fixed fee. Unlike UB, under SB, offers can be adjusted, so if a manufacturer makes a SB offer $\#t (w_t, F_t)$ with $w_t > c$, it is always possible to follow it up with another offer (w_{t+1}, F_{t+1}) so that $w_{t+1} < w_t$, $F_{t+1} \geq F_t$, $\pi_{t+1}^R > \pi_t^R$ and $\pi_{t+1}^M \approx \pi_t^M$. In other words, any offer to the retailer that involves $w_t > c$ can be improved—manufacturer can offer the retailer a *concession*. This can be done without making the manufacturer worse off because as w approaches the efficient level, both manufacturer and retailer can be made better off.³

Because Ho and Zhang (2008) already showed that wholesale prices under the TPT contract are significantly above production cost, the above argument implies that it would be reasonable for wholesale prices to decrease during Stage 1 of the SB treatment. Whether or not they do and how this affects contract performance are empirical questions that we intend to answer with our laboratory experiment. We summarize them below:

1. Will manufacturers under SB make concessions in Stage 1 and will retailers sometime accept Stage 1 offers?
2. Will average wholesale prices decrease during Stage 1 of SB?
3. Will SB exhibit higher efficiency than UB?

³ Consider the setting in which $c = 20$ and $q = 100 - p$ (parameters used in our experiments). Consider an offer $w = 60$ and $F = 0$. Retailer's best reply is $q = 20$ which results in $\pi_R = 400$ and $\pi_M = 800$. Now suppose retailer rejects this offer and manufacturer follows up with $w = 40$ and $F = 300$. Retailer best reply is $q = 30$ which results in $\pi^R = \pi^M = 900$. The retailer received a concession of $900 - 400 = 500$, and at the same time, the manufacturer is also better off, earning 900 instead of 800.

3. Experimental Setting and Design

3.1 Experimental Setting

Contracting arrangements in our experiments are the wholesale price contract (*WP*) and the two-part-tariff contract (*TPT*). In all treatments we set the manufacturer's production cost to be $c = 20$ and the demand function $q = 100 - p$ (i.e. $A = 100$). For these parameters, the first best production quantity of $q^* = 40$ results in the total channel profit of 1600 to be allocated between the manufacturer and the retailer. Assuming profit maximization and no errors, under the wholesale price contract, in equilibrium, the manufacturer charges $w_{WP} = 60$, the retailer orders $q_{WP} = 20$, which results in the manufacturer profit of 800, the retailer profit of 400, and the channel efficiency of $(800+400)/1600 = 75\%$. Under the TPT contract, the manufacturer should set $w_{TPT} = c = 20$ in equilibrium and charge $F = 1600$ to extract the entire channel profit. This would result in the retailer order of $q_{TPT} = 40$, manufacturer profit of 1600, retailer profit of 0, and channel efficiency of 100%.

To investigate the effect of the bargaining protocol on contract performance, we tested contracting mechanisms under both, UB and SB bargaining protocols. Under the UB protocol, the manufacturer makes an offer, the retailer either orders q or rejects the manufacturer's offer; in the latter case, the round ends with both players earning zero. Under the SB protocol, the players have 5 minutes, during which the manufacturer can make offers to the retailer. If the retailer accepts an offer, the round ends. If the retailer rejects it, the manufacturer can make another offer. This initial stage is called Stage 1.

If Stage 1 ends without an agreement, the bargaining process enters stage 2-- the ultimatum stage, which is identical to the UB protocol—the manufacturer makes one last and final offer and the retailer can either order q or declare the final impasse, which results in zero profits for both players.

3.2 Experimental design

The main experiment manipulates two factors, the contract mechanism (*WP* and *TPT*) and the bargaining protocol (*UB* or *SB*) for a 2×2 full factorial design. We manipulate all factors between subjects. Participants play for 15 rounds, keeping their roles (Manufacturer or Retailer) for the entire 15 rounds. In each round a retailer and a manufacturer are randomly matched with a

person in the other role within the same cohort. We conducted all treatments in cohorts of six people (three manufacturers 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 2. In total, 186 subjects participated in our study.

Table 2. Experimental design and sample sizes

Contract Mechanism	Bargaining Protocol	
	Ultimatum Bargaining	Structured Bargaining
Wholesale price (WP)	3 cohorts of 6	4 cohorts of 6
Two-Part-Tariff (TPT)	12 cohorts of 6	12 cohorts of 6 ⁴

We ran only three WP-UB cohorts because results of WP-UB experiments are well-established in the literature (Ho and Zheng 2008, Lim and Ho 2007, Loch and Wu, 2008, Katok and Pavlov 2009) and our data replicated existing results, so there was no reason to collect more observations in this treatment. We ran four cohorts of the WP-SB treatment because the analysis of the wholesale price contract is not the main focus of our paper. Even with this small number of independent observations of the WP contract, we will show in the next section that the effect of structured bargaining on profits and some of the contract parameters is statistically significant.

In all sessions, participants arrived at the computer laboratory 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 (see on-line Appendix A.3). After all participants had a chance to read the instructions, the experimenter read instructions to them aloud to ensure common knowledge. Participants then completed multiple rounds of the game, which was implemented in zTree (Fischbacher, 2007), and were paid their actual accumulated earnings, privately and in cash. Participants were not allowed to communicate during the experiment.

All sessions were conducted at an experimental laboratory at a public university in Texas. SB sessions lasted approximately 1.5 hours, UB sessions lasted approximately one hour. Average

⁴ In six of the twelve cohorts, the manufacturer was restricted to making offers that were not worse in terms of the optimal retailer profit than the last offer on the table. For those cohorts, this restriction was in effect in both stages. This “good faith” bargaining restriction turned out to be innocuous (see Table 7 in Appendix A.2) and all the results we report hold if we use either the data with or without this restriction. Therefore, we present the results based on the pooled data.

earnings, including a \$5 participation fee, were \$27. Participants were students recruited through a web-based recruitment system, with cash being the only incentive offered. The majority of our participants were graduate students in business and engineering.

4. Results

4.1 Overall Negotiation Outcomes

We organize the first part of the results section to correspond to research questions that we summarized at the end of Section 2. We then proceed to propose a simple behavioral model that qualitatively organizes our data, and report on structural estimation of this behavioral model.

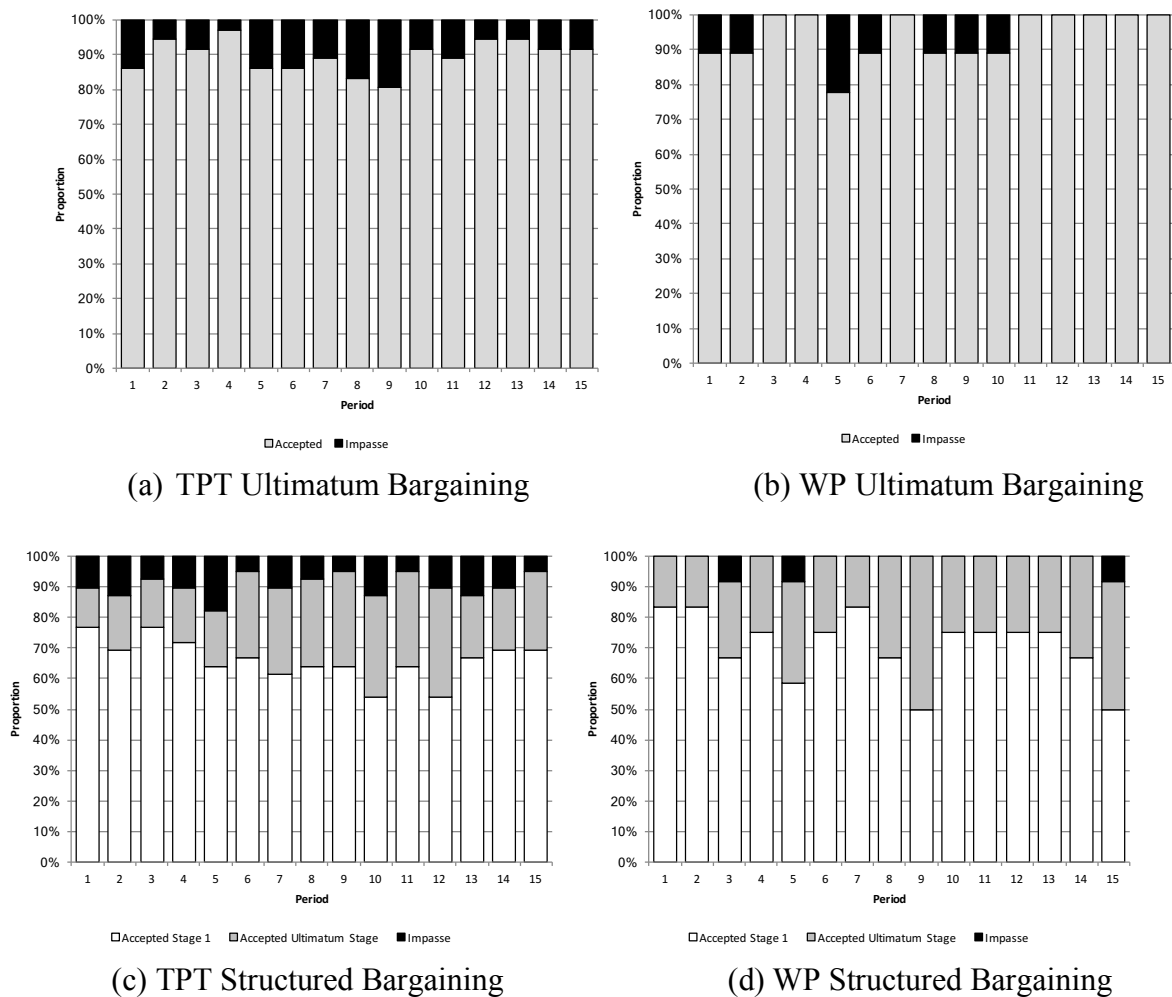


Figure 1. Proportion of impasses and offers accepted in the two SB stages.

Figure 1 shows the outcomes of negotiations in the four treatments of our study. For UB treatments in Figure 1(a) and Figure 1(b), we show the proportion of negotiations that ended in

agreement and impasse. For SB treatments in Figure 1(c) and Figure 1(d), we show the proportion of negotiations that were successfully completed in Stage 1, in the ultimatum stage, and the proportion of negotiations that ended in impasse. Overall, the bargaining protocol does not affect the proportion of negotiations that ended in impasse, which is quite low (see Table 3 for exact proportions). Under SB, over 70% of negotiations were completed in Stage 1 (69.1% in TPT and 70.5% in WP).

4.2 Manufacturers' Concessions under the SB protocol

The first question we ask is whether manufacturers will make Stage 1 offers as well as concessions under SB, and whether retailers will sometimes accept these offers. We already showed that Stage 1 offers are both made, and often accepted. Next, we examine the extent to which manufacturers make concessions. We define a manufacturer concession as the difference between an offer and the previous offer in terms of retailer profit (not utility), conditional on the retailer placing the profit-maximizing order. In other words, it is the amount by which a retailer profit conditional on placing the optimal order increases from one offer to the next.

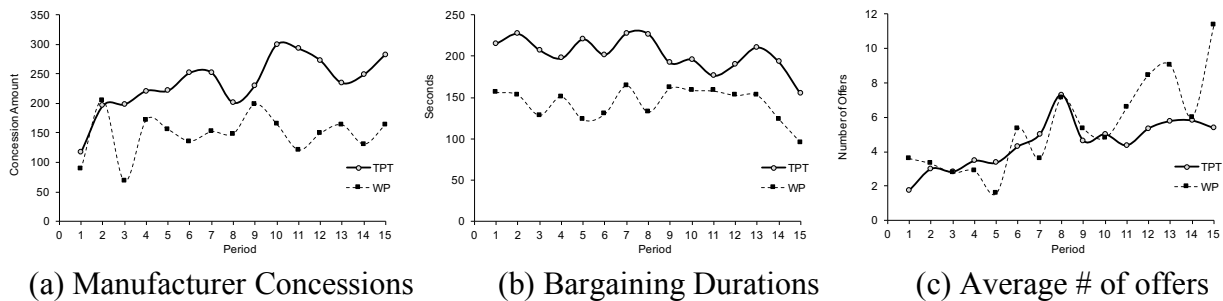


Figure 2. Bargaining behavior under the SB protocol.

Figure 2 shows how the bargaining process evolves over the 15 periods of the experiment. The average amount manufacturers concede (Figure 2 (a)) appears to be stable. From Figure 2 (b), the average duration is shorter than the maximum of 300 seconds. Toward the end of the session, participants make about 5-10 concessions in less than 200 seconds. This observation that manufacturers learn to make smaller concessions is reminiscent of results reported in experimental economics. For example, McKelvey and Palfrey (1995) estimate logistic QRE and show that errors are smaller in later rounds, and de Bruyn and Bolton (2008) report the same phenomenon in the context of bargaining. Correctly accounting for the error structure is particularly important in stochastic choice models (Wilcox, 2011, pp 99-100).

4.3 Contract Performance

Table 3 presents descriptive statistics for all four treatments, along with first and last round averages for SB treatments.⁵ The table also displays the results of comparisons between the first and last round SB averages, and between the UB and last round SB averages. Unless otherwise noted, all statistical tests use cohort average as the unit of analysis, and p -values are two-sided and refer to the Wilcoxon rank-sum test.

Table 3. Summary of contract parameters and performance.

	Two-Part-Tariff (TPT)			Wholesale Price (WP)		
	SB: First offer	SB: Last offer	UB	SB: First offer	SB: Last offer	UB
w	41.96 (2.36)	35.63* (2.05)	43.03 ^{††} (2.27)	57.40 (5.14)	50.41 (1.75)	57.11 [†] (2.14)
F	478.44 (38.02)	463.13 (52.20)	418.74 (77.10)	NA		
Optimal order (q^*)	29.01 (1.16)	31.32 (1.47)	28.45 (1.21)	21.30 (2.57)	24.79 (0.88)	21.69 [†] (1.23)
Actual order (q) for accepted offers	NA	32.02 (1.03)	28.58 [†] (1.13)	NA	24.14 (0.91)	19.25 ^{††} (1.68)
Proposal retailer profit	428.43 (46.75)	600.81** (33.63)	448.74 ^{††} (32.55)	486.82 (95.19)	622.81 (42.60)	473.31 [†] (45.98)
Proposal manufacturer profit	985.74 (34.99)	887.99* (24.59)	924.36 (60.55)	730.56 (22.10)	732.65 (16.03)	768.79 (6.63)
Offer Efficiency	0.8839 (0.0187)	0.9305* (0.0104)	0.8868 ^{††} (0.0140)	0.7609 (0.0692)	0.8485 (0.0171)	0.8039 (0.0290)
Final Efficiency	NA	0.8653 (0.0196)	0.7845 ^{††} (0.0256)	NA	0.8111 (0.0243)	0.6637 ^{††} (0.0192)
Impasse Rate		0.0574 (0.0143)	0.1019 [†] (0.0176)		0.0167 (0.0107)	0.05926 (0.0323)

Note: Standard errors are in parentheses. Average difference between the first and the last offer according to Wilcoxon rank-sum tests: * $p < 0.10$, ** $p < 0.05$. Average difference between the last round SB and UB according Wilcoxon rank-sum tests: [†] $p < 0.10$, ^{††} $p < 0.05$.

First, considering the differences between the first and last round of SB treatments: wholesale prices decrease, the fixed fees under TPT do not change, the best reply order increases, as does the retailer's profit and offer efficiency. The manufacturer profit decreases under TPT.

⁵ We also did the same analysis for the last five rounds, and results in terms of statistical significance are very similar except that differences in manufacturer profit are weakly significant in the last five rounds when comparing TPT-SB and TPT-UB. When comparing WP-SB to WP-UB, the difference in offer efficiency becomes significant.

All differences except in the fixed fee between the first and last offers are statistically significant under TPT and are not significant under WP.

Next, we consider the differences between the UB treatments and the last offer of SB treatments: Here again, wholesale prices are higher, fixed fees (in the TPT treatments) are not different, both conditionally-optimal and actual UB orders are slightly lower than corresponding SB orders. Retailer profits, and offer efficiency, are significantly lower under UB, while impasse rates under UB are slightly higher (weakly significant). As a result, overall efficiency under UB is significantly lower. At the same time manufacturer profits are not different under the two protocols.

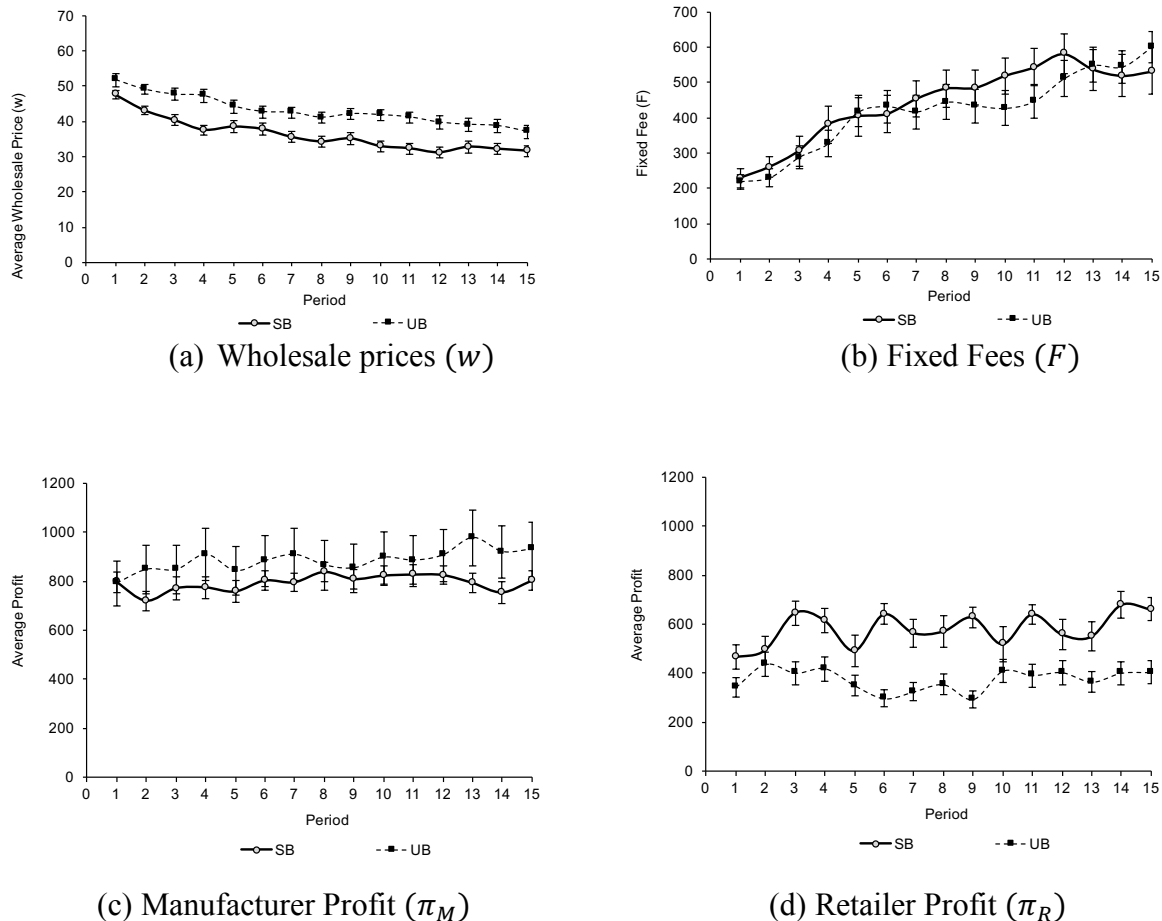


Figure 3. Average contract parameters and profits over rounds under the TPT contract. Vertical bars represent standard errors.

The story that emerges is that under SB, our participants improve efficiency through a series of concessions that manufacturers make by lowering wholesale prices while keeping the

fixed fees mostly fixed. This allows the manufacturers to improve retailers' profits while not substantially lowering their own.

Average contract parameters evolve over the 15 rounds of the session under the TPT contract (Figure 3(a) and (b)). The figure shows that, over the course of the session, average TPT wholesale prices appear to decrease and average TPT fixed fees appear to increase, but the resulting profits appear to be quite stable (Figure 3(c) and (d)).⁶

To check whether the better performance under SB can be attributed to the fact that manufacturers make, and retailers observe, more offers, we conducted an additional long TPT-UB treatment in which participants interacted for 40 rounds instead of 15. All other aspects of the experiment were the same as in the TPT-UB treatment, and 48 participants (8 cohorts of 6) were included in the long TPT-UB treatment. We provide a detailed analysis in on-line Appendix A.1 that shows that learning in the long TPT-UB treatment lasts for the first 10-15 rounds only.

5. Behavioral Model and Estimation

5.1. Probabilistic Choice and Reference-Dependent Utility

The purpose of the behavioral model we estimate in this section is to explain the dynamic patterns of Stage 1 offers in terms of w and F . Theoretically, with TPT, whether under UB or SB protocols, the manufacturer and the retailer should be able to reach full channel efficiency (coordination) and distribute the surplus between themselves. This full efficiency outcome occurs when the manufacturer charges a wholesale price equal to the marginal cost of 20.

While no behavioral model has hitherto been proposed for the dynamic framework of structured bargaining (SB) we investigate here, we can begin with a model that has been proposed for the static ultimatum bargaining (UB) TPT setting, and extend it to the dynamic setting. Our goal is to have a model that is *parsimonious* and can fit both static and dynamic settings with one set of parameters, while capturing the qualitative differences between the static and dynamic settings.

⁶ Under the wholesale price contract, SB wholesale prices appear to be quite stable, while UB wholesale prices appear to show a slight downward trend over rounds. Profits do not appear to show much trend over rounds. A plot analogous to Figure 3 for wholesale price contract is in Figure 7 in the online Appendix A.1.

We build on the model of Ho and Zhang (2008), who proposed a reference-dependent utility model in which the retailer has disutility from the fixed fee, and also makes random errors. If the retailer places a profit maximizing order $q^* = (A - w)/2$, then her monetary profit from offer (w, F) is $\pi_R^* = (p - w)q^* - F$. Let the retailer's utility from (w, F) under UB condition be

$$u_R(w, F) = \pi_R(w, F) - \beta F + \varepsilon \quad (1)$$

where β is the reference dependence parameter with respect to the fixed fee. Note that equation (1) is equivalent to the retailer's utility function in Ho and Zhang (2008).

Like Ho and Zhang (2008), we model retailers as making random errors (McKelvey and Palfrey 1995). Under the assumption that the error terms follow a Type I extreme value distribution, we consider the retailer's probability of accepting some offer that results in utility of $u_R(w, F)$ versus rejecting this offer and earning utility of 0. Under the UB protocol, we specify this probability P of acceptance as a logit:

$$P = \frac{\exp(\tau u_R(w, F))}{1 + \exp(\tau u_R(w, F))} \quad (2)$$

where τ is the rationality parameter. If a player never makes errors then $\tau = \infty$, and the probability of accepting any offer for which $u_R(w, F) > 0$ is 1. If $\tau = 0$, then any offer is accepted with probability 0.5. For any intermediate values of τ , *we can say that the probability of acceptance is monotonically increasing in the expected utility to the retailer.*

5.2. Reciprocal Concessions Model

Recall that SB differs from UB only by the presence of Stage 1, which gives manufacturers an opportunity to make *concessions*. Cialdini et al. 1975 show that concessions increase the likelihood of agreement, which they explain by the idea that a concession activates a social norm to reciprocate. Cialdini et al. (1975) summarize this reciprocal concession idea by observing that "...the likelihood of a concession by one party is positively related to the occurrence of a concession by another party." (p. 207). Cialdini et al. (1975) report on a set of experiments in which subjects are more likely to agree to a small favor when the request was preceded by another

request for a much larger favor that was rejected. Much work has followed the original experiments, and the idea has been refined. For example, Fern et al. (1986) and O’Keefe and Hale (1998) found that the likelihood of acceptance does not seem to depend on the size of the initial concession, so Hale and Laliker (1999) proposed that the reciprocity norm becomes activated as long as the concession size exceeds some threshold. The full review of this literature is beyond the scope of our paper and we refer interested readers to a review by Cialdini and Goldstein (2004). But our main point is that there exists a parallel between this work and our setting in that retailers are more likely to accept a Stage 1 offer when it was a result of one or more concessions.

We model the idea that concessions matter by including the concession amount into retailer’s utility function. Therefore, we will add the subscript t to represent Stage 1 offer number t . A Stage 1 offer t , (w_t, F_t) results in the retailer’s potential profit of

$$\pi^R(w_t, F_t) = (p_t - w_t)q_t^* - F_t,$$

Let Δ_t be concession from offer t . For the first offer, $\Delta_1 = 0$, and after the first offer, the concession from Stage 1 offer t is

$$\Delta_t = \pi^R(w_t, F_t) - \pi^R(w_{t-1}, F_{t-1}). \quad (3)$$

If the retailer were to place the profit-maximizing order $q_t^* = (A - w_t)/2$, then we can write the retailer’s utility from offer t as

$$u_R(w_t, F_t) = \pi^R(w_t, F_t) + \sum_{i=1}^t \theta_i \Delta_i - \beta F_t + \varepsilon. \quad (4)$$

The difference in retailer’s utility from *accepting* an offer under the UB and SB protocols is that in the SB, Δ_t affects the maximization of the manufacturer’s profit in the previous period because w_{t-1} and F_{t-1} appear in the expressions for the concessions. The utility from concessions captures the fact that the retailer reacts to past offers. We let the probability that retailer accepts offer number t be $P_t = \Pr(\text{Accept}|w_{t-1}, F_{t-1}, w_t, F_t)$ specified as

$$P_t = \frac{\exp(\tau u_{Rt}(w_t, F_t, w_{t-1}, F_{t-1}))}{\exp(\tau \psi) + \exp(\tau u_{Rt}(w_t, F_t, w_{t-1}, F_{t-1}))} \quad (5)$$

where $u_{Rt}(w_t, F_t)$ is defined by equation (4) and denotes the retailer’s utility form accepting the offer number t , and ψ denotes the retailer’s belief about what would be his expected utility from potential future offers that will come if offer t is rejected. In this formulation ψ is a constant, and

we will show in Section 5.5 that the value of ψ has a minor effect on the optimal contract parameters.

5.3 Exploring Retailer Behavior

In this section we explore retailers' behavior in order to further examine the factors that we conjecture to affect the likelihood that retailers accept offers. First, to provide the reader with a sense of how retailer's probability of accepting an offer depends on the fixed fee and the concession amount, we categorized all offers according to the retailer's profit from this offer ($\pi_R \leq 530$, $530 \leq \pi_R < 640$, and $\pi_R > 640$), the magnitude of the fixed fee ($F \leq 350$, $350 < F \leq 700$ and $F > 700$) and the magnitude of the concession ($\Delta \leq 15$, $15 < \Delta \leq 60$, and $\Delta > 60$).⁷ In Figure 4(a) we plot the probability of acceptance for the nine π_R/F pairs. In Figure 4(b) we plot the probability of acceptance for the nine π_R/Δ pairs.




Figure 4. Relationship between offer acceptance, fixed fee and concession amount.

It is apparent from the figure that on aggregate, and controlling for profit, higher fixed fees decrease probability of acceptance, while higher concessions increase probability of acceptance. To show this formally, we estimate a logit model with fixed effects for individuals with the dependent variable 1 if an offer was accepted and 0 if it was rejected. In the baseline model (Model

⁷ We chose the bin classifications so that each has approximately the same number of observations.

1) we include only independent variables from equation (4): $\pi^R(w_t, F_t)$, $F_t \Delta_t$. In Model 2 we add Δ_{t-1} in order to check the extent to which retailers pay attention to the concession history, in addition to the latest concession. To see how retailer behavior evolves over time, we add, in Model 3, the period number, the offer number, and an interaction variable between the period number and the fixed fee. We report estimates of the logit model in Table 4.

Table 4. Logit model for the likelihood of accepting an offer.

Parameter	Description	Model 1	Model 2	Model 3
$\pi^R(w_t, F_t)$	Retailer profit from offer t conditional on optimal order	0.0080*** (0.0006)	0.0080*** (0.0006)	0.0084*** (0.0006)
F_t	Fixed fee from offer t	-0.0008*** (0.0002)	-0.0008*** (0.0002)	-0.0013*** (0.0005)
Δ_t	Latest concession	0.0078*** (0.0002)	0.0080*** (0.0009)	0.0073*** (0.0009)
Δ_{t-1}	Previous concession		0.0020** (0.0099)	0.0014 0.0010)
<i>Period</i>	Period number (1-15)			-0.2121*** (0.0272)
t	Offer number			0.1101*** (0.0178)
$Period \times F_t$	Period number \times Fixed fee from offer t			0.0001** (0.00005)
Average of Fixed Effects		-4.0791*** (0.2264)	-4.1242*** (0.2292)	-3.1627*** (0.2834)
Obs. (N)		3091		
LL		-846	-844	-791

Note: *** $p < 0.01$, ** $p < 0.05$.

Consistent with the behavior that we formalized in equations (4) and (5), we see that the likelihood of an offer being accepted increases with the size of the offer and with the size of the latest concession, while it decreases with the size of the fixed fee. The coefficient for the previous offer's concession (Δ_{t-1}) is also positive and significant. Therefore, we conclude that equations (4) and (5) capture the main behavioral drivers. The effect of concession history exists, but is of

secondary importance.⁸ We estimated Model 3 to gain insight into dynamics. It shows that the propensity to accept offers decreases over rounds, and increases with the offer number. Additionally, the decrease in the likelihood of offer being accepted due to the size of the fixed fee, becomes smaller over rounds (positive and significant coefficient for $Period \times F_t$), suggesting that retailers' sensitivity to the fixed fee may be transitory.

5.4 Behavioral Model Estimation

Now that we have evidence that concessions matter to retailers, we proceed to estimate retailers' behavioral parameters using maximum likelihood estimation. For this estimation we estimate the effect on the latest concession only. We specify the log likelihood to be maximized in equation (6), where $Accept_t$ is an indicator variable that takes on the value of 1 if retailer accepted offer number t , and 0 otherwise.

$$\ln L = \sum_R \sum_t [Accept_t \ln(P_t) + (1 - Accept_{Rt}) \ln(1 - P_t)] \quad (6)$$

The estimation entails a binary logit using only retailer's acceptance/rejection decision.

We report in Table 5 separate estimates for SB and UB in columns labeled "Behavioral." We also estimate and display in Table 5 for comparison, a model in which we restrict $\beta = \theta = 0$, which we label "noise only." All estimates in Table 5 include individual effects.

The procedure for testing whether the β and τ parameters for UB and SB treatments are statistically different is as follows: we estimated the model on the pooled data and then compared the Bayesian Information Criteria (BIC) for the unrestricted model (separate parameters for UB and SB) and the restricted model (some or all parameters are the same across the two treatments). Using this method, we cannot reject the null hypothesis that β and τ estimates are the same for UB and SB, and β estimates are positive. We also conclude that the behavioral model provides a much better fit than the model with noise only.

⁸ We also estimated a model in which we included Δ_{t-2} that yielded nearly identical estimates for the Model 2 variables, and a positive and marginally significant Δ_{t-2} coefficient ($p = 0.096$). So the positive effect of concessions and the negative effect of the fixed fee appear to be robust.

Table 5. Estimates for Equation (5) for SB and UB separately

Parameter	Structured Bargaining (SB)						Ultimatum Bargaining (UB)					
	Noise only			Behavioral			Noise only			Behavioral		
Period	1-5	6-10	11-15	1-5	6-10	11-15	1-5	6-10	11-15	1-5	6-10	11-15
τ	0.007* (0.0005)			0.007* (0.0005)			0.015* (0.002)			0.015* (0.002)		
	0.006* (0.001)	0.010* (0.001)	0.015* (0.002)	0.006* (0.001)	0.010* (0.001)	0.015* (0.002)	0.057* (0.024)	0.028* (0.009)	0.066* (0.034)	0.057* (0.024)	0.046* (0.018)	0.166* (0.100)
β	NA			0.108* (0.024)			NA			0.077* (0.041)		
				0.219* (0.076)	0.027 (0.034)	0.037 (0.024)				0.070 (0.109)	0.135* (0.030)	0.0543* (0.036)
θ	NA			1.090* (0.139)			NA					
				0.804* (0.274)	0.804* (0.180)	0.892* (0.171)						
N	3,091						540					
$-2\ln L$	1966.0			1853.0			157.0			153.6		
AIC	2034.0			1925.03			231.0			229.6		
BIC	2239.2			2142.3			389.8			392.7		

Notes: Standard Errors are in parenthesis; * $p < 0.01$. BIC formula: $BIC = -2\ln L + k\ln(N)$; For the reference dependent model $k=40=4$ parameters + 36 individual intercepts for SB, $k=38=2$ parameters + 36 individual intercepts for UB. For the noise only model $k=37$.

Our aggregate estimates of the β parameter, which captures the aversion to the fixed fee, is close to the analogous estimate of Ho and Zhang (2008)⁹. The interpretation of the positive and significant estimate of θ is that retailers derive utility from receiving concessions in addition to the utility they derive from profit.

In addition to estimating models for the data aggregated over all 15 rounds, we examined whether retailers' behavioral parameters evolve over time. We did this by estimating the model for rounds 1-5, 6-10, and 11-15. We show these estimates in Table 5 immediately below the aggregate estimates. We note two observations about the estimates of behavioral parameters over rounds. First, under SB, the aversion to the fixed fee appears to disappear after about 5 rounds,

⁹ Ho and Zhang (2008) express the retailer's utility as $(p - w)(100 - p) - \lambda F$, while our analogous expression can be written as $(p - w)(100 - p) - F - \beta F$. Therefore, $\beta + 1 = \lambda$, and our β estimates are equivalent to λ in the 1.1 range, which is slightly lower than the Ho and Zhang 2008 estimate of 1.37 (experiment 1) and 1.27 (experiment 2) (see p. 696). This is because, unlike us, Ho and Zhang (2008) use not only retailer's accept/reject decision but also manufacturer's offers. The evidence suggests that manufacturers make offers as if they believe that the retailers are far more sensitive to F than the retailers actually are.

while under UB it does not. Second, the preference for concessions persists virtually unchanged over rounds.

5.5. Manufacturer's Behavior

We would like to model the manufacturer as a fully rational expected-profit maximizer, who is aiming to maximize his expected profit and is reacting to the retailer's response function. In our experimental setting, the manufacturer is not constrained in terms of the number of offers that she can make, but there is a time limit in Stage 1, so as a simplification, let us assume that the maximum number of offers is T . Then the manufacturer's optimization problem at time 1 is:

$$\max_{w_t, F_t} \sum_{t=1}^T \left[\prod_{i=1}^{t-1} (1 - P_i) \right] P_t \pi_t^M \quad (7)$$

Where P_t is defined in equation (5) and $\pi_t^M = (w_t - c) \left(\frac{A - w_t}{2} \right) + F_t$.

An important issue we must consider is manufacturer's *beliefs* about retailer's behavioral parameters (β , θ and τ). The natural assumption would be for the manufacturer to have rational expectations, and at least on average, correctly consider retailer's behavioral parameters as we estimated them in Table 5. We can check whether manufacturers (on average) have rational expectations with respect to the loss aversion parameter β because, Ho and Zhang (2008) showed that there is a one to one relationship between the wholesale price the manufacturer offers under the TPT contract, and the manufacturer's belief about retailer's loss aversion, namely:

$$\hat{\beta} = \frac{w - A}{2w - (A + c)} - 1 \quad (8)$$

Using equation (8) and the data from our TPT-UB treatment, we compute manufacturers' belief about the retailer's loss aversion to be $\hat{\beta} = 0.66$ (standard error is 0.195), which is much higher than the actual retailers' fixed fee sensitivity¹⁰. In general, manufacturer's optimization

¹⁰ We computed average w for each participant in the role of manufacturer in the TPT-UB treatment. Given our parameters, equation (8) is not valid for $w \geq 60$, so we use $\beta = 0$ for three participant who did not satisfy this restriction and used equation (8) to compute β 's for the remaining participants, and averaged them.

problem in equation (7) does not have a closed form solution, but there is a closed form solution for the wholesale price in the last period, w_T (see Appendix for details):

$$w_T^* = \frac{(1 + \theta + \beta)c + \beta A}{1 + \theta + 2\beta} \quad (9)$$

Solving for the other variables requires deriving the first order conditions. As an illustration, we also include in the Appendix, the first order conditions (FOC) derivations for the last period fixed fee, F_T .

There is no further insight that can be gained from analytical derivations for earlier periods $t < T$, so we omit them, and instead solve the problem numerically. We present computational results for the case of $\hat{\beta} = 0.66$ and parameters θ and τ come from Table 5.¹¹ We use exponentially declining weights on past concessions $\theta_1 \dots \theta_{t-1}$. There is also an additional set of constraints $\pi_t^R \geq \bar{\pi} \forall t$, restricting offers to ones that result in a minimum level of retailer profit. We set this minimum acceptable profit at $\bar{\pi} = 200$. These constraints are binding in early periods because without them, the optimal early offers include an arbitrarily high F_t because they maximize the size of the subsequent concessions, increasing retailer's utility and the probability of acceptance. We also set parameter $\Psi = 300$ (recall from equation (5) that Ψ denotes the retailer's belief about what would be his expected utility from potential future offers if offer t is rejected). We performed sensitivity analysis on these constants.¹²

In Figure 5 we display a distribution of the lengths of negotiations that we observe in our data. The distribution of lengths appears to lie in the range of $T = 1$ to $T = 15$, with average

¹¹ The TPT-UB treatment can be considered as simply a special case of $T = 1$, and solving equation (8) for this special case yields $w_t^* = 42.91$ and $F_t^* = 360.43$, comparable to observed averages in the TPT-UB treatment of $w = 43.03$ and $F = 418.74$.

¹² w_t^* and F_t^* change as Ψ changes, but qualitatively, the pattern of offers is quite consistent. For $\Psi = 0$, wholesale prices start at 36.2, gradually increase to 40.4 by offer 9, and then decrease consistent with our observed data to 35.6. At the same time, fixed fees start at 815 and gradually decrease to 520. For $\Psi = 600$, wholesale prices oscillate between 37 and 55 for the first 9 rounds, and then gradually decrease to 35.6, in line with our data, while fixed fees change little throughout the negotiation, but are about 100 lower than the average fixed fees we observe. So, the value $\Psi = 300$ can be considered a fitted parameter since we cannot observe the retailer's beliefs from our data. The parameter $\bar{\pi}$ has very little effect on the results when it is below about 300. Even for $\bar{\pi} = 0$ wholesale prices oscillate between 53 and 38 for the first 7 periods, and then gradually decrease to 36.6, and fixed fees are virtually unaffected. Setting $\bar{\pi}$ exceedingly high does change the dynamics. For example, for $\bar{\pi} = 600$, wholesale prices start at 24, and gradually increase to 35.6 instead of decreasing, while the fixed fee shows much more movement than it does in our data, starting at over 800 and gradually falling to about 300 by round 15. So, the value $\bar{\pi} = 200$ can also be considered a fitted parameter.

negotiation length of 5.72 and with only about 3.7% of negotiations having more than 15 offers. Therefore, we solved the problem for $T = 15$.¹³





Figure 5. Distribution of the lengths of negotiations observed in our data.


From this setup, we obtain several insights about how the manufacturer, faced with a behavioral retailer, would optimally adjust contract parameters. Figure 6Figure 7 graphically compares optimized and average observed offers from negotiations that lasted not more than 15 rounds. Figure 6Figure 7 shows that our model predicts bargaining dynamics well.

We conclude that manufacturers' behavior in terms of how they structure their offers is qualitatively close to expected profit maximization that accounts for the retailer's behavior and over-estimates the retailers' sensitivity to the fixed fee.


¹³ The WP-SB treatment is a special case of $F = 0$, and solving equation (8) for that special case, using a procedure identical to one we used for TPT, results in a sequence of wholesale prices that start at $w_1^* = 60$ and declines to $w_{15}^* = 49.43$, which is comparable to observed wholesale prices that decline from $w_1 = 57.4$ to $w_T = 50.41$.




(a) Wholesale prices (w).



(b) Fixed Fees (F).



(c) Retailer profits (π^R).



(d) Manufacturer profits (π^M).

Figure 6. Comparison between optimal and observed bargaining process.

5.6. Discussion

The behavioral model that we propose is one in which the retailer dislikes the fixed fee in the TPT contract, and likes concessions. While Ho and Zhang (2008) were the first to note the disutility from the fixed fee, we are the first to develop a quantitative model of dynamic bargaining that explicitly models concessions. While stylized and parsimonious, our model of retailer captures realistic human behavior, because when people negotiate, they start with aggressive offers, and then make multiple concessions to finally arrive at a compromise solution. Therefore,

our model captures the qualitative model of Cialdini et al. (1975) in a quantitative framework that we apply to contracting in supply chains.

Our model of manufacturer is also reasonable. Manufacturer in our setting tries to be rational and maximize his expected profit but overestimates the retailers' disutility from the fixed fee. Retailer's disutility from the fixed fee is measured by the behavioral parameter β , which we estimate to decrease from about 0.2 in the first 5 rounds to about 0.03 in subsequent rounds. Manufacturers, however, estimate β to be about 1.1 in rounds 1-5, 0.66 in rounds 6-10, and 0.48 in rounds 11-15. Manufacturer also treats the problem with indefinite end as if it was a finite, although a long, problem. We do not argue that our behavioral model fully captures the complex process of bargaining. To date, no model exists that does that. But the model we propose is parsimonious, trackable, and sufficiently realistic to capture several important features of the real bargaining process (aggressive initial offers and multiple concessions that improve offer efficiency), as well as organize the data in our experiment.

A limitation of our model is that the retailer is not fully strategic. A strategic profit-maximizing retailer, faced with a manufacturer who makes concessions, would reject all stage 1 offers—not observed in our data. A strategic profit-maximizing manufacturer, anticipating that the retailer would reject all stage 1 offers, would not make concessions. Thus, a model with two strategic profit-maximizing players does not have a hope of organizing our data.

6. Conclusions

We began our work with an observation that the ultimatum bargaining (UB) protocol has been commonly used in laboratory studies that tested analytical models of channel coordination. One contribution of our paper is to compare the UB protocol to the structured bargaining (SB) protocol and show that the protocol for bargaining used in the laboratory can make a difference to the conclusions regarding contract performance.

The SB protocol that we introduced has properties that imply an outcome that should be similar to the UB protocol, making it a useful step in exploring the effect of bargaining on contract

performance.¹⁴ While we cannot claim that a bargaining mechanism that allows manufacturers to make concessions is the true mechanism responsible for the treatment effect we observe, it is both, intuitively plausible and is consistent with our data.

An important feature of our setting, is that the negotiation between the manufacturer and the retailer is not a zero-sum game, but rather there is an opportunity to expand the channel profit through setting efficient contract parameters. The Two-Part-Tariff contract is an ideal vehicle to demonstrate the tradeoff between using a contract that appears intuitively attractive (one with high wholesale price and a low fixed fee) and using one that is efficient (one with low wholesale price and high fixed fee). The aspect of our paper that is both, novel, and managerially relevant, is that there is a complementarity between the structure of the contract and the process used for negotiating this contract. In other words, the structure of the contract is not independent from how specific contract parameters are negotiated. Under the two-part-tariff, the natural contract with a high wholesale price and a low fixed fee is inefficient. The need to make concessions in the SB condition forces the manufacturer to discover a way to expand the channel profit by lowering the wholesale price—an opportunity absent in the UB condition.

The insight that the contract structure is not independent of the bargaining process is likely a general one. Coordinating supply contracts are often complex, and this complexity is difficult for decision-makers to manage. For example, Kalkanci et al. (2011) explore a setting with stochastic demand and asymmetric information and compare the wholesale price contract to two different quantity discount contracts. Like our study (as well as the previously-cited studies of a deterministic demand setting by Ho and Zhang (2008), Lim and Ho (2007) and Katok and Pavlov 2013), as well as others cited in the review article by Chen and Wu (2018)) they find that coordinating contracts do not improve efficiency. But a bargaining process that would allow the buyer and the seller to share some of their private information may well reverse this conclusion.

¹⁴ It is important to emphasize that the SB protocol we investigate here is neither a positive nor a normative description of actual bargaining. That is, we do not propose that our SB protocol is a perfectly accurate reflection of the manner in which channel partners negotiate. It is merely a protocol that is closer to real bargaining than the UB protocol. Both protocols are abstractions of a very complex and nuanced process, but SB allows manufacturer to make concessions, which is a feature that is likely to be important to reaching agreements in many bargaining situations.

Therefore, a managerial take-away from our study is that an effective bargaining process should present opportunities for the bargainers to guide the evolution of contract parameters towards efficiency. How contract parameters evolve, and what is the best bargaining process to enable evolution towards an efficient outcome, depends on the contract structure. We showed that under the two-part-tariff, it is sufficient to allow concessions, and the need to make concessions naturally guides the contract towards efficiency. Davis and Hyndman (2017) show that with stochastic demand, the bargaining process that includes not only the wholesale price, but also the order quantity, helps counter the inefficiency that comes from sub-optimal newsvendor problem solution. And it is likely that in settings with asymmetric information, such as Kalkanci et al. (2011), a bargaining process that facilitates some information sharing would improve the efficiency of the quantity discount contract.

Appendix: Optimal Contract Parameter Characterization at the end of the time horizon

In period $t=T$ (terminal period), supplier maximizes w.r.t w_T and F_T :

$$\max_{w_T, F_T} Z_T = \left[(w_T - c) \left(\frac{A - w_T}{2} \right) + F_T \right] P_T$$

Where we use P_T to represent the probability of an offer that provides utility u_t to the retailer in the final period ($t = T$), and $u_t = \pi^R + \sum_{i=1}^t \theta_i \Delta_i - \beta F_t$ per equation (4).

Let us to rewrite retailer's utility from offer T as $u_T = (1 + \theta_T)\pi_T^R - \theta_T\pi_{T-1}^R - \lambda_T F_T + \text{Past concessions}$ by denoting¹⁵ $\lambda_t = 1 + \theta_t + \beta$ and $\text{Past concessions} = \sum_{i=1}^{T-1} \theta_i \Delta_i$. Also, the supplier profit from offer in period T is $\pi_T^S = (w_T - c) \left(\frac{A - w_T}{2} \right) + F_T$.

Necessary conditions are expressed by the FOCs: $\frac{\partial Z_T}{\partial w_T} = 0$; $\frac{\partial Z_T}{\partial F_T} = 0$. Solving, we get:

¹⁵ The new parameter λ_t is in effect just accounting for the parameter β . There is an F_t inside π_t^R . And that F_t is multiplied by $(1 + \theta_t)$. By setting $\lambda_t = 1 + \theta_t + \beta$, we are cancelling that redundant term, and leaving only β as a multiplier on F_t .

$$\frac{\partial Z_T}{\partial w_T} = \pi_T^S Pr_T' \frac{\partial u_{RT}}{\partial w_T} + Pr_T \frac{\partial \pi_T^S}{\partial w_T} = 0$$

$$\frac{\partial Z_T}{\partial F_T} = \pi_T^S Pr_T' \frac{\partial u_{RT}}{\partial F_T} + Pr_T \frac{\partial \pi_T^S}{\partial F_T} = 0$$

Combining the two: $\frac{\partial u_{RT}/\partial w_T}{\partial u_{RT}/\partial F_T} = \frac{\partial \pi_T^S/\partial w_T}{\partial \pi_T^S/\partial F_T}$

$$\text{Where } \frac{\partial u_T}{\partial w_T} = (1 + \theta_T) \frac{w_T - A}{2}; \quad \frac{\partial u_{RT}}{\partial F_T} = -(1 + \theta_T + \beta); \quad \frac{\partial \pi_T^S}{\partial w_T} = \frac{A+c-2w_T}{2};$$

$$\frac{\partial \pi_T^S}{\partial F_T} = 1$$

$$\text{So: } (1 + \theta) \frac{w_T - A}{2} = -\lambda_T \frac{A+c-2w_T}{2}$$

Which comes out to: $w_T^* = \frac{\lambda_T(A+c) - (1+\theta)A}{(2\lambda_T - (1+\theta))}$ and can be rewritten as equation (9).

To derive F^* , we plug w^* into the FOC, and the expression simplifies to

$$1 + \exp(\tau u_{RT}) = \tau \lambda \pi_T^S.$$

Note that the *Past concessions* constant is important. It changes the u_T and therefore the F_T determined above.

Incidentally (not needed in the computation but used for verification), prior to combining the FOCs, the FOC w.r.t. w is:

$$\frac{\partial Z_T}{\partial w_T} = \tau \pi_T^S Pr_T (1 - Pr_T) \frac{w_T - A}{2} + Pr_T \frac{A+c-2w_T}{2} = 0$$

Acknowledgments

The authors gratefully acknowledge the financial support of the Deutsche Forschungsgemeinschaft through the DFG-Research Group "Design & Behavior", and its members for useful comments. Katok also gratefully acknowledges the support of the National Science Foundation award 1243160. We thank seminar participants at the University of Wisconsin, Georgia Institute of Technology, University of South Carolina, Koç University, Bilkent University, and the Hong Kong University of Science and Technology.

References

Bartling, B, Schmidt K (2015) Reference Points, Social Norms, and Fairness in Contract Renegotiations, *Journal of the European Economic Association*, 13 (1): 98–129.

Biyalogorsky, E, Koenigsberg, O (2010), Ownership Coordination in a Channel: Incentives, Returns, and Negotiations, *Quantitative Marketing and Economics* 8(4): 461-490

Becker-Peth, M., Katok, E, Thonemann, U W (2017) How to Benefit from Manufacturer-Retailer Negotiations: An Example of the Buyback Contract, Working Paper University of Cologne.

Bolton, G. E., and E. Karagözoğlu (2014), On the Interaction of Hard and Soft Bargaining Leverage. A test of Schelling's hypothesis and a modified Zeuthen-Harsanyi model, *Working Paper*.

Cachon G.P. (2003) Supply chain coordination with contracts. *Handbooks in operations research and management science* 11, 227-339.

Cachon, G. P. and M. Fisher (2000), Supply Chain Inventory Management and the Value of Shared Information, *Management Science*, 46(8), 1032-1048

Chen, K.Y and Wu, Y. (2018) "Buyer-Supplier Interactions" in *The Handbook of Behavioral Operations*, K. Donohue, E. Katok and S. Leider eds. Wiley.

Cialdini, R. B. and N. J. Goldstein (2004) Social Influence: Compliance and Conformity, *Annual Review of Psychology* 55, 591-621

Cooper, D., J. Kagel. (2017). Other Regarding Preferences: A Selective Survey of Experimental Results. in *The Handbook of Experimental Economics*, 2., J. Kagel, A. Roth, eds. Princeton.

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

Davis, A.M. and K. Hyndman (2017) Multidimensional Bargaining and Inventory Risk in Supply Chains: An Experimental Study, Working paper, Cornell University.

Davis, A. M. and S. G. Leider. 2016. Capacity investment in supply chains: Contracts and the hold-up problem. Working Paper, Cornell University.

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

Desai, P., Koenigsberg O., and Purohit D., (2004), "Strategic Decentralization and Channel Coordination," *Quantitative Marketing and Economics*, 2(1), 5-22.

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

Fern E.F., Monroe K.B., Ramon A.A. (1986). Effectiveness of multiple request strategies: a synthesis of research results. *J. Marketing Research* 23:144–52.

Fischbacher, U. 2007. z-Tree: Zurich Toolbox for Ready-made Economic experiments, *Experimental Economics*, 10(2), 171-178.

Gneezy, U., Haruvy, E., & Roth, A. E. (2003). Bargaining under a deadline: Evidence from the reverse ultimatum game. *Games and Economic Behavior*, 45(2), 347-368.

Gode and Sunder. 1993. Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality. *Journal of Political Economy*.

Güth, W., R. Schmittberger, B. Schwarze. 1982. An Experimental Analysis of Ultimatum Bargaining. *Journal of Economic Behavior and Organization* 3 (4), 367–388.

Hale, J L, Laliker, M (1999) Explaining the door-in-the-face: Is it really time to abandon reciprocal concessions? *Commun. Stud.* 50: 203–210

Ho, T H, Lim, N, Cui, T H (2010). Reference dependence in multilocation newsvendor models: A structural analysis. *Management Science*, 56(11), 1891-1910.

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

Iyer, A. V., and Bergen M. E., (1997), “Quick Response in Manufacturer-Retailer Channels,” *Management Science*, 43(4), 559–570.

Iyer, G., Narasimhan C., and Niraj R., (2007) Information and Inventory in a Distribution Channel. *Management Science*, 53(10), 1551–1561.

Iyer, G., and Villas-Boas M. J., (2003), “A Bargaining Theory of Distribution Channels,” *Journal of Marketing Research*, 40, 80–100.

Jeuland, A. P., S. M. Shugan. 1983. Managing Channel Profits. *Marketing Science* 2(3): 239-272.

Kalkanci, B K-Y Chen and F Erhun, (2011) Contract Complexity and Performance under Asymmetric Demand Information: An Experimental Evaluation, *Management Science*, 57(4) 689-704.

Katok, E, Pavlov, V (2013) Fairness in Supply Chain Contracts: A Laboratory Study, *Journal of Operations Management* 31: 129-137.

Katok, E. and L. Tan (2017) Procurement in the Presence of Supply Disruptions, Working Paper University of Texas at Dallas.

Leider, S. and W. S. Lovejoy. (2016) Bargaining in supply chains. *Management Science* 62(10): 3039– 3058.

Lovejoy, W. S. 2010. Bargaining chains. *Management Science* 56(12) 2282–2301.

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

Lin, F., Y. Sung and Y. Lo (2005), Effects of Trust Mechanisms on Supply-Chain Performance: A Multi-Agent Simulation Study, *International Journal of Electronic Commerce*, 9(4), 91-112

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

McKelvey, R. D., and T. R. Palfrey (1995). Quantal response equilibria for normal form games. *Games and economic behavior*, 10(1): 6-38.

Moorthy, S. 1987. Managing Channel Profits: Comment. *Marketing Science*, 6, 375–379.

Ochs, J. and A.E. Roth, 1989. An Experimental Study of Sequential Bargaining, *American Economic Review*, 79(3): 355-84.

O’Keefe D J, Hale S L (1998) The door-in-the-face influence strategy: a random-effects meta-analytic review. In *Communication Yearbook*, ed. ME Roloff, 21:1–33. Thousand Oaks, CA: Sage

Padmanabhan, V., and Png I. P. L., (1997), “Manufacturer’s Return Policies and Retail Competition,” *Marketing Science*, 16(1): 81–94.

Pavlov, V., E. Katok. (2010) Fairness and coordination failures in supply chain contracts. Working paper, Penn State University.

Roth, A. E. 1995. Bargaining Experiments. In *The Handbook of Experimental Economics*, John H. Kagel, A. E. Roth, editors, Princeton University Press, 253-348.

Roth, A.E. and X. Xing. (1994) Jumping the Gun: Imperfections and Institutions Related to the Timing of Market Transactions. *American Economic Review*, 84, (September): 992-1044.

Rudi, N., Kapur, S., & Pyke, D. F. (2001). A Two-Location Inventory with Transshipment and a Local Decision Making. *Management*, 47(12), 1668–1680.

Schwartz, Alan and Robert E. Scott (2007) Precontractual Liability and Preliminary Agreements. *Harvard Law Review* 120(3): 661-707.

Su, X. (2008) Bounded rationality in newsvendor models. *Manufacturing & Service Operations Management* 10(4): 566-589.

Swaminathan, J. M. and S. R. Tayur (2003) Models for Supply Chains in E-Business, *Management Science*, 49(10): 1387-1406

Tirole, J. (1988) *The Theory of Industrial Organization*. MIT Press, Cambridge, MA.

Villa, S. and E. Katok (2017) Negotiating Transshipments Prices for Improving Supply Chain Coordination, Working Paper, University of Texas at Dallas.

Wilcox, N.T. (2011) Stochastically more risk averse: A contextual theory of stochastic discrete choice under risk. *Journal of Econometrics* 162(1): 89-104

On-Line Appendix

A.1 The effect of learning and additional repetitions

There appears to be little learning under the wholesale price contract, as can we show in Figure 7 below.

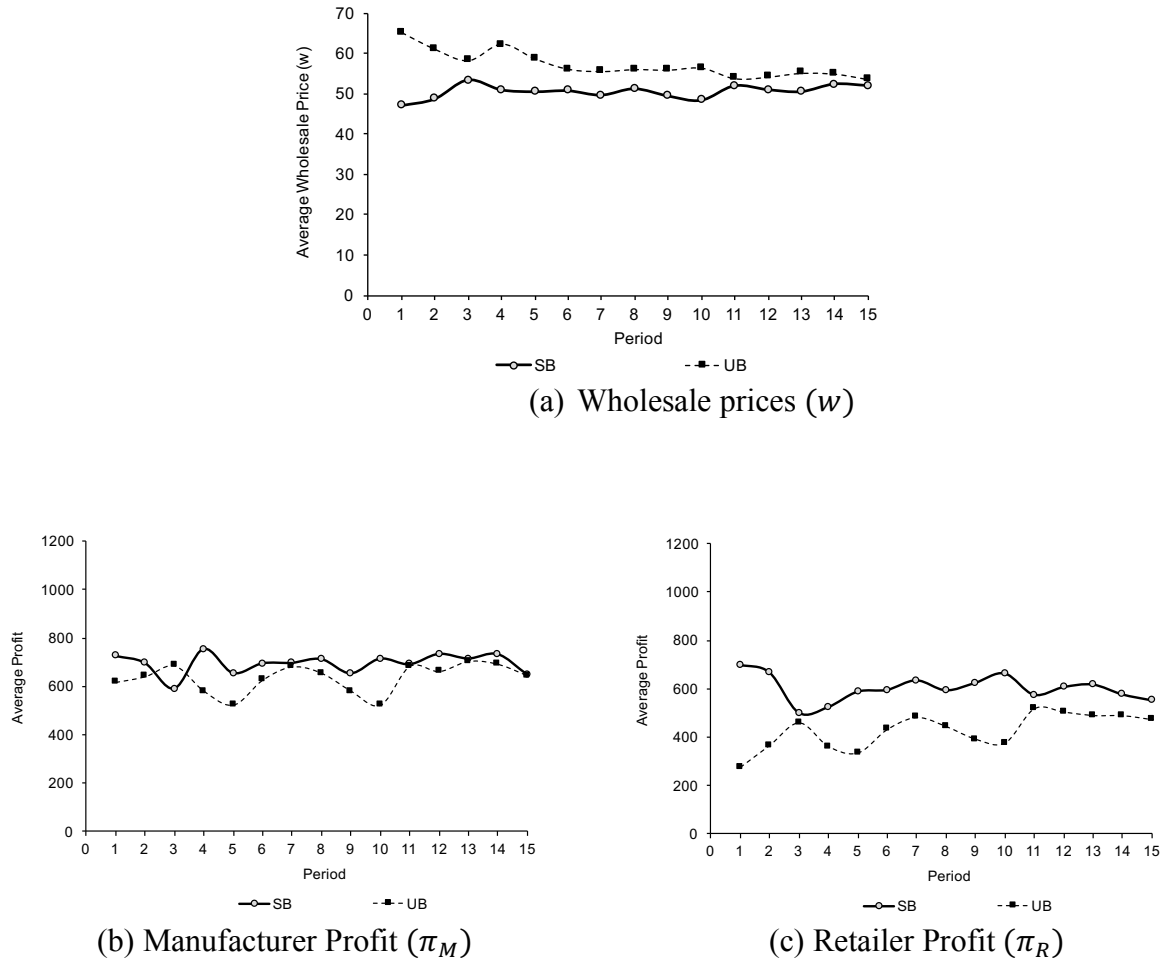


Figure 7. Average contract parameters and profits over rounds under the WP contract.

In contrast, there appeared to be a substantial amount of learning under the TPT contract, as we report in Figure 3 in the main body of the paper. Here we further explore the effect of additional repetition on the TPT contract.

Each round under SB involves multiple decisions, while each UB round involves one decision. In order to test a conjecture that improvement under SB can be fully attributed to repetitions, we conducted an additional TPT-UB-long treatment, in which participants played for

40, rather than 15 rounds. The procedures were identical in all other regards. This treatment included 8 cohorts of 6 participants. We are comparing it with the TPT-SB treatment that included 12 cohorts of 6 participants.




Figure 8. Average contract performance over periods (vertical bars represent standard errors).

Table 6. Averages and standard errors by period blocks in the TPT-UG-long treatment

Avg. (SE.)	Round blocks			
	1-10	11-20	21-30	31-40
w	46.19*** (2.964)	36.74 (3.917)	34.86 (4.35)	34.66 (4.002)
F	310.77** (86.649)	592.17 (145.001)	646.85 (156.059)	651.55 (153.883)
Offer Eff.	0.86*** (0.021)	0.91 (0.021)	0.92 (0.022)	0.93 (0.019)

Matched pair t-test (two-sided) comparing averages between adjacent blocks. *** $p < 0.001$, ** $p < 0.05$.




Figure 8 plots average wholesale prices, fixed fees and offer efficiency over rounds. The figure shows that under ultimatum bargaining, the learning due to repetition takes place at the start of the session, and after period 15 there is very little additional learning that takes place. To demonstrate this formally, we calculated average wholesale price, fixed fee, and offer efficiency, for rounds 1-10, 11-20, 21-30 and 31-40, and display these averages and the corresponding standard errors in Table 6. We computed standard errors using cohort averages as the unit of analysis. The table also reports the results of a matched pair t-test that compares averages between adjacent blocks.

Analysis in Table 6 shows that learning is evident at the start of the session (differences between periods 1-10 and periods 11-20 are highly significant) but tapers off after the first 10-20 periods. No learning is evident in the second half of the session.

A.2 The Effect of Good Faith Restriction (TPT-SB Treatment)

Overall, in the six TPT-SB cohorts without the good faith restriction there were 1376 concessions made, and of those 119, or about 8%, violated the restriction. These violations did

not have any significant effect on average contract parameters, profits, orders, efficiency metrics, or rejections. Table 7 shows the averages for the first and last offer broken down by treatments with and without the restriction.

Table 7. Means, Standard Errors, and comparisons of sessions with and without the good faith restriction.

	Good Faith Restriction (first offer)			Good Faith Restriction (last offer)		
	No	Yes	p-value	No	Yes	p-value
Wholesale Price (w)	42.46 (3.96)	41.49 (2.83)	0.8449	36.35 (2.77)	35.64 (3.16)	0.8791
Lump Sum (F)	450.65 (59.35)	506.22 (50.28)	0.4913	444.54 (125.45)	474.07 (71.01)	0.83
Retailer Profit (π_R)	455.86 (73.56)	400.99 (62.51)	0.5824	615.47 (29.2)	598.89 (43.16)	0.7848
Manufacturer Profit (π_M)	939.23 (43.71)	1032.38 (50.72)	0.1939	870.53 (42.57)	902.51 (28.98)	0.5358
Order (q)				31.64 (1.59)	31.31 (1.88)	0.9042
Optimal Order (q^*)	28.77 (1.98)	29.26 (1.42)	0.8448	32.19 (1.42)	32.15 (1.57)	0.9855
Offer Efficiency	0.87 (0.03)	0.89 (0.02)	0.5474	0.93 (0.01)	0.94 (0.01)	0.6522
Overall Efficiency				0.87 (0.02)	0.88 (0.03)	0.8743
Rejection Rate				0.06 (0.02)	0.05 (0.02)	0.7689

Note: H_0 : no difference due to good faith restriction.

A.3: Compendium Instructions

Welcome

You are about to participate in a decision-making experiment. If you follow these instructions, you can earn a considerable amount of money. Your earnings depend on your choices as well as the choices of others. In the end of the session, your profit, expressed in the experimental tokens, will be converted to US dollars at rate 600 tokens for one US dollar and added to the show-up fee of \$5. You will be paid in cash.

The Experiment

In this experiment you will be assigned into one of two roles -- a **manufacturer** or a **retailer**. Your role is currently displayed on the screen in front of you. Half the people in this room are assigned into each role and these roles will remain fixed for the duration of the experiment. The experiment will last 15 periods. In each period, a manufacturer and a retailer are matched. The matching will change every period.

The manufacturer

A manufacturer produces a product at *cost* 20 tokens per unit.

The manufacturer proposes a *wholesale price*, which is paid per unit, and a *lump sum*, which is to be paid by the retailer regardless of the number of units ordered.

If the retailer accepts the proposal, then the manufacturer will earn the number of units the retailer orders, multiplied by the (*wholesale price* – *cost*) plus the *lump sum* that the retailer pays.

If the retailer rejects the proposal, the manufacturer can make another proposal (unless negotiations reached “Impasse” stage; see below).

The retailer

The retailer evaluates the manufacturer’s proposal and determines if the proposal is acceptable.

If so, the retailer decides on the *order quantity* to buy from the manufacturer at the *wholesale price*.

The retailer may reject the manufacturer’s proposal and await a new proposal (unless negotiations reached “Impasse” stage; see below).

Participant Earnings

If a proposal is accepted, the retail price will be:

$$\text{Retail Price} = 100 - \text{Order Quantity}$$

The retailer’s profit is then:

$$\text{Retailer's Profit} = (\text{Retail Price} - \text{Wholesale Price}) \times \text{Order Quantity} - \text{Lump Sum}$$

The manufacturer’s total profit from the trade is:

$$\text{Manufacturer's Profit} = (\text{Wholesale Price} - 20) \times \text{Order Quantity} + \text{Lump Sum}$$

The Impasse stage

If no proposal is acceptable to the retailer within 5 minutes, the parties will enter an impasse stage. The **manufacturer** makes one final offer. The retailer can either accept the proposal or declare an impasse. If the retailer decides to declare an impasse, no trade takes place and each party earns zero profit in this period.

The Interface

Your computer screens, as shown below, have a “Calculate” button that allows both manufacturers and retailers to try different choice scenarios without actually submitting a decision.

Below is a screenshot of the calculator interface available to the manufacturer and the calculator interface available to the retailer.

Period
1 of 2

You are a **Supplier**

Please choose the wholesale price and the lump sum.

Enter the wholesale price Enter the Lump Sum Payment

Wholesale price	Lump Sum	Retailer's best order if accepts	Your Profit if the Retailer accepts	Your Profit if the Retailer rejects	Retailer's highest profit if accepts	Retailer's Profit if rejects
80.0	50	10.0	650.0	0.0	50.0	0.0

Period

1 of 2

You are a **Retailer**

The supplier offered you the following contract:
Wholesale price of 80.0 per unit and lump sum payment of 50 units.
Please choose your order quantity:

Enter your order

Reject Contract **SUBMIT** **CALCULATE**

Wholesale price	Your Order	Your Profit if You order this amount	Your Profit if You Reject	Supplier's Profit if You order this amount	Supplier's Profit if You Reject
80.0	10.00	50.00	0.0	650.00	0.0
80.0	2.00	-14.00	0.0	170.00	0.0
80.0	30.00	-350.00	0.0	1850.00	0.0
80.0	20.00	-50.00	0.0	1250.00	0.0
80.0	10.00	50.00	0.0	650.00	0.0

The manufacturer's interface allows manufacturers to try different combinations of *wholesale price* and *lump sum* without actually making an offer.

After the manufacturer submits an offer, it appears on the retailer's screen and the retailer can then try different *order quantities* before reaching a decision.

For a manufacturer, the calculator will assume the order quantity maximizes retailer's profits and show both parties' profits assuming that the retailer orders this amount.

You can use the "Calculator" button as much as you need. Whenever you are ready to submit your decision, click on the "Submit" button.

Example 1.

Suppose the manufacturer offers Wholesale price = 40 and Lump Sum = 333.

Suppose the retailer accepts the offer and orders 50 units. Then

$$\text{Manufacturer's profit} = (40 - 20) \times 50 + 333 = 20 \times 50 + 333 = 1000 + 333 = 1333$$

When the retailer re-sells these 50 units on the market, the market price is

$$\text{Retail Price} = 100 - 50 = 50.$$

The retailer's profit will be

$$\text{Retailer's profit} = (\text{Retail Price} - \text{Wholesale price}) \times \text{Order Quantity} - \text{Lump Sum}$$

$$= (50 - 40) \times 50 - 333 = 10 \times 50 - 333 = 500 - 333 = 167.$$

If no trade takes place

If the retailer does not accept the contract, each party earns nothing and the game proceeds to the next period.

Example 2.

Suppose the manufacturer offers Wholesale price = 80 and Lump Sum = 1000.

If the retailer accepts

Suppose the retailer accepts the offer and orders 10 units. Then

$$\text{Manufacturer's profit} = (80 - 20) \times 10 = 60 \times 10 = 600 + 1000.$$

When the retailer re-sells these units on the market, the market price is

$$\text{Retail price} = 100 - 10 = 90.$$

The retailer's profit will be

$$\begin{aligned} \text{Retailer's profit} &= (\text{Retail price} - \text{Wholesale Price}) \times \text{Order Quantity} - \text{Lump Sum} \\ &= (90 - 80) \times 10 - 1000 = 10 \times 10 - 1000 = 100 - 1000 = -900. \end{aligned}$$

Thus, for the retailer, accepting this contract and ordering 10 units results in a loss.

If no trade takes place

In fact, it turns out that at wholesale price = 80, the retailer's best order quantity should be 10. However, given the *lump sum* = 1000, it is simply not possible for the retailer to make profit with this contract. If the retailer rejects this offer, no trade takes place and each party earns zero profit.