Equilibrium Behavior in a Model of Multilateral Negotiations*
PRELIMINARY DRAFT

Charles J. Thomas
Clemson University

September 3, 2010

Abstract

This paper characterizes equilibrium behavior in multilateral negotiations, an exchange process in which a buyer attempts to purchase from one of several sellers with privately known production costs. The buyer receives an initial set of price offers, then can try to play the sellers off one another to obtain concessions. Consistent with recent experimental evidence, expected prices in the negotiations equal those in first-price auctions, despite differences in the time path of price offers. However, the buyer's inability to commit not to haggle can lead to costly delay, and being more patient can weaken the buyer's bargaining position.

1 Introduction

Negotiations often involve one party trying to consummate a transaction with one of several potential trading partners. For example, a soon-to-be-minted Ph.D. who receives multiple tenure-track job offers might try to leverage one school's offer to obtain concessions from other schools, along dimensions like salary, startup funds, or teaching load. Similarly, Iberia Airlines' purchase of new planes in 2003 involved a "months-long dogfight" between Airbus and Boeing over terms including price, maintenance, crew training, and the number of seats.[1] Even corporate takeover battles can take this form. Kraft's recent takeover of Cadbury was dogged by the prospect of counteroffers by Hershey and Ferrero,[2] while Fisher and Merck recently fought for control of the laboratory supplies maker Millipore.[3]

The preceding examples feature a process of voluntary exchange that consists of initial offers to an individual from multiple potential trading partners who have private information about the magnitude of gains from trade, followed by the prospect of additional rounds of offers that continue until an agreement is reached or the negotiations are abandoned. Importantly, the individual can hold an offer while shopping around for a better one, and can communicate with its potential trading partners about where they stand during the negotiations.

*Email: cjt@clemson.edu; Mail: 222 Sirrine Hall, Clemson University, Clemson, SC 29634. I have received helpful advice from participants in the Junior Faculty Brownbag and the IO workshop at Clemson University, and from seminar participants at the University of Alicante.
JEL: C78, D44, D82
Keywords: negotiations, bargaining, auctions, dynamic games, incomplete information

[1] "Airbus and Boeing Duke it Out to Win Lucrative Iberia Deal," Wall Street Journal, March 10, 2003.
[2] "Ferrero rules out Cadbury counterbid," *Financial Times*, January 26, 2010, p. 20.
[3] "Merck in $7.2bn Millipore deal," *Financial Times*, March 1, 2010, p. 17.

Exchange conducted in this fashion arguably embodies what many economists would consider to be basic elements of competition, but the economics literature has given relatively little attention to exchange exhibiting all of these features. Analyses of bargaining tend to emphasize bilateral settings, consider exogenously specified outside options rather than endogenously determined ones, or require one party to abandon current negotiations before beginning talks with a new potential trading partner. Analyses of auctions focus on multilateral settings, but lack the communication and interplay that are inherent features of the negotiations described above.

A recent exception is experimental research by Thomas and Wilson [2002, 2005, 2010] that investigates an exchange process they call "multilateral negotiations." In the experiments, a buyer seeking to purchase from one of several sellers receives initial offers from them. The buyer may accept one offer immediately, or try to use the initial offers to play the sellers off one another to obtain pricing concessions. The negotiations are conducted via "chat" technology over a computer network, and are extremely unstructured in terms of the content of the communication. In every scenario considered in the experiments, the sellers have privately known production costs that affect the magnitude of the gains from trade.

The preceding experiments were conducted before the development of theoretical models that closely approximate the environment studied, but this has become a standard approach for heuristic investigations in settings in which theoretical arguments are limited by the environment's strategic complexity.[4] In fact, one hope is that such an approach "can help stimulate the construction and study of new (and useful) game-theoretic models of bargaining."[5]

In this paper I develop a simple model of multilateral negotiations that contains the relevant structural features: multiple sellers, with privately known costs, who can make multiple offers that can be held while the buyer seeks concessions from rival sellers. In particular, I assume there are two sellers, their costs can be one of two levels, and each seller has two opportunities to make offers. Given the strategic complexity that soon will become evident, it seems sensible to use such a simple model to gain some initial insights and to provide guidance about fruitful approaches for enriching the strategic environment.

One goal of modeling this setting is to evaluate theoretical issues like the effect of introducing another seller into a bilateral bargaining setting, or the effect of the parties' patience on negotiation outcomes. Another goal is to see if the theory predicts two experimental results in the papers by Thomas and Wilson. The first is that multilateral negotiation outcomes can depend on the buyer's ability to reveal credibly to a seller the price offers made by rival sellers. The second is that the negotiation outcomes may closely resemble the outcomes from different auction formats.

Several results emerge from the equilibrium analysis, and I highlight the following ones because they relate to some commonly held perceptions about negotiation outcomes. First, negotiated prices on average are the same as those in first-price and second-price auctions, but negotiated outcomes may entail costly delay. Second, being more patient can weaken the buyer's bargaining position when it can credibly reveal rivals' offers, which contrasts with standard intuition from bilateral negotiation models. Impatience lets the buyer credibly commit not to haggle, which induces the sellers to make aggressive initial offers that are accepted immediately. Third, when the buyer cannot credibly reveal rivals' offers, claims about superior offers may not induce price cuts. In the relevant equilibria, a seller will not cut its price if the buyer rejects its initial offer. Instead the seller holds firm at its initial offer, and refuses to be drawn into offering potentially needless discounts. The first result is consistent with recent experimental evidence, but the second and

[4] For example, see Smith [1982].
[5] Muthoo [1999, p. 342]

third have not been evaluated experimentally.

2 Background Literature

Earlier I wrote that existing theoretical research has not combined all of the essential features of multilateral negotiations in a single model. To provide context for my analysis, below I explain more fully the variety of approaches that consider some elements of the strategic environment I investigate.[6] All of them are drawn from the literatures on bargaining or auctions.

In terms of general overview, Kennan and Wilson [1993] and Ausubel, Cramton, and Deneckere [2002] provide extensive background information on the bargaining literature, with an emphasis on settings with two parties and incomplete information. McAfee and McMillan [1987] and Milgrom [1989] likewise survey the auction literature, which has as a central consideration transactions that involve a single auctioneer facing multiple bidders who have private information about the gains from trade.

Fudenberg and Tirole [1983] provide an early treatment of sequential bilateral bargaining with incomplete information that has formed the basis for much subsequent research.[7] They use a two-period model with one buyer and one seller, each of whom has private information that can take on only one of two values. Among other things, they find that agreement may be delayed even if it would occur for sure with only one bargaining round, and that a player is not necessarily harmed by being more impatient. My model closely follows their variant with one-sided incomplete information, but introduces a second seller to compete for the buyer's business.

Shaked and Sutton [1984] model bilateral bargaining with complete information and alternating offers. The buyer bargains with one seller at a time, but can switch to a different seller by incurring a commonly known cost. Offers from one seller are void upon switching to another seller, and all sellers have the same commonly known production 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. The Walrasian outcome is supported as the switching cost goes to zero, and it is identical to the outcome if the buyer instead conducted a first-price auction between the two sellers. My model incorporates incomplete information, and lets the buyer bargain simultaneously with both sellers.

Fudenberg, Levine, and Tirole [1987] also consider the ability to switch to a new trading partner, but in a setting with incomplete information. Specifically, the uninformed party does not know the magnitude of the gains from trade with the current trading partner, and can switch to another trading partner at any time. The uninformed party makes all of the offers, and offers are void upon switching trading partners. The authors find that the ability to change partners has a much different effect than in the complete information setting of Shaked and Sutton [1984]. My model considers simultaneous bargaining with multiple trading partners, and has the informed parties make all of the offers.

Reinganum and Daughety [1991, 1992] consider a model in which a buyer can secure multiple offers by engaging in costly search. They assume that all sellers have the same commonly known production cost, and that competition yields a price equal to marginal cost if the buyer obtains two offers. My model assumes

[6] For the purpose of distinguishing terminology it is worth noting that the multilateral negotiations considered here differ from the multilateral bargaining examined in papers such as Krishna and Serrano [1996]. Multilateral bargaining involves more than two agents bargaining over the division of a common surplus. This is distinct from the present analysis, in which the buyer is interested in dealing with only one seller, and the size of the pie with different sellers can vary in a privately known way.

[7] For example, see Cramton, Sobel and Takahashi [1983], Chatterjee and Samuelson [1987], and other references in Ausubel, Cramton, and Deneckere [2002].

the sellers have privately known production costs, explicitly analyzes the competition that ensues when the buyer obtains offers from multiple sellers, and ignores the search aspect.

Vincent [1992] considers the effect of adding an additional seller to a bilateral bargaining setting, in what seems to be the closest analysis to the present one. Assuming that the two sellers have the same commonly known production cost, he finds that the buyer may not secure all gains from trade for certain institutional rules governing exchange. This result stems from the sellers' ability to support collusive equilibria through the prospect of potentially an infinite number of bargaining periods. My model differs from his by assuming incomplete information and a finite horizon.

Chatterjee and Lee [1998] examine a bargaining scenario with complete information in which the buyer can hold an offer from one seller while it incurs a cost to acquire an offer from another seller. The significant difference from my model is that the authors assume that the competing offer is a draw from a known distribution, rather than being the outcome of strategic interaction with another seller.

McAfee and Vincent [1997] model an auction environment in which the buyer imposes an announced reserve price, but cannot commit not to auction the item in the future if the current reserve price is not met. They find that the auction can take several rounds to complete. The main difference between our models is that in theirs the reserve price is a single number rather than a function of all of the offers. In multilateral negotiations, the buyer makes its decision to accept or reject a current offer by using the set of received offers to forecast future offers.

3 The Model

Consider a setting in which two risk-neutral sellers produce homogeneous products and compete to fulfill an indivisible contract for one risk-neutral buyer. Each seller's privately known production cost is a random draw from a commonly known binomial distribution. A seller is **low-cost** (c_L) with probability $\alpha \in [0,1]$, and is **high-cost** (c_H) with probability $1-\alpha$, where $0 = c_L < c_H$. The buyer's commonly known value for a single unit of the sellers' products is $c_H + \gamma$, where $\gamma > 0$ denotes the minimum extent of the gains from trade. The buyer and sellers discount future payoffs according to the discount factors $\{\delta_B, \delta_S\} \in [0,1]^2$.

Trade is conducted through multilateral negotiations, which begin in period 1 with each seller simultaneously submitting an initial price offer after it learns its production cost. If the buyer accepts one of the initial offers, then the negotiations conclude. If the buyer rejects both initial offers, then in period 2 each seller submits an offer that must be weakly lower than its initial offer.[8] If an offer is accepted, then the transaction price is the price offered by the winning seller. A transaction that occurs in period $t \in \{1,2\}$ at price p yields payoffs $\delta_B^{t-1}(c_H + \gamma - p)$ and $\delta_S^{t-1}(p - c_w)$ for the buyer and the winning seller, respectively, where c_w is the winning seller's cost. The losing seller receives 0. If no transaction occurs in either period, then each party's payoff is 0.

The experimental results in Thomas and Wilson [2005] suggest that the buyer's ability to reveal credibly to a seller the offers it has from other sellers influences the negotiation outcomes. For that reason I analyze both **verifiable** and **nonverifiable** multilateral negotiations, which differ based on the buyer's ability to reveal sellers' offers credibly. With verifiable offers, each seller learns its rival's offer if the game continues to period 2. With nonverifiable offers, each seller learns only that the buyer is seeking new offers from the sellers.

[8] This constraint embodies the idea that "rejection" of an offer, in the sense of seeking better offers, does not mean that the current offer is void. The constraint on future prices also brings another consideration into the seller's decision about its initial price offer, because any initial offer constrains future offers should the buyer reject the current offer.

The players' strategies consist of the following elements. In period 1 each seller makes a nonnegative initial price offer weakly less than $c_H + \gamma$. The buyer accepts one of the offers, or rejects them both. If period 2 is reached, then each seller makes a nonnegative subsequent price offer that can be contingent upon its available information, and that must be weakly less than its initial offer. The buyer accepts one of the offers, or rejects them both.

The solution concept used is Perfect Bayesian Equilibrium (PBE). To derive a PBE of this game, one determines each seller's initial offer, the buyer's rejection decision regarding the initial offers, beliefs about a seller's cost based on what is observed after period 1, each seller's subsequent offer, and the buyer's rejection decision regarding the subsequent offers. I examine PBEs that are pooling or separating, depending on the nature of the initial offers.

Before proceeding it is worth describing behavior in bilateral bargaining and in a first-price auction. Comparing equilibria in the multilateral negotiations to the equilibria from these benchmarks illustrates the effects of multiple sellers and the prospect of multiple rounds of offers. The following two results are well-known.

Result 1 (Bilateral Bargaining) *Consider a variant of the multilateral negotiation model in which there is only one seller. The expected price, the seller's expected payoff, and the buyer's expected payoff are*

$$\overline{P}^{bb} = c_H + \gamma \qquad \overline{\pi}_1^{bb} = \alpha c_H + \gamma \qquad \overline{\pi}_B^{bb} = 0.$$

Trade occurs without delay.

Because the parties with private information make all of the offers in the multilateral negotiations, the relevant comparison is to a bilateral bargaining model in which the informed party makes all of the offers. In this case the uninformed buyer earns no surplus, so comparing these bilateral negotiations to multilateral negotiations measures the benefits that accrue to the buyer from increasing the intensity of competition amongst the informed sellers.[9]

Result 2 (First-Price Auction) *Consider a variant of the multilateral negotiation model in which the two sellers each can make only one offer. The expected price, the sellers' expected payoffs, and the buyer's expected payoff are*

$$\overline{P}^{fpa} = \left(1-\alpha^2\right) c_H \qquad \overline{\pi}_1^{fpa} = \overline{\pi}_2^{fpa} = \alpha\left(1-\alpha\right) c_H \qquad \overline{\pi}_B^{fpa} = \alpha^2 c_H + \gamma.$$

Trade occurs without delay.

Result 2 is based on McAfee [1994], and Krishna [2002] argues that mechanism gives the buyer the highest expected payoff of all possible exchange mechanisms, when attention is restricted to mechanisms that yield trade with probability 1. One implication is that the buyer must be weakly worse off if it cannot commit not to haggle with the sellers, as is the case in the multilateral negotiations.

When the sellers are symmetric in terms of the distribution of their production costs, two other exchange mechanisms yield the same expected price and payoffs as does the first-price auction. One is a second-price auction, and the other is a version of the multilateral negotiation model with complete information about the sellers' costs. However, both relationships break down when the sellers are asymmetric.

[9] Result 1 assumes the buyer has no bargaining power. One also could consider how the buyer's best outcome with one seller compares to the outcome with two sellers, along the lines of the analysis in Bulow and Klemperer [1996].

4 Multilateral Negotiations with Verifiable Offers

Consider a multilateral negotiation setting in which the buyer is able to reveal credibly to a seller the offer from the other seller. In this case a seller does not have to worry that the buyer is lying about the existence of a superior rival offer, but does have to worry that having its rival see its offer may lead to more aggressive offers in period 2, should the buyer reject both initial offers.

Rather than specify a "communication" game describing how the buyer informs each seller about its rival's offer, I simply assume that all sellers' offers are publicly revealed after period 1. This assumption makes common knowledge each party's beliefs about the sellers' privately known production costs.

4.1 Pooling Equilibria

In a pooling equilibrium, each seller's initial offer reveals no information about its cost. Thus, the buyer uses its prior beliefs when forecasting its expected future payoff from rejecting both initial offers. Similarly, a seller's offer in period 2 is based on its prior beliefs about its rival's cost.

Proposition 1 *With verifiable offers, there exist pooling equilibria if and only if the buyer is sufficiently patient that*

$$\delta_B \geq \frac{c_H + \gamma - \delta_S (1-\alpha) c_H}{c_H + \gamma - \delta_S (1-\alpha) c_H + \delta_S \alpha^2 (1-\alpha) c_H}.$$

The equilibrium path of each pooling equilibrium features the following behavior: Each seller makes an initial offer that is independent of its cost, drawn from the range $[c_H, c_H + \gamma]$ according to some probability distribution that may be seller-specific. The buyer rejects both offers with probability 1. In period 2, the sellers make offers as they would in a first-price auction. In all pooling PBEs the expected price, the sellers' expected payoffs, and the buyer's expected payoff are

$$\overline{P}^v = \left(1-\alpha^2\right) c_H \qquad \overline{\pi}_1^v = \overline{\pi}_2^v = \delta_S \alpha (1-\alpha) c_H \qquad \overline{\pi}_B^v = \delta_B \left[\alpha^2 c_H + \gamma\right].$$

There exist $\delta_B < 1$ such that pooling PBEs exist if and only if $\alpha \in (0,1)$ and $\delta_S \in (0,1]$.

There is a multiplicity of pooling PBEs, all of which involve non-serious initial offers of c_H or higher, and all of which the buyer rejects. With the exception of delay, outcomes in the pooling equilibria are the same as outcomes in a first-price auction. Moreover, for there to be a pooling equilibrium the buyer's discount factor must be fairly high. Therefore, the difference between the two exchange methods is not too great, in terms of the buyer's expected payoff. However, the sellers can be much worse off in these equilibria, relative to a first-price auction. If δ_S is quite low while δ_B is high enough to support a pooling equilibrium, then the delay is very costly to the sellers.

While the buyer must be sufficiently patient for there to exist pooling PBEs, Proposition 1's constraint on δ_B is derived from the constraint preventing a low-cost seller from profitably deviating to an initial offer low enough to induce the buyer to accept. The issue for a deviating low-cost seller is that a rejected low initial offer constrains the second offers from that seller and from a low-cost rival, which benefits the buyer. The more patient is the buyer, the more willing it is to wait until period 2 after observing a given deviation. For a sufficiently patient buyer, a deviating low-cost seller must offer such a significant concession to win in period 1 that it prefers to pool with its initial offer and wait to compete until period 2.

It is instructive to consider parameter values for which pooling PBEs do not exist. If $\delta_S = 0$, then a low-cost seller will deviate to an initial offer just above 0 that is sure to be accepted, because the anticipated payoff in period 2 is worth nothing in present value terms. Similarly, if $\alpha = 1$, then a low-cost seller knows the price in period 2 will be 0, and hence will deviate to an initial offer above 0 that is sure to be accepted. If $\alpha = 0$, then the buyer knows the price in period 2 will be c_H. A high-cost seller expects a payoff of 0 in period 2, and so will deviate to an initial offer just above c_H that is sure to be accepted. If high-cost sellers instead are making initial offers of c_H in any candidate pooling PBE, then the buyer will purchase in period 1 rather than wait to get the same price in period 2.

4.2 Separating Equilibria

In a separating equilibrium, each seller's initial offer reveals its cost. Based on this information, the buyer must decide whether the price it anticipates paying in period 2 is sufficiently attractive that it is worth rejecting the lower initial offer. This choice depends on the entire set of initial offers, not simply the best one. For this reason the analysis of multilateral negotiations differs from the analysis of optimal reserve prices over sequences of auctions performed by McAfee and Vincent [1997]. In their analysis the reserve price that must be satisfied is not contingent on the entire set of offers received.

Proposition 2 *With verifiable offers, there exist separating equilibria for all parameter values. The equilibrium path of each separating equilibrium features high-cost sellers making initial offers of c_H, and low-cost sellers making initial offers that depend on the buyer's discount factor. If the buyer is sufficiently impatient that*

$$\delta_B \leq \frac{\gamma}{c_H + \gamma},$$

then low-cost sellers make initial offers identical to those from a first-price auction. The buyer accepts the lower initial offer with probability 1. If the buyer is of "intermediate" patience such that

$$\frac{\gamma}{c_H + \gamma} \leq \delta_B \leq \frac{\alpha c_H + \gamma}{c_H + \gamma},$$

then low-cost sellers make initial offers that are drawn from the range $[(1-\alpha)\, c_H, (1-\delta_B)\,(c_H+\gamma)]$ according the distribution used in a first-price auction, or that equal $c_H - \epsilon$. The buyer accepts the lower initial offer unless both offers equal $c_H - \epsilon$, in which case both offers are rejected and the price in period 2 is 0. If the buyer is sufficiently patient that

$$\frac{\alpha c_H + \gamma}{c_H + \gamma} \leq \delta_B,$$

then low-cost sellers make initial offers equal to $c_H - \epsilon$. The buyer accepts the lower initial offer unless both offers equal $c_H - \epsilon$, in which case both offers are rejected and the price in period 2 is 0. In all separating PBEs the expected price and the sellers' expected payoffs are

$$\overline{P}^v = \left(1-\alpha^2\right) c_H \qquad \textit{and} \qquad \overline{\pi}_1^v = \overline{\pi}_2^v = \alpha\,(1-\alpha)\, c_H.$$

The buyer's expected payoff varies with its discount factor, with

$$\overline{\pi}_B^v = \begin{cases} \alpha^2 c_H + \gamma & \textit{for } \delta_B \in \left[0, \frac{\gamma}{c_H+\gamma}\right] \\ \alpha^2 c_H + \gamma - \frac{(1-\alpha)^2[\delta_B c_H - (1-\delta_B)\gamma]^2}{(1-\delta_B)(c_H+\gamma)} & \textit{for } \delta_B \in \left[\frac{\gamma}{c_H+\gamma}, \frac{\alpha c_H+\gamma}{c_H+\gamma}\right] \\ \left(1-\alpha^2\right)\gamma + \alpha^2 \delta_B\,(c_H+\gamma) & \textit{for } \delta_B \in \left[\frac{\alpha c_H+\gamma}{c_H+\gamma}, 1\right] \end{cases}$$

For the separating equilibria in Proposition 2, the expected price and the sellers' expected payoffs are identical to those from the first-price auction described in Result 2, but the actual pricing behavior may look radically different. The differences in pricing behavior affect only the buyer's expected payoff, and they emerge as follows.

If the buyer is sufficiently impatient, then low-cost sellers' initial offers are drawn from the same probability distribution as in a first-price auction. These PBEs rely upon the sellers' anticipation that the buyer will accept the lower initial offer along the equilibrium path of play, so effectively there is only one round of offers. This anticipation is fueled either by the gains from trade being sufficiently large (γ large or c_H small), or by the buyer being sufficiently impatient (δ_B small). In both cases the buyer has no incentive to reject equilibrium initial offers to receive lower offers in period 2.

If the buyer is sufficiently patient, then it rejects the initial offers if both are from low-cost sellers and are sufficiently high. In fact, in equilibrium there is a cutoff price that depends on δ_B, below which low-cost sellers draw their initial offers from the same distribution as in a first-price auction. Prices above the cutoff are accepted only if the rival is revealed to be high-cost. Thus, a low-cost seller considering an initial offer above the cutoff has an incentive to set a price arbitrarily closely below c_H. Lower prices above the cutoff win in precisely the same circumstances, but yield a lower payoff, so the seller is better off abandoning the middle of the price-setting distribution in favor of the extremes. Once δ_B is sufficiently high, low-cost sellers completely abandon low prices. Instead they make high initial offers of $c_H - \epsilon$ that are accepted only if the rival is revealed to be high-cost.

Proposition 2 yields four results that contrast with the standard intuition from bilateral bargaining models that your being more patient is good for you and bad for your rival. The first two results are that the buyer's expected payoff is highest when it is least patient, and that the buyer's expected payoff is not monotonic in its discount factor δ_B. The buyer's expected payoff is invariant to its discount factor for low values of δ_B. Once δ_B reaches a value that induces low-cost sellers to begin avoiding mid-range prices, the buyer's expected payoff begins to decline. Although delay is less costly as δ_B increases, that positive effect is swamped by the greater likelihood of delay as δ_B goes up. Once δ_B reaches a level where low-cost sellers' initial offer exclusively is $c_H - \epsilon$, then the buyer's expected payoff begins to increase as δ_B increases. At this point, increasing δ_B does not increase the likelihood of delay, so the only effect of increasing δ_B is to increase the present value of the buyer's payoff. Finally, when δ_B reaches 1 the buyer's expected payoff returns to the level that could be achieved with low values of δ_B.

The next two results are that the sellers' payoffs are invariant to the buyer's discount factor δ_B, and that all parties' payoffs are invariant to the sellers' discount factor δ_S. For all δ_B, high-cost sellers receive 0, and low-cost sellers get an expected payoff equal to what they would receive in a first-price auction. Low-cost sellers receive positive payoffs only in period 1, so their discount factor δ_S has no effect on the present value of their payoffs.[10] Finally, the sellers' pricing behavior is unaffected by their discount factor, and hence the buyer's expected payoff also is unaffected by δ_S.

5 Multilateral Negotiations with Nonverifiable Offers

Considering nonverifiable offers introduces an additional complication into this dynamic game of incomplete information. If the initial offers are rejected, then each seller now must infer its rival's cost not based on

[10] Also note that the nature of the separating equilbrium is invariant to δ_S. This suggests that if the sellers each had a different discount factor, a seller's payoff would be unaffected by changes in its or its rival's discount factor.

the rival's initial offer, but based on the buyer's rejection of both offers. The basic problem this difference creates can be seen in the following simple example.

Consider a candidate equilibrium in which the buyer rejects the initial offers if both sellers are revealed to be low-cost, and accepts the best initial offer otherwise. A low-cost seller whose serious initial offer is rejected anticipates that its rival also is low-cost, and hence will make a subsequent offer of 0. However, even if the buyer receives only one serious initial offer, it may be worth rejecting simply to induce the single low-cost seller to so significantly, and unnecessarily, improve upon its initial offer. Without verifiability, a seller must be wary that a rejection of its initial offer may be part of the buyer's attempt to get the seller to compete against itself.

As in the setting with verifiable offers, equilibria with nonverifiable offers may be distinguished by whether the initial offers are pooling or separating. I examine each type of equilibrium in turn.

5.1 Pooling Equilibria

The lack of verifiability eliminates the possibility of pooling PBEs.

Proposition 3 *With nonverifiable offers, there do not exist pooling equilibria.*

To understand why pooling equilibria do not exist with nonverifiable offers, recall that they exist with verifiable offers because the buyer's rejection of a low-cost seller's deviating offer below c_H affects prices in period 2. This effect is absent with nonverifiable offers, because a non-deviating rival seller expected its initial offer would be rejected. Consequently, its subsequent offer is unaffected by its rival's deviation, because it does not realize that a deviation occurred. In any candidate pooling PBE, the sellers' offers are at least c_H, and are rejected with probability 1. The buyer expects a price of $\left(1-\alpha^2\right)c_H$ in period 2, and a low-cost seller expects a payoff of $(1-\alpha)\,c_H$. A low-cost seller can profitably deviate to an initial offer of $(1-\alpha)\,c_H$. If that offer is rejected, the seller wins for sure in period 2 by setting the same price. Winning with this price in period 1 is better than receiving the same expected payoff in period 2, and for the buyer the deviating initial offer is the same price it anticipates receiving in period 2 if it rejects both initial offers. Therefore, low-cost sellers have an incentive to deviate from any candidate pooling equilibrium, so there can be no such equilibria.

5.2 Separating Equilibria

There exists a separating equilibrium that exhibits behavior identical to that in a first-price auction, but without conditions on the buyer's degree of impatience, like those present with verifiable offers.

Proposition 4 *With nonverifiable offers, for all δ_B there exists a separating equilibrium in which the sellers make initial offers identical to those from a first-price auction, and the buyer accepts the lower initial offer with probability 1. In this separating equilibrium the expected price, the sellers' expected payoffs, and the buyer's expected payoff are*

$$\overline{P}^n=\left(1-\alpha^2\right)c_H \qquad \overline{\pi}_1^n=\overline{\pi}_2^n=\alpha\,(1-\alpha)\,c_H \qquad \overline{\pi}_B^n=\alpha^2 c_H+\gamma.$$

Trade occurs without delay.

In the equilibrium in Proposition 4, sellers expect the buyer to purchase in period 1. Consequently, their offers are the same as in a first-price auction. If the buyer unexpectedly rejects both initial offers, each seller can hold any beliefs about its rival's likelihood of being low-cost. If the sellers' beliefs remain unchanged from their prior beliefs, then each seller has no incentive to reduce its price below its initial offer. Therefore, the sellers stand firm at their initial offers in the face of rejection without verifiability.

The separating PBE with nonverifiable offers does not require sufficient impatience on the part of the buyer, unlike the similar separating PBE with verifiable offers. The difference with nonverifiable offers is that the buyer is unable to induce price cuts by the sellers, even if both of them are low-cost. Consequently, it is rational for the buyer to accept any initial offer on the equilibrium path, for any discount factor.

6 Conclusion

This paper develops a model of multilateral negotiations that reveals a close relationship between these negotiations and first-price auctions. In every equilibrium, expected prices are the same in the negotiations and in the auctions, and expected payoffs are weakly lower in the negotiations. Lower expected payoffs arise from delay in consummating a transaction, and all parties are not necessarily affected in the same way.

The model uses simplifying assumptions that may affect both quantitative and qualitative aspects of the results. In particular, one might naturally consider continuously distributed private information, an infinite horizon, and more sellers. However, the model introduced here should be considered a first step toward developing a better understanding of this commonly observed means of exchange.

7 Appendix

7.1 Pooling Equilibria with Verifiable Offers

The sellers' strategies in any pooling PBE with verifiable offers take the following form. Let $P_i \subseteq [c_H, c_H + \gamma]$ denote seller i's set of initial pooling offers, and let $\sigma_i(p)$ denote seller i's probability distribution over P_i. Both P_i and $\sigma_i(p)$ depend on seller i's identity, but not on seller i's realized cost.

Let $\widehat{\alpha}_i(p_i)$ denote the updated belief that seller i is low-cost, based seller i's initial offer p_i. I assume that the buyer and both sellers form the same beliefs $\widehat{\alpha}_i(p_i)$.

Proof of Proposition 1: Let p_1 and p_2 denote the initial offers made by sellers 1 and 2, and denote the lower and higher initial offers by $p_L = \min[p_1, p_2]$ and $p_H = \max[p_1, p_2]$. Consider the following beliefs and strategies.

Beliefs: Beliefs take the following form:

$$\begin{aligned}
\widehat{\alpha}_i(p_i) &= \alpha, \text{ if } p_i \in P_i \\
\widehat{\alpha}_i(p_i) &\geq \alpha, \text{ if } p_i \notin P_i,\ p_i \geq c_H \\
\widehat{\alpha}_i(p_i) &= 1, \text{ if } p_i \notin P_i,\ p_i < c_H
\end{aligned}$$

A seller whose initial offer is drawn from its set of pooling offers is believed to be low-cost with the same probability as the prior belief. A seller whose initial offer is not drawn from its set of pooling offers, but weakly exceeds c_H, is believed to be low-cost with at least as great a probability as the prior belief. Finally, a seller whose initial offer is strictly less than c_H is believed to be low-cost with probability 1.

Buyer's strategy: With respect to the initial offers p_1 and p_2, there are six distinct cases to consider:

1. If $p_i \in P_i$ and $p_j \in P_j$, then reject both initial offers.

2. If $p_i \in P_i$ and $p_j \notin P_j$, with $p_j \geq c_H$, then reject both initial offers.

3. If $p_i \in P_i$ and $p_j \notin P_j$, with $p_j < c_H$, then reject both initial offers if

$$\frac{(1-\delta_B)(c_H+\gamma)}{1-\delta_B(1-\alpha^2)} \leq p_j,$$

and otherwise accept p_j.

4. If $p_i \notin P_i$, with $p_i \geq c_H$, and $p_j \notin P_j$, with $p_j \geq c_H$, then reject both initial offers.

5. If $p_i \notin P_i$, with $p_i \geq c_H$, and $p_j \notin P_j$, with $p_j < c_H$, then reject both initial offers if

$$\frac{(1-\delta_B)(c_H+\gamma)}{1-\delta_B\left(1-\widehat{\alpha}_i^2(p_i)\right)} \leq p_j,$$

and otherwise accept p_j.

6. If $p_i \notin P_i$, with $p_i < c_H$, and $p_j \notin P_j$, with $p_j < c_H$, then reject both initial offers if

$$(1-\delta_B)(c_H+\gamma) \leq p_L,$$

and otherwise accept p_L.

In period 2, accept the lower offer if it is less than or equal to $c_H + \gamma$, and otherwise reject both offers.

High-cost seller i's strategy: Draw initial offer p_i from P_i according to $\sigma_i(p_i)$. In period 2, make an offer equal to the smaller of c_H and p_i.

Low-cost seller i's strategy: Draw initial offer p_i from P_i according to $\sigma_i(p_i)$. In period 2, there are nine distinct cases to consider.

1. If $p_i \in P_i$ and $p_j \in P_j$, then draw offer in period 2 from

$$F_1(p) = \frac{p-(1-\alpha)c_H}{\alpha p}.$$

2. If $p_i \in P_i$ and $p_j \notin P_j$, with $p_j \geq c_H$, then draw offer in period 2 from

$$F_2(p) = \frac{p-(1-\alpha)c_H}{\alpha p}.$$

3. If $p_i \in P_i$ and $p_j \notin P_j$, with $p_j < c_H$, then draw offer in period 2 from

$$F_3(p) = \frac{p-(1-\alpha)p_j}{\alpha p}.$$

4. If $p_i \notin P_i$, with $p_i \geq c_H$, and $p_j \in P_j$, then draw offer in period 2 from

$$F_4(p) = \frac{p-(1-\alpha)c_H}{\widehat{\alpha}_i(p_i)\,p}.$$

5. If $p_i \notin P_i$, with $p_i \geq c_H$, and $p_j \notin P_j$, with $p_j \geq c_H$, then draw offer in period 2 from

$$F_5(p) = \frac{p - (1 - \min[\widehat{\alpha}_i(p_i), \widehat{\alpha}_j(p_j)])\, c_H}{\widehat{\alpha}_i(p_i)\, p}.$$

6. If $p_i \notin P_i$, with $p_i \geq c_H$, and $p_j \notin P_j$, with $p_j < c_H$, then draw offer in period 2 from

$$F_6(p) = \frac{p - (1 - \widehat{\alpha}_i(p_i))\, p_j}{\widehat{\alpha}_i(p_i)\, p}.$$

7. If $p_i \notin P_i$, with $p_i < c_H$, and $p_j \in P_j$, then draw offer in period 2 from

$$F_7(p) = \frac{p - (1 - \alpha)\, p_i}{p}.$$

8. If $p_i \notin P_i$, with $p_i < c_H$, and $p_j \notin P_j$, with $p_j \geq c_H$, then draw offer in period 2 from

$$F_8(p) = \frac{p - (1 - \widehat{\alpha}_j(p_j))\, p_i}{p}.$$

9. If $p_i \notin P_i$, with $p_i < c_H$, and $p_j \notin P_j$, with $p_j < c_H$, then offer 0 in period 2.

To determine whether or not the preceding strategies and beliefs form a PBE, it is necessary to check that no player has an incentive to deviate, given its beliefs, and to check that the beliefs are consistent with the strategies and with Baye's Rule.

Consistency of beliefs: Given the sellers' strategies, observing $p_i \in P_i$ provides no information about seller i's cost. Thus, the belief $\widehat{\alpha}_i(p_i) = \alpha$ is consistent with the strategies and with Baye's Rule. Observing $p_i \notin P_i$ is a zero-probability event. Hence, any belief is consistent with the strategies and with Baye's Rule.

Optimality of buyer's strategy: First, consider the buyer's decision in period 2. Accepting the lower offer is preferable to rejecting it and earning 0 whenever the lower offer yields a non-negative payoff. This occurs whenever the lower offer is less than or equal to $c_H + \gamma$.

Now consider the buyer's decision to accept an initial offer, based on the six distinct cases of both initial offers.

1. If $p_i \in P_i$ and $p_j \in P_j$, then both sellers are believed to be low-cost with probability α. If the buyer rejects both initial offers, then it anticipates a price of $(1 - \alpha^2)\, c_H$ in period 2. Thus, it is rational for the buyer to reject both offers if and only if

$$\underbrace{c_H + \gamma - p_L}_{\text{payoff from accepting}} \leq \underbrace{\delta_B \left[c_H + \gamma - (1 - \alpha^2)\, c_H\right]}_{\text{payoff from rejecting}} \Longleftrightarrow \frac{c_H + \gamma - p_L}{\alpha^2 c_H + \gamma} \leq \delta_B.$$

2. If $p_i \in P_i$ and $p_j \notin P_j$, with $p_j \geq c_H$, then sellers i and j are believed to be low-cost with respective probabilities α and $\widehat{\alpha}_j(p_j) \geq \alpha$. If the buyer rejects both initial offers, then it anticipates a price of $(1 - \alpha^2)\, c_H$ in period 2. Thus, it is rational for the buyer to reject both offers if and only if

$$\underbrace{c_H + \gamma - p_L}_{\text{payoff from accepting}} \leq \underbrace{\delta_B \left[c_H + \gamma - (1 - \alpha^2)\, c_H\right]}_{\text{payoff from rejecting}} \Longleftrightarrow \frac{c_H + \gamma - p_L}{\alpha^2 c_H + \gamma} \leq \delta_B.$$

3. If $p_i \in P_i$ and $p_j \notin P_j$, with $p_j < c_H$, then sellers i and j are believed to be low-cost with respective probabilities α and 1. If the buyer rejects both initial offers, then it anticipates a price of $\left(1-\alpha^2\right) p_j$ in period 2. Thus, it is rational for the buyer to reject p_j if and only if

$$\underbrace{c_H + \gamma - p_j}_{\text{payoff from accepting}} \leq \underbrace{\delta_B \left[c_H + \gamma - \left(1-\alpha^2\right) p_j\right]}_{\text{payoff from rejecting}} \Longleftrightarrow \frac{\left(1-\delta_B\right)\left(c_H+\gamma\right)}{1-\delta_B\left(1-\alpha^2\right)} \leq p_j.$$

4. If $p_i \notin P_i$, with $p_i \geq c_H$, and $p_j \notin P_j$, with $p_j \geq c_H$, then sellers i and j are believed to be low-cost with respective probabilities $\widehat{\alpha}_i\left(p_i\right) \geq \alpha$ and $\widehat{\alpha}_j\left(p_j\right) \geq \alpha$. If the buyer rejects both initial offers, then it anticipates a price of $\left(1-\min\left[\widehat{\alpha}_i\left(p_i\right), \widehat{\alpha}_j\left(p_j\right)\right]^2\right) c_H$ in period 2. Thus, it is rational for the buyer to reject p_L if and only if

$$\underbrace{c_H + \gamma - p_L}_{\text{payoff from accepting}} \leq \underbrace{\delta_B \left[c_H + \gamma - \left(1-\min\left[\widehat{\alpha}_i\left(p_i\right), \widehat{\alpha}_j\left(p_j\right)\right]^2\right) c_H\right]}_{\text{payoff from rejecting}} \Longleftrightarrow \frac{c_H+\gamma-p_L}{\min\left[\widehat{\alpha}_i\left(p_i\right), \widehat{\alpha}_j\left(p_j\right)\right]^2 c_H+\gamma} \leq \delta_B.$$

5. If $p_i \notin P_i$, with $p_i \geq c_H$, and $p_j \notin P_j$, with $p_j < c_H$, then sellers i and j are believed to be low-cost with respective probabilities $\widehat{\alpha}_i\left(p_i\right) \geq \alpha$ and 1. If the buyer rejects both initial offers, then it anticipates a price of $\left(1-\widehat{\alpha}_i^2\left(p_i\right)\right) p_j$ in period 2. Thus, it is rational for the buyer to reject p_j if and only if

$$\underbrace{c_H + \gamma - p_j}_{\text{payoff from accepting}} \leq \underbrace{\delta_B \left[c_H + \gamma - \left(1-\widehat{\alpha}_i^2\left(p_i\right)\right) p_j\right]}_{\text{payoff from rejecting}} \Longleftrightarrow \frac{\left(1-\delta_B\right)\left(c_H+\gamma\right)}{1-\delta_B\left(1-\widehat{\alpha}_i^2\left(p_i\right)\right)} \leq p_j.$$

6. If $p_i \notin P_i$, with $p_i < c_H$, and $p_j \notin P_j$, with $p_j < c_H$, then both sellers are believed to be low-cost with probability 1. If the buyer rejects both initial offers, then it anticipates a price of 0 in period 2. Thus, it is rational for the buyer to reject p_L if and only if

$$\underbrace{c_H + \gamma - p_L}_{\text{payoff from accepting}} \leq \underbrace{\delta_B \left[c_H + \gamma - 0\right]}_{\text{payoff from rejecting}} \Longleftrightarrow \left(1-\delta_B\right)\left(c_H+\gamma\right) \leq p_L.$$

Optimality of high-cost seller i's strategy: First, consider high-cost seller i's decision in period 2, given the initial offers p_i and p_j. If $p_i \geq c_H$, then setting a price of c_H is a best-response to the strategy of either a high-cost or a low-cost rival. Similarly, if $p_i < c_H$, then setting a price of p_i is a best-response to the strategy of either a high-cost or a low-cost rival. Next, consider high-cost seller i's decision about its initial offer, p_i. Given the buyer's strategy and the rival seller's strategy (either high-cost or low-cost), there is no incentive to deviate from an initial offer drawn from P_i according to $\sigma_i\left(p_i\right)$. With an initial offer $p_i < c_H$, the best possible outcome is for the rival to win in period 2, and the worst is for seller i to win and incur a loss, in either period. An initial offer $p_i \geq c_H$, with $p_i \notin P_i$, will not be accepted, and will not affect the seller's expected payoff in period 2.

Optimality of low-cost seller i's strategy: First, consider low-cost seller i's decision in period 2, given the initial offers p_i and p_j. For each of the nine distinct cases, period 2 effectively is a first-price auction with beliefs as specified above. Such an auction has a unique Nash equilibrium, with pricing for seller i as specified in its strategy. Next, consider low-cost seller i's decision about its initial offer. Given the buyer's strategy and the rival seller's strategy, there is no incentive to deviate to an initial offer $p_i \notin P_i$, with $p_i \geq c_H$. If $p_i > p_j$, seller i's best-case scenario is that both initial offers are rejected, which provides an expectation of expected profits in period 2 that are identical to what they would be with an initial offer $p_i \in P_i$. If $p_i < p_j$ and is rejected, then seller i's expected payoff in period 2 is identical to what it would be with an initial offer $p_i \in P_i$. An initial offer strictly less than c_H that is rejected serves only to decrease the seller's expected payoff in period 2 from what it would be with an initial offer equal to c_H. Thus, a deviating low-cost seller will deviate only to a price that is sure to be accepted. From the arguments above, the optimal price for a deviation is the highest initial offer that will be accepted in this case,

$$p_L = \frac{(1-\delta_B)(c_H+\gamma)}{1-\delta_B(1-\alpha^2)}.$$

In order for this deviation to be less profitable than making an initial offer drawn from P_i, it must be the case that

$$\underbrace{\frac{(1-\delta_B)(c_H+\gamma)}{1-\delta_B(1-\alpha^2)}}_{\text{payoff from optimal deviation}} \leq \underbrace{\delta_S(1-\alpha)c_H}_{\text{payoff from pooling}} \Longleftrightarrow \frac{c_H+\gamma-\delta_S(1-\alpha)c_H}{c_H+\gamma-\delta_S(1-\alpha)c_H+\delta_S\alpha^2(1-\alpha)c_H} \leq \delta_B.$$

This constraint on δ_B is tighter than is the constraint that prevents the buyer from accepting an initial offer equal to c_H, so for a pooling PBE to exist requires

$$\frac{c_H+\gamma-\delta_S(1-\alpha)c_H}{c_H+\gamma-\delta_S(1-\alpha)c_H+\delta_S\alpha^2(1-\alpha)c_H} \leq \delta_B.$$

If $\alpha \in (0,1)$ and $\delta_S > 0$, then the critical value of δ_B is strictly less than 1. Hence, for those parameter values there exists a pooling PBE. ■

7.2 Separating Equilibria with Verifiable Offers

The sellers' strategies in any separating PBE with verifiable offers take the following form. Let $P_i^H \subseteq [c_H, c_H+\gamma]$ and $P_i^L \subseteq [0, c_H+\gamma]$ respectively denote the sets of initial price offers for high-cost and low-cost seller i, with $P_i^H \cap P_i^L = \varnothing$. Let $\sigma_i^H(p_i)$ and $\sigma_i^L(p_i)$ respectively denote seller i's probability distribution over the sets of initial offers when it is high-cost or low-cost. Both P_i^k and $\sigma_i^k(p)$ depend on seller i's identity, and on seller i's realized cost.

Proof of Proposition 2: As a first step I demonstrate that P_i^H contains only c_H, and that $P_i^L \subseteq [0, c_H]$. Letting p_L and p_H denote the lower and higher initial offers, consider the following possible combinations of initial offers.

1. If both sellers are revealed to be high-cost, then the buyer will buy today provided that

$$\underbrace{c_H+\gamma-p_L}_{\text{payoff from accepting}} \geq \underbrace{\delta_B[c_H+\gamma-c_H]}_{\text{payoff from rejecting}} \Longleftrightarrow c_H+(1-\delta_B)\gamma \geq p_L.$$

2. If both sellers are revealed to be low-cost, then the buyer will buy today provided that

$$\underbrace{c_H+\gamma-p_L}_{\text{payoff from accepting}} \geq \underbrace{\delta_B[c_H+\gamma-0]}_{\text{payoff from rejecting}} \Longleftrightarrow c_H-\delta_B c_H+(1-\delta_B)\gamma \geq p_L.$$

3. If one seller is revealed to be low-cost and one is revealed to be high-cost, then the buyer will buy today provided that

$$\underbrace{c_H + \gamma - p_L}_{\text{payoff from accepting}} \geq \underbrace{\delta_B \left[c_H + \gamma - \min\left[c_H, p_L\right]\right]}_{\text{payoff from rejecting}}.$$

If $p_L \leq c_H$, then the preceding condition is $c_H + \gamma \geq p_L$, which is automatically satisfied. If $p_L > c_H$, then the preceding condition is $c_H + (1 - \delta_B)\gamma \geq p_L$, which is identical to the condition in (1).

Let $\underline{p}_i^k = \inf P_i^k$ and $\overline{p}_i^k = \sup P_i^k$, for $k \in \{H, L\}$. Based on the conditions in (1)-(3), $\overline{p}_i^k \leq c_H + (1 - \delta_B)\gamma$, for $k \in \{H, L\}$.

Consider low-cost seller i making an initial offer $p_i > c_H + (1 - \delta_B)\gamma$. If seller j also is low-cost, then seller i's payoff is 0. Either seller j's initial offer is accepted, or both initial offers are rejected and the price in period 2 is 0. If seller j is high-cost, then seller i's payoff either is 0 (because seller j's initial offer is accepted), or is $\delta_S (c_{\dot{H}} - \epsilon)$ (because seller i wins in period 2 at price c_H). Suppose seller i instead makes an initial offer of $c_H - \epsilon$. Against a low-cost seller j, seller i's expected payoff weakly increases. Against a high-cost seller j, seller i definitely wins today, so seller i earns at least $c_H - \epsilon > \delta_S (c_{\dot{H}} - \epsilon)$. Hence, low-cost seller i will not make an initial offer $p_i > c_H + (1 - \delta_B)\gamma$, so $\overline{p}_i^L \leq c_H + (1 - \delta_B)\gamma$.

Consider high-cost seller i making an initial offer $p_i > c_H + (1 - \delta_B)\gamma$. If seller j also is high-cost, then seller i's payoff is 0. Either seller j's initial offer is accepted, or both initial offers are rejected and the price in period 2 is c_H. If seller j is low-cost, then seller i's payoff is 0. Either seller j's initial offer is accepted, or both initial offers are rejected and seller i loses in period 2. Suppose seller i instead makes an initial offer of $c_H + \epsilon$. Against a high-cost seller j, seller i wins today with positive probability, so seller i's expected payoff is strictly positive. Hence, high-cost seller i will not make an initial offer $p_i > c_H + (1 - \delta_B)\gamma$, so $\overline{p}_i^H \leq c_H + (1 - \delta_B)\gamma$.

Consequently, in equilibrium no player makes an initial offer strictly greater than $c_H + (1 - \delta_B)\gamma$, so $\overline{p}_i^k \leq c_H + (1 - \delta_B)\gamma$, for $k \in \{H, L\}$.

Suppose $\overline{p}_i^L > \overline{p}_i^H$. Either $\overline{p}_i^L = \overline{p}_j^L$, or $\overline{p}_i^L = \overline{p}_j^H$.

Suppose $\overline{p}_i^L = \overline{p}_j^L$. If low-cost seller i makes an initial offer $p_i \to \overline{p}_i^L$, then that offer loses with probability 1 in period 1. If seller j is low-cost, then seller i's payoff is 0. Either seller j wins in period 1, or both sellers set a price of 0 in period 2. If seller j is high-cost, then seller j wins today, because $p_L \leq c_H + (1 - \delta_B)\gamma$, and the buyer expects a price of c_H in period 2. Seller i's expected payoff strictly increases if it deviates to $c_H - \epsilon$, because it wins today against a high-cost rival, and possibly against a low-cost rival. Consequently, in equilibrium it cannot be the case that $\overline{p}_i^L > \overline{p}_i^H$ with $\overline{p}_i^L = \overline{p}_j^L$.

Suppose $\overline{p}_i^L = \overline{p}_j^H$. If low-cost seller i makes an initial offer $p_i \to \overline{p}_i^L$, then that offer loses with probability 1 in period 1. If seller j is low-cost, then seller i's payoff is 0. Either seller j wins in period 1, or both sellers set a price of 0 in period 2. If seller j is high-cost, then seller j wins today, because $p_L \leq c_H + (1 - \delta_B)\gamma$, and the buyer expects a price of c_H in period 2. Seller i's expected payoff strictly increases if it deviates to $c_H - \epsilon$, because it wins today against a high-cost rival, and possibly against a low-cost rival. Consequently, in equilibrium it cannot be the case that $\overline{p}_i^L > \overline{p}_i^H$ with $\overline{p}_i^L = \overline{p}_j^H$.

Consequently, in equilibrium it cannot be the case that $\overline{p}_i^L > \overline{p}_i^H$. Therefore, $\overline{p}_i^H \geq \overline{p}_i^L$ for all i. Moreover, one can show straightforwardly that $\overline{p}_i^H = \overline{p}_j^H$ in equilibrium.

Suppose that $\overline{p}_i^H > c_H$. If high-cost seller i makes an initial offer $p_i \to \overline{p}_i^H$, then that offer loses with probability 1 in period 1. This yields a payoff of 0 against either a high-cost or a low-cost rival. Deviation to $c_H + \epsilon$ yields a positive probability of winning, and hence a strictly positive expected payoff.

Consequently, in equilibrium it must be the case that $\overline{p}_i^H = \overline{p}_j^H = c_H = \underline{p}_i^H = \underline{p}_j^H$. Therefore, high-cost sellers make an initial offer of c_H, and low-cost sellers set prices less than or equal to c_H.

(FYI-this is incomplete) ■

7.3 Pooling Equilibria with Nonverifiable Offers

Lemma 5 *In any pooling PBE with nonverifiable offers, along the equilibrium path every initial offer must be rejected with probability 1.*

Proof of Lemma:

Proof of Proposition 3: The Lemma has three consequences for any pooling PBE. A low-cost seller's expected payoff is $\delta_S (1-\alpha) c_H$, the expected price is $(1-\alpha^2) c_H$, and the buyer's expected payoff is $\delta_B [\alpha^2 c_H + \gamma]$. With these features, a low-cost seller can profitably deviate to a lower initial offer that the buyer will accept.

Suppose that a low-cost seller deviates to an initial offer of $(1-\alpha) c_H$. If the buyer accepts, then the deviating seller is better off than in the candidate equilibrium, because it gets the same expected payoff, but without discounting. The buyer will accept this deviating initial offer because it anticipates the price in period 2 will be the same as the deviating offer. With discounting, the buyer is better off accepting the price today.

The reason the buyer anticipates the price in period 2 also will be $(1-\alpha) c_H$ is because the rival seller expected the initial offers to be rejected. Therefore, the rival seller's price in period 2 will be c_H if it is high-cost, or will be drawn from the range $[(1-\alpha) c_H, c_H]$ if it is low-cost. In either case, the deviating seller is guaranteed to win in period 2 by setting its price equal to its deviating initial offer of $(1-\alpha) c_H$, so there is no reason to cut its price further.

Therefore, in any candidate pooling equilibrium, a low-cost seller has an incentive to deviate from the proposed strategies. Consequently, there can be no pooling equilibria with nonverifiable offers. ■

7.4 Separating Equilibria with Nonverifiable Offers

In this section I construct a separating PBE with nonverifiable offers. There may be others.

Proof of Proposition 4: Letting p_L and p_H denote the lower and higher initial offers, consider the following strategies and beliefs.

Buyer's strategy: If $p_L \le c_H$, then accept p_L. If $p_L > c_H$, then accept p_L if $p_L \le c_H + (1-\delta_B)\gamma$, and otherwise reject both initial offers. In period 2, accept the lowest offer if it is less than or equal to $c_H + \gamma$, and otherwise rejects both offers.

High-cost seller's strategy: Make an initial offer of c_H. In period 2, make an offer equal to the smaller of c_H and the initial offer.

Low-cost seller's strategy: Make an initial offer drawn from the distribution

$$F(p) = \frac{p-(1-\alpha) c_H}{\alpha p}.$$

In period 2, if the initial offer was strictly less than c_H, then make an offer equal to the initial offer. If the initial offer was greater than or equal to c_H, then make an offer arbitrarily closely below c_H.

Beliefs: A seller whose initial offer is greater than or equal to c_H is believed by the buyer to be high-cost with probability 1, while a seller whose initial offer is strictly less than c_H is believed by the buyer to be low-cost with probability 1. If both offers are rejected and a seller's initial offer was strictly less than c_H, then the seller believes its rival is low-cost with probability α. If both offers are rejected and a seller's initial offer was greater than or equal to c_H, then the seller believes its rival is high-cost with probability 1.

To determine whether or not the preceding strategies and beliefs form a PBE, it is necessary to check that no player has an incentive to deviate, given its beliefs, and to check that the beliefs are consistent with the strategies and with Baye's Rule.

First, consider the buyer's incentive to accept an initial offer, based on both initial offers.

If $p_L \geq c_H$, then each seller is believed to be high-cost with probability 1. This belief is consistent with the strategies and with Baye's Rule, because observing an initial offer strictly greater than c_H is a zero-probability event. If the buyer rejects both initial offers, then it anticipates a price of c_H in period 2. It is rational for the buyer to accept p_L if and only if

$$\underbrace{c_H + \gamma - p_L}_{\text{payoff from accepting}} \geq \underbrace{\delta_B \left[c_H + \gamma - c_H\right]}_{\text{payoff from rejecting}} \Longleftrightarrow p_L \leq c_H + (1 - \delta_B)\,\gamma.$$

If $p_H \geq c_H > p_L$, then the seller offering p_H is believed to be high-cost with probability 1, and the seller offering p_L is believed to be low-cost with probability 1. These beliefs are consistent with the strategies and with Baye's Rule, because observing initial offers greater than c_H is a zero-probability event. If the buyer rejects both initial offers, then it anticipates that the low-cost seller will submit the exact same offer in period 2. With discounting, it therefore is rational for the buyer to accept the lower initial offer.

If $c_H > p_H$, then each seller is believed to be low-cost with probability 1. This belief is consistent with the strategies and with Baye's Rule. If the buyer rejects both initial offers, then it anticipates a price of p_L in period 2. With discounting, it is rational for the buyer to accept p_L.

Next, consider a high-cost seller's decision about its initial offer. Given the buyer's strategy and the rival seller's strategy, there is no incentive to deviate from an initial offer equal to c_H. With a higher initial offer, the rival's initial offer will be accepted. With a lower initial offer, the best possible outcome is for the rival to win, and the worst is for the high-cost seller to win and incur a loss in either period 1 or 2.

Finally, consider a low-cost seller's decision about its initial offer. Given the buyer's strategy and the rival seller's strategy, there is no incentive to deviate to an initial offer strictly greater than c_H. With such an offer, the rival's initial offer is sure to be accepted. An initial offer strictly less than c_H is sure to be accepted. Given the rival seller's strategy, it is a best-response for a low-cost seller to use the distribution $F(p)$ defined above. If the initial offer is rejected and is in the range $[(1-\alpha)\,c_H, c_H]$, then the seller's expected payoff in period 2 is identical for any price between $(1-\alpha)\,c_H$ and its initial offer. Hence, it is a best-response to hold firm at its initial offer. If the initial offer is rejected and is strictly less than $(1-\alpha)\,c_H$, then the seller's expected payoff in period 2 is greatest by setting a price equal to its initial offer. ■

References

[1] Ausubel, L., Cramton, P., and Deneckere, R. (2002) "Bargaining with Incomplete Information," in Aumann, R., and Hart, S. (eds.) Handbook of Game Theory 3 (Amsterdam: Elsevier Science B.V.).

[2] Bulow, J., and Klemperer, P. (1996) "Auctions Versus Negotiations," *American Economic Review* 86(1), 180-194.

[3] Chatterjee, K., and Lee, C. (1998) "Bargaining and Search with Incomplete Information About Outside Options," *Games and Economic Behavior* 22(2), 203-237.

[4] Chatterjee, K., and Samuelson, L. (1987) "Bargaining with Two-Sided Incomplete Information: An Infinite Horizon Model with Alternating Offers," *Review of Economic Studies*, 54(2), 175-192.

[5] Cramton, P. (1984) "Bargaining with Incomplete Information: An Infinite-Horizon Model with Two-Sided Uncertainty," *Review of Economic Studies* 51(4), 579-593.

[6] Cramton, P. (1992) "Strategic Delay in Bargaining with Two-Sided Uncertainty," *Review of Economic Studies* 59(1), 205-225.

[7] Fudenberg, D., Levine, D., and Tirole, J. (1985) "Infinite-Horizon Models of Bargaining with One-Sided Incomplete Information," in Roth, A.E. (ed.) Game Theoretic Models of Bargaining (Cambridge: Cambridge University Press).

[8] Fudenberg, D., Levine, D., and Tirole, J. (1987) "Incomplete Information Bargaining with Outside Opportunities," *Quarterly Journal of Economics* 102(1), 37-50.

[9] Fudenberg, D., and Tirole, J. (1983) "Sequential Bargaining with Incomplete Information," *Review of Economic Studies* 50(2), 221-247.

[10] Kennan, J., and Wilson, R. (1993) "Bargaining with Private Information," *Journal of Economic Literature* 31, 45-104.

[11] Krishna, V. (2002) Auction Theory (San Diego: Academic Press).

[12] Krishna, V., and Serrano, R. (1996) "Multilateral Bargaining," *Review of Economic Studies* 63(1), 61-80.

[13] McAfee, R.P. (1994) "Endogenous Availability, Cartels, and Merger in an Equilibrium Price Dispersion," *Journal of Economic Theory* 61(1), 24-47.

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

[15] McAfee, R.P., and Vincent, D. (1997) "Sequentially Optimal Auctions," *Games and Economic Behavior* 18(2), 246-276.

[16] Milgrom, P. (1989) "Auctions and Bidding: A Primer," *Journal of Economic Perspectives* 3, 3-22.

[17] Muthoo, A. (1999) Bargaining Theory with Applications (Cambridge: Cambridge University Press).

[18] Reinganum, J., and Daughety, A. (1991) "Endogenous Availability in Search Equilibrium," *RAND Journal of Economics* 22(2), 287-306.

[19] Reinganum, J., and Daughety, A. (1992) "Search Equilibrium with Endogenous Recall," *RAND Journal of Economics* 23(2), 184-202.

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

[21] Thomas, C.J., and Wilson, B.J. (2002) "A Comparison of Auctions and Multilateral Negotiations," *RAND Journal of Economics* 33(1), 140-155.

[22] Thomas, C.J., and Wilson, B.J. (2005) "Verifiable Offers and the Relationship Between Auctions and Multilateral Negotiations," *Economic Journal* 115 (506), 1016-1031.

[23] Thomas, C.J., and Wilson, B.J. (2010) "Horizontal Product Differentiation in Auctions and Multilateral Negotiations," unpublished manuscript.

[24] Vincent, D. (1992) "Modelling Competitive Behavior," *RAND Journal of Economics* 23(4), 590-599.

[25] Waehrer, K., and Perry, M. (2003) "The Effect of Mergers in Open Auction Markets," *RAND Journal of Economics* 34(2), 287-304.