

WORKING PAPERS



A Comparison of Auctions and Multilateral Negotiations

Charles J. Thomas
Federal Trade Commission
and
Harvard Law School

and

Bart J. Wilson
Economic Science Laboratory
University of Arizona

WORKING PAPER NO. 231

July 2000

FTC Bureau of Economics working papers are preliminary materials circulated to stimulate discussion and critical comment. All data contained in them are in the public domain. This includes information obtained by the Commission which has become part of the public record. The analyses and conclusions set forth are those of the authors and do not necessarily reflect the views of other members of the Bureau of Economics, other Commission staff, or the Commission itself. Upon request, single copies of the paper will be provided. References in publications to FTC Bureau of Economics working papers by FTC economists (other than acknowledgment by a writer that he has access to such unpublished materials) should be cleared with the author to protect the tentative character of these papers.

BUREAU OF ECONOMICS
FEDERAL TRADE COMMISSION
WASHINGTON, DC 20580

*A Comparison of Auctions and Multilateral Negotiations**

Charles J. Thomas
Federal Trade Commission
Division of Economic Policy Analysis
600 Pennsylvania Ave., NW
Washington, D.C. 20580
E-mail: cjthomas@ftc.gov

and

Bart J. Wilson
Economic Science Laboratory
University of Arizona
Tucson, AZ 85721-0108
E-mail: bwilson@econlab.arizona.edu

July 2000

Abstract: We compare the well-known first-price auction with a common but previously unexamined exchange process that we term “multilateral negotiations.” In multilateral negotiations, a buyer solicits price offers for a homogeneous product from sellers with heterogeneous costs, and then plays the sellers off one another to obtain additional price concessions. Using experimental methods, we find that transaction prices are statistically indistinguishable in the two institutions with a sufficiently large number of sellers, but that prices are higher in multilateral negotiations than in first-price auctions as the number of sellers decreases. With fewer sellers, the institutions are equally efficient, but with more sellers, there is some evidence that multilateral negotiations are slightly more efficient.

* We thank Jeremy Bulow, Jim Cox, Vernon Smith, Todd Swarthout, and Keith Waehrer for helpful comments. We also thank the Economic Science Laboratory at the University of Arizona for providing financial support, and Todd Swarthout for skillfully writing the software. The data and a demonstration copy of the instructions and software are available upon request. This work is not intended to represent the position of the Federal Trade Commission or any individual Commissioner.

JEL: C78, D44, C9, L14, L40

Keywords: bargaining, auctions, e-commerce, merger analysis

1. Introduction

Economists have developed models of auctions and of bilateral bargaining to enhance our understanding of two of the most commonly observed exchange institutions. Auction models are used to examine the allocation of products from an auctioneer to a number of bidders. For example, a manufacturer may procure inputs by soliciting sealed bids that name each supplier's contract price. In addition to their use in industrial procurement, auctions have been used to allocate products as varied as art, produce, government securities, and offshore mineral rights.¹ More recently, several governments have used auctions to allocate such valuable resources as radio spectra, electric power, and pollution rights.² As these examples illustrate, the lessons of auction theory have been applied to several sectors of the economy.

Bilateral bargaining models are used to examine the allocation of surplus between two economic agents. For example, two traders in an Edgeworth Box economy presumably would reach an agreement on the contract curve through some sort of bargaining process. More concretely, bilateral bargaining models have been applied to settings such as wage negotiations between a firm and a union, and contracting for the purchase of specialized inputs.³ Bilateral bargaining is viewed by some economists as the most basic form of exchange, and it is considered to be a natural default in the event of a breakdown in some other exchange process. For example, if the previously described procurement auction did not yield acceptable contract terms to the manufacturer, then the manufacturer and one of the suppliers might negotiate a contract in a setting separate from the other potential suppliers. Bilateral bargaining models have been extended to allow for "multilateral bargaining," in which more than two agents bargain over the division of surplus common to all of the agents.⁴

In this paper we use experimental techniques to study a previously unexamined exchange process that combines features from both auction and bilateral bargaining models. Specifically, we study a setting in which a buyer solicits price offers from sellers, and then confronts each seller with claims about its rivals' price offers in order to elicit a more favorable offer. The buyer plays the sellers off one another until the expected value of future concessions is less than the incremental cost of obtaining those concessions. This process is pervasive in industrial procurement, with buyers extending formal Requests for Proposals (RFPs) and then haggling

¹ See Milgrom and Weber [1982] and McAfee and McMillan [1987].

² See McAfee and McMillan [1996], Wolfram [1998], and Cason [1995] for details, respectively.

³ See Horn and Wolinsky [1988].

⁴ See Krishna and Serrano [1996].

with suppliers after receiving the initial offers. This process is also common to many other transactions, including the securing of job offers and the purchasing of computers, contractors' services, and automobiles. We refer to this exchange institution as "multilateral negotiation" to distinguish it from the multilateral bargaining setting described above.

Due to the prevalence of multilateral negotiations as a means of exchange, it is important to understand them for the same reasons that it is important to understand other common price formation processes, such as auctions and bilateral bargaining. We are interested in such issues as the institution's efficiency, the effect of the number of sellers on the transaction price, and the effect of the agents' and the institution's characteristics on the agents' bargaining positions.

We also are interested in the relationship between the outcomes of multilateral negotiations and of various auction formats for two reasons. First, this relationship has important implications for institutional design. For example, a buyer attempting to procure services from potential sellers has incentives to use the most profitable means of exchange. The fact that some buyers in an industry use multilateral negotiations, while others use one-shot sealed-bid auctions, suggests either that the processes are outcome-equivalent or that there are factors that make one process more favorable than the other. Identifying these factors should lead to a more informed selection of an exchange process. Concerns about institutional design are particularly important in the emerging e-commerce field, where buyers and sellers are in the process of developing software agents to handle the procurement process online.⁵ For example, Su, Huang, and Hammer [2000] have implemented a prototype server for automated, Web-based negotiations between buyers and sellers in e-commerce exchange. More generally, a bevy of researchers and practitioners in computer science and management information systems are creating artificially intelligent mechanisms for negotiations and auctions. Regrettably, there is no empirical and little theoretical economic research comparing these institutions that could guide them in their work. This paper provides a first step in such a research agenda.

Second, the relationship between the outcomes of multilateral negotiations and auctions has important implications for antitrust analysis. Recently, auction models have been employed by U.S. competition authorities and private parties to evaluate the impact of proposed mergers.⁶ Even if transactions resemble multilateral negotiations more closely than auctions, an analyst

⁵ According to an article in *Business Week*, 1/17/2000, "By the end of next year, 91% of U.S. firms will use the Internet for procurement, compared with today's 31 percent."

⁶ For example, auction results were used to evaluate the recent merger between Rite-Aid and Revco. See Baker [1997]. Section 2.2.1 of the 1992 Merger Guidelines highlights the possible effect of mergers in auction markets.

might elect to study behavior by using an auction model because such models have been extensively studied, while at present there do not exist any formal multilateral negotiation models. If we find that the outcomes of auctions and multilateral negotiations are similar,⁷ then there should be less concern about using a modeling approach that does not precisely fit the market's characteristics. Otherwise, caution should be used in applying auction models in settings in which transactions more closely resemble multilateral negotiations.

We study the relationship between first-price auctions and multilateral negotiations by permitting fairly unstructured negotiation between a buyer and several sellers. Each experimental session anonymously matches a buyer with either two or four sellers, and consists of several periods of negotiations and first-price auctions. When multilateral negotiations are employed, the buyer can communicate electronically in real-time with the sellers, but the sellers cannot communicate with each other. When auctions are employed, the buyer plays a passive role, and none of the players can communicate with each other. We match sellers' costs across groups and institutions to study whether outcomes depend on which institution is used. Similarly, we vary the number of sellers to see how the outcomes change within an institution. Finally, we exploit within-group differences in the outcomes of the different institutions to test whether sellers' behavior later in the session depends upon their prior institutional experience.

With two sellers who have no prior experience with either institution, we find that the transaction prices in multilateral negotiations are statistically higher than the transaction prices in first-price auctions. However, with four sellers, we cannot statistically distinguish between the transaction prices in the two institutions. We also find an experience effect in the two-seller treatment, in that sellers first exposed to multilateral negotiations set higher prices in first-price auctions than do sellers first exposed to first-price auctions. Moreover, examination of the within-group behavior with two sellers suggests there is virtually no difference between the two institutions. From these results we conclude that the number of sellers and their prior experience have an economically significant effect on the relationship between first-price auctions and multilateral negotiations. Specifically, with no prior experience with multilateral negotiations or first-price auctions, the transaction prices of first-price auctions and multilateral negotiations are statistically indistinguishable if there are sufficiently many sellers. Efficiency of multilateral negotiations is slightly higher. With a small number of sellers, multilateral negotiation prices are

⁷ This similarity could involve either equivalence in price levels or equivalence in percentage price changes as the number of sellers changes.

higher and efficiency is the same. In contrast, our within-group results suggest that a buyer who traditionally has used multilateral negotiations may not receive lower price offers by switching to an auction process. However, given that multilateral negotiations are costly in terms of the time spent determining the transaction price, our results suggest that buyers in this setting should prefer to employ first-price auctions rather than multilateral negotiations.

The paper is structured as follows. Section 2 describes the existing theory most closely relevant to multilateral negotiations, and explains the reasoning underlying our hypothesis that the outcomes of first-price auctions and multilateral negotiations will be similar. Section 3 describes the experimental design and the procedures we use to examine the relationship between the two institutions. Section 4 presents our results, while Section 5 briefly concludes and provides directions for future research.

2. Related Theoretical Background

The exact exchange mechanism we wish to study has not been formally modeled in the bargaining literature, presumably due to the difficulty in finding a tractable solution to a complex problem. Hence, the basis of our study is driven by our intuition about how first-price auctions and multilateral negotiations should be related. To begin, we first describe first-price auction theory and the theoretical setting we envision as being an appropriate starting point for formalizing multilateral negotiations. Next we describe existing work that relates to multilateral negotiations, and finally we explain our intuition about the relationship between auctions and multilateral negotiations.

Consider a setting in which S risk-neutral sellers compete to fulfill a contract for a single risk-neutral buyer. V_B is the commonly known value that the buyer places on having the contract fulfilled. Each seller's cost c is a privately known independent draw from the continuous distribution function G with density g that is strictly positive over the support $[\underline{c}, \bar{c}]$. In the auction literature, this is referred to as a symmetric independent private value (IPV) setting. The first-price auction proceeds with each seller simultaneously submitting a secret price offer. The firm offering the lowest price is awarded the contract at the price that it offered. All other firms receive nothing. The winning seller's profit is $p - c_w$, where p is the transaction price and c_w is the winning seller's cost. The buyer's profit is $V_B - p$, total surplus is $V_B - c_w$, and efficiency is $(V_B - c_w)/(V_B - c_l)$, where c_l is the lowest realized cost.

In contrast, the multilateral negotiation proceeds with each seller simultaneously making a secret price offer to the buyer in the first period of play. The buyer can accept one of the offers or reject them all. If the buyer accepts an offer, then the game concludes and the transaction price is the price p offered by the winning seller. As in the auction setting, the winning seller's profit is $p - c_w$, where c_w is the winning seller's cost, and the buyer's profit is $V_B - p$. If the buyer rejects all offers, then the buyer announces to each seller a (possibly different and not necessarily true) competing offer that is better than the seller's standing offer, and the period ends. In the next period of play, the sellers again make secret price offers, and the game continues as before. The buyer and sellers have $T > 0$ discrete periods of play to reach an agreement, where T can be infinite, but the players are penalized if the game continues to another period. The players' delay cost is represented either by a per-period cost $D \geq 0$, or by a discount factor $\mathbf{d} \in [0,1]$. That is, if the buyer purchases at price p in period \mathbf{t} , then the buyer's present discounted payoff, from the start of the game, is either $V_B - (\mathbf{t} - 1)D - p$ or $\mathbf{d}^{\mathbf{t}-1}(V_B - p)$. The sellers' payoffs are similar.

There is little existing theoretical research that is closely related to the preceding multilateral negotiation model, but there is work that addresses similar questions. One relevant approach incorporates multiple sellers into a bilateral bargaining framework. This approach is taken by Shaked and Sutton [1984], who model an alternating offer setting in which the buyer can switch to a different seller at some commonly known cost. In their setting, the buyer can bargain with only one seller at a time, and offers from one seller are void upon the buyer's switching to another seller. Assuming that it is commonly known that both sellers have the same cost, the authors find that the presence of another seller constitutes a credible threat that permits the buyer to obtain greater surplus than if switching were impossible. In fact, the Walrasian outcome is supported as the switching cost goes to zero. The Walrasian outcome is identical to that which would occur if the buyer instead arranges a first-price auction between the two sellers.⁸

Our approach differs from Shaked and Sutton [1984] in that we do not assume that the sellers' costs are common knowledge. We also relax their strict structure of alternating offers made by the buyer and the sellers. In the alternating offer setting, the buyer's bids and seller's offers tend to converge to the final transaction price. In our setting, we envision the buyer not

⁸ Fudenberg, Levine, and Tirole [1987] examine an alternative model in which the buyers' values are private information and the seller makes all of the offers.

proposing a price at which he is willing to buy, but instead obtaining price concessions by presenting a seller with claims about the discounts offered by rival sellers. Moreover, multilateral negotiations permit the buyer to hold multiple offers simultaneously, a feature distinct from the alternating offer structure.

Another relevant approach incorporates dynamic aspects of a buyer's ability to decline all offers in a procurement auction. This approach is taken by McAfee and Vincent [1997], who model a buyer's selection of a reserve price, when the buyer cannot commit not to solicit offers in the future if he declines all offers made today. They find that an equilibrium of the repeated first-price auction is revenue equivalent to the unique equilibrium of the repeated second-price auction. They also find that, as the time between auctions goes to zero, the expected price converges to that of a static auction with no reserve price.

Our approach differs from McAfee and Vincent [1997] in that we do not assume that the offers are made public.⁹ More significantly, we also do not assume that the buyer announces or commits to a reserve price. While the buyer in a multilateral negotiation likely uses a "reservation offer" to decide whether to accept a current offer, the decision likely is based on the set of received offers. This decision structure permits the buyer to use information revealed by the current offers to forecast future offers.

Within the auction literature, it has been hypothesized that the negotiations we describe bear some relation to second-price auctions.¹⁰ The general argument is that the buyer should be able to obtain concessions from a seller until the seller's offer is just equal to the seller's cost. Such an argument supposes that the buyer's bargaining ability greatly exceeds the seller's. We agree partly with this characterization, but our intuition is that the relationship between multilateral negotiations and various auction formats depends critically on the buyer's ability to credibly reveal to a seller its rivals' offers. In particular, if the offers cannot be credibly revealed, then the multilateral negotiation would be similar to a first-price auction.¹¹ To understand why the ability to credibly reveal offers might potentially play a critical role, consider the problem facing a seller in a multilateral negotiation when its rivals' offers cannot be credibly revealed. When the buyer tries to use a rival's offer to elicit a better offer from the seller, the

⁹ Public revelation of the bids has no effect on the bidders' behavior in their setting, because the only bids that reveal information about a player's type also satisfy the reserve price. Hence, whatever information the "serious" bids reveal cannot be used against the bidders, because the auction concludes upon receipt of at least one serious bid.

¹⁰ See Waehrer and Perry [1999].

¹¹ If the offers can be credibly revealed, then the multilateral negotiation should be similar to a second-price auction. This is the case because sellers should be willing to make concessions until the price reaches their cost.

seller must be concerned that the buyer is not being truthful about the terms or the existence of the rival's offer. Consequently, the seller must be aware of the danger that he could end up bidding against himself by offering price reductions that are undercut by fictitious discounts from a rival.

To explain our intuition more carefully, note that in many auction settings a seller must be concerned about bidding against itself, just as in multilateral negotiations. For example, in the "button auction" described by Milgrom and Weber [1982], the contract price starts at a high level, each potential seller depresses a button to signal its willingness to fulfill the contract at the current price, and the price decreases in continuous fashion. A seller signals its lowest offer by releasing its button at that price. A firm cannot depress the button once the button has been released, so a firm cannot exit and then re-enter the bidding. The winning firm is the last firm to release its button, and it is paid the price at which it released its button.

Suppose that sellers do not see the identity or price of firms that drop out, so they have no knowledge of whether any other firms are participating at any point in time. Thus, a seller must be concerned that it is continuing to depress its button after all other sellers have dropped out. The outcome of this game seemingly should be related to the outcome of a multilateral negotiation without credible revelation. However, this game is strategically equivalent to a Dutch clock procurement auction, in which the price starts at zero and increases in continuous fashion, and the winner is the first firm to depress its button. Moreover, the Dutch auction is strategically equivalent to a sealed-bid first-price auction. Thus, from a theoretical perspective we hypothesize that there should exist a relationship between first-price auctions and multilateral negotiations without credible revelation of offers.

For several reasons, the preceding relationship may not be exact, either empirically or theoretically. First, Coppinger, Smith, and Titus [1980] and Cox, Roberson, and Smith [1982] report that, in buying auctions, prices are higher in first-price auctions than in Dutch clock auctions in which the price clock starts at a high level. Cox, Smith, and Walker [1983] conclude that the non-isomorphism results from bidders mistakenly adjusting their beliefs that their rivals' values are lower than what they initially anticipate, when no one takes the item as the clock ticks down. They argue that the real-time nature of the Dutch clock auction leads to a failure in the isomorphism. Therefore, in a procurement auction, if multilateral negotiations are isomorphic to Dutch auctions, then we would expect multilateral negotiation prices to exceed first-price auction prices.

Second, in practice the outcomes of multilateral negotiations likely depend critically on the players' ability to haggle. For example, suppose that one seller has a low cost draw relative to his rivals, and that all of the sellers begin with high offers that they reduce in the course of the negotiations. When the high cost rivals stop offering discounts, the low cost seller may do the same if the buyer fails to report (i.e., does not lie about) further competing discounts from the other sellers. Thus, the transaction price may be higher with multilateral negotiations than with a one-shot first-price auction. However, a skilled buyer may be able to keep a seller offering discounts below his equilibrium first-price auction offers, because losing sellers in first-price auctions would be willing to lower their initial price offers in order to win the contract. Hence, the buyer might have the power to extract more favorable price offers in a multilateral negotiation than in a one-shot first-price auction.

Another reason that the relationship between multilateral negotiations and first-price auctions may differ behaviorally is that in the multilateral negotiations one must provide incentives for the sellers to make serious offers. That is, there is no reason for sellers to make an offer until the last possible moment, particularly if there are no delay costs and if they are concerned that serious initial offers will be used against them later in the negotiation.¹² In our experimental framework there exists a time limit on each negotiation period, and there clearly exist frictions that prevent the buyer from receiving infinitely many offers. Consequently, a seller might be concerned that it will be left out of the communication process if it does not make serious offers. Moreover, if the seller does not stay current with the state of play, then even if he tries to come in late in the negotiations he will not have a good sense of what the market price is.

Reputation effects may also prevent a buyer from pushing sellers too hard to obtain offers better than he could receive in first-price auction. For example, if the buyer purchases from a seller who refused to lower its price in the face of an alleged offer made by another seller, then the seller will know that the buyer has lied in the past about competing offers. Presumably, this will make the seller more wary about the veracity of future competitive offers alleged by the buyer. Also, relationships may develop between the buyer and the sellers, if the transactions are more appropriately considered to be the stage games of a repeated game.

3. Experimental Design and Procedures

Because of the difficulty in finding a tractable solution for what appears to be a complex theory of multilateral negotiations, we conducted a heuristic experiment (see Smith [1982], pp. 941-942) to compare and contrast market performance and behavior in first-price auctions and multilateral negotiations. In the absence of a formal theoretical analysis, we determined it is better to design our multilateral negotiation process to look more like real-world multilateral negotiations than the structure we outlined in Section 2. We focus on the case without credible revelation of rivals' offers because it seems to be the more empirically relevant case. Moreover, the conclusions regarding the relationship between auctions and multilateral negotiations seem less likely to hold in this setting than in the setting with credible revelation of offers, so the case we examine seems likely to be the more interesting one.

Our experimental design permits us to compare the outcomes of first-price auctions and multilateral negotiations, both across and within subject groups. Moreover, our design also permits us to compare the efficiency of the two institutions, and to compare the first-price auction outcomes with their theoretical predictions.

Using “ F ” to denote a sequence of first-price auctions and “ N ” to denote a sequence of multilateral negotiations, our initial intent was to use both FNF and NFN sequences of treatments, to permit across- and within-group comparisons of the institutions, and to test for experience effects. However, time constraints prevented this due to the time necessary to conduct the multilateral negotiations. Specifically, we wanted the N and F sequences that were paired across treatments to have the same number of transactions, but our pilot session indicated that the treatment with two sets of negotiations would have taken too much time for the number of paired auctions we wanted to conduct.¹³

To compensate, we derived an alternative design that gives us a sufficiently large number of observations to make the desired statistical comparisons of the institutions. Specifically, we pair two treatments, one with the sequence $NFFN$, and one with the sequence $FFNF$. The first and third sequences consist of 12 transactions; the second consists of 16, while the fourth

¹² Multilateral negotiation will be equivalent to a first-price auction if the buyer's per-period delay cost exceeds the maximum possible gain from obtaining additional offers, or if the discount factor is zero. In both cases the buyer will accept an offer in the first period.

¹³ We wanted a large enough number of first-price auction observations so that we could estimate each seller's price-setting function with a minimal 20 degrees of freedom.

consists of 6.¹⁴ We vary these two treatments by changing the number of sellers. One has two sellers per buyer, while the other has four sellers per buyer.

For each of the four treatments, $\{2 \text{ sellers}, 4 \text{ sellers}\} \times \{NFFN, FFNF\}$, we have four groups of subjects. Each subject is assigned a specific role in a specific group for the duration of the session. A seller's characteristics consist of 46 random draws from the Uniform distribution on $[0, 6.00]$. Of the eight groups with four sellers, seller i ($i = 1, 2, 3, 4$) has the same cost draws across groups. Of the eight groups with two sellers, seller i ($i = 1, 2$) has the same cost draws across groups. Moreover, the costs of sellers 1 and 2 in the two-seller treatment are the same as the costs of sellers 1 and 2 in the four-seller treatment. We maintained a constant support for the cost draws across the two- and four-seller treatments because we did not want the distance between the order statistics to affect the real-time negotiations. A larger dispersion in cost draws may make it easier or harder for sellers to extract more surplus from the buyer, and because we did not know if or how the variance of the cost draws mattered, we held it constant.

Our experiment consisted of a total of 736 first-price auctions or rounds of multilateral negotiation using sixty-four undergraduate student volunteers. Some students had participated previously in market experiments, but with significantly different trading institutions. No subject participated in more than one of the sessions reported in this experiment.

The instructions for the first-price auction sequences are based upon those used by Cox, Roberson, and Smith [1982] and Cox, Smith, and Walker [1983, 1988]. The instructions for the multilateral negotiation sequences are newly developed, as our experiment is the first to study this institution.

In addition to reading the self-paced instructions displayed by the software, the subjects followed along as the experiment monitor read out loud from a handout with both additional and review information. The public instructions explained (and made common knowledge) that the sellers' costs were assigned randomly each period and that the distribution of the draws was $U[0, 6.00]$. The instructions also revealed that the buyer's value was 6.00. While revealing the buyer's value may create incentives that are affected by the perceived financial benefit of being a buyer rather than a seller in the experiment,¹⁵ it is consistent with prior auction experiments. Most auction experiments use buying auctions rather than selling auctions, and bids are

¹⁴ Six trailing periods of negotiation were the most that could be comfortably run within a two-hour experiment.

¹⁵ For example, sellers may not compete as vigorously, if doing so can be supported, because they realize the large profits being enjoyed by the buyer. In most auction experiments, there is no subject playing the auctioneer and obtaining monetary prizes. This hypothesis is certainly testable; however, we do not pursue it in this paper.

constrained to be nonnegative. This implicitly informs the subjects that bids in excess of zero have a chance of being accepted. Revealing the buyer's value also helps to avoid the possibility of no trade in the first-price auctions, which would occur if the lowest offer exceeded the buyer's value.

The random cost draw for a given period was disclosed to the subject at the beginning of the period. In the first-price auction environment, after learning his cost each seller had a maximum of four minutes to submit his private offer to sell.¹⁶ The computer automatically awarded the item to the seller that submitted the lowest offer once all of the offers had been submitted, provided that the lowest offer was less than 6.00. At the end of the auction, the final market price was announced electronically to all market participants, after which the session proceeded to the next period.

In the multilateral negotiation environment, after learning his cost each seller had a maximum of 30 seconds in the first phase of the period to submit his initial offer to the buyer. The instructions indicated that the seller would then be able to lower his offer at any time in the second phase of the period. Once the buyer received all initial offers, the clock was reset to four minutes for the negotiation phase. At any time during the negotiation phase, a seller could (only) lower his offer, and the buyer could accept the offer of a single seller.¹⁷ Furthermore, a buyer and a specific seller could engage in nonbinding discussions concerning a deal. The sellers could only communicate with their buyer.¹⁸ However, the buyer could negotiate individually with any seller, but only one at a time, while retaining standing offers from the other sellers. This process is meant to parallel the naturally occurring process of a buyer formally soliciting RFPs and then negotiating in person or over the phone until a transaction price is agreed upon. A transcript of the discussions between the buyer and the seller remained on the screen for the duration of the period. The subjects only knew the laboratory identification number of the parties with whom they were communicating. Actual names were not disclosed in the discussions. Once the buyer accepted an offer, the final market price was announced electronically to all market participants,¹⁹ after which the session proceeded to the next period. The sellers had no information on the initial offers or the subsequent discounts made by the other sellers, unless the

¹⁶ The four-minute time limit was never binding.

¹⁷ We avoided the potentially loaded term "negotiation" by calling the negotiation phase a "discussion" phase.

¹⁸ The experiment monitor received numerous requests for the ability of a seller to communicate with other sellers, even though it was explicitly discussed in the instructions that a seller could only communicate with a buyer.

¹⁹ There were only two cases in which the buyer did not accept any offer. Both of these were in the two-seller treatment.

buyer revealed it to them in their discussions. However, the sellers could not verify this information.

The subjects were not told the number of trading periods in the session or in any institutional regime within the session. Moreover, the subjects did not know in advance the rules of any future trading institution, as the instructions for any particular institution were displayed only prior to commencing trade in that institution. It was public information that the same set of sellers was matched with the same buyer for the duration of the experiment.²⁰

Participants received \$5 for showing up on time, plus their salient earnings. In the IPV first-price auctions reported in Cox, Roberson, and Smith [1982], the subjects earned considerably less than the risk-neutral predictions. Hence, calibrating US\$ payoffs on the risk-neutral predictions is not helpful per se in targeting a salient earnings amount for a two-hour experiment. For this reason and because we were holding the support of the cost draws constant for all sessions, the subjects in the two-seller and four-seller treatments were given different exchange rates imputed from the US\$ payoffs of subjects in the aforementioned IPV first-price auction experiments.²¹ In the four-seller sessions, the buyers' exchange rate was US\$1 for 8 experimental dollars, and the sellers' exchange rate was US\$1 for 0.25 experimental dollars. In the two-seller sessions, the exchange amounts were 6 and 2 experimental dollars for each US\$1, respectively. To equalize the buyers' and sellers' earnings expectations, the exchange rates are more favorable to the sellers. This is the case because a buyer receives a payoff every period, but a seller only expects to win every two or four periods, depending on the number of sellers in the treatment. The average subject's earnings (in addition to the \$5 show up fee) for the experiment were \$17.12. The *FFNF* four-seller sessions and all two-seller sessions lasted one hour on average, whereas the *NFFN* four-seller sessions took almost a full two hours. The average US dollar payoffs by subject type and the first-price, risk-neutral Nash predictions (conditional on the set of random draws) are listed in Table 1. We discuss the systematic differences in the next section.

²⁰ Repeated auction play is a common feature of naturally occurring markets and has also been employed in the previous auction experiments discussed above. Those papers examine issues related to static games, as do we.

²¹ We observed that even though Cox, Roberson, and Smith [1982] adjusted the expected risk-neutral payoffs to be identical as the number of bidders increased, many bidders in the auctions with more bidders had lower payoffs.

4. Experimental Results

For each period of play, our data include the institution used, the transaction price, each seller's cost, the buyer's value, as well as each seller's initial and subsequent offers in the multilateral negotiations. We also have a verbatim record of the communications between buyers and sellers. While not used in the statistical analysis, the transcripts provide qualitative insights about the players' strategies and their beliefs about their rivals' strategies. The data permit us to compare the transaction prices and efficiency across institutions, and to examine whether there exist experience and numbers effects.

For clarity's sake, in what follows we summarize our results in a series of five findings. In addition to the qualitative results displayed in tables and figures, we analyze the data using a linear mixed-effects model for repeated measures.²² The results from estimating this model by the four regimes of 46 periods are given in Table 3 below. The results from nonparametric tests, which are reported in Appendix A, lend the same qualitative support. Two exceptions are noted in the evidence for Findings 2b and 3a. The dependent variable is the observed market price. The treatment effects (*Two Sellers* vs. *Four Sellers*, and *Negotiation* vs. *First-price Auction*) and an interaction effect from the 2×2 design are modeled as (zero-one) fixed effects, whereas the 16 independent sessions are modeled as random effects, e_i . Specifically, we estimate the model

$$Price_{ij} = \mathbf{m} + e_i + \mathbf{b}_1 \text{TwoSellers}_i + \mathbf{b}_2 \text{Negotiation}_i + \mathbf{b}_3 \text{TwoSellers}_i \times \text{Negotiation}_i + \mathbf{e}_{ij},$$

where the sessions are indexed by i and the repeated periods by j (e.g., $j = 1, 2, \dots, 12$, for the first regime of twelve periods).²³ We also accommodate heteroskedastic errors by session when estimating the model via restricted maximum likelihood.

We begin by assessing the effect of changing the number of sellers within a specific institution. This finding is a baseline result that establishes that the change in the number of sellers affects transaction prices in a manner consistent with the theoretical predictions from standard oligopoly models.

After comparing some of their observed bid functions with the risk-neutral predictions, we decided to increase the saliency of the four-seller treatment by lowering the exchange rate further.

²² See Laird and Ware [1982] and Longford [1993] for a description of this technique commonly employed in experimental sciences.

²³The linear mixed effects model for repeated measures treats each session as one degree of freedom with respect to the treatments. Hence, with four parameters, the degrees of freedom for the estimates of the treatment fixed effects are $12 = 16 \text{ sessions} - 4 \text{ parameters}$.

Finding 1: *For all regimes, the primary effect of reducing the number of sellers from four to two significantly increases transaction prices.*

Evidence: Table 2 reports the average transaction price for the first 12 periods, by institutional regime and by the number of sellers. The average price in the first-price auctions is 1.87 with four sellers and is 3.10 with two sellers, which is a 66% increase. This percentage price increase nearly matches the predicted theoretical increase of 65%.²⁴ The average price in the multilateral negotiations is 1.88 with four sellers and is 3.69 with two sellers, which is a 97% increase. The results in both institutions suggest that the transaction price increases as the number of sellers decreases. Similar price comparisons can be made for the remaining periods.

Figure 1 displays by treatment the transaction prices in each of the first twelve periods. As with the data presented in Table 2, visual examination of the average prices in the two-seller and four-seller *FFNF* treatments suggests that in each period the transaction prices are higher with fewer sellers. The same conclusion holds for the *NFFN* treatments as well, and similar comparisons can be made for the remaining periods. To formally test these conclusions, we refer to the estimates in Table 3 from linear mixed-effects model for repeated measures. The primary effect of the *Two Seller* treatment is highly significant in all four regimes, raising transaction prices by $\hat{b}_1 = 1.105, 1.009, 2.464,$ and 1.506 experimental dollars, respectively. ■

We now turn our attention to our primary finding that compares the transaction prices and the efficiency of first-price auctions and multilateral negotiations in the initial institutional regime.

Finding 2a: *In the initial institutional regime, transaction prices are statistically indistinguishable in first-price auctions and multilateral negotiations with four sellers. However, multilateral negotiation prices are statistically higher than first-price auction prices with two sellers.*

Evidence: The average prices for the first 12 periods that are reported in Table 2 support this finding. With four sellers, the average price is 1.87 in the first-price auctions and is 1.88 in the multilateral negotiations. With two sellers, the average price is 3.10 in the first-price auctions and is 3.69 in the multilateral negotiations.

²⁴ Conditional on the cost draws, the Nash prediction is that the average price is 2.49 with four sellers and 4.12 with two sellers, or that prices rise by 65%.

Reference to Figure 1, but now comparing the graphs from left to right, illustrates the across-institution differences per period. The price patterns in the four-seller treatments are strikingly similar, while the price patterns in the two-seller treatments are noticeably higher. However, as we discussed in Section 2, ex ante we had no hypothesis whether the multilateral negotiation would grant either the buyers or the sellers more power than they have in a first-price auction. Because the baseline treatment in the linear mixed-effects model is four sellers in first-price auctions, the estimate of the *Negotiation* coefficient ($\hat{\mathbf{b}}_2$) represents the amount by which the *Negotiation* treatment affects transaction prices vis-à-vis first-price auctions, holding the number of sellers constant at four. Table 3 reports that the point estimate for *Negotiation* is nearly zero and is highly insignificant (p -value = 0.9856). This is not too surprising, given our visual examination of the four-seller transaction prices in Table 2 and Figure 1. However, for two sellers the negotiation treatment significantly raises transaction prices by $\hat{\mathbf{b}}_2 + \hat{\mathbf{b}}_3 = -0.004 + 0.884 = 0.880$ experimental dollars above the level for two sellers in first-price auctions (p -value = 0.0038). With two sellers in first-price auctions, the estimated transaction price is $2.98 = \hat{\mathbf{m}} + \hat{\mathbf{b}}_1 = 1.878 + 1.105$, and so the *Negotiation* treatment raises two-seller transaction prices by 29.5% to $3.862 = \hat{\mathbf{m}} + \hat{\mathbf{b}}_1 + \hat{\mathbf{b}}_2 + \hat{\mathbf{b}}_3 = 1.878 + 1.105 - 0.004 + 0.884$. ■

Finding 2b: *In the initial institutional regime, changing the number of sellers does not significantly affect the level of efficiency. However, there is some evidence that with four sellers, multilateral negotiations are slightly more efficient than first-price auctions.*

Evidence: Table 4 reports the average efficiency for the first 12 periods, by institutional regime and by the number of sellers. The observed high efficiency levels are consistent with those reported in previous auction experiments.²⁵ The results from a linear mixed-effects model for the efficiency levels are reported in Table 5. Once again, the baseline treatment is four sellers in first-price auctions. In the auction treatment, reducing the number of sellers from four to two has no effect on efficiency ($\hat{\mathbf{g}} = 1.57$, p -value = 0.2154). A similar result emerges in the negotiation treatment, with the *Two Seller* primary effect offset by the *Two Sellers* \times *Negotiation* interaction effect ($\hat{\mathbf{g}} + \hat{\mathbf{g}} = 1.57 - 3.36 = -1.79$). With four sellers the *Negotiation* treatment significantly raises efficiency levels by $\hat{\mathbf{g}} = 2.98$ percentage points (p -value = 0.0134). However, with two

²⁵ For example, see Cox, Roberson, and Smith [1982] and Cox, Smith, and Walker [1983].

sellers the *Negotiation* primary effect of increasing efficiency is more than offset by the *Two Sellers* \times *Negotiation* interaction effect ($\hat{\beta} + \hat{\beta} = 2.98 - 3.36 = -0.39$). Non-parametric tests do not reject the null hypothesis of equal efficiency levels between the first-price and multilateral negotiation institutions with four sellers or with two sellers. Given that both institutions are highly efficient and that there are several instances in which both institutions are 100% efficient, it is not surprising that a (low power) nonparametric test would fail to detect a mere 3-percentage point difference in the *Negotiation* treatment with four sellers. ■

Findings 2a and 2b report that transaction prices with two sellers are higher in multilateral negotiations than in first-price auctions, but that efficiency is the same. From these two findings, we infer that the change in the trading rules simply results in a transfer of surplus from the buyer to the more efficient of the two sellers. Because the first-price auction price-setting behavior is more aggressive than the risk-neutral Nash prediction, Finding 2a also indicates that the use of multilateral negotiations moves prices toward the risk-neutral predictions for first-price auctions. Findings 4 and 5 explore this last observation at the level of the individual subject. With four sellers, transaction prices are the same across institutions, but the level of efficiency increases, implying that the winning seller receives the same price in either institution, but the lowest cost seller is more likely to transact with multilateral negotiations. We do note, however, that this effect on efficiency is economically small (3 percentage points).

One possible explanation for why prices are higher with multilateral negotiations than with first-price auctions is that the former institution may foster tacit collusion. However, this is not a likely explanation for the observed institution effect because (a) there is no discernable evidence of bid rotation or increasing starting offers, (b) efficiency remains relatively high, and (c) the negotiation transcripts indicate that the sellers are concerned about how competitive their offers are. McAfee and McMillan [1992] show that effective collusion without side payments in first-price auctions requires an inefficient allocation of contracts. It seems reasonable that collusion in multilateral negotiations also requires inefficiency. Therefore, because these two institutions are equally efficient, it is unlikely that the higher transaction prices in the multilateral negotiations can be attributed to collusion. In fact, excluding two periods for which a buyer rejected all offers and purchased nothing, the mean efficiency is actually greater with multilateral negotiations (98.5% versus 96.8%). In Appendix B, selections from the real-time ordered

transcripts from each of the four two-seller *NFFN* discussions indicate that the sellers are attentive to how competitive their offers are.

Our third finding illustrates how the history of trading with different institutions affects market performance. As procurement moves onto the Internet, agents that adopt new institutions will have had experience with other institutions. Firms that employ negotiations in traditional procurement settings may move towards auctions on the Internet, and vice versa. Hence, the impact of historical experience with a particular institution is an interesting question to investigate. We report this finding in three parts, the first of which compares the common 16 periods of first-price auctions in periods 13 through 28. The second part compares the behavior for periods 29 through 40, and the third part assesses the return to the original institution in periods 41 through 46, just prior to which all subjects have equal experience with the institutions. The quantitative support is drawn from Table 3.

Finding 3a: *With two sellers, sellers in first-price auctions who only have prior experience with multilateral negotiations transact at significantly higher prices than sellers who only have prior experience with first-price auctions. With four sellers, there is no difference in transaction prices.*

Evidence: The average prices for periods 13 through 28 that are reported in Table 2 support this finding. With two sellers, the average price is 3.01 for those sellers only with prior auction experience and is 4.01 for those sellers only with prior negotiation experience, which is a 33% increase. With four sellers, the average price is 2.01 for those sellers only with prior auction experience, and it is 2.12 for those sellers only with prior negotiation experience, which is only a 5% increase.

Figure 2 displays by treatment the transaction prices for periods 13 through 28. The sellers in the *FFNF* sessions have had 12 periods of experience with first-price auctions, while the sellers in the *NFFN* sessions have had 12 periods of experience with multilateral negotiations. Recall that the cost draws are common for all bidders of the same seller identification. Visual examination of the average prices illustrates the across-institution differences. The price patterns in the two-seller treatments are quite different, while the price patterns in the four-seller treatments look very similar. The alternative to the null hypothesis that the transaction prices are identical is that a history of negotiation may have a permanent effect such that first-price auction prices are higher in the *NFFN* sessions. We employ a one-sided test for the two-seller treatment because we have already observed that *N* prices are statistically

higher than F prices in the initial 12 periods, but we employ a two-sided test for the four-seller treatment. Referring to Table 3, for two sellers the negotiation treatment significantly raises transaction prices by $\hat{\mathbf{b}}_2 + \hat{\mathbf{b}}_3 = 0.108 + 0.965 = 1.073$ experimental dollars above the level for two sellers in first-price auctions (p -value = 0.0002). However, we find no evidence of a difference in the transaction prices in the four-seller treatment ($\hat{\mathbf{b}}_2 = 0.108$, p -value = 0.6513). A nonparametric test yields the same conclusion for two sellers, but not with four. ■

Finding 3b: *With two sellers, sellers in first-price auctions who have prior experience with multilateral negotiations and first-price auctions transact at higher prices than do sellers in multilateral negotiations who only have prior experience with first-price auctions. With four sellers, there is no difference in transaction prices.*

Evidence: The average prices for periods 29 through 40 that are reported in Table 2 support this finding. With two sellers, the average multilateral negotiation price is 3.22 for those sellers only with prior auction experience, and the average first-price auction price is 4.07 for those sellers with prior negotiation and auction experience. With four sellers, the average price is 1.69 for those sellers only with prior auction experience, and is 1.74 for those sellers with prior negotiation and auction experience.

Figure 3 displays by treatment the transaction prices for periods 29 through 40. The sellers in the *FFNF* sessions have had 28 periods of experience with first-price auctions, while the sellers in the *NFFN* sessions have had 12 periods of experience with multilateral negotiations and 16 periods of experience with first-price auctions. Because of the potentially offsetting effects of the history of the *NFFN* sellers and treatment effect of the multilateral negotiations in the *FFNF* sessions, ex ante we have no clear prediction how prices might differ across the two institutions. Therefore, we employ a two-sided test for equivalence of transaction prices. Referring to Table 3, for two sellers the history of negotiation significantly reduces transaction prices by $-\hat{\mathbf{b}}_2 - \hat{\mathbf{b}}_3 = 0.068 + 0.733 = 0.801$ experimental dollars below the level for two sellers in first-price auctions (p -value = 0.0139). However, we find no evidence of a difference in the transaction prices in the four-seller treatment ($\hat{\mathbf{b}}_2 = -0.068$, p -value = 0.8205). ■

Finding 3c: *With equal amounts of differently ordered experience in both first-price auctions and multilateral negotiations, there is an across-group “return to baseline effect” of the institution treatment. That is, transaction prices are statistically indistinguishable in the first-price auctions and the multilateral negotiations with four sellers, but multilateral negotiation prices are higher than first-price auction prices with two sellers.*

Evidence: The average prices for the final 6 periods that are reported in Table 2 support this finding. With four sellers, the average price is 1.60 in the first-price auctions and is 1.54 in the multilateral negotiations. With two sellers, the average price is 3.01 in the first-price auctions and is 3.52 in the multilateral negotiations.

Figure 4 displays by treatment the transaction prices in each of the final six periods of the sessions. At the start of period 41, all subjects have the same amount of experience in both institutions, but in a different order. At this point in the session each subject has experience with 12 multilateral negotiations and with 28 first-price auctions. Again, we employ a one-sided test for the two-seller treatment because we have already observed that N prices were statistically higher than F prices in the initial 12 periods, but we employ a two-sided test for the four-seller treatment. Referring to Table 3, for two sellers the negotiation treatment significantly raises transaction prices by $\hat{b}_2 + \hat{b}_3 = -0.095 + 0.596 = 0.501$ experimental dollars above the level in first-price auctions (p -value = 0.0371). However, with a two-tailed test we find no evidence of a difference in the transaction prices in the four-seller treatment ($\hat{b}_2 = -0.095$, p -value = 0.7644). ■

The three parts of Finding 3 indicate that the order in which subjects are exposed to the two institutions has a lasting effect on the transaction prices. With two sellers, initial exposure to multilateral negotiations leads to higher prices throughout the session than does initial exposure to first-price auctions. One explanation for these results is that the competition induced by the one binding offer of the first-price auction overwhelms the opportunity for the sellers later to keep prices higher with multilateral negotiations. Another explanation is that the nature of the competition induced by the initial institution may permanently affect how sellers formulate their bidding/negotiating strategies conditional on their cost draws. Regardless of the reason, the experience effect in the two-seller treatment has a nontrivial effect on the subjects’ earnings. The data reported in Table 1 indicate that a Seller 1 first exposed to the two-seller multilateral negotiations earns on average US\$8.49 more (or 76%) than its counterpart first exposed to first-

price auctions, and a Seller 2 earns *twice* as much as its counterpart (US\$14.29 versus US\$7.14). All of this comes at a cost to the buyers, who earn US\$6 less when first exposed to multilateral negotiations.

Finally we turn our attention to assessing the *within-subject* effects of the multilateral negotiation and first-price auction institutions. We focus on the two-seller treatment because, as the above findings indicate, the four-seller treatment yields very similar and competitive pricing results across institutions (as Figures 1 through 4 depict). Figures 5 and 6 display every first-price offer (blue diamonds) and each winning negotiated contract price (red circles) for all 16 sellers in the two-seller treatment. These figures also contain the equilibrium risk-neutral offer function (in yellow) and the 95% confidence intervals for the mean response from the OLS regression of the first-price offer on c , c^2 , c^3 and an intercept ($n = 28$).²⁶ We estimate a cubic offer function because Cox and Oaxaca [1996] find that with a cubic bidding function 999 out of 1,000 first-price bids are consistent with the log-concave model of Cox, Smith, and Walker [1988]. Cox and Oaxaca [1996] also find that only 0-10% of their subjects behave consistently with the linear risk-neutral model.

Finding 4: *The initial institution has a permanent effect on the sellers' behavior in the two-seller treatment. For a given seller, winning multilateral negotiation prices are not greater than the confidence and prediction intervals of the seller's first-price offer function.*

Evidence: It is clear from Figures 5 and 6 that the change in the institution does *not* have much of an impact on the individual (within-subject) behavior. For the multilateral negotiation cost draws in the *NFFN* treatment, only 4 out of 46²⁷ N contract prices are greater than the upper bound of the subjects' prediction intervals.²⁸ (Note that *NFFN* session 2 differs from the other sessions in that the first-price offers *increase* for both subjects *after* experience with multilateral negotiations.) Similarly, only 3 out of 48 N contract prices are greater than the upper bound of the subjects' prediction intervals in the *FFNF* treatment. ■

²⁶ Three obvious outliers, one of which appears to be due to an absent decimal, were omitted from three of the regressions: *FFNF* session 2, subject 2 ($c = 0.40$, $offer = 185$), *FFNF* session 3, subject 1 ($c = 0.21$, $offer = 35$), and *FFNF* session 2, subject 2 ($c = 3.95$, $offer = 1.05$). The respective R^2 's after omitting these data points are 96.7%, 97.4%, and 96.4%.

²⁷ cf. 19.

²⁸ The prediction intervals are slightly larger than the confidence intervals displayed in Figures 5 and 6. The confidence intervals, which are plotted, indicate the capability of the model to explain the relationship between offers and costs for purposes of interpolation (repeated observations for the same cost draw in the first-price offer function). Prediction intervals reveal the first-price offer function's performance for extrapolation for new costs draws in the multilateral negotiation environment.

From Finding 2 we conclude for the two-seller treatment that transaction prices for the two institutions across subjects are initially different (N prices exceed F prices in the first 12 periods). However, as the examination of the individual behavior in Finding 4 indicates, the within-subject behavior is largely unaffected by the change in the institution. The *NFFN* and *FFNF* sequences take separate paths based upon the initial prices, such that N and F prices *at the subject level* are statistically the same for the remaining 34 periods. Hence, we find that both initial institutions have a permanent effect on the individual's offering and negotiation behavior. Finding 5 further examines this effect on the offer functions.

Finding 5: *With two sellers, the first-price offer function at low cost draws is closer to the risk-neutral prediction for sellers with a history of multilateral negotiations than for sellers without a history of multilateral negotiations.*

Evidence: Inspection of Figures 5 and 6 provides clear support for this finding. For 7 of the 8 subjects in the *NFFN* treatment, the confidence intervals for the mean response include or exceed the risk-neutral prediction for low cost draws. Except for the obvious and highly variable seller 1 in *FFNF* session 2, the *FFNF* confidence intervals for the offer functions noticeably lie in the risk-averse area. Table 6 reports the predicted offers from the subject-specific first-price offer functions for two low cost draws, 1.00 and 2.00 (2.00 is the expected value of the minimum order statistic of two $U[0, 6.00]$ draws). One-sided Wilcoxon rank sum tests for the offer predictions with costs of 2.00 and 1.00 lead us to reject the null hypothesis in favor of the alternative that the *NFFN* predicted first-price offers are greater than the *FFNF* predicted first-price offers (for $c = 2.00$: $W = 87$, $n_1 = 8$, $n_2 = 8$, $p\text{-value} = 0.0249$, and for $c = 1.00$: $W = 90$, $n_1 = 8$, $n_2 = 8$, $p\text{-value} = 0.0103$). ■

One explanation for Findings 4 and 5 is that the subjects' institutional experience influences their risk preferences in the first-price auctions.²⁹ An alternative explanation is that initial exposure to an institution affects the subjects' formulation of their strategies in both institutions, but does not influence their risk preferences.

²⁹ See Berg, Dickhaut, and McCabe [1996] for an experiment that finds that individuals' risk preferences are not stable across institutions.

5. Conclusion

In this paper we study an extremely common but previously unexamined exchange process that we term “multilateral negotiations.” In this process, a buyer solicits price offers for a homogeneous product from sellers with heterogeneous costs, and then plays the sellers off one another to obtain additional price concessions. This process is observed in industrial procurement, haggling over new car prices, and wage negotiations, to name but a few examples. Using experimental techniques, we evaluate subject behavior in multilateral negotiations, and we compare the outcomes of multilateral negotiations to the outcomes of first-price auctions.

When subjects have no experience with either institution, we find that transaction prices are statistically indistinguishable in the two institutions with a sufficiently large number of sellers. We also find that prices are higher in multilateral negotiations than in first-price auctions as the number of sellers decreases. However, with two sellers the institutions are equally efficient. With four sellers there is some evidence that multilateral negotiations are slightly more efficient. We also find an experience effect that has remarkable implications regarding the subjects’ risk preferences. With two sellers in first-price auctions, subjects with prior experience only with multilateral negotiations transact at higher prices than do subjects with prior experience only with first-price auctions. Moreover, this ranking of transaction prices across the paired groups is maintained when the second group switches to multilateral negotiations, despite our first finding that multilateral negotiation prices exceed first-price auction prices with inexperienced subjects. Perhaps more surprisingly, we find that for a given seller there is virtually no difference in behavior across the two institutions. Finally, we find that the price-setting behavior in first-price auctions of subjects initially exposed to multilateral negotiations cannot be distinguished from risk-neutral price-setting behavior. This finding is in sharp contrast both to the behavior of subjects initially exposed to first-price auctions, and to the behavior of subjects in virtually all other studies of first-price auctions. We hypothesize that the initial exposure to multilateral negotiations provides the subjects with some insight in formulating their strategies in both institutions, or that the initial institution affects the subjects’ risk preferences.

Our results have implications for institutional design and for antitrust analysis. First, we find that an initial exposure to multilateral negotiations with a small number of sellers leads to a permanently higher path for transaction prices than an initial exposure to first-price auctions, even in subsequent first-price auctions. With a large number of sellers, there is virtually no across-institution difference in the price paths that is related to the subjects’ historical

experience. Hence, if a buyer and his suppliers traditionally have engaged in multilateral negotiations off the Internet, then switching to an online auction may not result in lower prices.

Second, we find that negotiated prices are never less than first-price auction prices, both across and within subject groups. This result suggests that buyers in our setting should prefer to employ first-price auctions rather than multilateral negotiations, given that multilateral negotiations are costly in terms of the time spent determining the transaction price. Of course, this conclusion raises the question of why first-price auctions are not observed more frequently in common transactions. One explanation is that reputation effects create a barrier for buyers trying to implement first-price auctions. For example, a car buyer is a short-run player in the market for new automobiles, and hence is likely to be unconcerned about maintaining a reputation. If the car buyer approaches several dealers and tells them that he wants their best offer, as in a first-price auction, then the sellers would be foolish to actually submit their first-price offers. If the buyer thought he had received first-price offers, then he still would want to haggle with the dealers. Moreover, the dealers might be willing to make concessions if asked, because each knows that if he currently has the second highest offer, then he may yet get a profitable sale by reducing his price. Thus, the buyer's inability to commit to the procurement format likely inhibits his use of what appears to be the preferred institution.

Third, in the initial regime we find that the equivalence of the institutions depends on the number of sellers. Not only can the price level differ substantially across the institutions for certain numbers of sellers, but the percentage price change caused by changing the number of sellers also differs. These findings suggest that caution be used in applying auction models in merger analyses in which transactions more closely resemble multilateral negotiations. For example, our results suggest that if a pair of mergers reduced the number of firms in a market from four to two, then an analysis using an auction model would understate the transaction's price effect compared to an analysis using multilateral negotiations.³⁰

While our results have important implications for institutional design and for antitrust analysis, they are limited by the scope of our experiment and would benefit from further research. First, it would be useful to extend our analysis to settings with different numbers of sellers. At this point, we do not know for what number of sellers multilateral negotiations and

³⁰ Concurrent analysis of two mergers is not merely hypothetical. For example, see *Federal Trade Commission v. Cardinal Health, Inc.*, 12 F. Supp. 2d 34, 1998 U.S. Dist. LEXIS 11778, (D.D.C. 1998) (Sporkin, S.) and *Federal Trade Commission v. McKesson Corp.*, 12 F. Supp. 2d 34, 1998 U.S. Dist. LEXIS 11778, (D.D.C. 1998) (Sporkin, S.). In these two merger cases, the four largest drug wholesalers attempted to consolidate into two firms.

first-price auctions become indistinguishable, or whether the price-concentration relationship is actually linear. In the latter case, multilateral negotiation prices would be less than first-price auction prices for five or more sellers. Second, from the perspective of merger policy, it would be useful to extend our analysis to settings with changes in the number of sellers and with differences in the cost distributions across sellers. To date, these issues have not been explored experimentally in the auction setting, much less in our multilateral negotiation setting. Finally, it would be interesting to let the buyer select his preferred institution, or be unable to commit not to haggle upon receiving the sellers' initial offers. As these examples illustrate, our initial experiment suggests a wide range of interesting research possibilities.

References

- Baker, J. (1997). "Unilateral Competitive Effects Theories in Merger Analysis," *Antitrust* 11, 21-26.
- Berg, J., Dickhaut, J., and McCabe, K. (1996). "Risk Preference Instability Across Institutions: A Dilemma," unpublished manuscript.
- Cason, T. (1995). "An Experimental Investigation of the Seller Incentives in the EPA's Emission Trading Auction," *American Economic Review* 85, 905-922.
- Cohn, L. Brady, D., and Welch, D. "B2B: The Hottest Net Bet Yet?" *Business Week*, 1/17/00, No. 3664, p. 36.
- Cox, J., Roberson, B., and Smith, V. (1982). "Theory and Behavior of Single Object Auctions," Vol. 2, Research in Experimental Economics, General Editor, V. Smith, (Greenwich: JAI, Press).
- Cox, J., Smith, V., and Walker, J. (1988). "Theory and Individual Behavior of First-Price Auctions," *Journal of Risk and Uncertainty* 1(1), 61-99.
- Cox, J., Smith, V., and Walker, J. (1983). "A Test That Discriminates between Two Models of the Dutch-First Auction Non-Isomorphism," *Journal of Economic Behavior and Organization* 4, 205-219.
- Cox, J., and Oaxaca, R. (1996). "Is Bidding Behavior Consistent with Bidding Theory for Private Value Auctions?" Vol. 6, Research in Experimental Economics, General Editor, R. Mark Isaac, (Greenwich: JAI, Press).
- Fudenberg, D., Levine, D., and Tirole, J. (1987). "Incomplete Information Bargaining with Outside Opportunities," *Quarterly Journal of Economics* 102, 37-50.
- Horn, H, and Wolinsky, A. (1988). "Bilateral Monopolies and Incentives for Merger," *RAND Journal of Economics* 19, 408-419.
- Krishna, V., and Serrano, R. (1996). "Multilateral Bargaining," *Review of Economic Studies* 63, 61-80.
- Laird, N. M. and Ware, J. H. (1982). "Random-Effects Models for Longitudinal Data," *Biometrics* 38, 963-974.
- Lehmann, E. L. (1975). Nonparametrics: Statistical Methods Based on Ranks. (San Francisco: Holden and Day).
- Longford, N. T. (1993). Random Coefficient Models. (New York: Oxford University Press).

McAfee, R.P., and McMillan, J. (1987). "Auctions and Bidding," *Journal of Economic Literature* 25, 699-738.

McAfee, R.P., and McMillan, J. (1992). "Bidding Rings," *American Economic Review* 82, 579-599.

McAfee, R.P., and McMillan, J. (1996). "Analyzing the Airwaves Auction," *Journal of Economic Perspectives* 10, 159-175.

McAfee, R.P., and Vincent, D. (1997). "Sequentially Optimal Auctions," *Games and Economic Behavior* 18, 246-276.

Milgrom, P., and Weber, R. (1982). "A Theory of Auctions and Competitive Bidding," *Econometrica* 50, 1089-1122.

Siegel, S. and Castellan, N.J. (1988). Nonparametric Statistics for the Behavioral Sciences. (Boston: McGraw-Hill).

Shaked, A., and Sutton, J. (1984). "Involuntary Unemployment as a Perfect Equilibrium in a Bargaining Model," *Econometrica* 52, 1351-1364.

Smith, V.L. (1982). "Microeconomics as an Experimental Science," *American Economic Review* 72, 923-955.

Su, S., Huang, C., and Hammer, J. (2000). "A Replicable Web-Based Negotiation Server for E-Commerce." *Proceedings of the 33rd Hawaii International Conference on System Sciences*.

Waehrer, K., and Perry, M. (1999). "The Effect of Mergers in Open Auction Markets." unpublished manuscript.

Wolfram, C. (1998). "Strategic Bidding in a Multiunit Auction: An Empirical Analysis of Bids to Supply Electricity in England and Wales," *RAND Journal of Economics* 29, 703-725.

Table 1. Average and Predicted Session Payoffs (US\$) by Sequence and Player Type*

	Average Payoff		Nash Prediction
	<i>FFNF</i>	<i>NFFN</i>	
<i>Four Sellers</i>			
Buyer	23.95	23.68	19.83
Seller 1	19.98	25.63	56.20
Seller 2	9.99	15.32	44.78
Seller 3	19.55	22.99	62.21
Seller 4	11.22	10.70	48.32
<i>Two Sellers</i>			
Buyer	22.30	16.24	14.52
Seller 1	11.18	19.67	23.83
Seller 2	7.14	14.29	19.74

*These payoffs do not include the \$5 payment for showing up on time.

Table 2. Average Transaction and Predicted Nash* Prices by Regime

Sequence	<i>Two Sellers</i>		<i>Four Sellers</i>	
	Observed	Nash Predictions	Observed	Nash Predictions
<i>FFNF</i>				
Periods 1-12: <i>F</i>	3.10	4.12	1.87	2.49
Periods 13-28: <i>F</i>	3.01	4.01	2.01	2.79
Periods 29-40: <i>N</i>	3.22	4.11	1.69	2.43
Periods 41-46: <i>F</i>	3.01	4.07	1.60	2.31
<i>NFFN</i>				
Periods 1-12: <i>N</i>	3.69	4.12	1.88	2.49
Periods 13-28: <i>F</i>	4.01	4.01	2.12	2.79
Periods 29-40: <i>F</i>	4.07	4.11	1.74	2.43
Periods 41-46: <i>N</i>	3.52	4.07	1.54	2.31

*The predicted Nash prices are conditional on the draws and the assumption of risk-neutral price-setting. The ex ante Nash predictions are a price of 4.00 with two sellers and 2.40 with four sellers.

Table 3. Estimates of the Linear Mixed-Effects Model

$$Price_{ij} = \mathbf{m} + e_i + \mathbf{b}_1 \text{TwoSellers}_i + \mathbf{b}_2 \text{Negotiation}_i + \mathbf{b}_3 \text{TwoSellers}_i \times \text{Negotiation}_i + \mathbf{e}_{ij},$$

$$\text{where } e_i \sim N(0, \mathbf{s}_1^2), \mathbf{e}_{ij} \sim N(0, \mathbf{s}_{2,i}^2).$$

	Estimate	Std. Error	Degrees of Freedom*	t-statistic	p-value
<i>Regime 1: Periods 1 – 12</i>					
m	1.878	0.166	174	11.302	0.0000
<i>Two Sellers</i>	1.105	0.268	12	4.117	0.0014 [†]
<i>Negotiation</i>	-0.004	0.233	12	-0.018	0.9856
<i>Two Sellers</i> × <i>Negotiation</i>	0.884	0.339	12	2.612	0.0227
			190 Obs.		
				$H_a: \mathbf{b}_2 + \mathbf{b}_3 \neq 0$	0.0038
<i>Regime 2: Periods 13 – 28</i>					
m	2.012	0.167	240	12.018	0.0000
<i>Two Sellers</i>	1.009	0.246	12	4.107	0.0008 [†]
<i>Negotiation History</i>	0.108	0.233	12	0.463	0.6513
<i>Two Sellers</i> × <i>Negotiation History</i>	0.965	0.313	12	3.079	0.0048 [†]
			256 Obs.		
				$H_a: \mathbf{b}_2 + \mathbf{b}_3 > 0$	0.0002 [†]
<i>Regime 3: Periods 29 – 40</i>					
m	1.757	0.209	176	8.406	0.0000
<i>Two Sellers</i>	2.464	0.249	12	9.900	0.0000 [†]
<i>Negotiation</i>	-0.068	0.292	12	-0.232	0.8205
<i>Two Sellers</i> × <i>Negotiation</i>	-0.733	0.403	12	-1.817	0.0942
			192 Obs.		
				$H_a: \mathbf{b}_2 + \mathbf{b}_3 \neq 0$	0.0139
<i>Regime 4: Periods 41 – 46</i>					
m	1.605	0.227	80	7.064	0.0000
<i>Two Sellers</i>	1.506	0.278	12	5.421	0.0001 [†]
<i>Negotiation</i>	-0.095	0.309	12	-0.307	0.7644
<i>Two Sellers</i> × <i>Negotiation</i>	0.596	0.401	12	1.486	0.0816 [†]
			96 Obs.		
				$H_a: \mathbf{b}_2 + \mathbf{b}_3 > 0$	0.0371 [†]

*N.B. The linear mixed effects model for repeated measures treats each session as one degree of freedom with respect to the treatments. Hence, with four parameters, the degrees of freedom for the estimates of the treatment fixed effects are 12 = 16 sessions – 4 parameters.

[†]One-sided test.

Note: The linear mixed-effects model is fit by restricted maximum likelihood. The results do not differ in any meaningful way if the model is estimated with autocorrelated errors by session. For brevity, the session random effects are not included in the table.

Table 4. Average Efficiency by Regime

Sequence	Two Sellers	Four Sellers
<i>FFNF</i>		
Periods 1-12: <i>F</i>	96.4%	96.8%
Periods 13-28: <i>F</i>	99.5%	98.8%
Periods 29-40: <i>N</i>	97.1%	97.1%
Periods 41-46: <i>F</i>	98.3%	98.6%
<i>NFFN</i>		
Periods 1-12: <i>N</i>	94.3%*	96.8%
Periods 13-28: <i>F</i>	96.1%	98.7%
Periods 29-40: <i>F</i>	92.5%	98.9%
Periods 41-46: <i>N</i>	97.6%	98.6%

*This mean includes two observations for which the buyer rejected both final offers and purchased nothing in those periods, resulting in 0% efficiency. The statistical tests in Finding 2b include these two observations. Excluding these two observations, the mean efficiency is 98.5%.
 Note: Efficiency is defined to be $(6 - \text{winner's cost}) / (6 - \text{lowest cost draw}) \times 100\%$.

Table 5. Estimates of the Linear Mixed-Effects Model for Efficiency Levels

$$Efficiency_{ij} = \mathbf{m} + e_i + \mathbf{g}TwoSellers_i + \mathbf{g}Negotiation_i + \mathbf{g}TwoSellers_i \times Negotiation_i + \mathbf{e}_{ij},$$

where $e_i \sim N(0, \mathbf{s}_1^2)$, $\mathbf{e}_{ij} \sim N(0, \mathbf{s}_{2,i}^2)$.

	Estimate (in %)	Std. Error	Degrees of Freedom*	<i>t</i> -statistic	<i>p</i> -value
<i>Regime 1: Periods 1 – 12</i>					
<i>m</i>	97.02	1.03	176	94.45	0.0000
<i>Two Sellers</i>	1.57	1.20	12	1.31	0.2154
<i>Negotiation</i>	2.98	1.03	12	2.90	0.0134
<i>Two Sellers</i> × <i>Negotiation</i>	-3.36	1.58	12	-2.12	0.0552
192 Obs.					

Table 6. Subject-Specific Offer Predictions for the Two-Seller Treatment

Session, Subject	$c = 1.00$		$c = 2.00$	
	<i>FFNF</i>	<i>NFFN</i>	<i>FFNF</i>	<i>NFFN</i>
Session 1, Subject 1	1.67	2.26	2.72	2.68
Session 1, Subject 2	1.59	2.96	2.11	3.43
Session 2, Subject 1	4.61	3.82	4.89	3.97
Session 2, Subject 2	1.60	4.20	2.70	4.27
Session 3, Subject 1	1.48	3.52	2.43	3.83
Session 3, Subject 2	1.69	3.95	2.56	4.09
Session 4, Subject 1	2.37	4.12	3.61	4.51
Session 4, Subject 2	2.69	3.82	3.24	4.22
Average	2.21	3.58	3.03	3.87
Risk Neutral Prediction	3.50		4.00	

Appendix A

This appendix reports nonparametric tests as support for Findings 1, 2, and 3.

Finding 1:

We use the average per-period prices across the sessions shown in Figure 1 in a one-tailed Wilcoxon signed rank test for data paired by period.³¹ We reject the null hypothesis of equal transaction prices in favor of the alternative that the two-seller prices are higher than the four-seller prices in both institutions (for first-price auctions: $V = 77$, $p\text{-value} = .0005$, $n = 12$, and for multilateral negotiations: $V = 77$, $p\text{-value} = .0005$, $n = 12$).

Finding 2a:

Again we use the average per-period prices across the sessions shown in Figure 1 in a two-tailed Wilcoxon signed rank test for data paired by period. We cannot reject the null hypothesis that transaction prices are equal in the first-price auction and multilateral negotiation environments with four sellers ($V = 38$, $p\text{-value} = .9096$, $n = 12$).

However, we reject the null hypothesis of equal transaction prices in the two institutions when there are only two sellers ($V = 72$, $p\text{-value} = .0068$, $n = 12$). In part, this rejection occurs because the Wilcoxon signed rank statistic accounts for the magnitude of the differences in prices, and for the two-seller test there are nine out of the twelve paired periods in which the average multilateral negotiation price exceeds the average first-price auction price. The differences in price for these nine periods range from 47 to 159 experimental cents, whereas the three exceptions range from 3 to 37 experimental cents.

Finding 2b:

We use the average per-period efficiency levels across the sessions in a two-tailed Wilcoxon signed rank test for data paired by period. We cannot reject the null hypothesis of equal efficiency levels between the first-price and multilateral negotiation institutions with four sellers ($Z = -1.354$,³² $p\text{-value} = .1759$, $n = 12$) or with two sellers ($Z = 0.119$, $p\text{-value} = .9053$, $n = 12$).

³¹ See Siegel and Castellan [1988].

Finding 3a:

We employ a one-sided test for the two-seller treatment because we have already observed that N prices are statistically higher than F prices in the initial 12 periods, but we employ a two-sided test for the four-seller treatment. The Wilcoxon signed rank tests using the per-period average prices shown in Figure 2 lead us to reject the null hypothesis in favor of the alternative for both the four-seller and two-seller treatments (for the four-seller treatment: $V = 114$, $p\text{-value} = .0155$, $n = 16$, and for the two-seller treatment: $V = 130$, $p\text{-value} = .0002$, $n = 16$).

Finding 3b:

For the two-seller treatment, we reject the null hypothesis in favor of the alternative of different prices ($V = 76$, $p\text{-value} = .0015$, $n = 12$). However, we find no evidence of a difference in the transaction prices in the four-seller treatment ($V = 42$, $p\text{-value} = .8501$, $n = 12$).

Finding 3c:

Using a one-sided Wilcoxon signed rank test we fail to reject the null hypothesis of equality in the four-seller sessions ($V = 8$, $p\text{-value} = .7188$, $n = 6$). However, we find a significant difference in the two-seller sessions ($V = 20$, $p\text{-value} = .0312$, $n = 6$).

³² Because there are ties for several periods where both treatments are 100% efficient, an exact probability for the test cannot be computed. The normal approximation given by Lehmann [1975, p. 130] is used in this case.

Appendix B

The following selections are taken from the real time ordered transcript for the four two-seller *NFFN* discussions indicates (bold added for emphasis).

Session 1, period 11:

(Seller 1's cost is 0.20 with an initial offer of 3.50. Seller 2's cost is 0.41 with an initial offer of 4.15.)

[Buyer to Seller 1]: can you go down to 3.00?
[Seller 1]: yes
[Buyer to Seller 2]: they are lower once again....is it possible to go down quite a bit?
[Seller 1]: will that work best for you?
[Seller 2]: tell me a price?
[Buyer to Seller 2]: what's the lowest you can possibly go?
[Seller 2]: Tell me there price
[Buyer to Seller 2]: can you beat 3.00?
[Seller 1]: if not I can sacrafice 2.94
[Buyer to Seller 2]: actually...they just went below that..can u go any lower?
[Seller 1]: do you have a better offer?
[Seller 2]: If we don't deal wuicker then we lose money. you arem aking money everytime so don't barter so much
[Buyer to Seller 1]: woah...they just went down a lot...can you go any lower? you are both real close.

The buyer accepted Seller 1's offer of 2.00.

Session 2, period 1:

(Seller 1's cost is 1.23 with an initial offer of 5.00. Seller 2's cost is 2.81 with an initial offer of 7.50.)

[Buyer to Seller 1]:
[Seller 1]: yes??
[Seller 2]: hi there
[Buyer to Seller 1]: seller 2 has offered me \$3. Can you beat that?
[Seller 1]: 2.90
[Seller 2]: is my price too high?
[Seller 1]: so?
[Buyer to Seller 2]: seller 1 just offered me 2.50 can you beat that
[Seller 1]: i give u 2.90
[Seller 2]: ha, nope. I'd lose a ton of money!
[Seller 2]: maybe next time
[Buyer to Seller 1]: seller 2 just countered with 2.25, can you beat that
(Note: The buyer is lying. Seller 2 never lowered his initial offer of 7.50.)
[Seller 1]: ok??

The buyer accepted Seller 1's offer of 2.20.

Session 3, period 7:

(Seller 1's cost is 2.37 with an initial offer of 3.75. Seller 2's cost is 2.25 with an initial offer of 4.50.)

[Seller 2]: give me an offer
[Buyer to Seller 2]: How about 3.50
[Seller 2]: make it 4.00
[Seller 2]: how is that
[Buyer to Seller 2]: sorry

The buyer accepted Seller 1's offer of 3.75.

Session 4, period 11:

(Seller 1's cost is 0.20 with an initial offer of 5.75. Seller 2's cost is 0.41 with an initial offer of 3.82.)

[Seller 1]: This is ridiculous...
[Seller 2]: This is a low price for ya
[Buyer to Seller 2]: sold at 2.75
[Buyer to Seller 1]: don't think you can do this round
[Seller 1]: where you at
[Buyer to Seller 1]: goota show me 3.25
[Seller 2]: How's 3.50
[Seller 1]: I'll show 4
[Buyer to Seller 2]: 3.0 is a deal
[Buyer to Seller 1]: down to 3.2
[Seller 2]: 3.15?
[Buyer to Seller 1]: 3.0 sells
[Buyer to Seller 2]: gotta compete
[Buyer to Seller 2]: show me 3.0
[Seller 1]: where you at
[Buyer to Seller 1]: i'm at 2.8
[Seller 2]: it's only 5 cents
[Seller 1]: what happened to 3
[Buyer to Seller 2]: gotta beat 3.0 now
[Buyer to Seller 1]: competition
[Seller 2]: 2.98
[Seller 1]: competition, hahahahaha
[Buyer to Seller 2]: show me 2.75 and i buy
[Buyer to Seller 1]: gotta beat 2.8
[Buyer to Seller 2]: gotta buy from the lowest bed
[Buyer to Seller 2]: beat 2.7

The buyer accepted Seller 2's offer of 2.65.

Seller 1's submitted offers were: 5.75, 5.50, 4.00, 3.50, 3.29, 3.00, 2.80, and 2.75.

Seller 2's offers were: 3.82, 3.50, 3.25, 3.05, 2.98, 2.85, and 2.65.







Figure 1. Transaction Prices by Treatment for Periods 1-12







Figure 2. Transaction Prices by Treatment for Periods 13-28

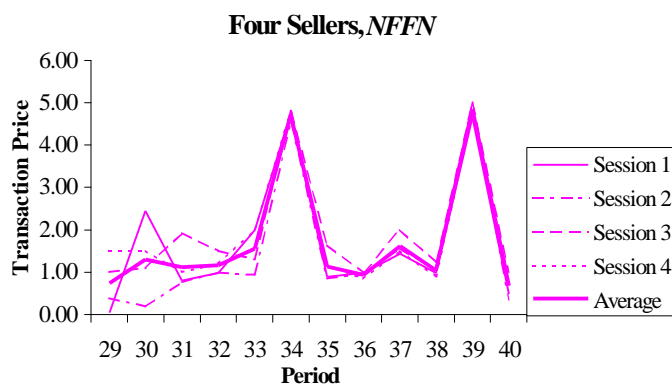





Figure 3. Transaction Prices by Treatment for Periods 29-40







Figure 4. Transaction Prices by Treatment for Periods 40-46

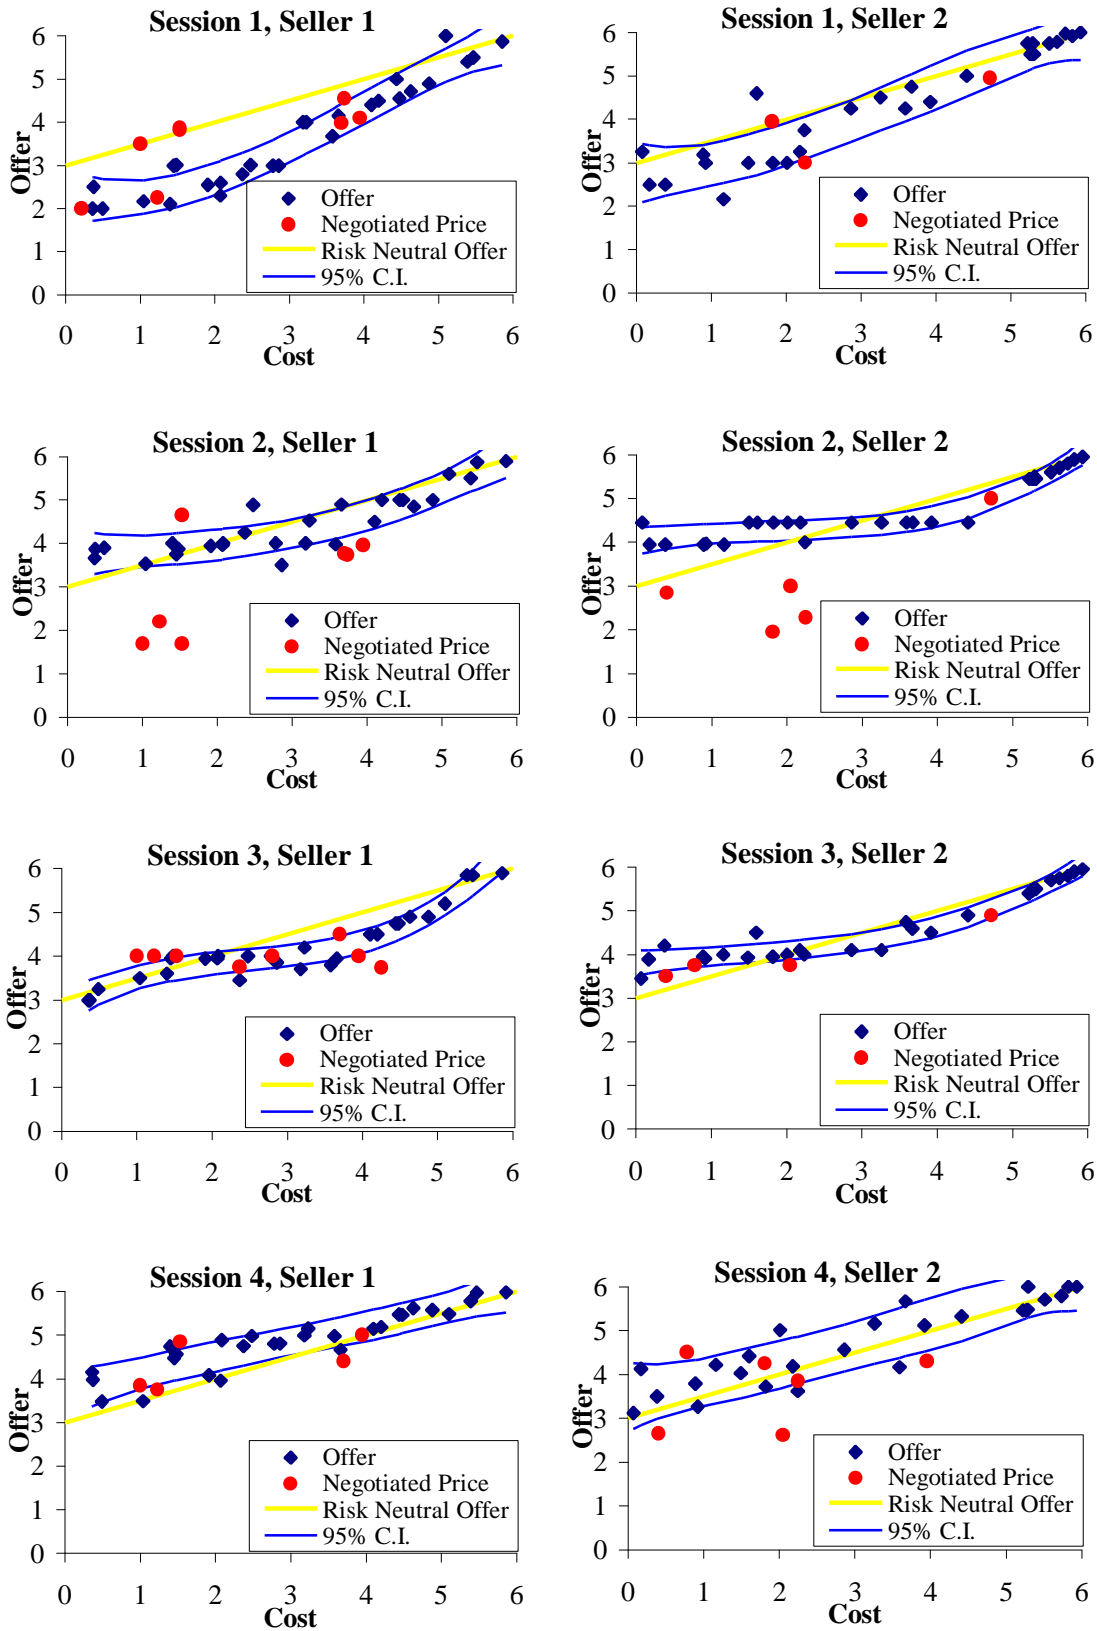


Figure 5. Negotiated Contract Prices and Offer Functions for *NFN* Two-Seller Treatment

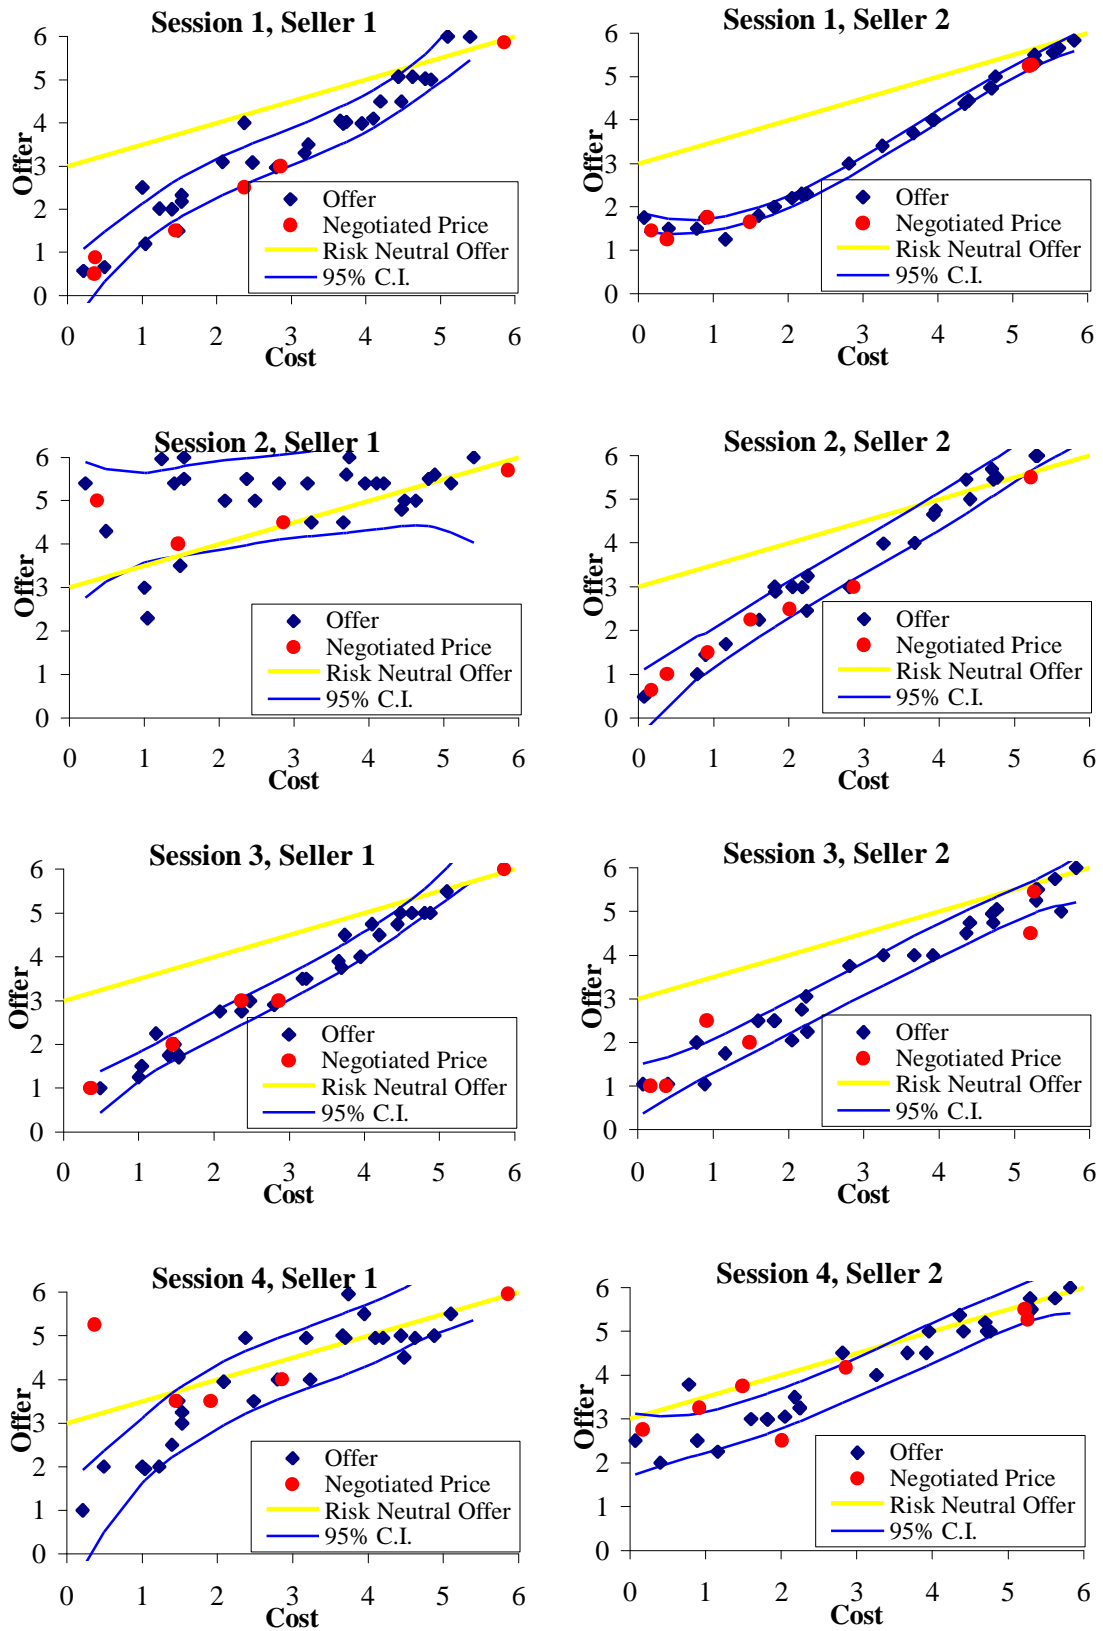


Figure 6. Negotiated Contract Prices and Offer Functions for *FFNF* Two-Seller Treatment