Decentralized Trading, Strategic Behaviour and the Walrasian Outcome

ARIEL RUBINSTEIN
The Hebrew University

and

ASHER WOLINSKY
Northwestern University

For a market with a finite number of agents, pairwise matching and bargaining, it is shown that, even when the market is frictionless, the equilibrium is not necessarily competitive. It depends on the amount of information agents use. If their behaviour is conditioned only on the sets of agents present and the time, the competitive solution is the unique subgame perfect equilibrium. If agents have full information and condition their behaviour on some of it, there are also noncompetitive equilibria in which behaviour depends on specific information such as the identity of the trading partner and past events.

1. INTRODUCTION AND DESCRIPTION OF THE MODEL

The competitive outcome is often motivated by a naive scenario in which price-taking agents respond to prices announced by an impartial auctioneer. However, the predictions of the competitive outcome are viewed as relevant for a much wider range of scenarios. In the present paper we consider a market for an indivisible good in which the trade process is decentralized, in the sense that prices are determined in direct contacts between pairs of agents. In such markets we still expect to get the competitive outcome, if certain conditions are satisfied. It is common to include in such conditions the existence of a "large" number of agents and certain informational assumptions. However, notice that we also expect the competitive outcome to emerge in environments which are not classified as competitive in the above sense. Thus, for example, we expect the competitive outcome even in the case of one seller of an indivisible unit who faces two buyers who bid simultaneously for the unit.

In this paper we investigate conditions under which the competitive outcome is the unique equilibrium in models of a market for an indivisible good in which the trade is decentralized. The approach is to model in detail some particular (albeit natural) trading processes as games in extensive form, to examine the appropriate noncooperative equilibrium outcomes and to sort out the circumstances under which they approximate the competitive outcome. This is then the same approach as adopted, for example, by Rubinstein-Wolinsky (1985), Binmore-Herrero (1988), and Gale (1987).

The basic model is perhaps the simplest framework that combines pairwise meetings and some form of strategic bargaining. It is related to the models of Binmore and Herrero (1988) and Gale (1987). The agents in the model are S identical sellers and B identical

buyers, and it is assumed throughout that $S < B$. Each seller has one unit of the good for sale and each buyer wishes to buy exactly one unit. The market operates over time. The time dimension is discrete and the periods will be labelled with $t = 0, 1, 2, \ldots$.

In each period the remaining sellers and buyers are matched pairwise in a manner that no agent meets more than one agent of the opposite type. The matching process is random. In each period at least one match takes place and all possible seller-buyer matches are equally probable. It should be mentioned that these are not the minimal assumptions necessary for the conclusions of the paper to hold. We adopt them for the sake of simplicity.

After a buyer and a seller have been matched, they go through a short bargaining process over the terms of a possible trade between them. First, one of the matched agents is selected randomly (with probability $1/2$) to propose a price between 0 and 1, and then the other agent responds by accepting the proposed price or rejecting it. If the proposal is accepted the parties will implement it and depart from the market. Rejection dissolves the match and the agents proceed to the next matching stage.

When a transaction is concluded with price p at time t the von Neumann–Morgenstern utility to the involved seller is p and to the buyer it is $(1 - p)$. A utility of 0 is assigned to an agent who never leaves the market. Notice that there is no impatience associated with delays in trade: an agent derives the same benefit from a transaction at price p independently of its timing.

Regarding the agents' *information*, we shall distinguish among different regimes which will be spelled out later. These informational assumptions will determine the notion of a history.

A *strategy* for an agent specifies an action after each possible history that ends at a decision point. After histories that end with the agent being selected to propose, the strategy specifies the price to be proposed. After histories that end with the agent facing a proposed price, the strategy specifies the agent's response (acceptance or rejection) to every possible offer.

The above together with specification of the agents' information constitute a complete description of a game. The solution concepts that will be applied are *subgame perfect* or *sequential equilibrium*, depending on whether the considered version is a game of perfect or imperfect information.

In the context of this model the *competitive solution* is such that all the available units are sold to buyers for the price 1. This solution is depicted by the familiar demand-supply diagram of Figure 1.

One may suppose that the frictionless trading process described in this model, where it is costless for agents to meet and exchange offers, already constitutes a sufficiently competitive environment which should result in the competitive solution. The idea is that, since $B > S$, at any price less than 1 there would be buyers who would agree to a slight increase in price rather than be left without the good, and since sellers are not impatient they will wait until the price is bid in this manner up to 1. This view is seemingly supported by the above cited work of Binmore–Herrero and Gale.

The message of the present paper is that such conditions of frictionless trading are not in themselves sufficient for the competitive solution to be the unique equilibrium. When there is perfect information, and in the absence of additional assumptions, the model has noncompetitive sequential equilibria which are interpretable, i.e. they can be described verbally without using technical terms specific to the model. In order to obtain the competitive solution as the unique equilibrium, one needs extra assumptions which limit the information that agents utilize for their market behaviour.

FIGURE 1

The family of noncompetitive equilibria that we construct for the case of perfect information to demonstrate the points mentioned above has the following interpretation. There is some predetermined price p^* and there is an implicit understanding that certain S buyers have the privilege of buying the units at that price. The equilibrium strategies are designed so as to prevent the sellers and nonprivileged buyers from circumventing this implicit understanding. Thus, even though the process permits active participation by all parties, at the equilibrium, the nonprivileged give way to the privileged. The source of the privileges and the manner in which the price p^* is determined are left unmodelled. Although one may think of a variety of social arrangements which have the flavour of instituting a noncompetitive "fair price" and allocating privileges by some nonmarket mechanism, it is not our purpose in this paper to study such arrangements. On the contrary, our purpose is to understand the conditions which isolate the competitive outcome as the solution.

The strategies that support the above described noncompetitive equilibria rely on some detailed information (e.g. a seller has to act differently depending on the identity of the buyer that he faces—whether or not this buyer is privileged). When attention is restricted to environments in which agents condition their behaviour only on sufficiently limited information, either because their observations are limited or because they disregard some information as irrelevant, these equilibria are ruled out. Specifically, when the relevant information on which agents condition their behaviour is only the time and the sets of sellers and buyers present in the market, then the competitive solution is the unique sequential equilibrium outcome.

One question suggested by these observations concerns the extent to which agents' information has to be limited in order to have the competitive solution as the unique equilibrium. We examine an intermediate case which lies between the two extremes described above and in which agents do not observe the offers exchanged between others. They observe only their own personal histories and the sets of the remaining agents. It is shown that, with this restriction, there are still sequential equilibria which give rise to noncompetitive outcomes of the type described above.

The paper concludes with a discussion of the role of two modelling assumptions: the absence of costs of time, and the random matching assumption. This discussion clarifies the relations between the present work and the related work by Binmore–Herrero and Gale. In particular, it enables us to point out how seemingly unimportant modelling decisions in these models might have an important qualitative effect to the extent that

the conclusions that these works derived from quite similar models sometimes differ markedly from our conclusions.

Finally, let us call attention to what we view as an interesting modelling point. In the description of the equilibrium strategies we use the language of finite automata. An agent's strategy is described by a set of states and rules of transition between them, where each state specifies the agent's behaviour (offers and acceptance policy) in it. This method not only provides an unambiguous description of the strategies, but also gives rise to a natural interpretation of the strategies in terms of "states of mind" of the agents. For the use of this tool in the context of repeated games see, for example, Rubinstein (1986).

2. FRICTIONLESS MARKET WITH FULL INFORMATION

In the model of the present section agents have a substantial amount of information in the sense that at the beginning of date t agents know everything that has happened in the market up to the end of date $t-1$. For concreteness assume that in a given period an agent receives the following pieces of information in the following order: (i) the identity of his match and the selection of the proposer; (ii) the proposal; (iii) the response; (iv) the information regarding this period's events throughout the market. The type of situation that seems to satisfy either these or qualitatively similar assumptions is an open trading situation which resembles an oral auction in that all traders are gathered together and offers are exchanged publicly.

Notice that this is a game of imperfect information because in a given period agents are not informed about events that take place simultaneously in other matches. All subgames of this game start just before the matching stage and after all agents are informed about past events. The concept of subgame perfect equilibrium is sufficient to ensure the optimality of the proposer's move in each of the matches, but it does not ensure the optimality of an agent's response to an out of equilibrium offer and hence we employ the concept of sequential equilibrium. The latter concept requires optimality of the responses, given the agent's beliefs about events that occur in the simultaneous matches. Recall that these beliefs are not completely arbitrary even after unexpected offers. They have to satisfy a consistency requirement (see Kreps and Wilson (1982)) which implies here that, after unexpected offers, the beliefs agree with the equilibrium plans of the other agents.

The intuitive arguments which suggest the competitive solution of price equal to 1 do not refer to any restrictions on the information of agents. Thus, it might appear that there is no reason to expect that the obvious competitive pressures in the model would be inhibited by the fact that agents can condition their behaviour on extensive information. However, the following proposition shows that when there is full information, there exist sequential equilibria which are not competitive.

Proposition 1. *For any price p^*, $0 \leqq p^* \leqq 1$, and any assignment $b(k)$ which fits a distinct buyer $i = b(k)$ to each seller k, there is a sequential equilibrium in which seller k's unit is sold to i at price p^*.*

The equilibrium is constructed around the idea that the trading process is not impersonal: there is one buyer $i = b(k)$ who for some reason has the privilege to buy seller k's good and this is known to everybody. This privilege means that seller k will demand price p^* only from i and will agree to this price offer only from i. Buyers other than i and the seller are prevented from circumventing this privilege and bidding the

price up by the following understandings: if the seller demands from some buyer a price p, $p^* < p < 1$, then upon refusal this buyer acquires the privilege; if a buyer other than i bids above p^*, the understanding is that buyer i gets the opportunity to match this offer.

Proof. Each equilibrium strategy will be described as a collection of states and rules of transition between them. A seller's strategy is based on the states RIGHT (i, p) $i = 1, \ldots, B$, $p \geq p^*$ and NORIGHT. The interpretation is that in state RIGHT (i, p) buyer i has the privilege of buying this seller's unit for p, while in state NORIGHT no buyer has this privilege. A buyer's strategy is based on the states RIGHT (k, p), $k = 1, \ldots, S$, $p \geq p^*$, and NORIGHT. In state RIGHT (k, p) the buyer has the privilege to buy seller k's good at the price p; in state NORIGHT the buyer has no such privilege.

The equilibrium strategies are described by Table I and the accompanying lists of transition rules. Each row corresponds to a state and describes the agent's behaviour in

TABLE I

	Seller k's Strategy	
State	Offer to buyer j	Accept from buyer j
RIGHT (i, p)	p if $j = i$ 1 if $j \neq i$	$q \geq p$ if $j = i$ none if $j \neq i$
NORIGHT	1	none

Transition Rules

Switch from NORIGHT to:
  RIGHT (j, p^*) If you (seller k) have just met buyer j who was in state NORIGHT.
  RIGHT (j, p) If you (seller k) have just received offer q, $q > p^*$ from buyer j who was in state RIGHT (l, p), for some $l \neq k$.

Switch from RIGHT (i, p) to:
  RIGHT (i, q) If buyer $j \neq i$ has just offered you (seller k) price q, $q > p$ or if another seller $l \neq k$ has just offered buyer i price q, $p^* \leq q < p$.
  RIGHT (j, p^*) If you (seller k) have just offered p, $1 > p > p^*$, to buyer j who was in NORIGHT.
  NORIGHT   If buyer i has just left the market or has just switched to RIGHT (l, p) for some $l \neq k$ and some p.

Otherwise stay in the same state.

	Buyer i's Strategy	
State	Offer to seller l	Accept from seller l
RIGHT (k, p)	p if $l = k$ p^* if $l \neq k$	any $q \leq p$ if $l = k$ any $q \leq p^*$ if $l \neq k$
NORIGHT	p^*	any $p \leq p^*$

Transition Rules

Switch from NORIGHT to:
  RIGHT (k, p^*) If you (buyer i) have just met seller k who is in NORIGHT or if seller k has just offered you p, $1 > p > p^*$.

Switch from RIGHT (k, p) to:
  RIGHT (k, q) If buyer $j \neq i$ has just offered seller k price q, $q > p$, or if seller $l \neq k$ has just offered you (buyer i) price q, $p^* < q < p$.
  RIGHT (l, p) If you (buyer i) have just offered q, $p^* < q < p$, to seller l who was in NORIGHT.
  NORIGHT   If seller k has left the market or has just switched to RIGHT $(j, \cdot)$ for some $j \neq i$.

Otherwise stay in the same state.

this state in the events that he is selected to propose or to respond. The transition between states takes place immediately after the relevant information has reached the agent. For example, if a seller's offer is supposed to change the buyer's state, then the buyer's response will already be based on the new state.

Now consider an initial assignment of states such that if $i = b(k)$, then seller k is in RIGHT (i, p^*) and buyer i is in state RIGHT (k, p^*) and if j is such that $j \neq b(l)$ for all sellers l, then buyer j is in state NORIGHT. Observe that this assignment of states together with the above strategies and with beliefs, according to which each agent sticks to the presumption that in the other contemporary matches agents follow their equilibrium strategies, constitute a sequential equilibrium in which buyer i gets seller k's good for the price p^*. To verify this notice that the deviations that might threaten the equilibrium are that, in a meeting between seller k who is in RIGHT (i, p^*) and some $j \neq i$ who is in NORIGHT, one of the parties would offer p, $1 > p > p^*$. However, seller k does not profit from offering such p since j will refuse and this would only result in switching the states of seller k, buyer j, and buyer i to RIGHT (j, p^*), RIGHT (k, p^*), and NORIGHT, respectively. Buyer j will not profit from offering p, $1 > p > p^*$ to seller k since this will only result in switching the states of seller k and buyer i to RIGHT (i, p) and RIGHT (k, p), respectively. ‖

The equilibrium constructed above captures a trading institution whose main feature is that some buyers are privileged. The source of this privilege is not modelled. It might be determined in the model, e.g. assigned to the buyer who happens to be the first to meet a seller, or it might be determined exogenously by factors which are not modelled. The nonprivileged buyers accept this arrangement. In any event all that such a buyer can accomplish by challenging the arrangement and trying to outbid a privileged buyer is bidding up the price to be paid by the privileged buyer.

Notice that it is not necessary that all units be sold for the same price. Using very similar strategies we can construct equilibria in which the different buyers get their units at different prices.

An important conclusion from Proposition 1 is that the basic conditions of excess demand and frictionless interaction are not sufficient to point out a unique solution and in order to obtain a unique solution we have to fill in more details about the nature of the interaction. The proposition itself describes a particular solution that arises when the trading process is guided by a convention that assigns to a particular buyer the privilege of purchasing the good at a predetermined price.

Notice that even if one is inclined to think that the solution described by Proposition 1 is uninteresting in itself, Proposition 1 still has a valid message as a comment on the competitive solution. It points out that the competitive solution requires some further assumptions which essentially guarantee that the information that agents refer to and condition their behaviour upon is sufficiently limited.

3. A SUFFICIENT CONDITION FOR THE COMPETITIVE SOLUTION

The strategies that support the noncompetitive equilibra of the previous section rely on the fact that agents are informed about events occurring throughout the market. The following proposition gives a sufficient condition under which the competitive outcome is indeed the unique sequential equilibrium. This condition restricts attention to environments in which agents' information in the beginning of period t includes only the set of agents who are present at time t and the time itself. Let $\bar{B}_t$ and $\bar{S}_t$ denote the sets of

buyers and sellers, respectively, who are present at time t. Under these assumptions on the information, a strategy is a price offer and acceptance policy as functions of $\bar{B}_t$, $\bar{S}_t$, and t.

Proposition 2. *If agents' information includes only $\bar{B}_t$, $\bar{S}_t$, and t, the unique sequential equilibrium outcome is such that the good is sold for the price* 1.

Proof. Obviously there exists an equilibrium in which sellers always demand 1. The proof that it is unique under the restriction is by induction on the number S of sellers in the market. Assume that the proposition is true if the number of sellers is smaller than S. Consider a particular equilibrium. Let $v_i(t)$ and $w_k(t)$ denote the expected utilities of buyer i and seller k, respectively, at the beginning of period t, provided that no transaction has taken place up to that time. Observe that $v_i(t)$ and $w_k(t)$ are well-defined, since in the absence of a transaction the sets of buyers and sellers are unchanged and strategies depend only on time.

Let m denote the infimum over all equilibria of $w_k(t)$. Of course, for all t, $\sum_{i=1}^{B} v_i(t) \leq S(1-m)$. Therefore, at any t there exists some i such that $v_i(t) \leq (1-m)S/B$. Consider a seller who adopts the strategy of demanding the price $1 - \varepsilon - (1-m)S/B$ and not agreeing to less as long as the original sets of agents are in the market. Either he will meet at some time t a buyer for whom $v_i(t+1) \leq (1-m)S/B$ who will then agree to that price, or some other seller will transact beforehand. In the former case the seller's utility will be $1 - \varepsilon - (1-m)S/B$, while in the latter case it will be 1 by the inductive hypothesis. Since a seller can always adopt this strategy, we have $m \geq 1 - \varepsilon - (1-m)S/B$ and hence $m \geq 1 - \varepsilon B/(B-S)$ for any $\varepsilon > 0$, which means that $m = 1$.

To complete the proof we have to show that the inductive hypothesis is true for $S = 1$. This is shown by repeating the above argument. As above, let m be the infimum of the seller's expected utility over all equilibria in the case $S = 1$. By always offering $1 - \varepsilon - (1-m)/B$ the seller must meet at some time a buyer who will agree. Hence $m \geq 1 - \varepsilon B/(B-1)$, for any $\varepsilon > 0$, which means $m = 1$. ‖

The assumption that agents' information includes only the commonly observed time, and the sets of buyers and sellers present in the market need not necessarily be interpreted as lack of information or weak memory. Notice that if agents have full information (as in Section 2), then Proposition 2 can be modified to say that the outcome of any sequential equilibrium in which the strategies depend only on $\bar{S}_t$, $\bar{B}_t$, and t is the competitive outcome (the proof is identical). Thus, the result can be interpreted as referring to a situation in which all agents believe that any other information is irrelevant for the trading. This seems to be what is often meant by anonymity assumptions. The conclusion then is that if the process is characterized by a sufficient degree of anonymity in the above sense, then the equilibrium is necessarily competitive.

4. FRICTIONLESS MARKETS WITH IMPERFECT INFORMATION

The model of the previous section represents one extreme in that agents rely just on small amounts of information. The model of Section 2 represents another extreme: each agent knows everything that has happened and the equilibrium strategies prescribe responses to offers that were exchanged between other agents. The latter information regime seems appropriate for analyzing environments which resemble an oral auction in that the exchange of offers is in the open, or environments in which a faithful record of the offers is kept. However, for analyzing another important class of environments, it is more

plausible to assume that agents know only their personal histories. The purpose of this section is to inquire whether the insights of the previous section are still valid in an intermediate information regime in which each agent's information is restricted to his personal history. That is, does this restriction destroy the ability to maintain arrangements that keep the price below the competitive level? As we shall see, the answer to this question is mixed. In order to make the point we wish to make, it suffices to consider the case of $S = 1$, although one can prove the appropriate version of the following proposition for any $S < B$.

Proposition 3. *Suppose that $S = 1$, that agents know their personal histories, and that the seller can identify the different buyers. For any price p^*, $0 \leq p^* \leq 1$, and any buyer i there exists a sequential equilibrium in which i gets the unit for price p^*.*

Proof. The strategies use the states described in Proposition 1. The seller's strategy is based on the states RIGHT (j, p), $j = 1, \ldots, B$, $p = p^*, 1$. The buyers' strategies are based on the states NORIGHT and RIGHT (p), $p = p^*, 1$. The behaviour in these states is as described in the proof of Proposition 1 and since there is only one seller we write RIGHT (p) rather than RIGHT $(1, p)$ and ignore the instructions which refer to other sellers.

Transition rules for the seller:
Switch from RIGHT (i, p^*) to:

RIGHT (j, p^*) after offering buyer $j \neq i$ price p, $1 > p > p^*$.
RIGHT $(j, 1)$ after receiving from $j \neq i$ an offer p, $1 > p \geq p^*$.

Otherwise stay in the same state.

Transition rules for buyer i:
Switch from NORIGHT to:

RIGHT (p^*) after receiving an offer p, $1 > p > p^*$.
RIGHT (1) after offering the seller price p, $1 > p > p^*$.

Otherwise stay in the same state.

The initial assignments of states is such that the seller and some buyer i are in states RIGHT (i, p) and RIGHT (p^*), respectively, and all other buyers are in NORIGHT. As the interaction proceeds, agents observe only their personal histories, they may not know the state of other agents. Therefore, they will base their behaviour on beliefs regarding the state of others. It is assumed that the beliefs of buyer i correspond to his state: when he is in state RIGHT (p^*), RIGHT (1), or NORIGHT, he believes accordingly that the seller is in RIGHT (i, p^*), RIGHT $(i, 1)$, or RIGHT (j, p^*) for some $j \neq i$. Thus the initial buyer's beliefs are consistent with the initial assignment and they are updated whenever the buyer switches from one state to another, as described by the above transition rules.

To verify that the evolution of a buyer's beliefs is consistent with his information (i.e. satisfies Bayes rule whenever it is applicable), observe from the above transition rules that the only instances in which buyer i updates his beliefs are when he switches from NORIGHT to RIGHT (p^*) or to RIGHT (1). In both cases the change of beliefs follows a meeting with the seller and is consistent with the events that happened in this meeting in the sense that the seller will indeed switch to RIGHT (i, p) or to RIGHT $(i, 1)$, respectively. In all other cases a buyer does not update his beliefs. Now, the buyer's beliefs remain unchanged either following a meeting with the seller or following a period

in which he was not matched. In the former case this is again consistent with the information obtained in the match, for then the buyer knows that the seller remains in the same state as well. In the latter case observe that, conditional on the information that the seller did not meet buyer i in the last period and has not reached an agreement, the seller's state has not changed with probability one.

Now, it is a matter of a routine inspection to verify that the above described strategies together with the beliefs constitute a sequential equilibrium in which the privileged buyer i receives the unit for the price p^*. ‖

Thus, the restriction on information does not rule out the possibility of sustaining the noncompetitive arrangement as an equilibrium.

Recall that the equilibrium of Proposition 1 was based on the understanding that, when a buyer has the privilege to a unit, he is committed to match competing offers. This required that agents be informed about offers exchanged between other traders. Since this information is not available in the present case, the equilibrium is sustained by different behaviour: a nonprivileged buyer who tries to bid for a unit does not raise its price for the privileged buyer, but instead is expected to purchase the good at the price 1 that leaves him with no surplus.

In our opinion, this feature of the equilibrium makes it somewhat non-robust, since it depends on the property that the buyers are indifferent between purchasing the good at price 1 and not purchasing the good at all (this point was suggested to us by J.-P. Benassy). However, it is not necessary to have this feature in order to support an equilibrium of this type. Using the idea developed in Proposition 6 below, it can be shown that essentially the same outcomes (in RIGHT (i, p) buyer i will get the unit with a probability arbitrarily close to 1 rather than 1) can be sustained by equilibria in which buyers never buy for their reservation prices.

5. THE ROLE OF THE RANDOM MATCHING AND THE NO-DISCOUNTING ASSUMPTIONS

This section discusses the role and the meaning of two modeling assumptions which were employed throughout. These are the random matching assumption that sellers and buyers are matched randomly for one period at a time, and the assumption that there are no costs of delay. This discussion reveals the extent to which the qualitative results are robust. It also clarifies the relations with the above cited work of Binmore-Herrero and Gale.

Recall that the model does not specify costs of time: agents behave as if they do not attribute importance to whether they reach agreement immediately or after a finite number of rounds. Let us introduce the cost of time by assuming a common discount factor δ. Thus, if the good is sold to a buyer after t model-periods at price p, the utilities to the seller and the buyer will be $\delta^t p$ and $\delta^t(1-p)$, respectively.

It is important to notice that, in the presence of discounting, some details of the matching technology, which are of no consequence in the absence of discounting, affect the structure of the transaction costs in a significant way. Consider, for example, the pattern of behaviour described in Proposition 1 where buyer $i = b(k)$ has the privilege of buying seller k's good. With the matching technology considered above, it will take a random number of periods before buyer i and seller k will meet. Therefore, with discounting, this behaviour will be relatively more costly for the seller than settling on the same price p^* with the first buyer he meets. In contrast, with an alternative matching

technology which allows buyer i and seller k to meet whenever they wish to do so, without waiting for the random process to match them, a transaction between k and i is obviously not more costly than a transaction between k and a randomly chosen buyer.

The differences in the transaction costs between the two scenarios are reflected in the corresponding sets of equilibria. This is demonstrated by the following proposition.

Let $x(k)$ and $y(k)$ be defined as the (x, y) solution to the system,

$$y = \delta(x+y)/2 \tag{1}$$

$$1 - x = \delta(1 - x + 1 - y)/2k. \tag{2}$$

Observe that $x(1)$ and $y(1)$ describe the two-person bargaining equilibrium characterized in Rubinstein (1982), where the equilibrium strategies are for the seller to always demand price $x(1)$ and for the buyer to offer $y(1)$.

Proposition 4. *Consider the case of* $S = 1$.

(i) *Under the random matching procedure of sections 1-4 there is a unique subgame perfect equilibrium. The equilibrium agreement is reached immediately and the price is* $x(B)$ *or* $y(B)$ *according to whether the seller or the buyer was selected to propose.*

(ii) *Under a voluntary matching procedure where the seller can choose in each period with which buyer to talk, there is a range of subgame perfect equilibria. For any buyer* i *and any* x, $1 \geqq x \geqq x(1)$, *there exists a subgame perfect equilibrium such that the good is sold to buyer* i *and the price is* x *or* $\delta x/(2 - \delta)$ *according to whether the seller or buyer* i *was selected to propose in the meeting between them.*

Proof. The proof of part (i) is a nontrivial extension of the proof for the two-person case. Since it is rather lengthy, it is deferred to the Appendix. To prove part (ii), let x and y be such that $1 \geqq x \geqq x(1)$ and (x, y) satisfy equation (1) above. Consider the following strategies. The seller always offers x and agrees to accept y or more. A buyer always offers y and agrees to x or less. In the first period the seller picks buyer i and in case of disagreement the seller will choose to continue the bargaining with the same buyer only if the buyer did not deviate. If a buyer offered less than y or rejected an offer of x or less than x, then the seller chooses to discontinue the bargaining with him and picks a new buyer.

To verify that these strategies constitute a subgame perfect equilibrium observe that, since y and x satisfy equation (1) and since the R.H.S. of (1) is the seller's expected utility of rejecting an offer, it is indeed optimal for the seller to accept any price above y and to reject any price below y. Since $1 \geqq x \geqq x(1)$ it follows from (2) that

$$0 \leqq 1 - x \leqq \delta(1 - x + 1 - y)/2. \tag{3}$$

The R.H.S. of (3) captures the buyer's expected utility in the event that he rejects a price offer higher than x, since then the buyer will remain matched and the bargaining will continue to the next period. The L.H.S. of (3) captures the fact that the buyer's expected utility is zero if he rejects a price offer of x or higher, since then the seller will return to the matching process. Therefore, inequality (3) establishes that the buyer's strategy regarding acceptance of offers is optimal: from the R.H.S. of (3) it pays to reject offers above x and from the L.H.S. of (3) it pays to accept price x or lower.

Given the acceptance strategies of the buyer and the seller, respectively, it is optimal for the seller to propose price x and for the buyer to propose price y. This completes the proof that the suggested strategies form a subgame perfect equilibrium.   ‖

Notice that part (ii) of the proposition can be trivially extended to the case of $S > 1$. However, we have been unable to extend part (i) to cover this case.

Part (ii) of the proposition is closely related to the ideas developed in discussions of the "outside option" principle in Shaked and Sutton (1984) and in Binmore (1985). In our notation Binmore has shown that $(x(1), y(1))$ is the unique equilibrium in a version of the model in which the buyers can always respond to a seller's offer with a counter-offer (i.e. a seller can replace a partner only after hearing his counter-offer). An extensive discussion of the strategic role of outside options, which is the basic element here, can be found in Shaked (1987).

Observe from (1) and (2) that $x(B) = (B - \delta)(2 - \delta)/[B(2 - \delta) - \delta]$ and $y(B) = \delta(B - \delta)/[B(2 - \delta) - \delta]$. It follows that for $B > 1$, $\lim_{\delta \to 1} y(B) = \lim_{\delta \to 1} x(B) = 1$. Thus, the random matching technology of Sections 1-4 together with discounting result in a unique equilibrium and, when δ is close to 1, the equilibrium price is close to the competitive price of 1. We do not view this observation as contradictory to the earlier conclusion that the competitive solution requires some assumption concerning the impersonality of the trade process. Recall that the equilibria that support the noncompetitive outcome involved special relationships between sellers and buyers, with the interpretation that the buyer in such a relationship has the privilege to buy the seller's unit. When the agents discount time and the matching is random, the special relationships of this type are costly to maintain, since time is required to pass between meetings. Proposition 4(i) shows that these costs, regardless of their magnitude, prevent the emergence of such special relationships. However, this is not the only effect of discounting. The other effect is that, in any match, the proposer has some amount of monopoly power over the responder. The latter effect depends on the magnitude of δ and this is the reason that only when δ is close to 1 (but less than 1) we get approximately the competitive outcome.

Indeed, when the structure of the transaction costs is modified by adopting the matching technology which allows a seller and a buyer to meet voluntarily, then the special relations of Proposition 1 continue to play a role. As part (ii) of the proposition shows, for any selection of a privileged buyer i and any price between the two-party bargaining equilibrium (approximately $\frac{1}{2}$ since $\lim_{\delta \to 1} x(1) = \lim_{\delta \to 1} y(1) = \frac{1}{2}$) and the competitive solution, there exists an equilibrium in which i gets the unit for that price.

The conclusion is that the introduction of discounting in itself does not destroy the qualitative results of Section 1. It is only that with discounting one has to be more careful in the specification of the matching technology. If one wants to model a situation in which agents condition their behaviour on the identities of their potential partners, then the matching technology should already incorporate the features that agents recognize each other and have discretion over whom to meet. A random matching technology of the type considered throughout might provide an acceptable model for a situation in which the trade is completely impersonal, but is obviously not a very sensible model for other situations in which identities matter and delays are costly.

The above observations clarify the difference in conclusions between the present paper and the work of Binmore-Herrero and Gale. The models analyzed in these papers resemble the present model, but they combine the random matching assumption with discounting of future gains and with a continuum of agents. As argued above, the combination of random matching and discounting alone already amounts to assuming

that the seller's option of continuing with an ongoing bargaining process is more costly than turning to another buyer and this rules out the noncompetitive equilibria. The continuum of agents assumption strengthens this effect still further, since even in the absence of discounting the combination of random matching and continuum rules out the possibility of developing any continuing seller-buyer relationships. Thus, given the random matching assumption, both the discounting and continuum assumptions can be interpreted as assumptions about the anonymity of the interaction. The continuum assumption implies complete anonymity, while the discounting assumption implies effective anonymity in the sense that alternative behaviour is more costly.

6. THE CASE OF HETEROGENOUS AGENTS

This section extends the result of Proposition 1 to the case of nonidentical agents. Besides pointing out that the earlier result is not peculiar to the case of identical agents, the analysis here sharpens the cutting edge of that result. With the heterogeneity of agents, it is possible to show that it is not only that the equilibrium need not be competitive, it also need not be efficient in the sense that the goods need not end up at the hands of the agents who value them the most.

To expose the main ideas, it suffices to consider the simple cases of one seller vs. two non-identical buyers and of two non-identical sellers vs. two non-identical buyers. It will be evident from the discussion that, as before, the results extend immediately to arbitrary numbers of buyers and sellers. As well, to avoid repetition on the considerations described in Section 2, we shall not provide detailed proofs of the following propositions, but rather outline the main arguments which are required for the extension.

Consider a market with one seller and two buyers referred to as h and l. The seller's reservation price is zero, while the buyers' reservation values are b_h and b_l, respectively, where $b_h > b_l$. Assume that all agents have full information as in Section 2.

Proposition 5. *For all p^*, $0 \leqq p^* \leqq b_h$ and for all i such that $b_i \geqq p^*$, there exists a subgame perfect equilibrium in which the seller sells the good to b_i, $i = h, l$, for the price p^*.*

Proof. As in the proof of Proposition 1, the system starts in state RIGHT (i, p^*). In this state, the seller demands p^* from buyer i, demands b_j from buyer $j \neq i$, rejects prices below p^*, and agrees to price offers p^* or above from buyer i; buyer i offers p^* and agrees to price p^* or less; buyer $j \neq i$ offers $p = 0$, and rejects any price above p^*.

If the seller deviates and demands from $j \neq i$, price p, $b_j > p > p^*$, the system switches to RIGHT (j, p^*). If $j \neq i$ offers $p > p^*$, the system switches to RIGHT (h, b_h), i.e. to the equilibrium in which b_h gets the unit for price $p = b_h$. ‖

Note that the equilibrium can support an inefficient allocation such that buyer l ends up with the unit, while buyer h, whose valuation is higher, remains without it.

Observe that in such an equilibrium, buyer h cannot benefit from bidding above b_l, since such a bid results in the equilibrium in which h is expected to purchase the unit at the price b_h that leaves him with no surplus. This feature of the equilibrium may seem somewhat unattractive, since h participates in his own "punishment" by buying at the price b_h, though he is indifferent between buying and not buying. The following proposition demonstrates that essentially the same outcomes can be sustained by equilibria that avoid this feature.

Proposition 6. *Suppose that buyers never buy for a price equal to their reservation value, then*

(i) *for all p^*, $0 \leqq p^* < b_h$, there is a SPE in which buyer h gets the unit for p^*.*
(ii) *for all p^*, $0 \leqq p^* < b_l$, and all $\alpha < 1$, there is a SPE in which buyer l gets the unit for price p^* with probability higher than α.*

Proof. The equilibrium is based on states RIGHT (i, p, t). When the system is in state RIGHT (i, p, t) the interpretation is that buyer i has the privilege to buy the unit at price p during the next t periods. If this privilege is unlimited in time, we shall write RIGHT (i, p) instead of RIGHT (i, p, ∞). In states RIGHT (i, p, t) or RIGHT (i, p) the seller's strategy is:

Demand from buyer j p if $j = i$ and b_h if $j \neq i$.
Accept from buyer j $q \geqq p$ if $j = i$ and only b_h if $j \neq i$.

Buyer i's strategy is:

Offer p and accept any $q \leqq p$.

Buyer j's, $j \neq i$, strategy is:

Offer 0 and accept any $q \leqq \min [p, b_j]$.

To present the rules of transition between states, let $\beta(t)$ denote the probability that the seller and buyer l will meet at least once in a stretch of t periods.

Transition Rules
Switch from RIGHT (l, p, t) to:

RIGHT (h, p) If l has just deviated or if $t = 0$.
RIGHT $(h, \tilde{q})$ if buyer h has just offered $q > p$, where $\tilde{q}$ is such that $b_h - \tilde{q} < (b_h - p)(1 - \beta(t))$ and $\tilde{q} > q$.
RIGHT (h, p) if the seller demanded $q > p$ from h.
RIGHT $(l, p, \tilde{t})$ if the seller has just demanded $q > p$ from l, where $\tilde{t} > t$ and satisfies $b_l - q < \beta(\tilde{t})(b_l - p)$.
RIGHT $(l, p, t-1)$ if the seller has just met h.

Switch from RIGHT (h, p) to:

RIGHT (h, b_l) if l has just offered q, $b_l > q > p$.
RIGHT (l, p, t) if the seller has just demanded from l price q, $b_l > q > p$, where t is s.t. $b_l - q < \beta(t)(b_l - p)$.

Otherwise, the system remains in the same state.

Now for part (i) of the proposition, let the system start at RIGHT (h, p^*). If $p^* \geqq b_l$, buyer l is practically out of the game and the result is obvious. If $p^* < b_l$, the deviations that threaten the equilibrium are that buyer l would offer $p > p^*$ or that the seller will demand from l, price p, $p^* < p < b_l$. Observe that, in the former case, the system switches to RIGHT (h, p) so that the seller has no incentive to accept l's offer. In the latter case, the system switches to RIGHT $(l, p^*, \tilde{t})$, where $\tilde{t}$ is chosen to assure that l will reject the seller's demand because the expected utility after rejection, $\beta(\tilde{t})(b_l - p^*)$, is greater than that of accepting, $b_l - p$.

For part (ii) of the proposition, let the system start at state RIGHT (l, p^*, t), where t is such that $\beta(t) > \alpha$. Observe that if the strategies are followed, then l gets the unit for p^* with probability $\beta(t) > \alpha$. To verify that this is a SPE consider the following possible deviations in state RIGHT (l, p^*, t). Buyer h will not benefit from offering $p > p^*$, since then the system switches to RIGHT $(h, \tilde{q})$ so that h gets $b_h - \tilde{q}$ which is lower than h's expected utility, $(b_h - p^*)(1 - \beta(t))$, of following the equilibrium (i.e. waiting in the hope that the seller and l will not meet in the next t periods). The seller will not profit from offering to h price p, $b_h > p > p^*$, because h will reject it and get the privilege to purchase the good for p^*. The seller cannot profit from demanding from l price $p > p^*$, since when l rejects it, the horizon t is adjusted so as to make the rejection beneficial for l. Finally, the seller has no incentive to reject l's offers of p^*, since when l's privilege runs out the system will switch to RIGHT (h, p^*) and the seller will not get more than p^*. ‖

Next consider a market with two sellers and two buyers: seller h, seller l, buyer h, and buyer l, with reservation values $S_h = 3$, $S_l = 0$, $b_h = 4$, and $b_l = 1$. It follows almost directly from Proposition 5 that this game has equilibria in which seller l sells to buyer l at a price between 0 and 1 and seller h sells to buyer h at a price between 3 and 4. To see this consider the subgame in which buyer h and seller h remain alone in the market, and fix the equilibrium price in this subgame to be $p = 3$, after any history. Notice that, in the full game, this effectively changes the reservation value of buyer h to 3, Therefore, the game between seller l and the two buyers is essentially the one described in Proposition 5 and hence has an equilibrium such that seller l sells to buyer l for a price between 0 and 1.

Observe that in this case an efficient trade will involve only the transfer of the unit from seller l to buyer h. In this sense the equilibrium outlined above involves too much trade—more than the competitive amount.

7. CONCLUSION

Conditions of a frictionless trading process are not sufficient to yield the competitive solution as the unique equilibrium. When the participants possess full information then there are plausible noncompetitive equilibria.

To single out the competitive solution as the unique equilibrium, one needs to limit the information that agents utilize in their market behaviour. For example, when agents condition their behaviour only on state variables such as the sets of active sellers and buyers and time, then the competitive solution is the unique equilibrium.

APPENDIX

Proposition 4, part (i). *Assume $S = 1$ and the random matching technology. There is a unique subgame perfect equilibrium. The equilibrium agreement is reached immediately and the price is $x(B)$ or $y(B)$ according to whether the seller or the buyer was selected to propose.*

Proof. Let us rewrite the appropriate version of system (1)–(2),

$$y = \delta(x + y)/2 \tag{B.1}$$

$$1 - x = \delta(1 - x + 1 - y)/2B. \tag{B.2}$$

Define the sets A_s and A_b as follows.

$A_s = \{x \mid x$ is a perfect equilibrium payoff to a seller in a subgame
starting with the seller's offer$\}$.

$A_b = \{x \mid x$ is a perfect equilibrium payoff to the seller in a
subgame with a buyer's offer$\}$.

Let $m_i = \inf A_i$, $M_i = \sup A_i$, $i = b, s$. The method of the proof is to show that both $(x, y) = (M_s, M_b)$ and $(x, y) = (m_s, m_b)$ are solutions for system (B.1)-(B.2). Since this linear system has a unique solution, the above implies that $m_s = M_s$ and $m_b = M_b$. Hence, the sets A_s and A_b are singletons, so that the equilibrium is unique. Thus, the idea of the proof is similar to the method introduced by Shaked and Sutton (1984) for the equilibrium analysis of the two-person bargaining game.

It turns out that routine arguments establish that (M_s, M_b) indeed solves system (B.1)-(B.2) and that,

$$m_b = (m_s + m_b)\delta/2. \tag{B.3}$$

Lemma 1. *If* $x \in A_s$ *and* $y \in A_b$, *then* $Z = \delta(x + y)/2 \in A_b$.

Proof. Consider the following strategies in the subgame starting with a buyer's offer. In the first period the buyer offers price Z and the seller agrees to Z and any price above it. If the game reaches the next period, then all players follow the equilibrium strategies that support seller's payoff x or y, according to whether in the next period it is the seller or some buyer who is selected to propose.

Now, since $Z = \delta(x + y)/2$, the seller will not profit from rejecting Z, but will not accept less than Z. The buyer cannot profit from offering less than Z, since the offer will be rejected and he will get at most $\delta(1 - x + 1 - y)/2 < 1 - Z$. Therefore, the strategies are at equilibrium. The equilibrium payoff is Z and hence $Z \in A_b$. ‖

Lemma 2. $M_b = \delta(M_s + M_b)/2$.

Proof. By Lemma 1, for all $x \in A_s$ and $y \in A_b$, $M_b \geqq \delta(x + y)/2$ and hence $M_b \geqq \delta(M_s + M_b)/2$. It is impossible to have $Z \in A_b$ which is strictly above $\delta(M_s + M_b)/2$. This is because, in the perfect equilibrium that supports Z, the buyer whose offer starts the play of this equilibrium does not get more than $1 - Z$. But since $Z > \delta(M_s + M_b)/2$ and since the seller will agree to any price offer above $M_b = \delta(M_s + M_b)/2$, the buyer can deviate profitably by offering a price between Z and $\delta(M_s + M_b)/2$. ‖

Lemma 3. $1 - M_s \leqq \delta(1 - M_s + 1 - M_b)/2B$.

Proof. Almost identical arguments to those of Lemma 1 establish that if $x \in A_s$, $y \in A_b$ and $1 - Z = \delta(1 - x + 1 - y)/2B$, then $Z \in A_s$. From this observation the lemma follows immediately. ‖

Lemma 4. *In all perfect equilibria in a subgame that starts with the seller's offer to a certain buyer, this buyer's payoff is at least* $1 - M_s$. *In the perfect equilibria in a subgame that starts with a buyer's offer, the payoff to this buyer is at least* $1 - M_b$.

Proof. Consider a subgame that starts with the seller offering to buyer i. Suppose that there exists a perfect equilibrium in this subgame in which buyer i gets $1 - M_s - \varepsilon$. This perfect equilibrium must be such that there is no immediate agreement, for otherwise the seller's payoff would be above M_s in contradiction to the definition of M_s.

Let p be a price between M_s and $M_s + \varepsilon$, and consider the following candidate for an equilibrium in this subgame. The seller offers p and buyer i accepts it. If the seller demands more than p or if buyer i rejects this offer, all players continue as in the original perfect equilibrium. It is easy to verify that, given our initial hypothesis, this is indeed a perfect equilibrium and the seller's payoff is $p > M_s$, in contradiction to the definition of M_s. Thus, the initial hypothesis is false and so in all perfect equilibria of this subgame buyer i's payoff is at least $1 - M_s$.

The second statement of the lemma follows immediately from the definition of M_b. ‖

Lemma 5. $1 - M_s = \delta(1 - M_s + 1 - M_b)/2B$.

Proof. Given Lemma 3, it is sufficient to show that there is no $Z \in A_s$ such that $1 - Z < \delta(1 - M_s + 1 - M_b)/2B$. By Lemma 4, the R.H.S. of the inequality is the minimum payoff guaranteed to any buyer, if the game continues to the next period. This is because, with probability $1/B$ a particular agent will

meet the seller in the next period and then, by Lemma 4, this buyer will get at least $1 - M_s$ or $1 - M_b$ according to the selection of the proposer. Thus, in equilibrium, no buyer will offer or accept price p such that $1 - p$ is smaller than the R.H.S. of the inequality, so that there may not be an equilibrium payoff Z satisfying the inequality. ‖

Lemmas 1-5 prove that (M_s, M_b) is a solution for system (B.1)-(B.2). The validity of (B.3) can be proven by arguments in the same spirit as those presented in Lemmas 1-4. The difficulty is to prove that (m_s, m_b) satisfy equation (B.2) as well.

Let z denote the supremum of a buyer's expected utility in a subgame starting at the beginning of a period. Consider the following inequalities.

$$\frac{1}{\delta}(1 - m_s) \leqq z \leqq \left\{ \frac{1}{B}\left[\frac{1}{2}(1 - m_b) + \frac{1}{2}(1 - m_s) \right] + \frac{B-1}{2B}[1 - (m_b + 1 - M_b)] \right.$$

$$\left. + \frac{B-1}{2B}[1 - (m_s + 1 - M_s)] \right\}. \tag{B.4}$$

To verify the first inequality in (B.4) note that the buyer will accept any offer above δz and therefore $m_s \geqq 1 - \delta z$. Next observe that the R.H.S. of (B.4) gives an upper bound on z. The term $\frac{1}{2}(1 - m_b) + \frac{1}{2}(1 - m_s)$ bounds the buyer's expected utility in the event that he meets the seller in the next period and it is weighted by the probability, $1/B$, of that event. The term $1 - (m_b + 1 - M_b)$ bounds the expected utility of the considered buyer, say i, in the event that in the next period the seller will meet another buyer $j \neq i$ and that buyer will be selected to propose. In this event the seller's expected payoff will be, by definition, at least m_b and buyer j's expected payoff will be at least $1 - M_b$ (because of Lemma 4). Therefore, in this event, the payoff to buyer i is at most $[1 - (m_b + 1 - M_b)]$ and this is weighted by the probability of the event, $(B-1)/2B$. Similarly, the term $1 - (m_s + 1 - M_s)$ bounds buyer i's utility in the event that in the next period the seller will meet another buyer $j \neq i$ and the seller will be selected to propose. Then, by definition, the seller's payoff will be at least m_s and by Lemma 4 buyer j's payoff is at least $1 - M_s$. Therefore, in this event buyer i's payoff is at most $1 - (m_s + 1 - M_s)$ and this is weighted by the probability $(B-1)/2B$ of this event.

Using the fact that (M_s, M_b) is a solution for (B.1)-(B.2) we get

$$M_s = \frac{(B-\delta)(2-\delta)}{B(2-\delta) - \delta}; \qquad M_b = \frac{\delta(B-\delta)}{B(2-\delta) - \delta}. \tag{B.5}$$

Substituting these into (B.3) and (B.4) and solving for m_s we get

$$m_s \geqq \frac{(B-\delta)(2-\delta)}{B(2-\delta) - \delta} = M_s.$$

Thus $m_s = M_s = x(B)$ and $m_b = M_b = y(B)$ and obviously from (B.5) $\lim_{\delta \to 1} x(B) = \lim_{\delta \to 1} y(B)$. ‖

REFERENCES

BINMORE, K. (1985), "Bargaining and Coalitions", in Roth, A. (ed.) *Game Theoretic Models of Bargaining* (Cambridge: Cambridge University Press).
BINMORE, K. and HERRERO, M. (1988), "Matching and Bargaining in Dynamic Markets", *Review of Economic Studies*, **55**, 17-32.
GALE, D. (1987), "Limit Theorems for Markets with Sequential Bargaining", *Journal of Economic Theory*, **43**, 20-54.
KREPS, D. and WILSON, R. (1982), "Sequential Equilibria", *Econometrica*, **50**, 863-894.
RUBINSTEIN, A. (1982), "Perfect Equilibrium in a Bargaining Model", *Econometrica*, **50**, 97-109.
RUBINSTEIN, A. (1986), "Finite Automata Play the Repeated Prisoner's Dilemma", *Journal of Economic Theory*, **38**, 83-96.
RUBINSTEIN, A. and WOLINSKY, A. (1985), "Equilibrium in a Market with Sequential Bargaining", *Econometrica*, **53**, 1133-1150.
RUBINSTEIN, A. and WOLINSKY, A. (1986), "Decentralized Trading, Strategic Behavior and the Walrasian Outcome" (CARESS Working Paper No. 86-12, University of Pennsylvania).
SHAKED, A. and SUTTON, J. (1984), "Involuntary Unemployment as a Perfect Equilibrium in a Bargaining Model", *Econometrica*, **52**, 1351-1364.
SHAKED, A. (1987), "Opting Out: Bazaars versus Hi-Tech markets" (ICERD, LSE, Discussion Paper 87/162).