Discussion Paper No. 893

MODELLING COMPETITIVE BEHAVIOR*

by

Daniel R. Vincent
J.L. Kellogg Graduate School of Management
Northwestern University
Evanston, IL 60208

April 1990

* This paper is a revised version of Chapter Two of my thesis. I would like to thank my advisers, Hugo Sonnenschein, Sandy Grossman and Andy Caplin, for their advice and encouragement. Financial support from the Sloan Foundation and SSHRC of Canada is gratefully acknowledged.

Abstract

A single seller of an indivisible object wishes to sell the good to one of many buyers. The seller has zero value for the good, the buyers have a commonly known identical value of one. This paper attempts to determine strategic environments which ensure the seller's ability to exploit the competitive behavior of the buyers to extract all the surplus in the game. It is shown that in many simple dynamic games, there are subgame perfect equilibria which involve the seller giving up the good for free. Even if the seller has an informational advantage which allows him to keep bidders from learning the bidding behavior of their opponents, there still exist (perfect Bayesian) equilibria which involve a sale at the price of zero. However, in this case, a simple refinement in the spirit of sequential equilibria can be used to rule out such collusive behavior and to show that the unique equilibrium outcome satisfying this refinement involve a price of one.

Modelling Competitive Behavior

Recent studies of strategic bargaining games have examined how two agents reach an agreement on the price of a good. These games use aspects such as impatience and abilities to make price offers as proxies for strategic power. Intuition suggests that adding either one extra buyer or one extra seller should be a way of introducing competitive pressures into the trading environment. This paper analyzes a collection of simple many-person bargaining games to illustrate that the modelling of competitive forces in a strategic game is delicate matter. Games which appear to be highly competitive support as subgame perfect equilibria, outcomes that yield substantial surplus to the competitors.

The games examined in this paper all have a natural interpretation. A single seller attempts to sell a single indivisible object to two or more buyers. Buyers have a reservation value of one monetary unit for the object, the seller's reservation value is zero. One's expectation is that a seller in this situation should be able to take advantage of competitive behavior on the part of his potential purchasers to gain all of the surplus. The purpose of this paper is to see what non-cooperative game theory can say about characterizing such competition.[1]

---

[1] The correct way to characterize the effects of the presence of more than one buyer in sequential offer trading games is an unresolved issue. Bargaining game theory illustrates that the determination of questions such as who gets to make offers, how quickly can offers be made, who is more impatient, who gets to make the first and the last offers all may have important implications for the outcome of the game. It may be that these exogenous features of the game structure are what generate surplus for a seller rather than the opportunity to sell to many buyers. The games examined here are infinite horizon games where only buyers make 'offers'. Time costs will matter but emphasis will be placed on games where the moves can be made relatively rapidly. The study of such restrictive structures is intended to allow us to isolate the effect the presence of another

Section One examines some simple many-agent bargaining games and shows by example that it is not always possible for a seller to exploit competitive behavior to gain in a trading game. The examples show that predicting the outcomes of strategic pricing behavior may be difficult. It is often suggested that the equilibrium outcome of a first price auction in the above framework with one seller and many buyers would involve a price of one.[2] This paper indicates that it is important to pay attention to the explicit structure of the auction game. In particular, if a seller cannot commit to keeping the object if he does not get a sale in the current period, buyers may engage in implicitly collusive equilibrium behavior in which they split the surplus.

The games in Section One confirm Stigler's claim that "...collusion will always be more effective against buyers [in my model, sellers] who report correctly and fully the prices tendered to them (Stigler, 1964)." What, then, of the case in which the seller is successful in keeping past bids private? Section Two examines more closely a condition which may allow the seller to take advantage of competitive behavior. An implicitly collusive behavior arose in some games because bidders conditioned future bids on their opponents' past behavior. This conditioning and the existence of a high price equilibrium ensured the presence of a credible threat to support a low price equilibrium. Section Two asks whether a seller who is able to restrict the ability of his potential buyers from observing each others' bids can ensure a high price equilibrium. Such a game, with its

---

interested buyer has on a seller's ability to get a high price for his good.

[2] See, for example, Milgrom and Weber (1982), and Milgrom (1985).

elaborate and asymmetric information structure is quite complicated. If attention is limited to perfect Bayesian equilibria, it is shown that even in this case the collusive outcome cannot be ruled out. However, Theorem Two shows that if a plausible (and simple) belief restriction which rules out beliefs of correlated deviations is also imposed, then the unique equilibrium outcome is the competitive outcome.

The result sheds some light on the importance of the differences between perfect Bayesian equilibrium and sequential equilibrium. More importantly, it suggests that a seller may be more likely to gain surplus from competition if he can keep private the buyers' bidding behavior. Thus a truly competitive situation may occur not, as one might think, in an auction pit with many buyers vigorously shouting out prices, but instead in a quiet office where the seller can privately monitor competing offers from buyers and prevent them from observing each other's bids. Such a result has disturbing implications for the many markets in which the open exchange of price offers and other transactions information are prized.

Section One: Non-Cooperative Collusive Behavior

Since the work of Bertrand, it has been common to assume that price-setting competition will compete away profits. The intuition is compelling; if the competing buyers were enjoying positive profits there would always be an incentive to overbid one's competitor slightly in order to gain more trade. This analysis ignores the role of the strategic environment. In a dynamic context, other intuitive reasoning may be equally persuasive.

Consider the simple pricing game, G_1. There are two buyers and one seller.[3] Buyer one offers a price, p_1, in the first period which the seller can accept and gain utility p_1 or reject. If he rejects, the game continues to the next period where buyer two offers a price p_2. If the seller accepts, he gains utility δp_2.[4] If the seller rejects p_2, buyer one offers again, payoffs are discounted by δ^2 and so on for potentially an infinite number of periods. The game may be seen to confer an advantage to the first buyer but one might expect that, if offers could be made arbitrarily quickly, this advantage would become arbitrarily small. Theorem One shows the expectation to be misguided.

Theorem One: The unique subgame perfect equilibrium outcome to G_1 yields the object to buyer one in the first period at a price of zero for any $\delta \in (0,1)$.

Proof: Let $p^* > 0$ be the supremum of all subgame perfect equilibrium price offers after any history. Suppose h_{t-1} is a history in an equilibrium which prescribes the offer of p^* by buyer i in period t. Suppose buyer i offered $p' \in (\delta p^*, p^*)$ instead. Since the best a seller can do if he rejects is no more than δp^*, he must accept p'. But that contradicts the definition of p^*. Therefore $p^* = 0$ and the seller must accept any price greater than zero in any period. It is straightforward to show that an offer of zero in every period and acceptance of zero by the seller is, in fact, a subgame perfect equilibrium. ||

---

[3] For simplicity, the games are described with only two bidders. THe arguments all hold (potentially with higher values of δ) with many bidders.

[4] Buyer two gains $\delta(1-p_2)$. Buyer one gains nothing.

The strategic environment of G_1 is not competitive at all. For any positive surplus the seller might expect to gain in this game, the current bidder could offer somewhat less and force the seller to accept. This ability is commonly known at every stage of negotiation and, thus, forces the seller to yield at any stage.

The proof of Theorem One suggests that the result may be due to the specific structure -- the separation of the buyers in the bidding stages turns out to confer a striking strategic advantage. Consider, instead, games G_2 and G_3 which are more in the spirit of Bertrand pricing games. In G_2, two buyers offer simultaneously in each period and the seller chooses which price if any to accept. If he rejects them both, the process is repeated in the next period with payoffs discounted, and so on for an infinite number of periods. Game G_3 is like G_2 except that within a period buyers offer sequentially -- buyer one offers, buyer two observes the offer and then makes his offer and the seller chooses. The process is repeated if he rejects.

These games are dynamic versions of a Bertrand game where the seller can never commit to walking away if a sale does not occur and cannot commit to a tie-breaking rule in the case in which bids are equal. It is easy to show that a price of one and sale to either buyer in the first period can be supported by a subgame perfect equilibrium. The proof relies on precisely the Bertrand intuition.

However, the explanation for dramatically different behavior can be as compelling and is familiar from the large literature on repeated games. Buyers are aware that competition with each other is both rational and self-defeating. This awareness, in a dynamic game, can generate tacit collusion.

Buyers offer a low (zero) price in every period and the seller gives the good up to either buyer with equal probability. An attempt by one to outbid the other by offering a higher price signals to both the seller and the other buyer the onset of a price war. If the seller is patient enough (if $\delta > .5$ with two bidders), he will wait until the next period to reap the benefits of the competition. The putative attempt by a buyer to increase his surplus would fail. The threat of competition acts as sufficient discipline to force buyers to 'collude' and offer a profit-generating price.[5]

It might be objected that the collusive equilibria rely heavily on the seller's use of mixed strategies. Aumman's(1987) interpretation of mixed strategies as representations of the buyers' uncertainty about which of them the seller will decide to grant the good to seems equally valid and perhaps preferable. The seller might indeed follow a pure strategy but, since at a price of zero he is indifferent, neither buyer may know who will win the good but places equal probability on the likelihood that each will.

This type of implicit collusion supports a price of zero in both games G_2 and G_3 and, in fact, can be shown to apply to a large class of dynamic bidding games. It is true, of course, that for a fixed δ, there exists a large enough number, $N(\delta)$, such that if there are more than $N(\delta)$ buyers, collusive equilibria fail to survive. While this fact conforms, perhaps, with our intuition that more buyers means more competition, the size of $N(\delta)$

---

[5] Although the games examined here are dynamic games rather than repeated games, the collusive equilibria have much the same flavor as the equilibria underlying the Folk Theorem in repeated games (Fudenberg and Maskin(1986)). This type of equilibrium behavior was first analyzed in dynamic games of incomplete information by Gul (1986) and Ausubel and Deneckere (1986). Vincent (1987) shows that the results extend to a large class of dynamic games.

is enormous. At an annual interest rate of 10%, even if offers could only be made daily, $N(\delta) = 3561$. One may be almost enough for monopoly, but it would be a surprise if more than 3000 buyers is needed for competition.[6]

Section Two: Silent Bid Auction Games

Section One indicates that in the context of an infinite horizon game, the ability of a seller to count on his many potential purchasers to compete away each other's surplus is uncertain. One interpretation of the result given its restriction to the particular class of games is that some additional aspect of strategic power is required by the seller in order to take advantage of his position. One such source of power lies in the technology of commitment in the game. An ability to force players to wait a significant (or significantly costly) length of time before rebidding can be to the seller's advantage. This ability would be reflected in a low δ. A simple example is setting δ equal to zero in G_2. This yields the one-shot Bertrand game in which it is well-known that the unique equilibrium is at a price of one. Another source of strategic power which is obvious from bargaining theory is the ability to make offers. Allowing the seller to make all the offers clearly is sufficient to grant him all the surplus. Of course, this is true even if he faced only one buyer.

Strategic power could also be granted by giving the seller an

---

[6] It has been conjectured that stock prices at the Chicago Board of Trade may obtain systematically higher prices than the New York Stock Exchange for reasons similar to the above. In the NYSE, pre-established time and size priority rules determine who wins a trade in the event of a tie. In the Board of Trade, no such rule exists and, the suggestion is, bidders would rather tie each other and take a chance at trade at a lower price rather than to engage in a self-defeating bidding war. I am grateful to Sandy Grossman for bringing this remark to my attention.

informational advantage. Note that in the games examined above, the informational assumptions were very strong. At the close of each stage, all players knew the entire history up to that stage. If the seller was the only player able to observe the behavior of all bidders the particular type of history contingent strategies might no longer be supportable. Such a game, of course, has a much more complicated information structure and, in general, a more stringent equilibrium refinement than subgame perfection is required.

The rest of this section examines a specific example of a game where the seller enjoys some informational advantage about the history of the game. Game G_2 with two buyers is altered so that a buyer only observes his own past price offers. At any stage of the game after the first period, he must form conjectures about his opponent's bidding behavior. The seller, on the other hand, observes the complete history of the game. The game can be interpreted as a sealed bid auction game where a seller always has the option of refusing to sell in a given period. If he does so, he must put it up for sale in the next period. However, he is able to keep past bids secret. The following analysis shows that such a structure also allows him to extract all the surplus from trade. Silent bid auctions induce competitive behavior.

A silent bid auction game can be created by considering the game G_2 with one seller and two buyers. Change the information structure of the game so that while the seller still observes the complete history of the game, buyers only observe their own price offers and the fact that the seller has rejected. A buyer's strategy at time t is a price in the interval, [0,1]. No buyer observes his rival's price offer but knows his own history of offers. The seller observes all offers. Since the buyers' strategy sets are convex, we restrict attention to the case in which the buyers choose only pure strategies; the seller may mix. Thus, in period 0,

$g_o^j \in [0,1]$, and subsequently,

g_t^j: $X^j(t-1)$ to $[0,1]$, $j = 1,2$.

where $X^j(t)$ denotes the space of all possible histories of price offers by b_j to time t. A buyer's strategy is an infinite sequence of functions

$g^j = (g_o^j, g_1^j, g_2^j, \ldots)$, $j = 1,2$.

Note that each buyer b_j observes only his own price offers and rejections by the seller. At any stage t, his information set is characterized by $X^i(t-1)$, that is, the space of all possible price offers by his opponent to period t - 1. When a perfect Bayesian equilibrium is characterized, it will be necessary to specify beliefs over this set: each buyer must form conjectures about his opponent's behavior. Define $m^1(u(t-1))$, to be b_1's probability distribution over b_2's price offers to period t - 1 given that b_1 has offered a sequence of prices u(t-1).[7] A generic strategy of the seller is the same as in G_2 -- that is, it is a sequence of distributions over S = {0,1,2}, $g^s = (g^s{}_0, g^s{}_1, \ldots)$ each conditional on the history of the game up to the current period.

Let a sequence of pairs of price offers to period t be denoted by x(t) = (u(t),v(t)) and let $u(t) = (u(t-1),u_t)$, that is, subscripts denote a single period offer. The term $u^t(g^1)$ denotes a sequence of price offers by

---

[7] Restricting attention to pure strategies of buyers ensures that along the equilibrium path, the probability distributions m^j are point distributions. Furthermore, given a point distribution belief, the expected future price offers of an opponent are deterministic.

b_1 given that he has offered a sequence of prices, u(t), and follows strategy g^1 henceforth. Thus, $u^o(g^1)$ denotes the full path of price offers determined by strategy g^1.

Let $\sigma = (g^1, g^2, g^s)$ be a triple of strategies. After any history, σ determines a probability distribution over the continuation of the game. Let $U(u(t),v(t),g^1) = U[(u^t(g^1)),v^t(g^2),g^s(u^t(g^1)),v^t(g^2)]$ be the expected utility of a continuation path to buyer 1 when history (u(t),v(t)) has occurred and σ is followed thereafter. A perfect Bayesian equilibrium is a triple,σ, and a pair of belief sequences, (m^1,m^2) such that

$$\int U(u(t),v(t),g^1)dm^1(u(t)) \quad \geq \quad \int U(u(t),v(t),g'^1)dm^1(u(t))$$

for all u(t) and for all $g^{1'} \in S^1$ and for m^1 satisfying Bayes' rule whenever possible given the information structure of the game and the equilibrium strategies. Similar conditions must hold for b_2 and the seller.

Despite the seller's informational advantage, a perfect Bayesian equilibrium in which buyers get the object for free exists and can be supported in the following way. Each buyer's strategy is to offer a price of zero in every period as long as he has offered zero in each preceding period. If there exists a previous period in which he did not offer zero he reverts to offering a price of one in every period. The seller accepts any maximum price greater than δ. He also accepts a price of zero if all previous price offers were zero (and randomizes between the two buyers). He rejects zero otherwise and always rejects prices strictly between zero and δ. The buyer's reversion to a price of one after a deviation is justified as a best response because he has reached an out of equilibrium node at which beliefs may be freely specified. Without any restriction on beliefs , the buyer can believe not only that he deviated but that the other buyer did as well. Believing (albeit incorrectly) that his opponent is also offering a price of one, the buyer's own choice of one is a best response.

The buyer's belief that his own deviation is correlated with an opponent's deviation even though neither can condition on the other's strategy would be ruled out by stronger equilibrium concepts. In particular, sequential equilibrium which is defined for finite games rules out such beliefs by the consistency restriction. The application of the consistency requirement to games with continuous action spaces is not straightforward, though, which is why perfect Bayesian equilibrium has tended to be more frequently used in such games. The approach taken here, is simply to rule out beliefs which involve correlated deviations whenever observed histories are enough to explain an agent's arrival at an out-of-equilibrium information set.

Assumption One (A1): Fix an equilibrium triple (g^1, g^2, g^s). Let $u(t)$ be any sequence of price offers by b_1. Let $x(t) = (u(t), v^o(g^2(t))$. If g^2 and g^3 are such that $\text{Prob}[g^s(x(i)) = (0)] > 0$ for all $i = 0,1,..t$, then $m^1(u(t))$ is such that $\text{Prob}[v(t) = v^o(g^2(t)] = 1$. A similar condition holds for b_2, $v(t)$ and $m^2(v(t))$.[8]

Remark: A1 ensures that if any buyer deviates and arrives at an information set which is explained by his deviation and the equilibrium behavior of his opponents, then he must believe that his opponents have continued to follow their equilibrium strategies. Appendix Two shows that, for finite games, A1

---

[8] A1 restricts beliefs at nodes which Kreps and Ramey (1987) denote 'reachable' nodes given a strategy profile, σ. The assumption itself corresponds to condition i) in their definition of both a structurally consistent assessment and of a convex structurally consistent assessment.

is an implication of sequential equilibrium.

Theorem Two shows that the unique perfect Bayesian equilibrium outcome to this game under A1 is a price of one in the first period. The idea of the proof is as follows. Suppose that a price less than one can be supported as an equilibrium outcome. Let the history leading to it be (u(t),v(t)). Lemma One shows that if some buyer could get certain and sole acceptance by offering just a bit more than $z_t = \max(u_t, v_t)$, the buyer would do so. Therefore, the seller's equilibrium strategy must involve some positive probability of rejection of a price higher than z_t.

If such a rejection is a best response, there are some restrictions that are implied for the continuation of the game following a deviation by a buyer offering the higher price. It is clear that if the seller rejects, he must expect to get a better deal in the continuation game -- it is also clear that this better price can only come from the deviating buyer. Lemma Two classifies the set of possible price paths following deviations by b_1 and shows that for any past history of buyer one-only deviations, buyer one never has a best response that involves offering more than the equilibrium price offer of his opponent, b_2. The combination of Lemmas One and Two yield Theorem Two.

Let x(t) = (u(t),v(t)), a price path ending with positive probability of acceptance, be supported by a perfect Bayesian equilibrium, σ.

Lemma One: Suppose $1 > v_t \geq u_t$ and let $\delta > .5$. There exists an $\epsilon^* > 0$ such that if the seller accepted $p' = v_t + \epsilon$, $0 < \epsilon < \epsilon^*$ with probability one, b_1 would be better off offering p' than offering u_t.

Proof: See Appendix One. ||

A consequence of Lemma One is that $\text{Prob}[g^s((u(t-1),p'),v(t)) = (1)] < 1$ for all $p' \in (v_t, v_t + \epsilon^*)$. That is, if there is an equilibrium in which a price less than one is accepted, the seller's equilibrium strategy must be to reject some higher (and out of equilibrium) price offer with some positive probability.

Note, also that if x(t) is an equilibrium outcome with positive probability of acceptance in period t, then

$$v_t \geq \max \delta^j v^t(g^2(t+j))_{t+j}, \text{ for all } j > 0,$$

otherwise the seller would reject for sure in period t and wait for b_2 to offer a higher price.

Fix a perfect Bayesian equilibrium, $\sigma = (g^1, g^2, g^s)$. The set X' denotes the set of possible outcomes of the game when b_2 and the seller follow their equilibrium strategy but b_1 may not:

$$X' = \{x(t)=(u(t),v(t)):v(t)=v^o(g^2(t))_{|t} \text{ and } \mathrm{Prob}[g^s(x(i)=(0)] > 0, \text{ for all } i=0,1,2,\ldots,t, \text{ for all } t\}^9.$$

Define j*,p* as

$$(j^*,p^*) = \operatorname{argmax} \delta^j p_j, \text{ such that}$$

$$p_j = v_j, \text{ for } (u(j),v(j)) \in X', \text{or}$$

$$p_j = g_j^1\ (u(j-1)) \quad \text{for } (u(j-1),v(j-1)) \in X'.$$

If (j*,p*) are not unique choose the pair with the highest j*. For $p > 0$ this exists since p is bounded by one and $\delta < 1$.[10] The idea is to assume that buyer two follows his equilibrium path but that buyer one may deviate

---

[9] The notation $_{|t}$ indicates a price path truncated to time t.

[10] If $p^* = 0$, then it is easy to see that at least one buyer will have an incentive to deviate and offer a slightly positive price.

any number of times. Find the highest discounted price offered either by buyer two or by buyer one assuming that b_1 uses his proposed equilibrium strategy (conditional on the history) to choose the offer he makes. Recall that a strategy describes behavior for all possible histories of the game, so this process is well-defined.

Lemma Two: $p^* = v_{j^*}$.

Proof: Clearly $p^* \geq v_{j^*}$. Suppose $p^* = g_{j^*}^{1}(u(j^*-1)) > v_{j^*}$. Note that A1 ensures that for all $x(j-1) \in X'$, b_1 believes that $x(j-1) = (u(j-1),v(j-1))$ and so he expects his opponent to offer v_{j^*} in period j^*. If the seller accepted for certain a lower price than p^* in period j^*, then b_1 would do better to offer it. Therefore, it must not be the case that the seller accepts such an offer with probability one. However, suppose b_1 offered p' less than p^* but close enough that p' exceeded both v_{j^*} and the best a seller can get in a continuation game if he rejects. Since the resulting outcome is again an element of X' by definition of p^* such a p' exists and the best the seller can expect is less than p^*/δ in the next period. He would have done better to accept p' immediately, contradicting the definition of $p^* = g_{j^*}^{1}(u(j^*-1))$. That is, the prescribed equilibrium strategy choice determined by g after a history $x(j^*-1)$ would not be p^* but a price less than p^*. Therefore, no matter what the history of buyer one deviations in X', buyer one will never make a maximum price offer higher that the offer he expects buyer two to make. ||

Theorem Two: Let δ be greater than .5. The unique perfect Bayesian outcome is (1,1)with acceptance in the first period.

Proof: Let $(u(t),v(t))$, $1 > v_t \geq u_t$ be a pbe outcome. Suppose b_1 deviates and makes an offer p' greater than v_t and suppose that rejecting p' is part of the seller's equilibrium strategy. The continuation path now is in X' and the highest discounted offer the seller can expect if he rejects p' is $\delta^{j*}p*$ $= \delta^{j*}z_{j*} \leq \delta^t v_t < \delta^t p'$. The first equality is by Lemma Two and the inequality follows by the assumption that buyer two follows his equilibrium strategy. Therefore rejecting $p' > v_t$ can not be a best response for the seller. But Lemma One shows that there is a $p' > v_t$ which, if accepted for sure, b_1 would prefer to offer. This implies $v_t = 1$, i.e. any equilibrium outcome occurs at a price of 1. Now suppose that $t > 0$. This implies that both first period price offers are less than δ since the seller will always accept a price greater than δ. Along the equilibrium path, no buyer gains a positive utility from trade. Each buyer would do better, then, to offer $p \in [\delta, 1)$ and gain acceptance in the first period rather than to wait for the later equilibrium price of one. ||

Section Three: Conclusion

The result of Section Two may seem counter-intuitive but in light of the preceding Section it is less surprising. It is commonly thought that more information rather than less is advantageous for a competitive environment. However, when the behavior of other agents can be observed the possibility of enforcing collusive behavior arises. The games of Section One show that such 'collusion' can occur even without an explicit agreement -- that is, apparent opponents can coordinate their actions to split the surplus while following equilibrium behavior in a non-cooperative game. The mere presence of additional interested buyers is not sufficient to ensure the seller's ability to capture all the surplus -- the strategic environment may frustrate this ability. Section Two shows that with a simple refinement in the spirit of sequential equilibrium, if agents can no longer observe each other's actions, then a seller can ensure himself the competitive price.

Appendix One

Proof of Lemma One: Without loss of generality let b_1 be the buyer with the lowest expected utility when x(t-1) is the history and (u_t, v_t) are the proposed equilibrium offers in period t.

Case I: Suppose the seller's strategy is such that $Prob[g^s(x(t)) = (0)] = w = 0$; that is, he accepts with certainty. In period t, the most b_1 can expect by offering u_t is

$$.5(1 - u_t) < (1 - u_t - \epsilon) \text{ for any } \epsilon \text{ such that } 0 < \epsilon < \epsilon^* = .5(1 - u_t).$$

Case II: Suppose the seller rejects x(t) with some positive probability, i.e. $w > 0$. For this to be a best response of the seller the value of the continuation game to the seller must be no lower than v_t. The best b_1 can do from x(t-1) is

$$\begin{aligned} .5w(1 - v_t) + \delta(1 - w)(1 - v_t/\delta) &< (.5w + \delta(1 - w))(1 - v_t) \\ &< \delta(1 - v_t) \text{ for } \delta > .5, \\ &< 1 - v_t - \epsilon \text{ as in Case I.} \end{aligned}$$

||

Appendix Two

This appendix shows that for finite extensive form games assumption A1 is implied by the consistency requirement in the definition of sequential equilibrium. Notation in this appendix is borrowed directly from that used in Kreps and Wilson (1982).

Lemma A1: Let (μ^*,π^*) be a sequential equilibrium of a finite extensive form game with perfect recall. Let h be an information set belonging to player i(h) and σ^i be the collection of pure strategies of i leading to h. Suppose there exists a $y \in h$ such that $\mathrm{Prob}(y:\pi^*_{-i},\sigma_i) > 0$. Then for all $x \in h$ such that $\mathrm{Prob}(x : \pi^*_{-i},\sigma_i) = 0$, $\mu^*(x : h) = \mu^*(x) = 0$.

Proof: Let Φ^0 be the set of assessments (μ,π) such that π is a profile of fully mixed strategies and μ are beliefs formed over the information sets derived via Bayes' rule given π. From Kreps and Wilson (1982), $\mu(x)$ is defined by

$$\mu(x) = P^\pi(x)/P^\pi(h(x)) \text{ where}$$

$$P^\pi(x) = \rho(p_{l(x)}(x)) \prod_{l=1}^{l(x)} \pi^{i(p_l(x))}(\alpha(p_{l-1}(x)))$$

and $P^\pi(h(x)) = \sum_{x'\in h} P^\pi(x')$.

Factor out the strategies belonging to i(h).

$$P^\pi(x) = \rho(p_{l(x)}(x))\, M(\pi^i,x) \prod_{\substack{l=1 \\ i(p_l(x))\neq i(h)}}^{l(x)} \pi^{i(p_l(x))}(\alpha(p_{l-1}(x)))$$

where

$$M(\pi^i,x) = \prod_{\substack{l=1 \\ i(p_l(x))=i(h)}}^{l(x)} \pi^{i(p_l(x))}(\alpha(p_{l-1}(x)))$$

Perfect recall ensures that only one strategy path for player i(h) leads to information set h. Thus for any $x' \in h$

$$M(\pi^i, x) = M(\pi^i, x') = M(\pi^i).$$

Setting $P'^{\pi}{-i}(x) = P^{\pi}(x)/M(\pi^i)$, yields

$$\mu(x) = P'^{\pi}{-i}(x)/(\sum_{x' \in h} P'^{\pi}{-i}\ (x')),$$

that is, $\mu(x)$ does not depend on the strategy choice of player i. Suppose there exists a $y \in h$ such that $P'^{\pi*}{-i}\ (y) > 0$. By hypothesis, x is such that $P'^{\pi*}(x) = 0$. By the continuity of $P'^{\cdot}$ in π, for all π converging to $\pi*$, $\lim P'^{\pi}(x) = 0$. The existence of y ensures that the denominator is positive so $\mu*(x) = 0$.

The implication of Lemma Al is that if i reaches an information set h and can explain his arrival at that set by assuming that $\pi*_{-i}$ was followed, consistency requires that he believe $\pi*_{-i}$ was, in fact, played.

Appendix Three

This appendix shows that changing the game structure to allow the seller some opportunity for offering at stages of the game may not qualitatively alter the results of Section III. Let G4 be the game in which there are two buyers, b1 and b2 and a seller. In periods 4i, buyer b1 makes offers which the seller may accept or reject, in periods 2 + 4i buyer b2 makes offers and in periods 1 + 2i the seller makes an offer to the buyer whose offer he just rejected.

In this game it is easy to show that a strategy profile in which each buyer offers δ in their offering period and accepts any price less than or equal to one in their responding period and the seller offers price of one and accepts any price greater than or equal to δ forms a perfect equilibrium.

Consider a slight change to the structure of G4. The game G5 is the same as G4 with the exception that the seller makes the first offer. Thus in the first period the seller offers a price to buyer b1, if it is rejected b1 then offers to the seller, if it is rejected, the seller offers to b2, b2 offers to the seller if he rejects and so on. With a common discount factor δ, a proof similar to that of Rubinstein's proof in the one buyer-one seller game shows the existence of a unique perfect equilibrium outcome at exactly the Rubinstein price for the two player game. That is, the unique perfect equilibrium outcome has the seller offering a price of $1/(1 + \delta)$ and b1 accepting in the first period. The same result holds for any number of buyers in the game; the seller does not gain from the presence of other buyers.

References

Aumann, Robert. Correlated equilibrium as an expression of Bayesian rationality. Econometrica 55. 1987. Pp. 1-18.

Ausubel, Larry and Deneckere, Ray. One is almost enough for monopoly. Discussion Paper 669. Northwestern University. December, 1985.

Freixas, X., Guesnerie, R., Tirole, J. Planning under incomplete information and the ratchet effect. Review of Economic Studies LII, 1985. Pp. 173-191.

Fudenberg, D. and Maskin, E.. The Folk Theorem in repeated games with discounting or with incomplete information. Econometrica 54, no. 3. May, 1986. Pp. 533 - 554.

Gul, Faruk. Bargaining foundations of dynamic oligopoly. Ph.D thesis. Princeton University. 1986.

Kreps, D. and Ramey, G.. Structural consistency, consistency, and sequential rationality. Econometrica 55. 1987. Pp. 1331-1348.

Kreps, D. and Wilson, R.. Sequential equilibria. Econometrica 50. 1982. Pp. 863 - 894.

Milgrom, Paul. Auction theory. World Congress of the Econometric Society. 1985.

Milgrom, P. and Weber, R. A theory of auctions and competitive bidding. Econometrica, vol. 52, no. 5. September, 1982. Pp. 1089 - 1122.

Rubinstein, A.. Perfect equilibrium in a bargaining model. Econometrica 50. 1982. Pp 97 - 109.

Stigler, George. A theory of oligopoly. Journal of Political Economy, LXXII. 1964.

Vincent, Daniel R.. Strategic interaction in dynamic trading games. Ph.D. thesis. Princeton. 1987.