

Working Paper

Bargaining in Supply Chains (Long Version)

Stephen Leider

Stephen M. Ross School of Business
University of Michigan

William S. Lovejoy

Booth School of Business
University of Chicago

Ross School of Business Working Paper
Working Paper No. 1259
March 2015

This work cannot be used without the author's permission.
This paper can be downloaded without charge from the
Social Sciences Research Network Electronic Paper Collection:
<http://ssrn.com/abstract=2538547>

Bargaining in Supply Chains (Long Version)

Stephen Leider

Ross School of Business, University of Michigan, Ann Arbor, MI 48104, leider@umich.edu

William S. Lovejoy

Ross School of Business, University of Michigan, Ann Arbor, MI 48104, wlovejoy@umich.edu

We study experimentally bargaining in a multiple-tier supply chain with horizontal competition and sequential bargaining between tiers. Our treatments vary the cost differences between firms in tiers 1 and 2. We measure how these underlying costs influence the efficiency, negotiated prices and profit distribution across the supply chain, and the consistency of these outcomes with existing theory. We find that the structural issue of cost differentials dominates personal characteristics in explaining outcomes, with profits in a tier generally increasing with decreased competition in the tier and increasing with decreased competition in alternate tiers. The Balanced Principal model of supply chain bargaining does a good job explaining our data, and outperforms the common assumption of leader-follower negotiations. We find a significant anchoring effect from a firm's first bid but no effect of the sequence of those bids, no evidence of failure to close via escalation of commitment, and mixed results for a deadline effect. We also find an interesting asymmetry between the buy and sell sides in employed bidding strategy. The buy side makes predominantly concessionary offers after the initial anchor, but a significant number of sell side firms engage in aggressive anti-concessionary bidding, a strategy that is effective in that it increases prices while not compromising closure rates.

Key words: bargaining, behavioral experiments, multiechelon supply chains, efficiency and profitability

1. Introduction

In the Operations Management literature, supply chain management as a field of study evolved from multi-echelon inventory and control theory via the recognition of the parochial interests of each firm in the chain. That is, the coordination of the activities of multiple independent firms to maximize total social value is not automatic, due to the private profitability interests of each. Firms recognize that by cooperating on material and information flows they can create value for society, but each also wishes to capture as much of that value as possible for their private use. The mechanisms by which this tension is managed vary broadly, from detailed legal contracting to more informal relationships. Managing this issue of social value versus private gain is central to supply chain management research. Arshinder et al. (2008) catalog representative papers in this area, and

Cochon (2003) reviews coordination through choosing the appropriate inter-firm contract form. These, and the references there, provide an overview of current perspectives and approaches.

Scholarly analyses of supply chains focus on issues of efficiency (are chain-wide profits maximized by the choices made by the independent firms?) and distribution (how are the chain-wide profits distributed along the chain?). The former is important from a social perspective (are resources appropriately allocated?) and the latter is important from a firm perspective (understanding the profitability consequences of alternative actions is necessary for decision makers). The answers to these questions remain unclear. Indeed, in many supply chain contexts of practical importance inter-firm negotiations can best be described as small numbers bargaining, an enduringly difficult yet fundamentally important economic context. Yet we will not really understand supply chains and their efficiency and distributional characteristics without understanding how inter-firm negotiations determine which firm(s) get the contract(s) and at what prices. This paper contributes to our body of knowledge by experimentally exploring these questions in the context of one common supply chain structure, but for which theory is new and no behavioral evidence yet exists.

Our supply chain setting

Supply chains formed for different purposes generally differ in their structural dependencies. For example, assembly operations in which an OEM must contract with at least one supplier for each of several components differ from retailers who can choose which portfolio of products to stock, and both of these differ from an OEM who contracts with a single tier 1 aggregator who in turn contracts with tier 2 supplier(s), etc. Our experimental set-up is designed to represent one common multi-tier supply chain context. An OEM designs a new product and wishes to bring it to market, but does not have ownership or control over all of the resources required to make that happen. The product is sufficiently new that the OEM is, at least temporarily, a monopolist vis-a-vis its customers. The OEM will send out a request for quote (RFQ) to multiple tier 1 suppliers. The tier 1 suppliers turn around and negotiate with their (tier 2) suppliers to get a better idea of their possible supply costs. The tier 1 - tier 2 negotiations end with an understanding of what they will do if they get the business from the OEM. Once their supply costs are known, the tier 1 firms respond to the OEM's RFQ, and then potentially further OEM - tier 1 negotiations ensue. The chain forms when the OEM selects a tier 1 supplier to do business with at an agreed upon price, and the tier 1 supplier selects a tier 2 supplier to do business with at their agreed upon price. We assume that there are sufficient economies of scale in supply that a single tier 1 supplier will be chosen by the OEM, and likewise a single tier 2 firm will be selected by the active tier 1 firm, so that the final chain consists of three firms (and implicitly the suppliers to tier 2, as described below).

This situation, with an OEM selecting a single tier 1 partner from several options via an RFQ and subsequent negotiations, and the tier 1 firms behaving similarly vis-a-vis tier 2, is common in practice when the downstream tiers in the supply chain are performing product-specific activities. Somewhere upstream in the chain (in our experiments this is after tier 2) inputs become more generic, competition more perfect, and firms can source inputs at something approximating a common market competitive price. This situation, with sole-sourcing downstream and competitive markets upstream, is representative of, but not limited to, the high tech, consumer products and services, entertainment, food, furniture, large complex engineered products, and automotive industries (see Lovejoy 2010a). The specific structure we analyze is shown in Figure 1. For ease of interpretation we label the OEM interfacing with the public as the *retailer*, the tier 1 firms as *manufacturers* and the tier 2 firms as *suppliers*, which intuitively signals the appropriate chain relationship of supplier to manufacturer to retailer. In our experiments there are two suppliers and two manufacturers, all with potentially different costs. c_s^1 and c_s^2 are the supply costs for suppliers 1 and 2, respectively (these include their costs of upstream supply and value adding cost). c_m^1 and c_m^2 are the value-adding costs for (tier 1) manufacturers 1 and 2, respectively. Their supply costs from tier 2 will be determined by negotiations. R is the revenue (net of any firm-specific value adding cost) at which the retailer can sell the (indivisible) finished item on the market. Like the manufacturers, the retailer's supply cost from tier 1 will be determined by negotiations. We assume complete information about value-adding costs, which is an abstraction in many applied settings but not unreasonable in others. Firms go to great lengths to understand their suppliers' costs, because that information helps them in negotiations. In practice these efforts include reverse engineering, cost modeling based on historical data, backing out component costs from competitors' published prices for different product configurations, direct inspection of suppliers, open books agreements, and other tactics.

Our contribution

This paper contributes to two broad categories in the supply chain and bargaining literatures, one regarding outcomes and one regarding the process by which those outcomes are reached.

In the first category, we use our experiment to ask the following questions: (a) will the efficient firms become active in the final contracts, (b) what will the distribution of profits be throughout the chain, and (c) what features drive this? To the authors' knowledge this is the first time these questions have been addressed in a three-tier supply chain experiment. Furthermore, existing theoretical models of supply chain bargaining (such as leader-follower models and the Balanced

Principal solution, described below) differ in their predictions about the impact of horizontal competition on the profit distribution, and we experimentally compare their predictions to outcomes.

In the second contribution category, we test several conventional wisdoms about the dynamics of negotiations, including framing and anchoring, escalation of commitment, concessionary versus anti-concessionary bidding, and deadline effects.

Additionally, we make a methodological contribution by studying *free-form* bargaining between firms in the supply chain. Most of the existing behavioral operations experiments that study supply chain contracting (e.g. Lim and Ho 2007, Loch and Wu 2008, Ho and Zhang 2008) use ultimatum style bargaining - however Haruvy et al. (2014) show that supply chain coordination improves substantially when firms can use a more flexible bargaining format. Unstructured bargaining is both more realistic and allows greater scope for structural factors such as cost differences to impact the resulting outcomes (e.g. a firm's ability to make and receive multiple offers over time provides greater opportunities to push the firms in the other tier to match offers).

Overview of our results

In the category of supply chain bargaining outcomes we find that supply chain efficiency is high in all cost treatments, but as is usual in bargaining experiments we find a higher level of non-closure than one would expect using purely rational economic reasoning. We find that horizontal competition significantly influences the distribution of profits within the supply chain. Generally, with minor exceptions, the profits that accrue to the eventually active firm in each tier will decrease as horizontal competition in that tier increases, and increase as horizontal competition in other tiers increases. Structural issues (e.g. cost structures) dominate individual characteristics (e.g. risk aversion) in determining these outcomes. The Balanced Principal solution explains the data quite well. The differences between treatments and the relative profits among tiers largely match its predictions, however the point predictions assign too much profit to the retailer. In predicting the outcomes of the retailer-manufacturer negotiations the Balanced Principal solution outperforms the assumption of either retailer leadership (i.e. assuming the retailer makes a "take it or leave it" offer) or manufacturer leadership. This lends further support to existing evidence that in small numbers bargaining situations, the popular leader-follower frameworks for analysis underperform relative to more bargaining-based frameworks that predict a less extreme distribution of wealth. If, despite this, one adopts a leader-follower model in a multiple seller, single buyer supply chain, we find the common agency format outperforms the more common designation of buyer as leader.

In the category of bargaining dynamics, we find a significant anchoring effect in that the first bids anchor the negotiations and the final price tends to be midway between the initial bids. However, we find no first mover advantage, so the first bid by each party in the negotiations matters, but not the sequence in which they make them. We see a significant deadline effect in stage 2 negotiations but much less so in stage 1, suggesting a more complex relationship to time in multi-party bargaining wherein some firms can be excluded. However, we find no relationship between when a deal is struck and final prices (which we would expect to see if impatience or anxiety-to-close put negotiators at a disadvantage). The actual sequence of bids leading to closure largely follows an intuitive path (anchor bid followed by concessions) on the buying side but not on the selling side, where anti-concessionary tactics are used to good effect. There also appears to be a psychological construct at play that grants the buyer a position of power unexplained by structural characteristics.

2. Experimental Design

We developed a laboratory Supply Chain Game with free-form bargaining to study how horizontal competition affects the efficiency and distribution of profits in a multi-tier supply chain. We studied a $2 \times 2 \times 1$ supply chain, consisting of two suppliers, two manufacturers and a retailer. Subjects were randomly assigned to one of these three roles, which they keep throughout the experiment. The retailer needed to establish a supply chain with one manufacturer and one supplier to bring a single unit of a good to the retail market. If the retailer could form a supply chain he received a fixed revenue R of \$40. Each supplier and manufacturer had a cost of \$5, \$15 or \$25 for their value-adding activities. In order to form a supply chain the retailer needed to negotiate a transfer price for a supply contract with a manufacturer, who in turn needed to have a supply contract with a supplier. Based on conversations with supply chain managers, we chose to have the supply chain game begin with the negotiations between the suppliers and manufacturers, followed by the negotiations between the manufacturers and the retailer (See Figure 1).

2.1. Supply Chain Game Preliminaries

Subjects played a total of six periods of the Supply Chain Game. In each period subjects were randomly and anonymously matched in groups of five, consisting of two subjects with the supplier role, two subjects with the manufacturer role, and one subject with the retailer role. Suppliers and manufacturers were randomly assigned to be Supplier 1 or Supplier 2 (Manufacturer 1 or Manufacturer 2) in each period. Subjects were also informed of each player's cost or revenue.

2.2. Supplier-Manufacturer Bargaining

The two suppliers and two manufacturers simultaneously bargained, with the manufacturers attempting to negotiate a supply contract with a supplier. Each manufacturer negotiated separately with each supplier until he reached an agreement with exactly one of them. Each manufacturer could contract with just one supplier, but it was possible that a single supplier could end up supplying both manufacturers. See the Appendix for a screenshot of the negotiation stage.

The subjects had 6 minutes in the first period (4 minutes in later periods) to negotiate.¹ They could make numerical price offers² at any time, and could also send chat messages. Only the recipient could see an offer or a chat message, although subjects were free to reveal that information using the chat window if they wished. An agreement was reached if a manufacturer accepted the last price offer from a supplier, or if a supplier accepted the last price offer from a manufacturer. However, subjects were required to wait until 2 minutes had elapsed before accepting an offer.³ If one manufacturer struck an agreement the other manufacturer could continue negotiating with both suppliers.

2.3. Manufacturer-Retailer Bargaining

At the conclusion of the Supplier-Manufacturer negotiations, all the subjects were shown the agreed upon transfer prices and the new total costs for each manufacturer. The manufacturers then negotiated with the retailer for 6 minutes in the first period (4 minutes in later periods). A manufacturer could only participate in this negotiation stage if he came to an agreement with one of the suppliers. As in the previous stage, subjects could make numerical price offers or send chat messages at any time, and could accept an offer after the first two minutes. An agreement was struck when a manufacturer accepted the last offer from the retailer, or the retailer accepted the last offer from one of the manufacturers.

At the conclusion of the second bargaining stage all five subjects were informed whether a complete supply chain was formed, which firms were part of the chain, and what the negotiated prices were between the retailer and manufacturer (p_{rm}) and between the manufacturer and supplier (p_{ms}).

¹ We gave subjects 6 minutes in the first period to allow them to get comfortable with the computer interface and the bargaining procedures. The 4 minute deadline in later periods was sufficient to allow most groups to negotiate without extensive time pressures while avoiding indefinite stalling.

² Subjects were only allowed to make or accept offers that would give them non-negative profits.

³ That is, during the first two minutes subjects can make price offers and send chat messages, but cannot accept an offer. We included this restriction based on earlier pilot sessions where we found that subjects would race to be the first to accept an offer rather than attempting to chat or negotiate. We felt that this time pressure was not reflective of typical negotiations, and was not our primary focus.

Each subject's period payoff was calculated based on the results of the Supply Chain Game. If a subject was not part of the final supply chain their period payoff was \$0. For subjects in the supply chain their period payoff was calculated as follows:

$$\text{Retailer: } \pi_R = R - p_{rm}$$

$$\text{Manufacturer: } \pi_M = p_{rm} - p_{ms} - c_m$$

$$\text{Supplier: } \pi_S = p_{ms} - c_s$$

2.4. Additional Tasks

After the Supply Chain Game subjects performed two additional tasks to measure individual risk and social preferences. The first task (based on Dohmen and Falk 2011) asked subjects to make fifteen choices between a 50-50 lottery between \$0 and \$4 or a fixed payoff that varied between choices (ranging from \$0.25 to \$3.75). The number of times the subject chose the fixed payoff provides a measure of risk aversion (where choosing the sure payoff more often indicates higher risk aversion).

The second task (based on Andreoni and Miller 2002 and Andreoni and Vesterlund 2001) involved five unilateral allocation decisions. For each decision subjects were asked to divide 50 tokens between themselves and another randomly selected anonymous participant. The five decisions differed in how much a token was worth to the allocator and to the recipient (ranging from \$0.05 to \$0.15). Based on the allocation decisions we can identify behavior consistent with several forms of social preferences, including selfishness, altruism, inequity aversion and social welfare maximization.

2.5. Experimental Treatments

We examined five different between-subjects cost treatments. In all cases the most efficient firms in a tier (i.e. Manufacturer 1 and Supplier 1) had a cost of \$5. In our Base treatment we set the second firms to have a cost of \$15 - hence they were at a \$10 disadvantage relative to the most efficient firms. We then varied the level of competitiveness within a tier by increasing to \$20 or decreasing to \$0 the cost difference in either the manufacturer or supplier tier. This yielded five cost profiles, described in Table 1.

2.6. Experimental Procedures

We ran 14 sessions at the University of Michigan during 2011-2012, with a total of 210 subjects participating⁴. Participants were Michigan undergraduate students. Sessions lasted approximately 1.5 hours. Subjects earned on average \$19.53.

⁴ We planned to run two to three sessions in each treatment to target approximately 40 subjects, and oversampled the Base treatment since it provides the baseline for both the manufacturer and supplier competition manipulations.

3. Bargaining Outcomes: Existing Theory and Experimental Literature

3.1. Efficiency

The supplier-manufacturer negotiations over supply partner and price are 2×2 (2 sellers, 2 buyers) bargaining situations, and the subsequent manufacturer-retailer negotiations are either 2×1 or 1×1 situations (it is 1×1 if one of the manufacturers does not secure a supply partner and so cannot enter negotiations with the retailer). Theory for this type of small numbers bargaining context is incomplete. Indeed, early economists including Edgeworth (1881), Marshall (1890) and Bowley (1928) all viewed these contexts to be indeterminate, because there are many different outcomes that can support (what would later be called) a Nash equilibrium. In contrast, Stigler (1942) and Fellner (1949) argued that an efficient solution would arise that maximizes total chain profits, but that the distribution of that profit along the chain is indeterminate. Their reasoning was that from any inefficient position there exists an alternative efficient solution that increases total chain profits and therefore offers a potential distribution of the additional surplus that can make everybody better off, stimulating its adoption. Coase's (1960) seminal arguments also support this efficiency claim. In general, with complete information what theory exists supports the expectation that with free-form negotiations chains with positive profit potential will form and the efficient firms will be active in the eventual supply contracts.

3.2. Distribution

The two generic approaches to predicting wealth distributions in small numbers bargaining contexts are noncooperative and cooperative game theory. The noncooperative approach generally seeks analytical tractability by specifying a detailed extensive form of the game (who talks, in which sequence?) and adopting Nash equilibria or a variant as a solution concept. This literature is dominated by models with two tiers only, and most often with one firm per tier (bilateral monopoly). Rather than extending simple models to multiple tiers and multiple actors per tier, research has focused on enriching bilateral models at the player level (for example, investigating the consequences of incomplete information or risk aversion). The authors know of no results in the non-cooperative literature, with either complete or incomplete information, that predict the distribution of profits in a more general supply chain with more than two tiers and more than one player per tier. One reason is that even for the simpler models there can be many possible equilibria with only extra-model rationales for preferring one over the other, and the nature of the equilibria tend to be very sensitive to the particular extensive form adopted in the analysis. This is disquieting given the less structured manner in which real negotiations appear to unfold.

Despite this, the noncooperative approach has advantages over the cooperative alternative and these likely explain the dominance of this approach in the extant literature. In a typical noncooperative model the sets of feasible actions for two players are distinct and declared in a predetermined sequence. The manufacturer, for example, may choose wholesale price and the retailer cannot contest that decision, or make a counter-offer. All the retailer can do is to choose actions in her feasible set, order quantity for example, in response to the price declaration from the manufacturer. This approach has some intuitive appeal, being familiar in many personal, and some industrial, purchasing situations. In addition, the noncooperative approach can feature analytical tractability via a commonly accepted approach and solution concept. For example, sequential or Stackelberg games, with a subgame perfect solution concept, are well regarded and familiar to researchers and readers alike. In contrast, cooperative games and bargaining theory have no such commonly accepted solution concept. We refer to sequential or Stackelberg extensive form noncooperative games as “Leader-Follower” or LF games. We note that these include the familiar principal-agent formats that inform the large literature on auctions (c.f Krishna 2002) and mechanism design (c.f Myerson 1981), both of which have presence and relevance in the supply chain literature.

However, it is often the case in the downstream stages of supply chains that negotiations ensue more along the lines of our experiment, with a small number of bargainers in each tier responding to an RFQ. In that setting, it is not clear who should, or can, act as the leader or from whence such powers would derive. To the authors’ knowledge, there is no LF theory that informs those 2×2 negotiations⁵. Only recently has there been a cooperative approach that makes predictions in this setting, as described below.

The lack of a theoretical prediction from the LF perspective for 2×2 negotiations prevents a comparison of our experimental results to noncooperative theory. However, the stage 2 manufacturer-retailer negotiations (which will feature either 2×1 for 1×1 negotiations) fall within a class of models that has been studied extensively from the LF perspective and which offers some predictions about the negotiated wealth distribution. Being designated the leader confers substantial power and influence over the outcome of a game (the leader can anticipate the follower’s reactions and craft her offers to exploit that anticipatory understanding). For example, with complete information a single retailer as leader facing multiple manufacturers (followers) would make a take-it-or-leave-it offer to the most efficient manufacturer, at that manufacturer’s cost of supply, extracting all of the possible value in the chain for herself. In general, complete information LF models with a single

⁵ see Pratt and Rusticini 2003 for results in a related but different setting, and the extensive literature on double auctions, for example Smith 1962, which initially evolved to study the multilateral auction-trading mechanisms representative of stock, bond and commodities exchanges.

leader predict extreme distributional outcomes. In our stage 2 setting, 2×1 or 1×1 negotiations with the retailer as leader will predict that the retailer takes all of the available wealth and no manufacturer makes any profit.

There is a variant of the LF paradigm that predicts less extreme outcomes, in which multiple leaders are allowed to make simultaneous offers to a single follower (an $n \times 1$ “common agency” problem as in Bernheim and Whinston 1986). In that case each leader’s aggression in negotiations is checked by competition from other leaders. With complete information the noncooperative solution is for the most efficient leader to make a take-it-or-leave-it offer to the single follower at a price equal to the indifference point of the next most competitive leader. So, in our stage 2 setting with two viable manufacturers as leaders facing a single retailer as follower, the noncooperative prediction would be for the efficient manufacturer to get the contract from the retailer, and to reap profits equal to the difference between his total costs and that of the less efficient manufacturer. If only one manufacturer is viable (for example if only one manufacturer closes a deal with a supplier) then the situation conforms to a standard 1×1 model in which the manufacturer as leader will take all of the available surplus. In stage 2 of our experiments we test the predictive power of both forms of LF model.

The notion that one player is given the powers of a leader due to unstated, extra-model assumptions is not part of the general cooperative approach. Bargaining models do not, in general, grant extraordinary power to any player in way that is not driven directly by the model parameters reflecting the bargaining context. For example, the more attractive the outside option (that can be embraced in the event of a breakdown of negotiations) a player has, generally the better he will be predicted to do in negotiations.

While there are no clean predictions for our 2×2 negotiations in the noncooperative literature, there are some in the cooperative literature. We use the umbrella term “cooperative game theory” broadly to refer to approaches to bargaining and negotiations that do not rely solely on assumptions of self-interest and Nash equilibria. Rather than proposing a detailed extensive form for interactions and attempting to predict their end-point through a detailed analysis of the give-and-take over time, cooperative approaches tend to propose sets of conditions that a solution “should” satisfy, and then look for outcomes that satisfy these. This allows for some intuitively compelling outcomes that may not arise from an LF approach. For example if two equally powerful players negotiate over the division of a fixed amount of money cooperative approaches can predict that they divide it evenly. This outcome has intuitive appeal and experimental support (see the extensive summary of the experimental literature in Kagel and Roth 1995), but is more difficult to get from a noncooperative

model (although, it can be done in sequential games, see below). However, there is currently a lack of a commonly accepted solution concept in the cooperative literature, and this may account for its weak presence in the current supply chain literature. Published solution concepts include core, Van Neumann and Morgenstern (VNM) sets, Shapely values, Nash bargaining and its generalizations, and the Balanced Principal (BP) solution (see, for example, the discussion in Lovejoy 2010b). The role of experiments will be critical in identifying which solution concept is appropriate for which setting.

In the stage 1 (suppliers and manufacturers) negotiations, the Balanced Principal (BP) approach is the only extant theory that is reasonable for the supply chain context and provides testable hypotheses. The BP solution is a refinement of both the core and VNM concepts. Generalized Nash bargaining and Shapely values are ill-suited to a context where one or more actors will be shut out of negotiations (that is, by their structure these solution concepts grant each firm some value, no matter how uncompetitive it is). The BP solution was derived by explicitly considering the supply chain negotiation context (Lovejoy 2010b).

We note, in closing, that there are other approaches to negotiations between two players or tiers, that we do not explore. These models can produce predictions that are sympathetic with bargaining outcomes. For example, Rubinstein's (1982) complete information alternating-offers model predicts, with symmetry of character between the two players and an infinite horizon, that they divide the surplus between them. Two-sided auctions tend to have multiple equilibria but some of these (with the proper mechanism for mapping bids into outcomes) can predict a non-extreme distribution of wealth. For example, Chatterjee and Samuelson (1987) analyze a simultaneous bid model with incomplete information, in which the players divide the spoils based on an exogenously specified weighted average of their bids. If either the buyer or seller is given full weight the model reduces to an LF model with a single take-it-or-leave-it offer being made by the advantaged player. When the weights are the same, a situation that has intuitive appeal when all else is equal, the distribution of wealth is naturally more even. Consequently, the results here are not intended to pass a final judgment on the noncooperative versus cooperative approaches to supply chain outcomes. Rather, we wish to better understand those outcomes and their determinants, and test the predictive power of several distinct alternative existing theories.

3.3. Experimental Literature

There is an extensive history of experimental investigations of bargaining in both psychology (see Rubin and Brown 1975 and Bazerman et al. 2000 for surveys) and economics (see Roth 1995 for a

survey) dating back more than 50 years. Siegel and Fouraker (1960) provide an early and extensive study of free-form bilateral monopoly bargaining, and study in particular the effect of information on bargaining outcomes. They find that increasing the information subjects have about each other's payoffs increases both the efficiency of the outcomes as well as the frequency of prices that yield equal payoffs. They conclude that insights from both economics and psychology contribute to explaining their experimental outcomes. Experiments also explored structured bargaining (e.g. Kelley et al. 1967) and multi-issue integrative bargaining (e.g. Pruitt and Lewis 1975), with contextualized free-form multi-issue bargaining tasks becoming a common paradigm.

Bazerman et al.(2000) identify two major themes in the social psychology of bargaining in the 1960s and 1970s: individual differences and structural variables. They conclude that the evidence suggests that individual differences have a small effect on bargaining outcomes, and are generally outweighed by situational features (Ross and Nisbett 1991, Thompson 1990, Thompson 1998). A number of structural variables have been shown to influence bargaining outcomes including representation of constituencies (Druckman 1967), bargaining deadlines (Pruitt and Drews 1969), and mediation (Pruitt and Johnson 1972). However, Bazerman et al. argue that the psychology literature moved away from these structural factors because the effects were typically consistent with naive intuition and because these objective features of a negotiation are often beyond the control of an individual negotiator. Instead, more recent studies of negotiations in the behavioral literature focus on decision biases in negotiations and how individuals construe the negotiating environment, and emphasize strategies and guidance for individual negotiators.

Prominent negotiation biases include increased concessions for positive versus negative framings (Bazerman et al. 1985, Neale and Bazerman 1985), anchoring effects (Northcraft and Neale 1987), overconfidence about obtaining favorable outcomes (Bazerman and Neale 1982, Neale and Bazerman 1985), falsely assuming that the possible surplus is fixed and therefore missing out on Pareto improving agreements (Thompson and Hastie 1990), falsely assuming that the parties' interests are in conflict (Thompson and Hrebec 1996), and having a self-serving bias in judging the fairness of outcomes (Thompson and Loewenstein 1992).

The early experimental economic literature also heavily used free-form bargaining, particularly to test the predictions of axiomatic bargaining theories (Nydegger and Owen 1975, Roth and Malouf 1979) and Coasean bargaining (Hoffman and Spitzer 1982), as well as examining structural factors such as deadlines (Roth et al. 1988) and individual characteristics such as risk aversion (Murnighan et al. 1988). Many experiments found that under free-form bargaining agreements it was very common for outcomes to equalize payoffs (under complete information) or tokens (under

incomplete information about payoffs), see Roth and Malouf (1979), Roth et al. (1981), and Roth and Murnighan (1982).

With the introduction of the Ultimatum Game by Güth et al. (1982) experiments examining structured sequential bargaining became popular as well, with follow up experiments by Binmore et al. (1985), Güth and Tietz (1988), and Ochs and Roth (1989) (also see Bearden (2001) and Falk et al. (2003) for surveys of the literature on the Ultimatum Game). Güth et al. found that extreme offers were uncommon, and were frequently rejected, while fairer offers were most common and were generally accepted. Ochs and Roth found that with two-period and three-period bargaining and across a range of discount rates the bargaining outcomes differ substantially from the LF perfect equilibrium outcome, with individuals often rejecting unfair offers (and frequently making disadvantageous counteroffers that gave the rejecter a lower absolute payoff, but a higher payoff share). Ochs and Roth conclude that many subjects will reject insultingly low offers, and that bargainers overall adapt their offers to account for these minimum thresholds. In a follow up paper Bolton (1991) replicates the four key patterns of Ochs and Roth (a first mover advantage, average offers deviating from the perfect equilibrium in the direction of equal division, frequent rejections, and frequent disadvantageous counteroffers), and uses them to formulate a model of fairness in bargaining. Similarly, experiments such as Neelin et al. (1988) that consider sequential bargaining with more rounds of negotiation also find significant deviations from the perfect equilibrium prediction. A number of experiments have examined multi-party bargaining, including bargaining in networks. For example, in the psychology literature Thompson et al. (1988) and Mannix et al. (1989) examine agendas, decision rules and power balances in a three party negotiation, while Polzer et al. (1998) study coalition formation. Fréchette et al. (2003, 2005a, 2005b) test experimentally multilateral models of legislative bargaining, including the effects of selection rules and bargaining power. As in the bilateral bargaining case outcomes were frequently more egalitarian than noncooperative theory predicted. Charness et al. (2007) study negotiations in buyer-seller networks of different architectures. They find that the resulting bargaining outcomes are broadly consistent with the theoretical predictions about network structure, and that individual payoffs can change due to the addition of a distant link. Chakraborty et al. (2010) also examine the influence of network structure, and again find that an individual's payoff can be influenced by distant features of the network.

4. Hypotheses and Results

The various treatments in the experiments were designed to test theory-based hypotheses about efficiency and the distribution of profits in a supply chain resulting from negotiations between

tiers in the chain. Most of the hypotheses are inspired by the Balanced Principal theory, because it is the only extant bargaining-based theory that provides whole-chain profitability predictions. Although the first stage of negotiations is conducted with some uncertainty regarding the outcome of the second stage of negotiations, the BP predictions (which are based on simultaneous joint determination of all transfer prices) are tested to see if they extend into this more complex, but realistic, context. There is no LF model that makes a clean prediction in the (stage 1) context of two principals and two agents (see the literature review for a discussion), or for whole chain profit profiles.

The second stage of negotiations is conducted with full information for all parties, and benefits from both BP and LF predictions. When applying LF models, the prediction will depend on who is designated the leader. The LF predictions will give either all surplus to the leader, or in the case of multiple leaders (a common agency problem) the efficient firm in the leadership tier will extract as much as the competition allows.

The BP solution to a multi-tier bargaining chain is unique if the difference in value-adding costs between the two most efficient firms in tier k , call this Δc_k , is nonincreasing as we move upstream in the chain (as k increases, see Lovejoy 2010b). In such solutions downstream firms for which the Δc is so large that the second-best firm is not competitive function as a sequence of monopolists who divide their available surplus equally, and upstream firms (where Δc is smaller, indicating real competition) have profits that vary with Δc in that tier. That is, the theoretical prediction in the complex setting of multi-tier bargaining aligns with intuition, the efficient firm in a tier will get the contract and make profits that vary with the level of cost competition (Δc) in its tier.

We use five treatments (cases) to feature different combinations of Δc in the manufacturer and supplier tiers, as shown in Table 1. In all cases the revenue to the retailer (net of value-adding costs) equals 40. This experimental structure necessarily includes cases where Δc increases going upstream (that is, cases where Δc in the supply tier is greater than Δc in the manufacturing tier, as in treatments MDiff0 and SDiff20) which admit non-unique BP solutions. However, it is still possible that all BP solutions share some testable attributes, as we describe below.

4.1. Supply chain formation

With complete information and positive potential profits current economic theory based on (potentially constrained) self-interested behavior would predict 100% closure, as discussed above. Theories in the industrial organizations and mechanism design literature can generate non-closure when there is incomplete information (c.f. Vickrey 1961, D'Aspremont and Gerard-Varet 1979, Myerson

and Satterthwaite 1983), and in fact in bilateral negotiations closure becomes less likely as the mean beliefs of the buyer and seller get closer to each other (that is, there is less believed surplus in the system). However, these theories would also predict closure in the limit as information becomes complete. Non-closure in experimental bargaining games (such as rejection in the Ultimatum Game) are generally explained through non-pecuniary motivations, such as concerns for fairness or spiteful punishment of selfish individuals.

The hypothesis that the probability of closure $p = 1$ is clearly rejected as soon as we have any non-closure, which we do have here. So, not all negotiations will close even when there is complete information and positive profits available. In our experiments a key contingency seems to be asymmetry of bargaining power, where in stage 2 of the negotiations 91% of the 2x1 chains close, but only 73% of the 1x1 chains close (this situation obtains when there is a breakdown in stage 1 negotiations, resulting in one of the manufacturers having no supplier under contract). Overall, 89% of all trials ended up with a complete supply chain (see Figure 2). So, closure is significantly more likely than non-closure in all cases, as expected. However, the probability of closure for the chain overall is significantly reduced when at least one of the manufacturers fails to close on a supply contract in stage 1. This cannot be because the failed manufacturer is intransigent, because that manufacturer is not active in the stage 2 negotiations. Rather, it appears to be the symmetry of power in 1×1 stage 2 negotiations, relative to asymmetry in 2×1 negotiations, that drives failure to close. Further analysis would be required to confirm this. For the remainder of this paper, all results (proportions, etc.) are computed for completed chains only.

4.2. Supply chain efficiency

By a repeat of the above arguments on closure, existing economic theory with complete information would predict that the efficient firms will be active in the final chain. If any one of these firms is excluded from a proposed contract, they can always make a more attractive offer to a member of their opposite (supplier or buyer) tier. In our multi-tier setting, however, there are two different ways to perceive efficiency. The first is ex-ante efficiency, which means that the low ex-ante (before stage 1 negotiations) value-adding cost firms are active in the final chain. After stage one negotiations, it is possible that a manufacturer with lower ex-ante costs must pay, due to poor bargaining skills, a higher input price from his supplier and has become the higher cost firm ex-post. Ex-post manufacturer efficiency would mean that the lowest (ex-post) cost manufacturer gets the contract from the retailer, ex-post whole chain efficiency means that the ex-post efficient manufacturer and supplier are both active in the final chain.

There is one additional difference between stage 1 and stage 2 negotiations, and that is the information available to the negotiators. Stage 1 negotiations between the manufacturers and suppliers feature partial (incomplete) information, because they do not yet know the outcome of the stage 2 negotiations and therefore do not know the true value to a manufacturer for getting the contract. In contrast, the stage 2 negotiations between the manufacturer(s) and retailer features complete information.

Ex-ante efficiency

Hypothesis (Ex-ante efficiency, suppliers): The ex-ante efficient supplier will be active in completed supply chains.

Hypothesis (Ex-ante efficiency, manufacturers): The ex-ante efficient manufacturer will be active in completed supply chains.

Hypothesis (Ex-ante efficiency, whole-chain): Both the ex-ante efficient supplier and ex-ante efficient manufacturer will be active in completed supply chains.

We test these hypotheses by looking for significant statistical evidence that efficiency occurs at a rate greater than one would expect from random formation. For example to test ex-ante efficiency among suppliers we let x denote the fraction of completed chains in which the ex ante efficient supplier is active and test $H_0 : x = .5$ against the alternative $H_A : x \neq .5$.⁶ For ex-ante efficiency in the whole-chain we adopt as our null hypothesis independent random selection of firms in each tier. In cases where these are unique (MDiff and SDiff both > 0) we would expect 1/4 of them to be globally efficient, so letting x denote the fraction of completed chains that contain the ex-ante efficient firms in both the supplier and manufacturer tiers we test the null hypothesis $H_0 : x = .25$ against the alternative $H_A : x \neq .25$.

The results by treatment and tier are shown in Figure 3. The ex ante efficient suppliers are chosen with at least 80% probability in all treatments, and ex ante efficient manufacturers are chosen with at least 75% probability. In all cases efficiency is significantly higher than random selection would imply ($p \leq 0.01$). Whole-chain efficiency is at least 60% in all treatments, significantly higher than with random selection ($p < 0.01$). There are no significant differences between treatments in the frequency with which the efficient supplier is selected. The MDiff20 treatment has significantly higher manufacturer efficiency than the other three treatments with different manufacturer cost differences⁷ ($p < 0.01$), and similarly has higher joint efficiency ($p < 0.01$).

⁶ We conservatively use two-sided tests throughout our analysis.

⁷ We do not use the MDiff0 treatment in any of the statistical tests for ex ante manufacturer efficiency, since every chain is efficient trivially. Similarly we do not test for the ex ante supplier efficiency of the SDiff0 treatment.

Ex-post efficiency

After the stage 1 negotiations close, the total cost (input price + value adding cost) is known for each manufacturer, giving the stage 2 manufacturers-retailer negotiations clean theoretical predictions. In stage 2 if there is just one viable manufacturer in a chain that closes, then efficiency is trivially assured. We exclude these trivial cases from our analysis, so all results for ex-post efficiency are based on negotiations that result in two viable manufacturers after stage 1. Noting that the ex-ante and ex-post efficient suppliers are the same, we test

Hypothesis (Ex-post efficiency, manufacturers): The ex-post efficient manufacturer will be active in completed supply chains.

We test this by letting x denote the fraction of chains, among those that eventually close and for which the stage 2 negotiations include two viable manufacturers, for which the ex post efficient manufacturer is active, and test $H_0 : x = .5$ against the alternative $H_A : x \neq .5$.

Hypothesis (Ex-post efficiency, whole-chain): Both the ex-post efficient supplier and ex-post efficient manufacturer will be active in completed supply chains. Let x denote the fraction of chains, among those that eventually close and for which the stage 2 negotiations include two viable manufacturers, for which both the ex post efficient supplier and manufacturer are active. We test $H_0 : x = .25$ against $H_A : x \neq .25$.

The results are shown in Figure 4. Both of the above null hypotheses can be rejected ($p < 0.01$), so the ex-post efficient manufacturer and ex-post efficient supplier-manufacturer pairs are significantly more likely to be chosen in all treatments. There were some differences among treatments. The MDiff20 case was significantly more likely to result in the efficient manufacturer being chosen relative to the other cases ($p < .01$ for all). This might be expected, since the efficient manufacturer has the strongest bargaining advantage in this case. Also, in the whole-chain test MDiff20 is modestly statistically different than MDiff0 ($p = .06$) and Base ($p = .05$). Interestingly, SDiff20 is not as impacting as MDiff20, attesting to the importance of the M-R negotiations in the second stage of bargaining.

In summary, we expect the ex-ante efficient firms to be active in the final supply chain, even though the first stage negotiations feature bargaining with incomplete information. In the second stage of negotiations, featuring complete information bargaining, we expect the more efficient supplier-manufacturer pair to be active. There were some significant treatment effects, with a strong manufacturer value-adding cost advantage being more highly related to efficient outcomes than a strong supplier advantage.

4.3. Effect of competition on profit distribution

Our experiments investigated how profits will be distributed along the supply chain, whether this depends on the level of horizontal competition in each tier, and whether the results conform to theory. Figure 5 plots the average profits earned by each tier of the supply chain in each treatment. It is clear that there is substantial variation in the profit distribution between the treatments. Retailers tend to consistently capture the largest share of the profit (earning between 40% and 60% of the total profit). Supplier and manufacturer profits vary widely across treatments, with suppliers earning between 10% and 40%, while manufacturers earn between 15% and 30%. Additionally, note that even when two firms in a tier have equal costs, and are therefore highly competitive, the winning firm still earns positive profits. Hence, it does not appear that firms are bargained down to their reservation profit.

Comparative statics

The only extant theory that predicts the profit distribution in multi-tier supply chains with horizontal competition is Lovejoy’s (2010b) Balanced Principal (BP) model. This model predicts the outcomes shown in Table 2. A bracketed interval means that the BP solution is not unique. In those cases, every profit profile between the endpoints is a BP solution, with the endpoints preserving a total surplus of 30 (profits sum to 30). So, for example in MDiff0 all (supplier, manufacturer, retailer) profit distributions $\lambda \times (0, 0, 30) + (1 - \lambda) \times (10, 0, 20)$ as λ ranges from 0 to 1, are BP solutions. As described in Lovejoy (2010b), nonuniqueness results when there is a “profit bottleneck” in the chain, which is a stage where the horizontal value-adding cost differential (Δc) is large relative to its downstream neighbor (for example, when Δc in the supply tier is larger than Δc in the manufacturing tier). In our experiment cases were Δc in the manufacturing tier is smaller than that in the supplier tier (that is, MDiff0 and SDiff20) can generate non-unique BP solutions. In contrast, when Δc is nonincreasing going upstream in the supply chain, the BP solution is unique.

Several predicted trends in tier-specific profits as a function of horizontal competition (in the same and alternative tiers) are robust to non-uniqueness issues. For example, retailer profits can be expected to be decreasing (or at least non-increasing) as Δc in the manufacturing tier increases, because the prediction interval for MDiff0 is everywhere (except for a single point) above the predictions for the Base and MDiff20 treatments. Table 2 shows the theory-driven predictions of tier-specific profits as function of horizontal competition. Our specific hypotheses (we test each claim using a Cuzicks non-parametric test for trends) and the results, are as follows.

- Supplier profits increase with less supplier competition: supported ($p < .01$).
- Manufacturer profits increase with less manufacturer competition: supported ($p < .01$).
- Manufacturer profits do not change with less supplier competition: not supported. Manufacturing profits decrease significantly as horizontal competition in the supply tier decreases.
- Retailer profits decrease with less supplier competition: supported ($p < .01$).
- Retailer profits decrease with less manufacturer competition: supported ($p < .01$).

We also tested the single “no prediction” outcome from theory by testing for any significant trend (in any direction) of supplier profits as a function of horizontal competition in the manufacturing tier. The result was no statistically significant effect on supplier profits as a function of manufacturing competition ($p = .4$).

The predictions of the BP theory are largely supported for tier-wise profits as a function of horizontal competition in the same and alternative tiers. The exception is the prediction of indifference in the manufacturing tier as a function of competition in the supply tier. The intuition behind the theoretical prediction is that because the efficient supplier can contract with both manufacturers during the first stage, the two manufacturers should be able to strike the same agreement with the supplier. Therefore, whether the supplier ought to be able to capture a large or a small profit, this should add the same amount to the cost of both manufacturers, and therefore the cost difference between the manufacturers should be preserved. However, our experimental outcomes differ from this prediction. Manufacturers make more than predicted (also see Table 6 below) when supplier competition is high in SDiff0 and Base but the same as predicted when competition is low in SDiff20, while Retailers make less than predicted with high competition. Suppliers, by contrast, make approximately the predicted amount in each case. Hence, the departure from theory in the effect of supplier competition on manufacturer profits is not a function of the supplier-manufacturer negotiations, but of the subsequent manufacturer-retailer negotiations where the retailer bargains less aggressively than predicted. In summary, the BP predictions for the effect of tier-wise competition are generally supported, the sole exception being where a remote tier drives the results and does so by bargaining less aggressively than anticipated.

Regression analysis

As a further test of the prediction that tier-wise profits (π) are driven by Δc_m and Δc_s , we ran three regressions with (supplier, manufacturer, retailer) profit as dependent variables and indicators for manufacturer and supplier cost differences being 10 or 20:

$$\pi = \alpha + \beta_1 I_{[\Delta c_m=10]} + \beta_2 I_{[\Delta c_m=20]} + \beta_3 I_{[\Delta c_s=10]} + \beta_4 I_{[\Delta c_s=20]} + \beta_5 Period + \epsilon$$

We included data from all five treatments together in the same regression. We included the period in the regression to account for any learning or fatigue effects over time. Recall that BP theory (see Table 2) predicts that supplier profits will increase as horizontal competition in the supplier tier decreases (Δc_s increases), with no prediction as a function of Δc_m ; manufacturer profits will increase in Δc_m and remain the same in Δc_s , and retailer profits will decrease as either Δc_m or Δc_m increase.

The regression results are shown in Table 3. The period has no significant effect, so there is no evidence of learning or fatigue effects over time. The profits to the various tiers behaved as follows.

Supplier profit: Supplier profits increase significantly as the horizontal competition in the supplier tier decreases (i.e. as SDiff goes from 0 to 10 and from 10 to 20). Supplier profit does not exhibit a monotonic pattern with respect to manufacturer costs - profits decline significantly compared to the base case if MDiff = 10, but by less and not significantly if MDiff = 20. These results are consistent with BP predictions (Recall BP had no prediction for the influence of MDiff due to nonuniqueness issues).

Manufacturer profits: Manufacturer profits increase significantly as manufacturer competition declines (i.e. profits increase with MDiff), as predicted by theory. Manufacturer profits decrease significantly when SDiff = 20 and decrease but not significantly when SDiff = 10. Recall that theory predicts constant manufacturing tier profits in these treatments as a function of SDiff, so theory is not challenged when SDiff = 10 but is challenged when SDiff = 20.

Retailer profits: Retailer profits decreased significantly with lower levels of horizontal competition in either upstream tier (with MDiff=20 yielding retailer profits significantly below the base case, but only directionally lower than MDiff=10). These outcomes are consistent with theory. Interestingly, retailer profits are hurt more by decreased competition in the supply tier than by decreased competition in the manufacturing tier (compare the effect of SDiff going from 10 to 20 to the effect of MDiff going from 10 to 20). This highlights the importance of competitive realities in remote tiers of a supply chain, and the importance of the stage 1, tier 1 to tier 2, negotiations to the profits of the retailer.

4.4. BP Prediction: Whole-chain profit profile

One of the predictions of the multi-tier BP theory is the relative distribution of profits throughout the chain, i.e. who gets the available surplus? Even in the non-unique cases there are profitability

patterns that are robust to the range of predicted outcomes. For example, in the MDiff0 case the predicted supplier profits are in the range $[0,10]$, the predicted manufacturer profit equals 0, and the predicted retailer profit is in the range $[20,30]$. All of these outcomes satisfy the inequalities $M \leq S < R$. We test this by considering the space of strict inequalities⁸, in which there are six possible events ($M < S < R$, $S < R < M$, etc.). We test the null hypothesis that the population proportion of outcomes with the BP predicted ordering significantly exceeds what one would expect from random outcomes. For example, let x equal the proportion of outcomes in the MDiff0 treatment with the predicted ordering $M < S < R$. We test the null hypothesis $H_0 : x = 1/6$ against the alternative $H_A : x \neq 1/6$. For the Base case and SDiff20 there is no strict three-tier ordering predicted by theory, but there are predicted pairwise orderings. For example for SDiff20 BP predicts $M \leq S$ and $M < R$. Again using a sample space of all possible 3-tier strict orderings, the two pairwise results in SDiff20 occur simultaneously in 2 of the six ordering, so the natural null hypothesis is $x = 1/3$ if the outcomes were random (and likewise for the Base case outcomes). Proceeding similarly for the remaining treatments, Table 4 shows the hypotheses and relative frequencies under random selection. The results of these tests are shown in Table 5.

The predicted profit profile occurred in approximately two-thirds of the outcomes for three cases (MDiff0, Base and SDiff0) - a significantly higher frequency than random ordering would predict ($p < 0.01$ for each). For MDiff20 and SDiff20, however, while the predicted pattern occurred directionally more often than the null the differences were not statistically significant. For MDiff20 the strong bargaining position of the Manufacturer seemed to disrupt negotiations in both stages; the retailer earned more than the manufacturer (as predicted) in only 52% of chains, and the manufacturer earned more than the supplier in only 63% of the chains. For SDiff20 there were a large number of cases where the supplier and manufacturer earned equal profits - as described in the footnote above this is consistent with the BP prediction but not accounted for in the conservative test. If we include these ties a majority of the outcomes match the BP order.

Point estimates for profitability

The BP predictions for profits to each tier, and the median experimental outcomes are shown in Table 6. While the median outcomes are consistent with BP's predicted relative profit profiles (as we expect given the results of the previous section), the point predictions are significantly different from the experimental outcomes in more than half the data⁹. In general, we find that

⁸ This is more conservative, and simplifies identifying the relevant null hypothesis. If we allow for ties where BP predicts, and use the same null hypothesis probabilities, all our results are the same except SDiff20 now matches the BP order 54% of the time ($p = 0.01$).

⁹ We use Wilcoxon rank-sum tests.

- Retailer profits are generally lower than predicted
- Manufacturer profits are sometimes higher than predicted
- Supplier profits are sometimes higher than predicted

As described above, retailers bargained less aggressively than theory predicts in the majority of cases, lowering their profits. This leaves more for the manufacturers and suppliers, and both do as well or better than predicted in all cases. We will return to this point below.

In summary, the BP predictions for the profit profile (the ordering of profits) across the three tiers of the supply chain were directionally supported in our experiments with statistical significance in 2/3 of the cases. The exceptions were MDiff20 and SDiff20 where strong bargaining positions in those tiers seemed to have a disruptive effect. In contrast, BPs point predictions for firm profits deviate from the data in many cases. This is due to a general feature of retailers bargaining less aggressively than theory suggests, leaving more surplus for the manufacturers and suppliers to divide among them. So the profile is maintained but upstream tier profits tend to be higher than predicted.

Horizontal competition is a significant driver of the profits in any tier. The regression results suggest that retailers are hurt by decreased competition in either upstream tier (losing between a fifth and a third of their profit), both manufacturers and suppliers are helped by decreased competition in their own tier (doubling their profits in the best case), suppliers show mixed effects of decreased competition in the manufacturing tier, and manufacturers are hurt by decreased competition in the supplier tier. All of these are consistent with theory except the last one, as discussed above.

4.5. A comparison of BP and LF supply chain models

There is no predictive LF model for our stage 1 (2×2) negotiations. However, in stage 2 we will have either a 2×1 or 1×1 model, for which both BP and LF solutions are available. Which model is more predictive of outcomes?

Comparisons of mean absolute errors

LF models need to identify a leader - i.e. who is endowed with the ability to make “take it or leave it” offers - and leadership can reside either in the manufacturing tier or the retailer tier. If we declare the retailer the leader (case RL), the subgame perfect equilibrium is for her to declare a price infinitesimally larger than the full cost of the (ex post) most efficient manufacturer, capturing all of the available value for herself. If we declare the manufacturing tier as the leader (case ML), there are two possibilities. If there is only one viable manufacturer (a 1×1 bargaining system)

the sole manufacturer will capture all of the value, leaving the retailer with nothing. If there are two viable manufacturers (a 2×1 bargaining system) we have a complete information common agency problem (Bernstein and Whinston 1986), which has the subgame perfect solution of the (ex post) efficient manufacturer getting the contract for a price equal to the full cost of his competitor, thereby capturing the full Δc_m for himself. The BP solution is for the efficient manufacturer to get the contract but get only one half of Δc_m for himself in the 2×1 case and half the total surplus in the 1×1 case. Figure 6a shows the mean absolute error between these model predictions and experimental results for manufacturer profits¹⁰.

Specifically, for each chain we computed the predicted manufacturer profit for BP and the two LF models (RL and ML) given the outcome of the first stage negotiations. We then computed the absolute value of the difference between the predicted and actual outcome. Finally, we took the mean of these values over all chains (the median errors are very similar). Visually, BP is the best model overall and RL is the worst. Rigorously, we used a pairwise, two tailed Wilcoxin signed rank test for the equality of medians with the following results. BP is more accurate than RL in all treatments ($p < 0.01$ for all); BP is statistically indistinguishable from ML in all cases except MDiff20, when BP is unambiguously better ($p < 0.01$); ML is better than RL ($p < 0.01$), except MDiff20 and SDiff20 for which ML and RL are indistinguishable. Hence, we find that the BP model is equal or better than any LF model in all treatments. Interestingly, using a standard principal-agent style of model with a retailer principal is the worst among these alternatives. If researchers want to adhere to LF models of some sort, these results suggest the common agency format (ML) is superior to the single-leader format (RL).

One complicating factor in comparing the models is that in between 15% and 35% of cases the ex post inefficient manufacturer is selected, which reduces the amount of available surplus. Neither BP nor either of the LF models predict this outcome, and hence looking at the profit levels can be a bit misleading in that one of the models may appear to have a small error not because it correctly predicted that a firm would receive a small portion of the full surplus, but because the whole surplus had shrunk. To account for this we also express the models' predictions as predicted profit shares, and compare the differences between the predicted manufacturer shares and the observed profit shares (shown in Figure 6b). This comparison gives a similar impression of the three models. The RL model is again the worst (with BP being significantly better in all five cases, and ML being better in four cases). BP is still arguably better than ML, although the comparison is closer:

¹⁰ Examining retailer profits yields the same comparison, except in cases where the inefficient manufacturer is chosen, reducing the total surplus from what the three models predict. To account for that we also examine the manufacturer's profit share, shown in Figure 6b and discussed below.

the three models are statistically indistinguishable in three cases, BP is much better than ML in MDiff20 (average error = 14% versus 40%, $p < 0.01$), while ML is slightly better than BP in SDiff0 (average error = 13% versus 18%, $p = 0.04$).

Best estimate for profits as a function of competitive context

We noted above that there is at least one situation where the manufacturer gets more than is theoretically predicted. We can generalize our prediction for the profit to the active manufacturer in stage 2 negotiations to $\alpha\Delta c_m$ in the 2×1 case, and $\alpha(r - c_1)$ in the 1×1 case, for some $\alpha \in [0, 1]$. The BP model sets $\alpha = .5$, an RL model would set $\alpha = 0$ and an ML model would set $\alpha = 1$. We know from above that the BP model is, generally, the best among these options, but we can also ask what the best α would be, based on our experimental results?

Table 7 presents the regression results for estimating the best-fitting α for our data. Columns (1) and (2) report the estimates for fitting α to manufacturer profit levels, while columns (3) and (4) fit manufacturer profit shares (to account for surplus changes due to selecting the inefficient manufacturer). Columns (1) and (3) fit a single α for all five treatments, while columns (2) and (4) fit a separate alpha for each treatment. Again, BP does quite well across these comparisons. The best-fitting α is not statistically different from the BP prediction of 0.5 for either profit levels or profit shares, and the point estimates are numerically very close. Similarly, the treatment-specific estimates of α are not significantly different from the BP assumption in two of the five treatments for both profit levels (MDiff0 and MDiff20) and profit shares (MDiff20 and SDiff20). In the other cases we find that the best fit gives somewhat more profit to the manufacturer - as we would expect from our previous analysis. By contrast, both the RL assumption of $\alpha = 0.0$ and the ML assumption of $\alpha = 1.0$ can be rejected for each profit measure, both overall and for each treatment.

In summary, BP is equal to or better than any LF model in all treatments using mean absolute error, and is equal to or better than any LF model in four of the five cases using predicted profit shares. If a researcher, for some reason, must use an LF model, an ML common agency choice is in most cases better than the more common RL choice. Estimating the best α shows that over all treatments the most predictive alpha is not statistically different from the BP choice of .5. Both of the LF choices ($\alpha = 0$ or 1) can be statistically rejected. The BP model does quite well as a benchmark theoretical model, and at most one might want to adjust the model slightly to give more profit to manufacturers.

4.6. Effect of individual (subject) differences

We next examine what effect (if any) negotiator characteristics have on the distribution of profits in a supply chain. In real supply chains these characteristics would be a combination of the values and culture of the firm and the individual negotiators, however in our experiments these are the same. We will refer to these as “firm” characteristics. We assess whether a firm’s risk attitudes or social preferences change the average level of its profits. We use the lottery choice task and the dictator allocation task (described in Section 2.4) to determine a subject’s preference type. We identify subjects as risk averse if they chose the “safe” option more than the median amount (8 out of 15 times). 29% of subjects were risk averse by this measure. We use the pattern of a subject’s choices across the five allocation decisions (with different token valuations) to identify four (not mutually exclusive) characteristics: selfishness, altruism, inequity aversion (“fairness”), and social welfare maximization (“efficiency”). Subjects were identified as being in the (approximate) top quartile of each preference type in the following manner:

1. Selfish: Subjects kept the full endowment in all decisions (30% of subjects)
2. Altruist: Subjects kept no more than 138 tokens in total (25% of subjects)
3. Fairness: Subjects allocations had a total absolute payoff difference of no more than \$7 (25% of subjects)
4. Efficiency: Subjects allocations created a total surplus of more than \$20 (17.5% of subjects)

To test for the effect of individual characteristics on the distribution of profits in the supply chain, we regress firm profits on the cost variables for the treatment, as well as an indicator variable X_i for each firm i in the supply chain that denotes whether the firm has the individual characteristic. Tables 8a-8c report the results of regressing supplier, manufacturer and retailer profits on indicator variables for each of these five individual characteristics (with each specification including indicators for one characteristic).

$$\pi = \alpha + \beta_1 I_{[\Delta c_m=10]} + \beta_2 I_{[\Delta c_m=20]} + \beta_3 I_{[\Delta c_s=10]} + \beta_4 I_{[\Delta c_s=20]} + \beta_5 X_r + \beta_6 X_m + \beta_7 X_s + \beta_8 Period + \epsilon$$

While we find a few significant coefficients, there do not appear to be overall consistent patterns. For example, while a manufacturer’s profits decrease with their own risk aversion, a retailer’s profits are not affected by risk aversion, while a supplier’s profits increase with the *retailer’s* risk aversion. Social preferences appear to matter mostly for the manufacturer-supplier relationship, with profits decreasing when the other party is selfish and increasing when the other party is fair. Furthermore, it is important to note both that we are testing a large number of coefficients, and that the estimated effects are small compared to the effects of competition (which are as much as six times as large). Hence, it appears that the effect of individual characteristics are small compared to structural factors (consistent with the social psychology literature discussed in section 3.3).

5. Bargaining dynamics

We have so far analyzed bargaining outcomes and which factors are important in driving those outcomes. Our data also allows us to investigate aspects of “bargaining dynamics,” participant behaviors during the bargaining process that affect outcomes (closure rate, the timing of closure and/or the negotiated prices). There is a rich literature on bargaining dynamics, see Kagel and Roth (1995), Bazerman and Neale (1992), and Bazerman and Moore (2013) and reference there. We have already discussed (see section 4.1 above) failure to close, which has more presence in our (and other) experiments than would be predicted by theory. Here we look at some additional issues suggested by the literature: anchoring and the efficacy of bidding first, escalation of commitment as a failure mode, deadline effects on the timing of offers, and the effect of impatience. We also look at the consistency of the intuitive process of converging via alternating concessions with our observed data, and we note some interesting behaviors that arise in our asymmetrical, multi-lateral setting that one would not see in bilateral negotiations.

Throughout we relate bids and results to the “Zone of Possible Agreement” or ZOPA, which is defined as the available surplus in a negotiation. We will refer to the retailer and her revenue as R , to suppliers 1 and 2 and their costs as $S1$ and $S2$, respectively, with the cost of $S1 \leq S2$, and to manufacturers and their costs as $M1$ and $M2$ with value adding costs for $M1 \leq M2$. The ZOPA in stage 1 is the interval $[S1, R - M1]$. The ZOPA in stage 2 is the interval from the lower of the two manufacturer’s supply costs (after contracting with suppliers) to the revenue R . We will refer to a ZOPA interchangeably as being the relevant interval and/or its width, relying on context to aid interpretation.

5.1. Some summary statistics and observations

In the first bargaining stage between suppliers and manufacturers there were on average 5.5 to 7.6 offers involving Manufacturer 1, and between 5.3 and 7.6 offers involving Manufacturer 2. In both cases there are significant differences between treatments (Kruskal-Wallis test: $p < 0.01$ and $p = 0.02$), with $SDiff0$ having the most offers in both cases. In the second stage of bargaining between manufacturers and the retailer there were between 5.8 and 6.8 offers, with no significant differences between treatments (Kruskal-Wallis test: $p = 0.90$).

In the first stage bargaining between suppliers and manufacturers (see Table 9) the buying tier (Manufacturers) asks for the substantial majority of the available surplus (between 71% and 79% by $M1$ and between 58% and 84% by $M2$), while the selling tier asks for less (between 21% and 51% and between 1% and 21%, respectively, for $S1$ and $S2$ negotiating with $M1$; between 29% to

71% and between 0% and 29% for S1 and S2 negotiating with M2), with S1 bargaining with M2 the only case where the supplier asks for more than the manufacturer. Interestingly, in many cases the supplier asks for approximately one third the overall surplus, consistent with an equal split between the three tiers. These initial offers are close together; those involving Manufacturer 1 leave between 5% and 25% of the ZOPA contested, while initial offers involving Manufacturer 2 leave between 6% and 29% contested. These initial offers tend to anchor the negotiations (investigated further below), with final accepted offers tending to be near the midpoint of the original offers.

The offers are generally consistent with competitive realities for the buyer and stronger seller. For example, M1's most aggressive offers are in SDiff0 and least aggressive in MDiff0, and S1's most aggressive offers are in SDiff20 and least aggressive in SDiff0. However, the weaker seller follows a less "rational" approach, being almost as aggressive in SDiff20 as they are in SDiff0.

We see similar behaviors in the second stage bargaining between Manufacturers and the Retailer (see Table 10). Again, the buyer's (Retailer's) first offers are quite aggressive, claiming between 70 and 88% of the ZOPA, with the most aggressive offers in MDiff0 and least aggressive in MDiff20. By contrast, M1s initial offers tend towards approximately equal splits. Also, sellers are acting rationally given competitive realities, with M1's aggressiveness increasing with the difference in real costs and M2's aggressiveness decreasing in that difference.

Again the initial offers tend to anchor the negotiations, so subjects are only bargaining over between 20 and 40% of the remaining contracting space. Table 11 reports the average offers of each firm in stage 2, as well as the average accepted offer. Both manufacturers and retailers on average make relatively small concessions compared to their initial offers, and accepted agreements are approximately halfway in between the initial offers of the Retailer and Manufacturer 1 (as well as being at approximately the midpoint of the average offers). With the initially more aggressive offers by the Retailer, the average agreement gives the Retailer between 55% and 75% of the possible surplus.

5.2. Bargaining dynamics in the literature and in our experiments

Here we further investigate some specific bargaining dynamics issues discussed in the literature and testable with our data.

Framing and anchoring

Final agreements tending to be midway between the initial offers can be interpreted via "framing" and "anchoring". Framing refers to how a negotiator perceives or interprets a deal. This goes

beyond the actual price paid and puts that price into a more complex cognitive context. There are a number of ways that negotiations can be framed. An individual may evaluate an offer as favorable or unfavorable based on their prior experience of accepted offers, for example. Alternately, they may consider structural factors such as the total available surplus along the chain, any individual firm's costs, or non-structural inputs like the initial bid. In this subsection we focus on that final way to frame a negotiation, it being anchored at a specific price so that any outcome is perceived not on its own terms, but relative to the anchor price.

Kahneman and Tversky (1981) note that the same outcome can be perceived as a gain or a loss, depending on the reference point or anchor it is compared to, and that these effects can influence behaviors in negotiations. So, it is often recommended in negotiations that one offer a first bid that establishes an aggressive position, because that establishes an anchor, and then anything less (even modestly) is perceived by the other as a favorable outcome. Galinsky and Mussweiler (2001) show that whoever makes the first offer does better, but that this advantage is reduced if the opponent focuses on the structural aspects of the negotiations rather than that anchor bid. There is also an argument that in a context of incomplete information one should let the opponent bid first, because it will reveal some information (after correcting for the anticipated aggressiveness). None of these conventional wisdoms have been tested in our supply chain context.

Table 12 shows the results of regressing the final price paid against the initial bids of each firm (as well as the total costs for Manufacturers 1 and 2 in the second stage negotiation). In stage 1 negotiations between M1 and the suppliers we find a significant impact of the initial offers from both the buyer and the low cost seller on final prices, and buyer and seller initial offers are equally influential on final price. For stage 1 negotiations between M2 and suppliers we see significant but unequal influences of the initial offers by the manufacturer and low cost seller.

In stage 2 we again find a significant and equal impact of the initial offers from both the buyer and the low cost seller on final prices. In the second stage negotiation there is also an effect of the manufacturer's total costs, however this may reflect actual bargaining power rather than anchoring, since the coefficients change substantially when we include treatment dummies.

So, at least for the advantaged player within a tier, one's first bid has a significant effect on outcome. But, is it advantageous to bid first? Table 13 reports the results of regressing the agreed upon price on an indicator for the buyer making the first offer. We find no significant effect on the final price, and the estimated coefficient in the specification with treatment controls is of small magnitude. Hence we find no first mover advantage or disadvantage.

Escalation of commitment

One source of inefficiency in negotiations is a failure to close even when closure can be beneficial to both parties and to society, and we see failure to close more frequently in experiments than would be anticipated by theory. One contributing mechanism (see the discussion in Bazerman and Moore 2013) could be “escalation of commitment” by which we mean the (seemingly irrational) allegiance to an initial position, once one has signaled that that position is firm. Escalation of commitment can result in failure to close if one party refuses to move off of his or her initial, but untenable, position. In our context, maintaining an initial position would be signaled by refusing to retreat from an extreme initial offer, even if that meant non-closure.

To test for this, we construct for each firm a measure of their “concession” - specifically the difference between their initial and final offer (measured as a percentage of the ZOPA). A larger concession then indicates that a subject changed their offer by a larger amount compared to their initial offer. We then regress an indicator for a closed agreement on the concessions of each firm, as well as the initial offers. The results are reported in Table 14. In columns (1) and (2) we report the results for the stage one bargaining where Manufacturer 1 is the “buyer” and the suppliers are the “sellers,” while columns (3) and (4) report the corresponding results with Manufacturer 2 as the “buyer.” Columns (5) and (6) report the results for the second stage bargaining with the Retailer as the buyer and the manufacturers as the sellers. Concessions generally have the anticipated directional effects but are either not significant or marginally significant in both stage 1 and stage 2 bargaining. That is, we see no evidence of closure inefficiencies due to escalation of commitment in our negotiations.

Convergence via concessionary offers

One intuitive impression of negotiations is of parties starting at different initial offers and then making a sequence of concessions as the negotiations converge to an agreement. We examine how often subjects make improving (i.e. concessionary) offers. Recall that subjects were not required to make monotonic offers - they could choose to rescind a generous offer and replace it with a less generous offer. However, intuitively we might expect bargainers to take turns making concessions. We find that this is largely true for both bargaining stages for firms on the “buying” side (i.e. Manufacturers in the first stage, and Retailers in the second stage). Almost all Manufacturer offers in the first stage are price increases (89% of offers for Manufacturer 1, 84% for Manufacturer 2), and almost all Manufacturers offer only price increases (79% for Manufacturer 1 and 74% for Manufacturer 2). In the second stage 75% of follow-up offers by Retailers are price increases, and 86% of Retailers who make more than one offer only make increasing price offers.

However, for firms on the “seller” side of the negotiation we see a more nuanced story. Most of the subjects on the selling side only make price decreasing offers (Stage 1: 77% of Supplier 1s and 78% of Supplier 2s; Stage 2: 73% of Manufacturer 1s and 73% of Manufacturer 2s). However, only about half of the offers are price decreases (Stage 1: 51% of Supplier 1 offers, 45% of Supplier 2 offers; Stage 2: 52% of offers from Manufacturer 1, 50% of offers from Manufacturer 2). This is because the non-concessionary seller firms make a substantial number of non-concessionary offers (55% and 42% for S1s and S2s in stage 1; 48% and 36% for M1s and M2s in stage 2). This is in contrast to non-concessionary buyer firms, who still mostly make concessionary offers (25% and 21% for M1s and M2s in stage 1; 14% for Retailers in stage 2). Hence across all four of the negotiations about one quarter of the subjects in the seller role are using a very aggressive bargaining strategy - making approximately as many anti-concessionary offers as concessionary offers.

Can we go further and identify the conditions under which a supplier firm uses anti-concessionary offers? In stage 1 negotiations there is no monotonic pattern by MDiff for whether suppliers are more likely to be anti-concessionary (trend with MDiff: $p = 0.97$). However, there is a significant trend with SDiff, with suppliers being more likely to be anti-concessionary when the difference in cost between suppliers is larger (SDiff0: 7%, Base: 26%, SDiff20: 19%, $p < 0.01$). In stage 2 there is no overall pattern for whether M1 or M2 is more likely to try being anti-concessionary ($p = 0.73$). There are some differences treatment by treatment, but the direction of the difference varies to an extent that no clear claim can be made. In summary, we observe more anti-concessionary behavior from the sell side than the buy side, and potentially when there is a significant cost disadvantage on the sell side, but that latter claim is speculative and requires more research.

What effect do these aggressive bargaining strategies have on outcomes? Table 15 reports the results of regressing the final price on an indicator for anti-concessionary behavior. In columns (1) and (2) this indicator equals 1 if the sell-side firm makes at least one anti-concessionary offer, while in columns (3) and (4) the indicator equals 1 if the buy-side firm makes an anti-concessionary offer. Columns (1) and (3) use Stage 1 bargaining, while Columns (2) and (4) use Stage 2 bargaining. We find that in both stages when the sell-side firm makes anti-concessionary offers this significantly increases the final price by more than a dollar. There is no corresponding change in final price when the (stronger) buy-side firm is anti-concessionary. We also run similar regressions testing for an effect on closure rates. For buy-side firms we find evidence that anti-concessionary offers reduce closure rates in first stage negotiations ($\beta = -0.160$, $p < 0.01$), but do not find a similar reduction in closure rates in the second stage ($\beta = 0.163$, $p = 0.16$). For sell-side firms we find no significant effect on closure in either stage, and in both cases the coefficient is very close to zero (Stage 1: $\beta = -0.048$, $p = 0.38$; Stage 2: $\beta = 0.035$, $p = 0.63$). Hence using an aggressive strategy

of making anti-concessionary offers is clearly beneficial for the sell-side firms: it raises the average price without leading to more impasses.

Deadline effect

Roth, Murnighan and Shoemaker (1988) observed a high concentration of agreements in the final 30 seconds of their bargaining experiments. Overall slightly less than half of all agreements occurred in the final 30 seconds of their experiments. The most common interpretation of this is that negotiators will probe for advantage as long as there is time remaining, but realize more acutely the need to close as time is expiring. Those results, however, were for bilateral bargaining where any one party could unilaterally keep the negotiations open. In our setting firms face the possibility of being closed out of the deal, which may inject pressure to close sooner rather than delay. Do we see a deadline effect in our setting?

The answer is yes, with some caveats. Agreement times during both stages are bi-modal, with a large fraction of agreements struck shortly after agreements become possible at 120 seconds, and then (typically) another mode toward the deadline. In stage 1 the deadline effect is not strong, with a significant difference between the last 30 seconds and previous 30 seconds in only two treatments for agreements involving Manufacturer 1 (Base: $p = 0.04$, SDiff20: $p < 0.01$) and only one treatment for Manufacturer 2 (MDiff20: $p < 0.01$). Hence, anxiety-to-close may be significant in stage 1.

However, there is a substantial deadline effect in stage 2. Figure 7 shows the distribution of agreement times in periods 2 to 6 (so as not to mix period lengths, see footnote 1). In each treatment except SDiff0 at least 25% of agreements are made in the last 30 seconds, with more than 50% of agreements closing in the last half minute in the SDiff20 treatment. This is between twice and five times as many agreements closing than in the previous 30 seconds. A test of proportions shows a significant increase in these four treatments (MDiff0: $p = 0.02$, Base: $p = 0.04$, MDiff20: $p = 0.06$, SDiff20: $p < 0.01$).

Timing of closure

One feature of our bargaining context (relative to simple bilateral bargaining) is the pressure to close a deal so as not to be left out (only one manufacturer is chosen by the retailer, and one supplier by each manufacturer). One might anticipate that this context would generate faster closure times, consistent with an anxiety-to-close interpretation driven by competitive realities. In addition it is intuitive, and a feature of some analytical models (c.f. Muthoo 2002 and references there) that impatience (as a personality trait) is disadvantageous in negotiations. If a negotiator feels that he or she has to reach a deal sooner rather than later, whether due to personality traits or structural

reality, she is unlikely to get the most out of the negotiations. We cannot disentangle personality traits from structural characteristics in driving a desire for earlier closing times, but we can test whether early closing times are related to structural characteristics and/or outcomes.

In stage 1 the earliest closing treatments for both manufacturers were MDiff0 (M1: 133 seconds, M2: 137 seconds) and MDiff20 (M1: 131 seconds, M2: 143 seconds). The latest closing treatment for Manufacturer 1 was the Base treatment (155 seconds), while the latest closing treatment for Manufacturer 2 was SDiff20 (164 seconds). In stage 2 negotiations we find that in periods 2 through 6 (with four minute bargaining periods) the average closure time for the Base (175 seconds) and SDiff20 (179 seconds) treatment is significantly later than for the MDiff0 (158 seconds) and SDiff0 (147 seconds) treatments ($p = 0.07$ and $p = 0.01$ for Base versus MDiff0 and SDiff0, $p = 0.05$ and $p = 0.02$ for SDiff20 versus MDiff0 and SDiff0). The MDiff20 treatment (159 seconds) is not significantly different from any treatment. These do not consistently line up with a structural interpretation. We would expect the pressure to close would be greatest when the cost differences are smallest. However, if we regress closing times in stage 2 on the absolute difference in total cost between the two manufacturers, while the sign is positive (as anticipated) it isn't significant (beta = .815, SE = .563, $p = 0.170$). If we include treatment dummies the effect of the cost difference is positive and significant (beta = 1.996, SE = .812, $p = 0.028$). However, the treatment dummies for closure time are still fairly large so the pressure to close isn't fully explaining the treatment differences. In summary, we see no consistent relationship between closing time and structural characteristics in our experiments.

Regardless of what is driving closing time, is it related to outcomes? Table 16 reports the results of regressing the share of the ZOPA that the buying firm earns under the agreement on the time of the agreement, where we track separately cases where the buying firm accepts a selling firm's offer and where a selling firm accepts a buyer's offer. We find no evidence for an impatience or anxiety effect for either firm - early agreements and late agreements yield very similar outcomes. We also find similar results using price rather than surplus share as the dependent variable. So, even if there was some feature driving a desire for earlier closing times, they do not appear to significantly affect outcomes.

5.3. Bargaining dynamics discussion

As is usual in bargaining experiments we find a higher level of non-closure than one would expect using purely rational economic reasoning. We find no evidence that escalation of commitment is the reason, but otherwise are unable with our data to identify the cause. The literature suggests issues of fairness and/or emotional concerns not typically part of economic models.

We find a significant anchoring effect in that the first bids anchor the negotiations and the final price tends to be midway between the initial bids. However, we find no first mover advantage. So, each party should consider their initial bid carefully but it does not matter in which sequence negotiators offer their initial bids.

We see a significant deadline effect in stage 2 but much less so in stage 1 negotiations, suggesting that with horizontal competition on both sides the desire to bargain to the end is tempered by a countervailing desire to close. However, we find no relationship between when a deal is struck and final prices (we would expect such a relationship if impatience or anxiety-to-close put negotiators at a disadvantage).

The actual sequence of bids leading to closure largely follows an intuitive path on the buying side (manufacturers in stage 1, retailer in stage 2), featuring an initial anchor bid and then concessions (price increases for buyers) leading to eventual closure. However, 25% of sell side firms use an aggressive strategy of making as many anti-concessionary offers as concessionary offers. Further, this is effective because it raises the final price without reducing the closure rate. There is no symmetrical benefit for the buying side. It may be that sell sides with significant cost differentials engage more in anti-concessionary tactics, injecting some irrationality into a context where rational analysis puts them at a disadvantage, but our data cannot resolve that definitively.

Another interesting buy-side versus sell-side asymmetry appears in our data. The original anchor bids and the lack of extreme reactions to them clearly indicate that both sides perceive the buy side to be in a stronger position than the sell side. The buy side demands a larger fraction of the possible surplus than the sell side, and that remains true through closure. This makes sense in the second stage were the monopolist retailer really is in a strong position. However, the justification for this behavior is not clear in the first stage of negotiations in MDiff0 and SDiff20 where the low cost seller should be in strong position. There appears to be a psychological construct at play that grants the buyer a position of power unexplained by the costs alone. We cannot diagnose this further with our data, but can suggest some possibilities. Stage 1 has the feature that no party knows how the retailer will behave, and the buyer has to carry that burden into stage 2. It may also be that in the more complex multilateral context of stage 1, boundedly rational negotiators fall back to intuitive, but not necessarily relevant, framings that give the buyer power they do not structurally deserve. Either way, negotiations in realistic supply chain contexts open up new questions that one may not encounter in more conventional bilateral bargaining experiments.

6. Summary and conclusions

Central questions in the study of supply chain performance are those of efficiency (is the total profit maximized?) and distribution (who gets the potential profits in the chain, driven by what contingencies?). In this paper we study these and other questions in the context of a three-tier, $2 \times 2 \times 1$ (suppliers - manufacturers - retailer) supply chain with varying levels of horizontal competition in the manufacturer and supplier tiers. Bargaining unfolds in a manner sympathetic to many real supply negotiations between a market-facing firm and its tier 1 suppliers, and between tier 1 and tier 2 suppliers. Despite its sympathy with industrial reality, to the authors' knowledge this is the first experimental study of a supply chain with more than 2 tiers and horizontal competition within tiers.

We find that chains form with high probability and supply chain efficiency is high across all treatments. Profits are influenced by the degree of horizontal competition in each tier, in a manner that is largely consistent with the Balanced Principal (BP) model of supply chain negotiated prices. Specifically, profits in a tier will generally increase with less competition (higher Δc) in the tier, and decrease with less competition in other tiers. Profits to a firm can depend significantly on competitive realities in a remote tier. Deviations from predicted outcomes tended to be in the direction of more equitable distributions of wealth, where (the retailer primarily) does not extract all of the value that she theoretically could demand.

There is no alternative theoretical prediction for our stage 1 negotiations known to the authors, but stage 2 negotiations can benefit from leader-follower (LF) as well as BP predictions. The BP model outperforms both types of LF model (with either the retailer or the manufacturing tier in the role of leader). If we restrict attention only to LF models, declaring the manufacturing tier as leader is best. This is interesting, given the tendency in the literature in $n \times 1$ models to declare the 1 player the leader. Our results suggest that the common agency approach (with the n -firm tier as leader) is more predictive of actual outcomes. Neither, however, is as effective as the BP model.

The BP model predicts that when there are two viable manufacturers negotiating with a single retailer, the efficient manufacturer will get the contract and enjoy profits equal to $.5\Delta c_m$. We generalize this to assuming profits equal to $\alpha\Delta c_m$ and estimate the best α from our experimental data. Pooling all treatments, α is not significantly different from .5 (its best fitting estimate equals .57), but in specific treatments the best α can differ significantly from .5. When this happens it is in the direction of more equitable distributions of wealth (the retailer does not bargain as aggressively as predicted).

The cost profile in the chain dominates personal negotiator characteristics (such as risk aversion, altruism, etc.) in influencing outcomes. This is consistent with current intuition that structural characteristics dominate interpersonal differences in these settings. However, the bargaining strategies that individuals use have significant impact on bargaining outcomes. The initial offers that both the buying and selling firm make significantly affect the final price agreement. Additionally, sell-side firms who choose to make non-concessionary offers achieve significantly higher prices.

References

- Andreoni, J., J. Miller. 2002. Giving According to GARP: An Experimental Test of the Consistency of Preferences for Altruism. *Econometrica*. **70**(2) 737-753.
- Andreoni, J., L. Vesterland. 2001. Which is the Fair Sex? Gender Differences in Altruism. *The Quarterly Journal of Economics*. **116**(1) 293-312.
- Arshinder, K., A. Kanda, S.G. Deshmuch. 2008. Supply chain coordination: Perspectives, empirical studies and research directions. *International Journal of Production Economics*. **115** 316-335.
- Bazerman, M.H., M.A. Neale. 1982. Improving Negotiation Effectiveness Under Final Offer Arbitration: The Role of Selection and Training. *Journal of Applied Psychology*. **67** 543-548.
- Bazerman, M.H., T. Magliozzi, M.A. Neale. 1985. The Acquisition of an Integrative Response in a Competitive Market. *Organizational Behavior and Human Decision Processes*. **34** 294-313.
- Bazerman, M.H., M.A. Neale. 1992. *Negotiating Rationally*. The Free Press, NY.
- Bazerman, M.H., J.R. Curhan, D.A. Moore, K.L. Valley. 2000. Negotiation. *Annual Review of Psychology*. **51** 279-314.
- Bazerman, M.H., D.A. Moore. 2013. *Judgment in Managerial Decision Making*. Wiley, NY.
- Bearden, J.N. 2001. Ultimatum Bargaining Experiments: The State of the Art. Mimeo, INSEAD.
- Bernheim, B., M. Whinston. 1986. Common Agency. *Econometrica*. **54**(4) 923-942.
- Binmore, K., A. Shaked, J. Sutton. 1985. Testing Noncooperative Bargaining Theory: A Preliminary Study. *American Economic Review*. **75** (5) 1178-1180. s
- Bolton, G. 1991. A Comparative Model of Bargaining: Theory and Evidence. *American Economic Review*. **81** 1096-1136.

- Bowley, A.L. 1928. On bilateral monopoly. *The Economic Journal* **38** 651-659.
- Cachon, G. 2003. Supply Chain Coordination with Contracts. A.G. de Kok, S.C. Graves, eds. *Handbook in Operations Research and Management Science, Supply Chain Management: Design, Coordination and Operation*. Elsevier, Amsterdam.
- Chakraborty, T., S. Judd, M.J. Kearns, J. Tan. 2010. A Behavioral Study of Bargaining in Social Networks. Mimeo, University of Pennsylvania.
- Charness, G., M. Corominas-Bosch, G.R. Frchette. 2007. Bargaining and Network Structure: An Experiment. *Journal of Economic Theory*. **163** 28-65.
- Chatterjee, K., L. Samuelson. 1987. Bargaining with Two-sided Incomplete Information: An Infinite Horizon Model with Alternating Offers. *Review of Economic Studies* **54**(2) 175-192.
- Coase, R. 1960. The Problem of Social Cost. *Journal of Law and Economics* **3** 1-44.
- D'Aspremont, C., L. Gerard-Varet. 1979. Incentives and Incomplete Information. *Journal of Public Economics*. **11** 25-45.
- Dohmen, T., A. Falk. 2011. Performance Pay and Multidimensional Sorting: Productivity, Preferences, and Gender. *American Economic Review*. **101**(2) 556-90.
- Druckman, D. 1967. Dogmatism, Prenegotiation Experience and Simulated Group Representation as Determinants of Dyadic Behavior in Bargaining Situations. *Journal of Personality and Social Psychology*. **6** 279-290.
- Edgeworth, F.Y. 1881. *Mathematical Psychics: An Essay on the Application of Mathematics to the Moral Sciences*. C. Kegan Paul & Co, Chesterland, OH.
- Falk, A., E. Fehr, U. Fischbacher. 2003. Reasons for Conflict: Lessons from Bargaining Experiments. *Journal of Institutional and Theoretical Economics*. **159**(1) 171-187.
- Fellner, W. 1949. *Competition Among the Few*. Knopf, NY.
- Fréchet, G.R., J.H. Kagel, S.F. Lehrer. 2003. Bargaining in Legislatures: An Experimental Investigation of Open versus Closed Amendment Rules. *American Political Science Review*. **97**(2) 221-232.
- Fréchet, G.R., J.H. Kagel, M. Morelli. 2005a. Nominal Bargaining Power, Selection Protocol, and Discounting in Legislative Bargaining. *Journal of Public Economics*. **89**(8) 1497-1517.

-
- Fréchette, G.R., J.H. Kagel, M. Morelli. 2005b. Gamsons Law versus Non-Cooperative Bargaining Theory. *Games and Economic Behavior*. **51**(2) 365-390.
- Galinsky, A., T. Mussweiler. 2001. First offers as anchors: the role of perspective taking and negotiator focus. *Journal of Personality and Social Psychology*. **81**(4) 657-669.
- Güth, W., R. Schmittberger, B. Schwarz. 1982. An Experimental Analysis of Ultimatum Bargaining. *Journal of Economic Behavior and Organization*. **75** 367-388.
- Güth, W., R. Tietz. 1988. Ultimatum Bargaining for a Shrinking Cake An Experimental Analysis. R. Tietz, W. Albers, R. Selton, eds. *Bounded Rational Behavior in Games and Markets*. Berlin: Springer. 111-128.
- Haruvy, E., E. Katok, V. Pavlov. 2014. Can Coordinating Contracts Improve Channel Efficiency? Mimeo, UT Dallas.
- Ho, T-H., J-J. Zhang. 2008. Designing Price Contracts for Boundedly Rational Customers: Does the Framing of the Fixed Fee Matter? *Management Science*. **54**(4) 686-700.
- Hoffman, E., M.L. Spitzer. 1982. The Coase Theorem: Some Experimental Tests. *Journal of Law and Economics*. **25** 73-98.
- Kagel, J., A.E. Roth, eds. 1995. *The Handbook of Experimental Economics*. Princeton University Press, Princeton, N.J.
- Kahneman, D., A. Tversky. 1981. The framing of decisions and the psychology of choice. *Science*. **211**(4481) 453-458.
- Kelley, H.H., L.L. Beckman, C.S. Fischer. 1967. Negotiating the Division of a Reward Under Incomplete Information. *Journal of Experimental Social Psychology*. **3** 361-398.
- Krishna, V. 2002. *Auction Theory*. Academic Press, NY.
- 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.
- Loch, C.H., Y. Wu. 2008. Social Preferences and Supply Chain Performance: An Experimental Study. *Management Science*. **54**(11) 1835-1849.
- Lovejoy, W.S. 2010a. Conversations with Supply Chain Managers. University of Michigan Ross School of Business working paper 1145.


- Lovejoy, W.S. 2010b. Bargaining Chains. *Management Science*. **56**(12) 2282-2301.
- Mannix, E.A., L.L. Thompson, M.H. Bazerman. 1989. Negotiation in Small Groups. *Journal of Applied Psychology*. **74**(3) 508-517.
- Marshall, A. 1890. *Principles of Economics*. MacMillan, London.
- Myerson, R. 1981. Optimal Auction Design. *Math of Operations Research*. **6** 58-73.
- Myerson, R., M. Satterthwaite. 1983. Efficient Mechanisms for Bilateral Trading. *Journal of Economic Theory*. **29** 265-281.
- Murnighan, J.K., A.E. Roth, F. Schoumaker. 1988. Risk Aversion in Bargaining: An Experimental Study. *Journal of Risk and Uncertainty*. **1** 101-124.
- Muthoo, A. 2002. *Bargaining Theory with Applications*. Cambridge University Press, Cambridge UK.
- Neale, M.A., M.H. Bazerman. 1985. The Effect of Framing and Negotiator Overconfidence on Bargaining Behaviors and Outcomes. *The Academy of Management Journal*. **28**(1) 34-49.
- Neelin, J., H. Sonnenshein, M. Spiegel. 1988. A Further Test of Noncooperative Bargaining Theory: A Comment. *American Economic Review*. **78**(4) 824-836.
- Northcraft, G.B., M.A. Neale. 1987. Experts, Amateurs and Real Estate: an Anchoring-and-Adjustment Perspective on Property Pricing Decisions. *Organizational Behavior and Human Decision Processes*. **39** 228-241.
- Nydegger, R.V., H.G. Owen. 1975. Two-Person Bargaining: An Experimental Test of the Nash Axioms. *International Journal of Game Theory*. **3** 239-249.
- Ochs, J., A.E. Roth. 1989. An Experimental Study of Sequential Bargaining. *American Economic Review*. **79** 355-384.
- Polzer, J.T., E.A. Mannix, M.A. Neale. 1998. Interest Alignment and Coalitions in Multiparty Negotiation. *Academy of Management Journal*. **41**(1) 42-54.
- Pratt, A., A. Rustichini. 2003. Games Played Through Agents. *Econometrica* **71**(4) 989-1026.
- Pruitt, D.G., J.L. Drews. 1969. The Effect of Time Pressure, Time Elapsed and the Opponents Concession Rate on Behavior in Negotiations. *Journal of Experimental Social Psychology*. **5** 43-69.

-
- Pruitt, D.G., D.F. Johnson. 1972. Mediation as an Aid to Face Saving in Negotiation. *Journal of Personality and Social Psychology*, **14** 239-246.
- Pruitt, D.G., S.A. Lewis. 1975. Development of Integrative Solutions in Bilateral Negotiation. *Journal of Personality and Social Psychology*, **31**(4) 621-633.
- Ross, L., R.E. Nisbett. 1991. *The Person and the Situation: Perspectives of Social Psychology*. New York: McGraw-Hill.
- Roth, A.E. 1995. Bargaining Experiments. J. Kagel, A.E. Roth, eds. *Handbook of Experimental Economics*. Princeton University Press, 253-348.
- Roth, A.E., M.W.K. Malouf. 1979. Game-Theoretic Models and the Role of Information in Bargaining. *Psychological Review*. **86** (6) 574-594.
- Roth, A.E., M.W.K. Malouf, J.K. Murnighan. 1981. Sociological versus Strategic Factors in Bargaining. *Journal of Economic Behavior and Organization*. **2** 153-177.
- Roth, A.E., J.K. Murnighan. 1982. The Role of Information in Bargaining: An Experimental Study. *Econometrica*. **50** 1123-1142.
- Roth, A.E., J.K. Murnighan, F. Schoumaker. 1988. The Deadline Effect in Bargaining: Some Experimental Evidence. *American Economic Review*. **78** 806-823.
- Rubin, J.Z., B.R. Brown. 1975. *The Social Psychology of Bargaining and Negotiation*. New York: Academic.
- Rubinstein, A. 1982. Perfect Equilibrium in a Bargaining Model. *Econometrica*. **50** 97-110.
- Siegel, S., L.E. Fouraker. 1960. *Bargaining and Group Decision Making*. New York: McGraw-Hill.
- Smith, V. 1962. An Experimental Study of Competitive Market Behavior. *Journal of Political Economy* **70** 111-137.
- Stigler, G. 1942. *The Theory of Price*. MacMillan, NY.
- Thompson, L. 1990. Negotiation Behavior and Outcomes: Empirical Evidence and Theoretical Issues. *Psychological Bulletin*. **108**(3) 515-532.
- Thompson, L. 1998. *The Mind and Heart of the Negotiator*. Upper Saddle River, N.J.: Prentice Hall.

- Thompson, L., R. Hastie. 1990. Social Perception in Negotiations. *Organizational Behavior and Human Decision Processes*. **47** 98-123.
- Thompson, L., D. Hrebec. 1996. Lose-Lose Agreements in Interdependent Decision Making. *Psychological Bulletin*. **120**(3) 396-409.
- Thompson, L., G. Loewenstein. 1992. Egocentric Interpretations of Fairness and Interpersonal Conflict. *Organizational Behavior and Human Decision Processes*. **51** 176-197.
- Thompson, L., E.A. Mannix, M.H. Bazerman. 1988. Group Negotiation: Effects of Decision Rule, Agenda and Aspiration. *Journal of Personality and Social Psychology*. **54**(1) 86-95.
- Vickrey, W. 1961. Counterspeculations, Auctions and Competitive Sealed Tenders. *Journal of Finance*. **16** 8-37.

Figures and Tables

Figure 1: Three Tier Supply Chain Set Up



Stage 1: Suppliers and Manufacturers negotiate simultaneously to determine transfer prices. Each Manufacturer agrees on a price with one of the Suppliers.

Stage 2: Manufacturers and the Retailer negotiate simultaneously to determine transfer prices. Retailer agrees on a price with one of the Manufacturers. Supply chain formed with one firm in each tier.

Table 1: Experimental Treatments

Treatment	Supplier Costs		Manufacturer Costs		Retailer Revenue
	S1	S2	M1	M2	
MDiff0	5	15	5	5	40
Base	5	15	5	15	40
MDiff20	5	15	5	25	40
SDiff0	5	5	5	15	40
SDiff20	5	25	5	15	40

Figure 2: Frequency of closure by treatment




Figure 3: Ex-ante efficiency by treatment and tier




Figure 4: Ex-post efficiency




Figure 5: Supply Chain Profit Distribution




Table 2: BP profitability predictions for each treatment and tier

Treatment	Supplier profit	Manufacturer profit	Retailer profit
MDiff0	[0, 10]	0	[20, 30]
Base	5	5	20
MDiff20	5	10	15
Predicted trend as Mfg ΔC increases	No prediction	Increases	Decreases
SDiff0	0	5	25
Base	5	5	20
SDiff20	[5, 15]	5	[10, 20]
Predicted trend as Supp ΔC increases	Increases	Stays the same	Decreases

Table 3: Tier-level Profit Distributions

VARIABLES	(1) Supplier	(2) Manuf	(3) Retailer
MDiff = 10	-2.649*** (0.834)	3.544*** (0.893)	-3.393** (1.226)
MDiff = 20	-0.971 (0.652)	5.161*** (0.643)	-4.579*** (0.973)
SDiff = 10	2.999*** (0.956)	-0.612 (1.145)	-3.842** (1.513)
SDiff = 20	6.961*** (1.086)	-2.764*** (0.837)	-6.730*** (1.245)
Period	-0.310* (0.158)	0.00543 (0.175)	0.283** (0.108)
Constant	6.603*** (1.180)	4.601*** (0.957)	19.05*** (1.776)
Observations	218	218	218
R-squared	0.261	0.177	0.306
M Diff 10 = 20?	p = 0.03	p = 0.12	p = 0.39
S Diff 10 = 20?	p = 0.00	p = 0.02	p = 0.04

Standard errors clustered at the session reported in parentheses. Dependent variables are Supplier profit, Manufacturer profit, and Retailer profit.

*** p<0.01, ** p<0.05, * p<0.1

Table 4: Whole-chain profit profiles

Treatment	Supplier Profit	Manuf Profit	Retailer Profit	Profit Order Tested	Relative frequency
M Diff 0	[0,10]	0	[20,30]	M < S < R	1/6
Base	5	5	20	S < R & M < R	1/3
M Diff 20	5	10	15	S < M < R	1/6
S Diff 0	0	5	25	S < M < R	1/6
S Diff 20	[5,15]	5	[10,20]	M < S & M < R	1/3

Table 5: Results of whole-chain profit profile hypothesis tests

Treatment	BP Order	% Match Order	Null: Random	Obs Random?
MDiff0	R > S > M	71.79%	1/6	p < 0.01
Base	R > S & R > M	66.04%	1/3	p < 0.01
MDiff20	R > M > S	20.83%	1/6	p = 0.44
SDiff0	R > M > S	65.12%	1/6	p < 0.01
SDiff20	R > M & S > M	45.71%	1/3	p = 0.15

Table 6: Profitability point predictions

Treatment	Predicted			Median Experimental Outcome		
	Supplier	Manuf	Retailer	Supplier	Manuf	Retailer
MDiff0	[0,10]	0	[20,30]	9.00	3.40***	17.00***
Base	5	5	20	5.00	7.00***	12.50***
MDiff20	5	10	15	8.25***	10.00	11.00***
SDiff0	0	5	25	2.00***	9.00***	17.00***
SDiff20	[5,15]	5	[10,20]	8.00	5.00	10.00

*** indicates that the median outcome is significantly different from the point prediction at p < .01 (using a rank-sum test). All other p-values are > 0.10.

Figure 6a: Model Error – Manufacturer Profit




Figure 6b: Model Error – Manufacturer Profit Share




Table 7: Best Empirical Fit of α

VARIABLES	(1) Manuf Profit	(2) Manuf Profit	(3) Manuf Profit Share	(4) Manuf Profit Share
α	0.573*** (0.0489)		0.566*** (0.0454)	
α (MDiff0)		0.493*** (0.0191)		0.579*** (0.0107)
α (Base)		0.732*** (0.0333)		0.746*** (0.0457)
α (MDiff20)		0.470*** (0.0389)		0.467*** (0.0387)
α (SDiff0)		0.746*** (0.0907)		0.702*** (0.0948)
α (SDiff20)		0.591*** (0.0322)		0.610*** (0.0760)
Observations	218	218	218	218
R-squared	0.787	0.822	0.782	0.813

Standard errors clustered at the session reported in parentheses. Dependent variables are Manufacturer profit in (1) and (2), and Manufacturer profit share in (3) and (4).

*** p<0.01, ** p<0.05, * p<0.1

Table 8a: Individual Characteristics – Supplier Profit

VARIABLES	(1)	(2)	(3)	(4)	(5)
<i>Individual Traits</i>	<i>Risk</i>	<i>Selfish</i>	<i>Altruist</i>	<i>Fair</i>	<i>Efficient</i>
Retailer has Trait	1.599*** (0.502)	0.941 (1.013)	0.704 (0.538)	-0.0458 (1.051)	0.259 (0.636)
Manufacturer has Trait	-0.140 (0.590)	-0.869** (0.349)	0.499 (0.917)	-1.055 (0.922)	1.718 (1.228)
Supplier has Trait	-0.423 (0.847)	0.0198 (0.453)	-0.327 (0.591)	-0.671 (0.583)	0.650 (0.629)
<i>Treatment Controls</i>					
MDiff10	-3.016*** (0.777)	-2.443** (1.039)	-2.460** (0.947)	-2.361** (0.852)	-2.212** (0.904)
MDiff20	-1.410** (0.612)	-1.237 (0.971)	-0.815 (0.702)	-0.584 (0.755)	-0.471 (0.849)
SDiff10	2.702* (1.261)	3.008* (1.524)	2.473 (1.552)	2.934 (1.843)	3.047 (1.879)
SDiff20	6.691*** (1.416)	6.899*** (1.517)	6.402*** (1.756)	6.838*** (1.839)	6.875*** (1.958)
Period	-0.247 (0.173)	-0.253 (0.166)	-0.238 (0.172)	-0.241 (0.173)	-0.239 (0.161)
Constant	6.752*** (1.511)	6.387*** (1.673)	6.502*** (1.980)	6.744*** (2.028)	5.509** (2.036)
Observations	190	190	190	190	190
R-squared	0.206	0.193	0.189	0.195	0.204

Standard errors clustered at the session reported in parentheses. This table represents five different regressions. In each, the dependent variable is Supplier profit. There are the usual treatment and period dummies (listed under “Treatment Controls”). “Retailer has Trait”, “Manufacturer has Trait” and “Supplier has Trait” are dummy variables equaling 1 if the specified firm has the indicated trait [is risk averse in (1), is selfish in (2), is an altruist in (3), is inequity averse in (4) or is surplus maximizing in (5)].

*** p<0.01, ** p<0.05, * p<0.1

Table 8b: Individual Characteristics – Manufacturer Profit

VARIABLES	(1)	(2)	(3)	(4)	(5)
<i>Individual Traits</i>	<i>Risk</i>	<i>Selfish</i>	<i>Altruist</i>	<i>Fair</i>	<i>Efficient</i>
Retailer has Trait	0.0134 (0.649)	0.539 (0.923)	0.710 (0.927)	-0.280 (0.898)	-0.257 (0.851)
Manufacturer has Trait	-1.261* (0.657)	-0.224 (0.685)	-0.0384 (0.741)	0.191 (0.786)	-0.796 (0.795)
Supplier has Trait	1.397* (0.722)	-1.362** (0.535)	0.160 (0.756)	0.954* (0.522)	-0.0611 (1.037)
<i>Treatment Controls</i>					
MDiff10	3.650*** (0.995)	3.226*** (0.733)	3.476*** (0.955)	3.267** (1.094)	3.185** (1.088)
MDiff20	5.699*** (0.818)	4.907*** (0.602)	5.279*** (0.678)	4.987*** (0.886)	4.962*** (0.857)
SDiff10	0.413 (0.931)	1.220* (0.642)	0.663 (0.955)	0.465 (0.778)	0.506 (0.897)
SDiff20	-1.627*** (0.284)	-0.798*** (0.260)	-1.381*** (0.235)	-1.444*** (0.356)	-1.426*** (0.357)
Period	-0.0583 (0.139)	-0.0834 (0.135)	-0.0967 (0.142)	-0.106 (0.136)	-0.0968 (0.130)
Constant	3.435*** (0.865)	3.779*** (0.804)	3.377*** (0.861)	3.779*** (0.899)	4.214*** (0.953)
Observations	190	190	190	190	190
R-squared	0.217	0.208	0.188	0.190	0.188

Standard errors clustered at the session reported in parentheses. This table represents five different regressions. In each, the dependent variable is Manufacturer profit. There are the usual treatment and period dummies (listed under “Treatment Controls”). “Retailer has Trait”, “Manufacturer has Trait” and “Supplier has Trait” are dummy variables equaling 1 if the specified firm has the indicated trait [is risk averse in (1), is selfish in (2), is an altruist in (3), is inequity averse in (4) or is surplus maximizing in (5)].

*** p<0.01, ** p<0.05, * p<0.1

Table 8c: Individual Characteristics – Retailer Profit

VARIABLES	(1)	(2)	(3)	(4)	(5)
<i>Individual Traits</i>	<i>Risk</i>	<i>Selfish</i>	<i>Altruist</i>	<i>Fair</i>	<i>Efficient</i>
Retailer has Trait	-0.161 (0.803)	-0.423 (0.704)	-0.286 (0.879)	1.057 (0.856)	-0.619 (0.873)
Manufacturer has Trait	1.062 (0.647)	0.315 (0.899)	0.295 (0.818)	1.658 (1.060)	-1.012 (1.028)
Supplier has Trait	0.200 (0.723)	0.766 (0.724)	0.503 (0.651)	-0.353 (0.814)	0.381 (0.758)
<i>Treatment Controls</i>					
MDiff10	-3.189** (1.456)	-3.307** (1.247)	-3.301** (1.321)	-3.562** (1.411)	-3.502** (1.238)
MDiff20	-4.503*** (1.071)	-4.415*** (0.853)	-4.669*** (1.060)	-4.873*** (1.453)	-4.716*** (1.228)
SDiff10	-5.587*** (1.309)	-5.998*** (1.172)	-5.516*** (1.270)	-5.973*** (1.378)	-5.688*** (1.264)
SDiff20	-8.463*** (0.809)	-8.895*** (0.915)	-8.572*** (0.796)	-8.921*** (0.887)	-8.631*** (1.046)
Period	0.246* (0.124)	0.251* (0.121)	0.256* (0.126)	0.263* (0.137)	0.259* (0.125)
Constant	20.48*** (1.818)	20.92*** (1.534)	20.73*** (1.627)	20.67*** (1.761)	21.45*** (1.787)
Observations	190	190	190	190	190
R-squared	0.347	0.344	0.338	0.367	0.346

Standard errors clustered at the session reported in parentheses. This table represents five different regressions. In each, the dependent variable is Retailer profit. There are the usual treatment and period dummies (listed under “Treatment Controls”). “Retailer has Trait”, “Manufacturer has Trait” and “Supplier has Trait” are dummy variables equaling 1 if the specified firm has the indicated trait [is risk averse in (1), is selfish in (2), is an altruist in (3), is inequity averse in (4) or is surplus maximizing in (5)].

*** p<0.01, ** p<0.05, * p<0.1

Table 9 – Stage 1 Bargaining Offers

Panel A: Offers with Manufacturer 1

Treatment	S1 First Offer		S2 First Offer		M1 First Offer		Avg Accepted Offer	
	Price	% of ZOPA	Price	% of ZOPA	Price	% of ZOPA	Price	% of ZOPA for M1
MDiff0	15.34	34%	19.79	1%	13.63	71%	15.64	65%
Base	15.48	35%	20.86	2%	10.84	81%	13.77	71%
MDiff20	16.02	37%	20.15	3%	12.09	76%	14.20	69%
SDiff0	11.39	21%	11.19	21%	7.06	93%	8.15	90%
SDiff20	20.32	51%	28.30	18%	12.66	74%	17.74	58%

Panel B: Offers with Manufacturer 2

Treatment	S1 First Offer		S2 First Offer		M2 First Offer		Avg Accepted Offer	
	Price	% of ZOPA	Price	% of ZOPA	Price	% of ZOPA	Price	% of ZOPA for M2
MDiff0	15.63	35%	19.78	2%	13.68	71%	15.42	65%
Base	14.92	50%	20.63	0%	11.00	70%	12.91	60%
MDiff20	16.25	56%	20.96	6%	10.88	71%	13.08	60%
SDiff0	10.84	29%	10.18	29%	8.14	84%	8.61	82%
SDiff20	19.26	71%	28.21	21%	13.49	58%	15.65	47%

Table 10 – Stage 2 Bargaining Initial Offers

Treatment	M1 Total Cost	M2 Total Cost	M1 First Offer		M2 First Offer		Retailer First Offer	
			Price	% of ZOPA	Price	% of ZOPA	Price	% of ZOPA
MDiff0	20.66	20.35	26.45	26%	26.73	31%	21.52	88%
Base	18.98	28.04	28.98	47%	32.83	24%	22.83	82%
MDiff20	19.22	38.10	32.02	62%	40.07	11%	25.52	70%
SDiff0	13.16	23.70	26.57	49%	30.46	26%	21.35	72%
SDiff20	22.74	30.65	31.89	50%	35.94	38%	26.44	81%

Table 11 – Stage 2 Bargaining Average Offers

Treatment	M1 Avg Offer	M2 Avg Offer	Retailer Avg Offer	Avg Accepted Offer	
				Price	% of ZOPA for R
MDiff0	25.28	25.67	21.40	23.70	75%
Base	28.88	32.52	24.01	27.32	61%
MDiff20	31.11	40.44	25.55	28.41	56%
SDiff0	24.79	28.69	21.97	23.31	63%
SDiff20	31.36	34.67	27.31	30.06	58%

Table 12 – Effect of Initial Offers on Final Price

VARIABLES	M1-S Bargaining		M2-S Bargaining		R-M Bargaining	
	(1)	(2)	(3)	(4)	(5)	(6)
Buyer's First Offer	0.458*** (0.130)	0.410*** (0.121)	0.673*** (0.0835)	0.635*** (0.0844)	0.231*** (0.0634)	0.239*** (0.0594)
Seller 1's First Offer	0.448*** (0.103)	0.408*** (0.0977)	0.172*** (0.0544)	0.141** (0.0547)	0.249*** (0.0758)	0.257*** (0.0751)
Seller 2's First Offer	0.0980* (0.0464)	-0.0119 (0.0383)	0.0681** (0.0241)	0.0193 (0.0349)	0.0587 (0.0544)	0.0465 (0.0609)
Seller 1's Total Cost					0.149** (0.0624)	0.0862 (0.0723)
Seller 2's Total Cost					0.120* (0.0622)	0.156* (0.0900)
Treatment Controls	No	Yes	No	Yes	No	Yes
Constant	-0.415 (1.023)	3.344** (1.257)	1.335 (0.951)	3.325** (1.295)	1.308*** (0.216)	1.211*** (0.176)
Observations	179	179	158	158	164	164
R-squared	0.668	0.699	0.734	0.748	0.070	0.056

Standard errors clustered at the session level reported in parentheses. Dependent variable is the agreed upon price. *** p<0.01, ** p<0.05, * p<0.1

Table 13 – Effect of Making the First Offer on Final Price

VARIABLES	M1-S Bargaining		M2-S Bargaining		R-M Bargaining	
	(1)	(2)	(3)	(4)	(5)	(6)
Buyer Made First Offer	-0.236 (1.330)	-0.315 (0.824)	0.0517 (0.937)	0.403 (0.592)	-0.702 (0.680)	-0.170 (0.667)
Treatment Controls	No	Yes	No	Yes	No	Yes
Constant	13.69*** (1.016)	15.73*** (0.437)	12.90*** (1.060)	15.26*** (0.408)	26.69*** (0.364)	23.74*** (0.665)
Observations	236	236	231	231	220	220
R-squared	0.000	0.352	0.000	0.270	0.004	0.297

Standard errors clustered at the session level reported in parentheses. Dependent variable is the agreed upon price. *** p<0.01, ** p<0.05, * p<0.1

Table 14 – Escalation of Commitment and Bargaining Closure

VARIABLES	M1-S Bargaining		M2-S Bargaining		R-M Bargaining	
	(1)	(2)	(3)	(4)	(5)	(6)
Buyer's Concession	0.146 (0.133)		0.247 (0.207)		0.100 (0.129)	
Seller 1's Concession	-0.224 (0.417)		-0.0381 (0.0529)		0.413* (0.220)	
Seller 2's Concession	0.0302 (0.0597)		-0.0838 (0.0557)		-0.0218 (0.0941)	
Total Concession		0.0315 (0.0847)		0.0352 (0.0575)		0.144 (0.124)
Buyer's Initial Claim (% ZOPA)	-0.204* (0.113)	-0.186* (0.0923)	-0.175 (0.127)	-0.148 (0.123)	-0.359 (0.225)	-0.296* (0.175)
Seller 1's Initial Claim (% ZOPA)	-0.375 (0.225)	-0.371* (0.208)	-0.355 (0.233)	-0.327 (0.244)	-0.415* (0.212)	-0.260 (0.160)
Seller 2's Initial Claim (% ZOPA)	-0.00907 (0.135)	-0.0481 (0.127)	0.221 (0.157)	0.160 (0.111)	0.0537 (0.0775)	0.0823 (0.0780)
Constant	1.262*** (0.147)	1.241*** (0.123)	1.197*** (0.138)	1.176*** (0.144)	1.308*** (0.216)	1.211*** (0.176)
Observations	186	206	171	197	164	164
R-squared	0.170	0.159	0.076	0.053	0.070	0.056

Standard errors clustered at the session level reported in parentheses. Dependent variable is an indicator for a bargaining agreement. *** p<0.01, ** p<0.05, * p<0.1

Table 15 – Anti-Concessionary Offers and Prices

	Stage 1 Bargaining	Stage 2 Bargaining	Stage 1 Bargaining	Stage 2 Bargaining
Variables	(1)	(2)	(3)	(4)
Sell-side is anti-concessionary	1.477* (0.793)	1.742** (0.786)		
Buy-side is anti-concessionary			0.294 (0.455)	-0.297 (0.761)
Negotiation Involving M2	-0.993*** (0.308)		-0.470 (0.460)	
Treatment Controls	YES	YES	YES	YES
Observations	323	144	273	100
R-squared	0.347	0.408	0.324	0.347

Standard errors clustered at the session level reported in parentheses. Dependent variable is the agreed upon price (conditional on coming to agreement). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Columns (1) and (3) report results for the first stage negotiations between the Manufacturers and the Suppliers (we pool together negotiations involving both Manufacturers). Columns (2) and (4) report results for the second stage negotiations between the Retailer and the Manufacturers. “Sell-side is anti-concessionary” is an indicator variable that equals 1 if the sell-side firm that is part of the final agreement made at least one anti-concessionary offer. “Buy-side is anti-concessionary” is an indicator variable that equals 1 if the buy-side firm made at least one anti-concessionary offer to the sell-side firm that was part of the final agreement. “Negotiation involving M2” is an indicator variable that equals 1 if the agreement is between Manufacturer 2 and one of the suppliers.

Figure 7 – Agreement Times by Treatment during Second Stage Bargaining

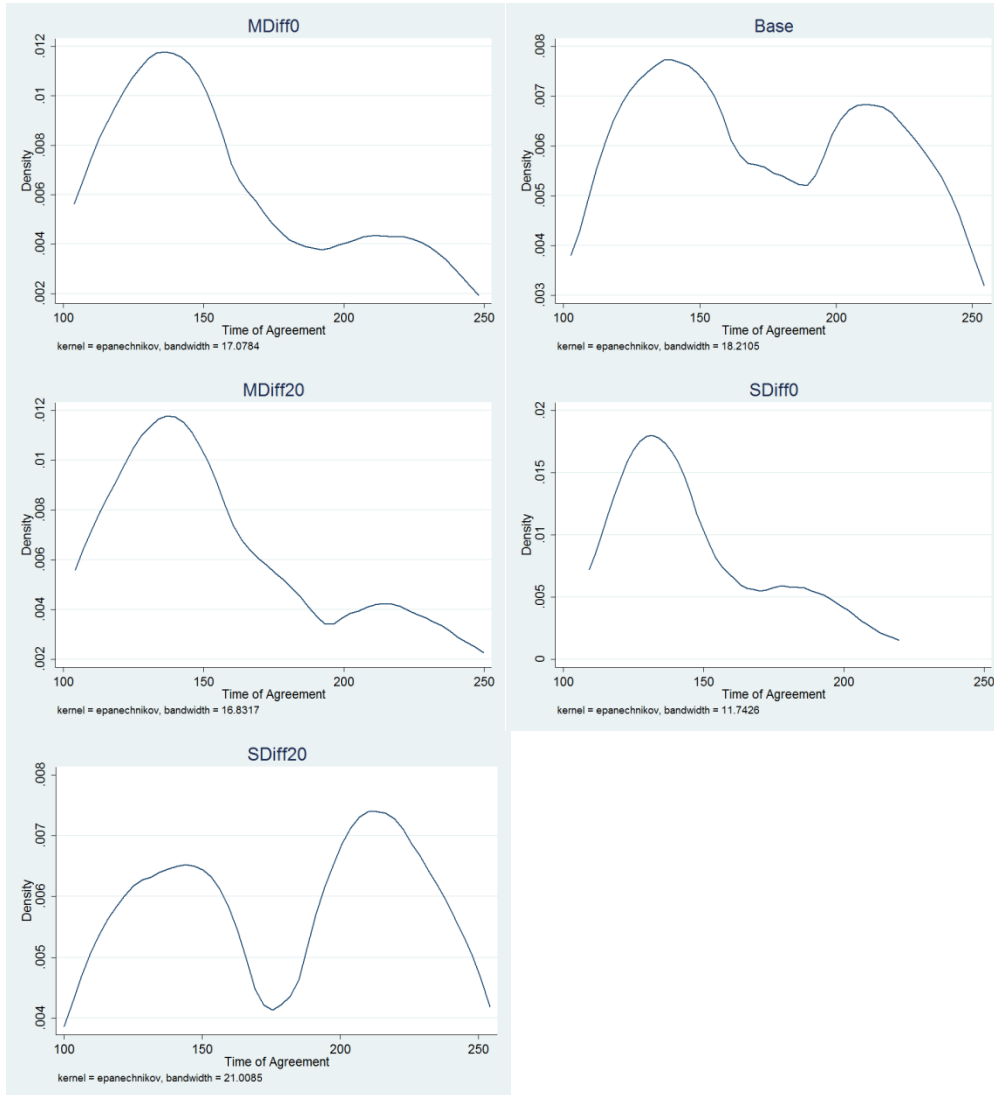


Table 16 – Effect of Agreement Timing of Buying Firm’s Share

VARIABLES	M1-S Bargaining		M2-S Bargaining		R-M Bargaining	
	(1)	(2)	(3)	(4)	(5)	(6)
Time Accepted by Buyer	-0.000366 (0.000503)	9.17e-05 (0.000363)	-0.000991 (0.000793)	-0.000186 (0.000505)	-0.000313 (0.000371)	-4.37e-05 (0.000373)
Time Accepted by Seller	-0.000172 (0.000525)	0.000220 (0.000306)	-0.000459 (0.000709)	0.000143 (0.000514)	-0.000464 (0.000335)	-0.000311 (0.000326)
Treatment Controls	No	Yes	No	Yes	No	Yes
Constant	0.762*** (0.0906)	0.690*** (0.0486)	0.779*** (0.107)	0.645*** (0.0811)	0.695*** (0.0580)	0.647*** (0.0674)
Observations	195	195	192	192	181	181
R-squared	0.009	0.440	0.046	0.362	0.016	0.220

Standard errors clustered at the session level reported in parentheses. Dependent variable is the fraction of the ZOPA allocated to the buyer in the agreement. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Appendix

Figure A1: Supplier-Manufacturer Negotiation Screen

Negotiations with Supplier 1		
Supplier 1's Last Offer		
Supplier	Manufacturer	Price
Supplier 1	Manufacturer 1	12.00
Accept Offer		
Your last offer to Supplier 1		
Supplier	Manufacturer	Price
Supplier 1	Manufacturer 1	8.00
Price to offer Supplier 1? <input type="text" value="8"/>		
Make Offer		
Supplier 1: Hi, let's make a deal Manufacturer 1: Hello		

You are: Manufacturer 1
Your cost is 5.00

Firm Costs and Revenue

S1's Cost	M1's Cost	R's Revenue
5.00	5.00	40.00
S2's Cost	M2's Cost	
15.00	15.00	

Negotiations with Supplier 2		
Supplier 2's Last Offer		
Supplier	Manufacturer	Price
Supplier 2	Manufacturer 1	16.00
Accept Offer		
Your last offer to Supplier 2		
Supplier	Manufacturer	Price
Supplier 2	Manufacturer 1	16.00
Price to offer Supplier 2? <input type="text" value="16"/>		
Make Offer		