Markov Equilibria in Dynamic Matching and Bargaining Games

Douglas Gale and Hamid Sabourian

April 2003

CWPE 0322

Markov Equilibria in Dynamic Matching and Bargaining Games

Douglas Gale
Department of Economics
New York University
269 Mercer Street
New York, NY 10003

Hamid Sabourian
Faculty of Economics and Politics
University of Cambridge
Austin Robinson Building
Sidgwick Avenue
Cambridge UK CB3 9DD

April 9, 2003

Abstract

Rubinstein and Wolinsky (1990) show that a simple homogeneous market with exogenous matching has a continuum of (non-competitive) perfect equilibria; however, the unique Markov perfect equilibrium is competitive. By contrast, in the more general case of heterogeneous markets, we show there exists a continuum of (non-competitive) Markov perfect equilibria. However, a refinement of the Markov property, which we call monotonicity, does suffice to guarantee perfectly competitive behavior: we show that a Markov perfect equilibrium is competitive if and only if it is monotonic. The monotonicity property is closely related to the concept of Nash equilibrium with complexity costs.

1 Introduction

The theory of competitive equilibrium provides an elegant and simple account of how markets work. By contrast, the strategic analysis of markets tends to be complex and intractable. Extensive-form market games have many equilibria, in which a variety of different kinds of behavior are sustained by threats and counter-threats

In a seminal paper, Rubinstein and Wolinsky (1990), henceforth RW, analyze a market for a single indivisible good. There is a finite number of buyers and sellers who are matched in pairs and bargain over the terms of trade. The trading process can be formalized as a simple, extensive-form game. RW show that this dynamic matching and bargaining game[1] possesses a large set of perfect-equilibrium outcomes,[2] a result reminiscent of the Folk Theorem for repeated games.[3] Most of these equilibria do not correspond to the perfectly competitive outcome.

RW also consider conditions under which perfect-equilibrium outcomes are competitive. For example, it is shown that any perfect equilibrium in which equilibrium strategies are Markovian (anonymous) is competitive. In a similar vein, Sabourian (2001), henceforth S, has investigated a refinement based on lexicographic minimization of the complexity costs of implementing strategies. In the context of RW's model, he shows that perfect equilibria satisfy this refinement only if they are competitive.

These results suggest that perfectly competitive behavior may obtain where agents are required to use simple strategies. Unfortunately, both RW and S restrict their attention to a very simple environment, a market for a single indivisible good consisting of B identical buyers and S identical

[1] In the sequel, when we refer to market games, we have in mind games of this general type.

[2] We use the term perfect equilibrium to embrace both subgame perfect equilibrium and sequential or perfect Bayesian equilibrium, as appropriate. In markets where a single pair of agents is matched at any time, it is sufficient to use the concept of subgame perfect equilibrium. Sequential or perfect Bayesian equilibrium is needed when simultaneous matching and bargaining are allowed. In the following discussion, where the technical differences are not important, we use the term perfect equilibrium to cover both cases.

[3] An important feature of RW is that it analyzes a market with a finite number of agents. The preceding literature (Rubinstein and Wolinsky (1985), Gale (1986a,b,c, 1987), Binmore and Herrera (1988a,b), McLennan and Sonnenschein (1991), Osborne and Rubinstein (1990)) assumes a non-atomic continuum of agents, each of whom has a negligible effect on equilibrium.

sellers, each of whom wants to trade at most one unit of the good. In a heterogenous market, where buyers (or sellers) have a range of valuations of the good, things turn out to be more complicated. In particular, the refinements proposed by RW are insufficient to guarantee competitive behavior.

Gale and Sabourian (2002a,b), henceforth GS, extend the approach of S to the case of a heterogeneous market. They show that in a heterogeneous market with deterministic matching and bargaining, every perfect equilibrium with complexity costs (PEC) corresponds to a competitive outcome. Although the general approach is similar to S, the analysis is more difficult. Further, as we shall see, it is necessary to impose substantive restrictions on the matching process in order to guarantee competitive behavior.

In the present paper, we take the Markov property as our starting point. Markov strategies play an important role in the analysis of dynamic games because of their simplicity and recursive structure and because the Markov property or some other stationarity assumption is often enough to reduce or eliminate the indeterminacy of equilibrium. Here are some examples:

- In the theory of repeated games, the Folk Theorem guarantees the existence of a large set of subgame pefect equilibria. However, the subgame perfect equilibria in Markov strategies correspond precisely to the Nash equilibria of the stage game (Masso and Rosenthal (1989)).

- In bargaining games under incomplete information, there are many non-stationary equilibria (Ausubel and Deneckere (1989), Sobel) but a unique Markov perfect equilibrium (Gul, Sonnenschein and Wilson (1986)).

- In bargaining games with more than two players and complete information, there are many subgame perfect equilibria but the Markov perfect equilibrium is unique (Shaked (1994), Herrero (1985)).

- And, as mentioned above, in a homogeneous market with exogenous matching, RW show that a market game has a continuum of (noncompetitive) perfect equilibria, whereas the Markov perfect equilibrium is unique and perfectly competitive.

The last result suggests that the Markov property might be sufficient for perfect competition in the present context. In fact, unlike the homogeneous case, a heterogeneous market has a continuum of non-competitive Markov-perfect equilibria. This surprising result raises the question of what is different about dynamic matching and bargaining models that makes stationarity (the Markov property) such a weak refinement. One motivation for the present study is to understand better the structure of games with random matching and bargaining and why they produce pathological Markov equilibria.[4] Of course, there are other reasons why equilibria fail to be competitive and we present examples of these in the sequel, but our main interest lies with the Markov property and the extra assumptions that are needed to get a determinate and competitive equilibrium.

Although it is not sufficient for perfect competition, the Markov property does play a role in characterizing competitive behavior. Our second result is to show that a simple strengthening of the Markov property, called *monotonicity*, is sufficient for the competitive outcome. This property is interesting in its own right and also helps us to see why different matching processes work in GS. In particular, if the matching process is deterministic, any Markov equilibrium (suitably defined to take account of time varying matches) automatically satisfies the monotonicity condition. Monotonicity also turns out to have a close relationship to lexicographic minimization of complexity costs.

The relationship between this paper and GS is somewhat complex. GS focuses on *deterministic* matching models, where the matching process can be exogenous or endogenous, and agents can be matched simultaneously or sequentially. In this paper we study *exogenous* and *random* matching processes, which for the most part are sequential, although in Section 7.2 we briefly consider the implications of *simultaneous* matching, which turns out to be problematic from the point of view of competitiveness and determinateness.

Another way of looking at the relationship between the two papers is to observe that GS analyzes cases in which the Markov equilibria are competitive whereas the present paper analyzes cases in which Markov equilibria are not necessarily competitive.

To sum up, the paper makes two main contributions. On the one hand, we show that the Markov property is not sufficient to induce a competitive outcome in markets with sequential, random matching or exogenous, simul-

[4]We also note that discounting does not help to reduce the non-competitiveness of equilibria.

taneous matching. On the other, we show that a simple strengthening of the Markov property (that can in turn be justified by appealing to complexity) is sufficient to induce a competitive outcome in markets with sequential random matching. Clearly, the success of the program of using complexity costs to select a competitive outcome depends on the matching and bargaining model used and on the precise definition of complexity used.

The rest of the paper is organized as follows. In Section 2, we explore the differences between homogeneous and heterogeneous markets and explain why the methods used in RW may not suffice to characterize competitive behavior in the richer, heterogeneous environment. The game is defined formally in Section 3. In Section 4, we present two examples that illustrate the types of non-competitive equilibrium behavior that can easily arise in heterogeneous markets.

In Section 5, we introduce the notion of monotonicity. In any Markov equilibrium, payoffs are non-increasing over time. More precisely, an agent will never trade if he expects a higher payoff in the continuation game. However, it may be that the agent's payoff will be higher with positive probability in some future subgame, as is the case in the examples of non-competitive Markov equilibria in Section 4. We call an equilibrium monotonic if every agent's payoff is non-increasing *with probability one* and not simply in expectation. Here we show that *monotonic, Markov equilibria with perfect responses*[5] *are precisely the competitive equilibria.*

In Section 6, we investigate the connection between the concept of monotonicity and a natural definition of complexity and show that lexicographic minimization of complexity costs in equilibrium implies monotonicity and hence competition.

Robustness and extensions are discussed in Section 7. We present an example of non-existence of perfect equilibria with complexity costs and an example of indeterminacy with simultaneous matching.

[5]Equilibrium with perfect responses is a mild refinement of Nash equilibrium: it requires that, in every match that occurs along the equilibrium path, the equilibrium strategy specifies an optimal response to every price offer, not just the offers that are made in equilibrium. This is a much weaker refinement than perfect equilibrium, which requires optimal responses in every subgame.

2 Dynamic matching and bargaining games

2.1 Homogeneous markets

RW study the following market game. There are S sellers and $B > S$ buyers. Each seller has one unit of an indivisible good and each buyer wants to buy at most one unit of the good. A seller's valuation of the good is 0 and a buyer's valuation is 1. Time is divided into discrete periods or dates indexed $t = 1, 2, ...$ At each date, the agents are randomly matched in pairs consisting of one seller and one buyer (each feasible configuration of pairwise matches has equal probability). One member of the pair is randomly chosen to be the proposer and the other is the responder. Each member has probability $1/2$ of being chosen as proposer. The proposer offers to trade at a price $p \in [0, 1]$. The responder accepts or rejects the offer. Unmatched buyers are forced to remain inactive throughout the period.

If agreement is reached, the two agents trade at the agreed price p and leave the market. The buyer receives a payoff $1 - p$ and the seller receives a payoff p. There is no discounting.

Agents have complete information about the past play of the game, but at the moment when they choose their actions they do not know the identity of the other matches or the actions simultaneously chosen by other agents. Simultaneous moves require the use of perfect Bayesian equilibrium as the solution concept.

The central result obtained for this model in RW is the following theorem.

Theorem 1 *For every price p^* between 0 and 1 and for every one to one function β from the set of sellers to the set of buyers there exists a perfect equilibrium in which seller s sells his unit to buyer $\beta(s)$ for a price of p^*.*

In other words, there is a continuum of perfect equilibrium outcomes. Here is the intuition behind the result for the case of a single seller ($S = 1$). One buyer b^* is identified as the intended recipient of the good at a price p^*. The equilibrium strategies require the seller to offer the good at a price of $p = 1$ whenever he is the proposer and is matched with a buyer $b \neq b^*$. Every buyer $b \neq b^*$ rejects the offer. Whenever buyer $b \neq b^*$ is the proposer he offers to buy the good at a price of $p = 0$ and the seller rejects. When the seller meets buyer b^*, whichever is chosen as the proposer offers a price $p = p^*$ and the responder accepts. These strategies clearly produce the required outcome and the payoffs of the seller and buyer b^* are p^* and $1 - p^*$, respectively.

To prevent a deviation, RW make use of the following punishment strategies. Suppose that the seller has deviated by proposing a price $p \neq p^*$. The responder rejects this offer and the game then moves into a subgame in which the rejecting buyer b^{**} becomes the intended recipient of the good and the selling price becomes $p^{**} = 0$. The strategies are the same as those given earlier with the price p^{**} in place of p^* and buyer b^{**} in place of buyer b^*.

Similarly, if one of the buyers deviates by offering a price $p \neq p^*$ then the seller rejects, another buyer $b^{**} \neq b^*$ is chosen to be the intended recipient and the price at which the unit is traded changes to $p^{**} = 1$.

Deviations from these punishment strategies can be treated in an exactly similar way.

These strategies are ultimately quite complicated, in the sense that there is no limit to the number of potential deviations and each additional deviation requires a tailor-made response that makes the play of the game more complicated. As RW point out, one can think of this construction as requiring a large amount of information for the players to execute the equilibrium strategies. As an alternative they consider a model in which the amount of information available to the agents is strictly limited. Specifically, the game satisfies the following assumption:

> (**Anonymity**) *At the beginning of each date t, all that the buyers and sellers know about the previous play of the game is the number of buyers B_t and the number of sellers S_t remaining in the game.*

Under this assumption, the proposer's strategy is a function of the number of buyers and sellers (B_t, S_t) and the date t. The responder's strategy is a function of the numbers of agents (B_t, S_t), the date t and the proposal p. RW show that, under the anonymity assumption, the only equilibrium outcome is the competitive one.

Theorem 2 *If each player's information consists only of B_t, S_t, and t, the unique perfect equilibrium outcome is such that the good is sold for a price $p = 1$.*

Two points are worth noting. First, Anonymity is weaker than the Markov assumption. Markov strategies are functions of a minimal set of payoff-relevant variables at each date. They cannot be conditioned on variables that do not directly affect the future payoffs of the game. Anonymous strategies, by contrast, are allowed to depend on t. Secondly, the assumption of anonymity has the immediate effect of preventing the agents from punishing a deviator, because deviations are not remembered after they occur.

Two further variations of this basic model are considered in RW. One is to introduce a common discount factor $0 < \delta < 1$ and study the perfect equilibrium for the case of a single seller and B buyers. Let $x(\delta, B)$ and $y(\delta, B)$ denote the unique solutions of the equations

$$\begin{aligned} y &= \frac{\delta(x+y)}{2} \\ 1-x &= \frac{\delta(1-x+1-y)}{2B}. \end{aligned}$$

For fixed values of δ and B, $0 < x(\delta, B) < y(\delta, B) < 1$; but $x(\delta, B)$ and $y(\delta, B)$ converge to 1 as $B \to \infty$ or $\delta \to 1$.

Theorem 3 *Suppose that $S = 1$ and that all agents discount the future using the common factor $0 < \delta < 1$. Then there is a unique perfect equilibrium. In this unique equilibrium trade takes place immediately and the price is $x(\delta, B)$ or $y(\delta, B)$ depending on whether the buyer or seller was chosen to propose.*

Theorem 3 can be interpreted as demonstrating that the multiple equilibria of Theorem 1 are not robust. However, RW show that the uniqueness found in Theorem 3 depends on the assumption of exogenous matching. They consider a further variant of the basic model in which a single seller can choose, at the beginning of each period, which buyer he wants to bargain with in that period. In the model with endogenous matching, the indeterminacy of equilibrium returns in a strong form.

Theorem 4 *If $S = 1$ and the seller can choose in each period the buyer with whom he wishes to bargain, there is a continuum of perfect equilibrium outcomes: for each buyer b and each price $(2-\delta)/2 \leq p \leq 1$, there is a perfect equilibrium in which buyer b receives the good and the price is either p or $\delta p/(2-\delta)$, according to whether the seller or buyer b is the proposer in the first meeting between them.*[6]

[6] Note that the degree of indeterminacy in Theorem 4 is independent of the number of buyers B but depends on the discount factor δ. However, the interval of possible prices p converges to $[1/2, 1]$ as the discount factor δ approaches 1.

The strong impression left by these results is that indeterminacy of equilibrium is a robust feature of market games and, in particular, there is no reason to expect the outcome to be perfectly competitive. However, as we have seen, the strategies required to support the family of equilibria in these models can be very complex. Indeed, RW suggest as much in the context of their Theorem 2. Gale (2000) argues that *bounded rationality*, in the form of a large but limited memory, also guarantees competitive behavior.

Following these suggestions, S undertakes a systematic analysis of the role of complexity in sustaining a multiplicity of non-competitive equilibria. More precisely, S introduces the notion of complexity costs into the market game of RW. Relative complexity is measured by a partial ordering on the agents' strategies. Roughly speaking, a strategy f'_k is said to be more more complex than a strategy f_k if there exists a set of histories H such that f_k and f'_k are identical except on H, f_k specifies the same action everywhere on H and f'_k specifies more than one action on H. (This concept of complexity costs was introduced in Chatterjee and Sabourian (2000)). More complex strategies are assumed to be more costly to use, but the costs are small relative to the payoffs. Therefore, an agent first maximizes his payoff and then minimizes the complexity cost. Formally, an agent *lexicographically minimizes complexity costs* if he chooses a strategy that is a best response and is no more complex than any other best response. A *perfect equilibrium with lexicographic complexity costs* (PEC) is a perfect equilibrium in which every agent lexicographically minimizes complexity costs.

Lexicographic minimization of complexity costs is quite effective at eliminating equilibria in RW's market game. In fact, S shows that there is a unique PEC, with and without discounting, with exogenous and endogenous matching, and with one or more sellers.[7] Furthermore, the equilibrium outcome selected is always competitive, in the sense that all trade occurs at the competitive price $p = 1$. The intuition for S's result *in the case of a single seller* is the following. In any non-competitive PEC in which the seller receives a payoff of less than 1 there cannot be an agreement at a price of 1 between a buyer and a seller after any history; otherwise some player can economize on complexity. (For example, consider what happens if, after some history, a buyer is offered the price $p = 1$ and he accepts; then by a complexity argument the buyer should accept $p = 1$ whenever it is offered.

[7]Note that some of the results depend on the precise definition of complexity used in S and on whether or not complexity costs are positive.

Accepting $p = 1$ whenever it is offered guarantees the seller an equilibrium payoff of one). This implies that for any non-competitive PEC *all* continuation payoffs of all buyers are positive; but this cannot be so in any SPE in which there are more buyers than sellers. (By competition, in any SPE with $B > S$, there must be a buyer with a zero continuation payoff after some history). The result for $S > 1$ is established by induction.

The conclusion suggested by S is that complex strategies are required to sustain a multiplicity of non-competitive outcomes and that a mild refinement of equilibrium eliminates most of them. This provides a rationale for the competitive equilibrium as the only sustainable outcome in a simple market game.

2.2 Heterogeneous markets

These are important results, but unfortunately they only apply to the case of a homogenous market consisting of B identical buyers and S identical sellers. A market consisting of *heterogeneous* buyers and sellers is quite different. For the purposes of the present paper, we define a heterogeneous market as follows. As in the homogeneous case, there is a single indivisible good that is exchanged for money and each agent wants to trade at most one unit of the good. Without loss of generality, we can assume that there are equal numbers of buyers and sellers.[8] Buyers are indexed by $i = 1, ..., n$ and sellers are indexed by $j = 1, ..., n$. Buyer i's valuation of the good is denoted by $v_i \geq 0$ and seller j's valuation is denoted by $w_j \geq 0$. We assume that buyers and sellers can be ordered so that

$$v_1 > v_2 > ... > v_n$$

and

$$w_1 < w_2 < ... < w_n.$$

These valuations define demand and supply curves that determine the competitive, market-clearing price(s) in the usual way. The marginal traders $i = j = m$ are defined by the conditions

$$v_m > w_m$$

[8]Sellers with extremely high valuations and buyers with extremely low valuations cannot trade in any case.

$$v_{m+1} < w_{m+1}.$$

The interval of perfectly competitive prices for the market is

$$\max\{v_{m+1}, w_m\} \leq p \leq \min\{v_m, w_{m+1}\}.$$

Exchange between any two inframarginal traders $i, j \leq m$ is efficient whereas exchange between an inframarginal trader and an extramarginal trader is inefficient. Exchange between two extramarginal traders is not individually rational.

Compared with the heterogeneous market, we can see that the homogeneous market is special in several respects.

- *Efficient trade*: In a heterogeneous market, trade between an inframarginal seller $j \leq m$ and an extramarginal buyer $i > m$ is always inefficient, but can be individually rational if $v_i > w_j$. Likewise, trade between an inframarginal buyer $i \leq m$ and an extramarginal seller $j > m$ is always inefficient, but can be individually rational if $v_i > w_j$. In a homogeneous market, by contrast, each buyer's valuation is greater than each seller's valuation, so trade is always efficient by definition. Inefficient trade is a possible deviation from perfect competition in a heterogeneous market that cannot occur in a homogeneous market.

- *Division of surplus*: In the homogeneous market, except for the special case $B = S$, the competitive equilibrium price is either 0 or 1 and all of the surplus goes to one side of the market. As was pointed out above, the analysis in S depends crucially on the fact that one side of the market receives a zero payoff in a competitive equilibrium. In a heterogeneous market, there will typically be agents receiving positive payoffs on both sides of the market. The characterization of equilibrium outcomes becomes much more complicated as a result.

- *Invariance of the competitive prices*: In the homogeneous market, the set of competitive prices remains constant, independently of the set of agents remaining in the market. For example, if $B > S$ then no matter how many pairs of agents have traded, the number of remaining buyers is greater than the number of remaining sellers and the competitive price remains equal to 1. In the heterogeneous market, this need not be so. For example, if the competitive interval is $[w_m, v_m]$ and the marginal buyer and seller trade first, the competitive interval becomes

$$[\max \{v_{m+1}, w_{m-1}\}, \min \{v_{m-1}, w_{m+1}\}],$$

which is strictly larger given the assumption that $v_{m-1} > v_m > v_{m+1}$ and $w_{m-1} < w_m < w_{m+1}$. Similarly, if an inframarginal buyer $i \leq m$ trades with an extramarginal seller $j > m$, or an extramarginal buyer $i > m$ trades with an inframarginal seller $j \leq m$, then the competitive interval changes. In some cases, the new competitive interval may not even intersect the old one. This leads to problems with the characterization of equilibrium, as we shall see.

For all these reasons, and a number of others, the analysis of a heterogeneous market is more complicated, conceptually, analytically, and substantively, than the analysis of a homogeneous market.

More important than the difficulty of analyzing a heterogeneous market is the existence of new kinds of non-competitive equilibria, not encountered in the homogeneous case, for example, equilibria in which trade is inefficient. Furthermore, the refinements proposed by RW cannot eliminate these non-competitive equilibria. For example, as we explained before, anomymity (a weak Markov property) is sufficient for perfect competition in a homogeneous market with exogenous random matching. In a heterogeneous market, it is easy to construct non-competitive Markov equilibria for a market with sequential random matching. The following example is studied in detail in Section 4.1.

> There are two buyers and two sellers and all agents are inframarginal: $w_1 < w_2 < v_2 < v_1$. In this simple example, perfect competition requires all trade to occur at a uniform price $p \in [w_2, v_2]$. We construct a continuum of Markov SPE in which different pairs of agents trade at different prices in the sequential random matching model. We can even choose these non-uniform prices to lie outside the competitive interval $[w_2, v_2]$.

Another refinement proposed by RW is the introduction of a small amount of discounting. In Section 4.2, we take the example above and introduce discounting.

> There is a common discount factor $0 < \delta < 1$ and the valuations are symmetric: $w_1 = 0$, $w_2 = w$, $v_2 = v$, $v_1 = v + w$. We

characterize the unique symmetric Markov SPE and show that it is not competitive: even in the limit as $\delta \to 1$, the equilibrium prices at which trade occurs are non-uniform.

Thus, discounting may eliminate indeterminacy (we have not established this) but it does *not* imply perfectly competitive behavior.

The Markov SPE of the examples in Sections 4.1 and 4.2 satisfy the conditions of PEC, for certain definitions of complexity. Whether the refinement proposed in S helps in selecting a competitive outcome in the heterogeneous case depends both on the definition of complexity and the nature of the matching process. In GS we use a natural extension of the definition of complexity in S and show that such a refinement selects a competitive outcome in deterministic models. This definition of complexity has a 'local' character — it defines a partial order on the set of strategies with reference to a given set of remaining agents. Thus, roughly speaking, in GS a strategy f'_k is said to be more complex than a strategy f_k if there exists a set of agents N such that f_k and f'_k are identical except on $H(N)$, where $H(N)$ is the set of histories such that N is the set of agents remaining in the market, f_k specifies the same action everywhere on $H(N)$ and f'_k specifies more than one action on $H(N)$. This definition of complexity has the property that Markov strategies are minimally complex. Since this implies that any Markov SPE is also a PEC, it follows that the examples of Section 4 constitute counter-examples to this refinement, given the definition of complexity used in GS. On the other hand, in Section 6 we use a different extension of the complexity concept introduced by S. This definition of complexity has a 'global' character and does not refer to the set of remaining agents. Roughly speaking, this definition says that a strategy f'_k is more complex than a strategy f_k if f_k and f'_k are identical everywhere except that, either as a proposer or as a responder to some price offer, f_k specifies the same action everywhere on the set of all histories H and f'_k specifies more than one action on H. Note that the Markov strategies are not necessarily minimally complex according to this definition of complexity. So the examples in Section 4 do not constitute counter-examples to the refinement using this definition of complexity. In fact, we show below that lexicographic complexity costs (according to this definition and applied to Markov strategies) can provide a justification for the competitive equilibrium.

These examples demonstrate that in order to characterize perfect competition in a heterogeneous market, stronger conditions will have to be placed on equilibrium strategies. In addition, substantial restrictions have to be placed on the *structure* of the game. In particular, if exogenous, simultaneous matching is allowed then a new set of non-competitive Markov-perfect equilibria can emerge. In Section 7.2 we consider an example of a heterogeneous market in which there are two buyers and two sellers but only one buyer and one seller are inframarginal.

> Suppose that $w_1 < v_2 < w_2 < v_1$. In a competitive equilibrium, only w_1 and v_1 can trade. However, if exogenous, simultaneous matching is allowed, it is easy to construct equilibria in which w_1 trades with v_2 and w_2 trades with v_1.

This will be true for any exogenous matching process, random or deterministic, in which, with positive probability, v_1 is matched with w_2 and v_2 is matched with w_1 at the first round. The intuition is similar to the Diamond corner (Diamond (1971)). If one pair of agents expects the other pair to trade at the first round, they are effectively in a two-person economy. The existence of the other pair provides no competitive pressure. In order to ensure that competition does occur, we have to ensure that not all trade can occur at once. We do that here by assuming that only one pair of agents is matched at a time.

These examples show, first, that in the random matching model, the Markov property is not sufficient for competition, secondly, that while discounting ensures uniqueness, at least within the class of symmetric equilibria, it does not deliver the competitive outcome as it did in the homogeneous case, and, thirdly, that exogenous, simultaneous moves are inconsistent with competition. Clearly, the heterogeneous case is quite different from the homogeneous case.

3 The market game

A heterogeneous market is defined, as in Section 2, by the valuations $v = (v_1, ..., v_n)$ and $w = (w_1, ..., w_n)$). The trading game is defined by the following rules:

- At each date, a pair of agents consisting of one buyer and one seller is chosen at random from the agents remaining in the market. One member of the pair is chosen at random to be the proposer; the remaining agent becomes the responder.

- The agent chosen to be proposer offers a price p. The responder must accept or reject this offer. If the offer is accepted, the good is traded at the agreed price and both agents leave the market. If the proposal is rejected, there is no trade and all agents begin the next period with the same endowments.

An agent's information at the beginning of date t consists of the matches, proposals and responses observed in all previous periods. During the period, all the agents observe the set of agents remaining in the market, the choice of proposer and responder, the price offered by the proposer, and the response. An agent's strategy maps all the available information into a choice of action at each date.

Given the market parameters $v = (v_1, ..., v_n)$ and $w = (w_1, ..., w_n)$, the game is defined as follows. Play occurs at a countable sequence of dates $t = 1, 2, ...$ At each date t, the set of players remaining in the game is denoted by N. The set N is *balanced*, that is, it contains an equal number of buyers and sellers. An ordered pair $\langle k, \ell \rangle$ is randomly selected from the set N, where the first agent k is the proposer and the second agent ℓ is the responder. We assume that:

- The pair $\langle k, \ell \rangle$ consists of a buyer and a seller;
- Each remaining buyer has an equal probability of being chosen and each remaining seller has an equal probability of being chosen;
- The buyer and seller chosen have equal probability of being chosen as proposer and responder.

We adopt this particular matching rule for simplicity. As we mentioned before, the matching probabilities are not important as long as each agent remaining in the game has a positive probability of being chosen and the matching probabilities are stationary, that is, they depend on the set of agents remaining in the market, but not on the date. The assumption that only one buyer and one seller are matched at each date is crucial, however, as the example in Section 7.2 shows.

Because we will later use arguments based on the complexity of strategies, we have to be somewhat pedantic about the description of the game. Let I denote the set of buyers, J the set of sellers, and $K = I \cup J$. At each date t, the play of the game consists of the choice of a matched pair $\langle k, \ell \rangle$, a proposal $p \in \mathbf{R}_+$, and a response $r \in \{accept, reject\}$. A finite history of the game consists of a finite sequence

$$\{(\langle k_1, \ell_1\rangle, p_1, r_1), ..., (\langle k_{t-1}, \ell_{t-1}\rangle, p_{t-1}, r_{t-1}), \langle k_t, \ell_t\rangle\}.$$

Let H^t denote the set of finite histories at date t and $H = \cup_{t=1}^{\infty} H^t$ the set of all finite histories. We use h to denote a generic finite history and also the initial segment of a history. Thus, $h' = (h, \langle k, \ell\rangle) \in H^t$ denotes a finite history $h' \in H^t$ with initial segment

$$h = [(\langle k_1, \ell_1\rangle, p_1, r_1), ..., (\langle k_{t-1}, \ell_{t-1}\rangle, p_{t-1}, r_{t-1})]$$

and a match $\langle k, \ell\rangle$ at the final date t.

For any finite history h, let $N(h)$ denote the set of remaining agents. We denote the set of histories after which agent k is the proposer by H_k^p, where

$$H_k^p = \{(h, \langle k, \ell\rangle) \in H\},$$

and the set of histories after which agent k is the responder by H_k^r, where

$$H_k^r = \{(h, \langle \ell, k\rangle) \in H\}.$$

A strategy for agent k is a function f_k defined on $H_k = H_k^p \cup (H_k^r \times \mathbf{R}_+)$ such that

$$f_k(h) \in \mathbf{R}_+, \forall h \in H_k^p$$

and

$$f_k(h, p) \in \{accept, reject\}, \forall (h, p) \in H_k^r \times \mathbf{R}_+.$$

Let F_k denote the set of strategies for agent k and let $F = \times_{k \in K} F_k$ denote the set of strategy profiles.

Given any strategy profile f, there is a unique (stochastic) outcome that determines the payoff $U_k(f)$ of each agent k. The market game $\Gamma = (K, F, U)$ is defined by the set of agents K, the set of strategy profiles F and the payoff function $U = \times_{k \in K} U_k$. A Nash equilibrium of Γ is a strategy profile f^* such that, for each k,

$$U_k(f^*) \geq U_k(f^*_{-k}, f_k), \forall f_k \in F_k..$$

The game described is a game of perfect information (and chance moves) so the appropriate equilibrium concept is usually perfect equilibrium. We will later introduce a weaker notion of equilibrium (to deal with some existence issues) but for the examples in the next section perfect equilibrium will do just fine. Informally, we define a subgame perfect equilibrium to be a strategy profile f^* such that for every subgame $(h, \langle k, \ell \rangle)$ or $(h, \langle k, \ell \rangle, p)$ the continuation strategies $f^*|_{(h,\langle k,\ell \rangle)}$ or $f^*|_{(h,\langle k,\ell \rangle,p)}$ form a Nash equilibrium for the subgames $\Gamma|_{(h,\langle k,\ell \rangle)}$ and $\Gamma|_{(h,\langle k,\ell \rangle,p)}$, respectively.

Here we are interested in *Markov* Nash equilibria, in which the equilibrium strategies have the Markov property, that is, for any finite histories $(h, \langle k, \ell \rangle)$ and $(h', \langle k, \ell \rangle)$ in H_k^p such that the set of remaining agents is $N = N(h) = N(h')$,

$$f_k^*(h, \langle k, \ell \rangle) = f_k^*(h', \langle k, \ell \rangle),$$

and for any finite histories $(h, \langle \ell, k \rangle)$ and $(h', \langle \ell, k \rangle)$in H_k^r and any price offer $p \in \mathbf{R}_+$,

$$f_k^*(h, \langle \ell, k \rangle, p) = f_k^*(h', \langle \ell, k \rangle, p).$$

For such an equilibrium, strategies can be treated as functions of information sets of the form $(N, \langle k, \ell \rangle)$ or $(N, \langle k, \ell \rangle, p)$. In what follows, we abuse notation by using the notation for finite histories to denote these information sets. In effect, we simply suppress the parts of the finite history that are not relevant for strategies. The set of information sets at which agent k controls play is denoted by H_k^M and defined by

$$H_k^M = \{ (N, \langle k, \ell \rangle) |\, k, \ell \in N, N \text{ balanced} \} \cup \\ \{ (N, \langle \ell, k \rangle, p) |\, k, \ell \in N, N \text{ balanced}, p \in \mathbf{R}_+ \} .$$

The set of Markov strategies is denoted by F_k^M and defined by

$$f_k \in F_k^M = \{ f : H_k \to \{accept, reject\} \cup \mathbf{R}_+ | \\ f(N, \langle k, \ell \rangle) \in \mathbf{R}_+,\ f(N, \langle \ell, k \rangle, p) \in \{accept, reject\} \} .$$

4 Examples

In this section we maintain the rules of the game defined in the preceding section. The first example shows the existence of a continuum of noncompetitive Markov perfect equilibria for a market with sequential random matching.

4.1 Random sequential matching

Suppose $n = 2$ and $v_1 > v_2 > w_2 > w_1$. Thus $m = n = 2$ and there is no possibility of inefficient trade. We construct a Markov perfect equilibrium (MPE) of the game as follows. Let p_{ij} denote the price at which buyer i and seller j will trade if they are the last pair of agents left in the market, where

$$w_j \leq p_{ij} \leq v_i$$

for $i, j = 1, 2$. In addition we assume that

$$p_{22} < p_{11}$$

and

$$p_{21} = \frac{1}{2}(p_{11} + p_{22}).$$

Let the numbers $(v_1^*, v_2^*, w_1^*, w_2^*)$ be defined implicitly as follows:

$$\begin{aligned} v_1 - v_1^* &= \frac{1}{2}(p_{11} + p_{21}) = w_1 + w_1^*, \\ v_2 - v_2^* &= \frac{1}{2}(p_{22} + p_{21}) = w_2 + w_2^*. \end{aligned}$$

Our strategy for proving the existence of a continuum of non-competitive equilibria is to treat the numbers v_i^* and w_j^* as if they represent the reservation utilities of the agents and then show that these are precisely the payoffs that agents achieve in equilibrium. Direct calculation shows that

$$v_1^* + w_1^* = v_1 - w_1 \quad \text{(1a)}$$

$$v_2^* + w_2^* = v_2 - w_2 \quad \text{(1b)}$$

$$v_1^* + w_2^* = v_1 - w_2 - \frac{1}{2}(p_{11} - p_{22}) < v_1 - w_2 \quad \text{(1c)}$$

$$v_2^* + w_1^* = v_2 - w_1 - \frac{1}{2}(p_{22} - p_{11}) > v_2 - w_1. \quad \text{(1d)}$$

So, in the early stages of the game (before any trade has occurred), each of three pairs of agents $(i, j) \neq (2, 1)$ can trade the good.[9]

[9] Trade between buyer 2 and seller 1 is not individually rational because, by construction, $v_2^* + w_1^* > v_2 - w_1$. Thus, we only have to consider the three subgames in which the first pair to trade is $(i, j) = (1, 1), (1, 2), (2, 2)$.

When a pair of agents $(i,j) \neq (2,1)$ is matched, if i is the proposer he offers a price $p = w_j^* + w_j$, which is accepted by j, and if j is the proposer he offers a price $p = v_i - v_i^*$, which is also accepted. When the pair $(i,j) = (2,1)$ is matched, no trade occurs. If the pairs $(1,1)$ or $(2,2)$ are the first to trade, the price does not depend on the identity of the proposer. If the pair $(1,2)$ is the first to trade, the price does depend on the identity of the proposer, but the mean price is clearly

$$\begin{aligned} \bar{p} &= \frac{1}{2}\left(v_1 - v_1^* + w_2 + w_2^*\right) \\ &= \frac{1}{4}\left(2p_{21} + p_{11} + p_{22}\right) \\ &= p_{21}. \end{aligned}$$

Each of the three cases is equally likely. The data are summarized in the following table.

Probability	First trade	Mean Price	Second Trade	Price
1/3	$(1,1)$	$v_1 - v_1^* = w_1 + w_1^*$	$(2,2)$	p_{22}
1/3	$(2,2)$	$v_2 - v_2^* = w_2 + w_2^*$	$(1,1)$	p_{11}
1/3	$(1,2)$	p_{21}	$(2,1)$	p_{21}

After the first trade occurs, there are three possible subgames that can occur on the equilibrium path, depending on which of three pairs traded first. If $(1,1)$ trade first, the subgame consists of buyer 2 and seller 2 who exchange the good at a price of p_{22}. If $(2,2)$ trade first, the subgame consists of buyer 1 and seller 1 who exchange the good at a price of p_{11}. If $(1,2)$ trade first, the subgame consists of buyer 2 and seller 1 who exchange the good at a price of p_{21}.

It is clear that these outcomes can be supported as MPE of the subgames (any individually rational trades are MPE outcomes for the two-person subgames). By inspection of the equations and inequalities in (1), the strategies in the first stages of the game, before the first trade occurs, are best responses, given the assumed payoffs from the continuation game. It remains to show that the assumed payoffs are achieved by these strategies. Direct calculation yields

$$\begin{aligned} v_1^* &= \frac{1}{3}v_1^* + \frac{1}{3}(v_1 - p_{11}) + \frac{1}{3}(v_1 - p_{21}) \\ &= v_1 - \frac{1}{2}\left(p_{11} + p_{21}\right), \end{aligned}$$

where the three terms in the first line correspond to first trades by $(1,1)$, $(2,2)$ and $(1,2)$, respectively. Similarly,

$$\begin{aligned} v_2^* &= \frac{1}{3}(v_2 - p_{22}) + \frac{1}{3}v_2^* + \frac{1}{3}(v_2 - p_{21}) \\ &= v_2 - \frac{1}{2}(p_{22} + p_{21}), \end{aligned}$$

$$\begin{aligned} w_1^* &= \frac{1}{3}w_1^* + \frac{1}{3}(p_{11} - w_1) + \frac{1}{3}(p_{21} - w_1) \\ &= \frac{1}{2}(p_{11} + p_{21}) - w_1, \end{aligned}$$

and

$$\begin{aligned} w_2^* &= \frac{1}{3}(p_{22} - w_2) + \frac{1}{3}w_2^* + \frac{1}{3}(p_{21} - w_2) \\ &= \frac{1}{2}(p_{22} + p_{21}) - w_2, \end{aligned}$$

as required.

So far we have defined the equilibrium path; it is clear how to define the complete strategies to deal with deviations from the equilibrium path (reject offers that are worse than the equilibrium payoff and accept those that are better).

This example shows that trade occurs at non-uniform prices and that some of these prices do not belong to the competitive interval (for some choices of p_{11} and p_{22}, seller 1 and buyer 1 trade at a price higher than v_2). Since the parameters p_{11} and p_{22} can be chosen arbitrarily within some small intervals, there exists a continuum of non-competitive, Markov-perfect equilibria.

Proposition 5 *Let $w_1 < w_2 < v_2 < v_1$ define a market. There exists a continuum of MPE with equilibrium payoffs $(v_1^*, v_2^*, w_1^*, w_2^*)$ defined by*

$$\begin{aligned} v_1 - v_1^* &= \frac{1}{2}(p_{11} + p_{21}) = w_1 + w_1^*, \\ v_2 - v_2^* &= \frac{1}{2}(p_{22} + p_{21}) = w_2 + w_2^*, \end{aligned}$$

where the parameters (p_{11}, p_{21}, p_{22}) *satisfy*

$$\begin{aligned} w_j &\leq p_{ij} \leq v_i, (i,j) \neq (1,2), \\ p_{22} &< p_{11}, \\ p_{21} &= \frac{1}{2}(p_{11} + p_{22}). \end{aligned}$$

Note also that the equilibrium payoffs change *non-monotonically* across subgames. For example, seller 1's payoff before any trade has occurred is w_1^*. If seller 2 and buyer 2 trade first, then seller 1 is in a subgame with buyer 1 and trade occurs at the price p_{11}. If seller 2 and buyer 1 trade first, then seller 1 is in a subgame with buyer 2 and trade occurs at a price p_{21}. By construction

$$w_1^* = \frac{1}{2}(p_{11} + p_{21}) - w_1,$$

where

$$p_{21} = \frac{1}{2}(p_{11} + p_{22})$$

and

$$p_{22} < p_{11}.$$

These relations imply that

$$p_{21} - w_1 < w_1^* < p_{11} - w_1.$$

Thus, if seller 1 is not one of the first agents to trade, his payoff may go up or down, depending on the identity of the first pair of agents to trade. This non-monotonic behavior of payoffs turns out to be a crucial feature of non-competitive equilibria.

The strategies that support the MPE depend only on the players who are matched and the identity of the proposer. It is clear that these strategies constitute a PEC, so lexicographic complexity costs (in the sense used by GS) do not guarantee a Walrasian outcome when there is random matching.

4.2 Discounting

In market games with exogenous matching, introducing a small amount of discounting is sometimes enough to ensure the uniqueness of perfect equilibrium. As mentioned before, under certain conditions, the unique perfect equilibrium with discounting converges to the competitive outcome as the discount rate converges to 1. The next example shows that this strategy does not work for markets with heterogeneous buyers and sellers.

Consider a symmetric market in which there are two buyers and two sellers and all agents are inframarginal. The valuations of the agents are denoted by $w_1 < w_2 < v_2 < v_1$, where without loss of generality we can assume that

$$\begin{aligned} w_1 &= 0 \\ w_2 &= w \\ v_2 &= v \\ v_1 &= v + w \end{aligned}$$

for some $0 < w < v$. Payoffs are discounted using a common discount factor $0 < \delta < 1$. The next result characterizes the symmetric MPE payoffs for this example.

Proposition 6 *Let* $(w_1, w_2, v_1, v_2) = (0, w, v+w, v)$ *define a symmetric market with common discount factor* δ*. For some* $\delta^* < 1$ *and all* $\delta^* < \delta < 1$*, there is a unique symmetric MPE of the game with discounting in which the equilibrium payoffs are defined implicitly by the equations:*

$$\begin{aligned} v_1^* &= w_1^* = \frac{(2+\delta)v + w - \delta v_2^*}{8 - 3\delta} \\ v_2^* &= w_2^* = \frac{(1+2\delta)v - \delta w - \delta v_1^*}{8 - 3\delta}. \end{aligned}$$

After the first pair of agents has traded, the remaining two agents are in a two-person bargaining game and we know from the bargaining literature (e.g., Rubinstein (1982)) that with equal discount factors they will split the surplus. More precisely, the division of the surplus will depend on who is the proposer but ex ante they have equal chances of being proposer and hence receive equal payoffs. If w_1 and v_1 (resp. w_2 and v_2) are left to trade in the two-person subgame, they each receive $(v+w)/2$ (resp. $(v-w)/2$) if w_1 and v_2 (resp. w_2 and v_1) are left to trade in the two-person subgame, they will each receive $v/2$. As $\delta \to 1$ the prices become independent of the proposer. If w_1 and v_1 (resp. w_2 and v_2) are left to trade in the two-person subgame, the price at which they trade will be approximately $(v+w)/2$; by contrast, if w_1 and v_2 (resp. w_2 and v_1) are left to trade in the two-person subgame, the price at which they trade will be approximately $v/2$ (resp. $w+v/2$). The competitive interval is $[w_2, v_2] = [w, v]$, which need not contain any of these prices.

In a symmetric MPE, v_1 and w_1 receive the same payoff and v_2 and w_2 receive the same payoff. Define the numbers v_1^*, v_2^*, w_1^*, and w_2^* implicitly by

$$\begin{aligned} v_1^* &= w_1^* = \frac{(2+\delta)v + w - \delta v_2^*}{8-3\delta} \\ v_2^* &= w_2^* = \frac{(1+2\delta)v - \delta w - \delta v_1^*}{8-3\delta}. \end{aligned}$$

Our strategy, as before, is to treat these numbers as if they represent the reservation utilities of the agents, define individual strategies accordingly, and then show that these are precisely the payoffs achieved in equilibrium. Substituting v_2^* into the first equation and simplifying yields

$$v_1^* = \frac{(8-3\delta)\left[(2+\delta)v + w\right] - \delta(1+2\delta)v + \delta^2 w}{(8-3\delta)^2 - \delta^2}.$$

Note that

$$\lim_{\delta\to 1} v_1^* = \lim_{\delta\to 1} w_1^* = \frac{1}{2}v + \frac{1}{4}w$$

and

$$\lim_{\delta\to 1} v_2^* = \lim_{\delta\to 1} w_2^* = \frac{1}{2}v - \frac{1}{4}w.$$

Thus,

$$\lim_{\delta\to 1} \delta\left(v_1^* + w_1^*\right) = v + \frac{1}{2}w < v + w \tag{2}$$

and

$$\lim_{\delta\to 1} \delta\left(v_2^* + w_2^*\right) = v - \frac{1}{2}w > v - w. \tag{3}$$

Thus, by (2) and (3), for δ close to 1, as long as there are four agents remaining, v_1 and w_1 must trade whenever matched and v_2 and w_2 cannot trade whenever matched. Summing the equations that define v_1^* and v_2^* we get

$$\delta\left(v_1^* + v_2^*\right) = \frac{\delta}{8-3\delta}\left[(2+\delta)v + w + (1+2\delta)v - \delta w\right] - \frac{\delta^2(v_1^* + v_2^*)}{8-3\delta}.$$

Thus,

$$\begin{aligned}
\left(8\delta - 3\delta^2 + \delta^2\right)\left(v_1^* + v_2^*\right) &= \delta\left((3+\delta)v + (1-\delta)w\right) \\
\delta(8-2\delta)\left(v_1^* + v_2^*\right) &= \delta\left((3+\delta)v + (1-\delta)w\right) \\
\delta\left(v_1^* + v_2^*\right) &= \frac{\delta}{(8-2\delta)}\left((3+\delta)v + (1-\delta)w\right) \\
&< \frac{\delta(4+\delta)v}{(8-2\delta)} < v.
\end{aligned}$$

Since $v_1^* = w_1^*$ and $v_2^* = w_2^*$ this establishes that, as long as there are four agents remaining, whenever v_1 and w_2 meet or v_2 and w_1 meet, they must trade.

The Markov strategies for the game with four agents remaining are as follows. If $(i,j) = (2,2)$ are matched, no individually rational trade is possible, so the proposer makes an offer that the responder must reject. Otherwise, if $(i,j) \neq (2,2)$ and i is chosen as the proposer, he proposes a price

$$p_{ij} = w_j + \delta w_j^*,$$

and w_j accepts any price greater than or equal to p_{ij} and rejects any other price. If j is chosen as the proposer, then he proposes a price

$$p_{ji} = v_i - \delta v_i^*$$

and v_i accepts any price less than or equal to p_{ji} and rejects any other price. The two-person subgames have unique MPE.

It remains to show that these strategies constitute a symmetric MPE. In the game with four agents remaining, there are four equally probable matches: (v_1, w_1), (v_1, w_2), (v_2, w_1), (v_2, w_2). The fourth results in no trade. The payoff to v_1 must satisfy

$$v_1^* = \frac{1}{4}\left(\frac{1}{2}\delta v_1^* + \frac{1}{2}\left(v + w - \delta w_1^*\right)\right) + \frac{1}{4}\left(\frac{1}{2}\delta v_1^* + \frac{1}{2}\left(v - \delta w_2^*\right)\right) + \frac{1}{4}\delta\frac{1}{2}v + \frac{1}{4}\delta v_1^*$$

and the payoff to v_2 must satisfy

$$v_2^* = \frac{1}{4}\left(\frac{1}{2}\delta v_2^* + \frac{1}{2}\left(v - \delta w_1^*\right)\right) + \frac{1}{4}\delta\frac{1}{2}(v - w) + \frac{1}{4}\delta\frac{1}{2}v + \frac{1}{4}\delta v_2^*.$$

Using symmetry, $v_1^* = w_1^*$ and $v_2^* = w_2^*$, these equations can be rewritten successively as

$$\begin{aligned}\left(1 - \frac{3}{8}\delta\right) v_1^* &= \frac{1}{8}(v + w) + \frac{1}{8}v + \frac{1}{8}\delta v - \frac{1}{8}\delta v_2^*,\\ \left(1 - \frac{3}{8}\delta\right) v_2^* &= \frac{1}{8}v + \frac{\delta}{8}(v - w) + \frac{1}{8}\delta v - \frac{1}{8}\delta v_1^*,\end{aligned}$$

or

$$\begin{aligned}(8 - 3\delta)\, v_1^* &= (2 + \delta)v + w - \delta v_2^*,\\ (8 - 3\delta)\, v_2^* &= (1 + 2\delta)v - \delta w - \delta v_1^*,\end{aligned}$$

which gives the definitions above.

This proves that we have constructed a symmetric, Markov-perfect equilibrium. To see that it is unique, one only has to note that any symmetric, Markov-perfect equilibrium will have payoffs $w_1^{**}, w_2^{**}, v_1^{**}$, and v_2^{**} that will uniquely determine the individual strategies. It can be further shown that, in any symmetric MPE, the pair (i, j) must trade in the initial four-person subgames if and only if $(i, j) \neq (2, 2)$. Then the preceding calculations show that the payoffs are unique and equal to w_1^*, w_2^*, v_1^*, and v_2^*.

The equilibrium described in the proposition is non-competitive because (a) trade takes place at non-uniform prices and (b) these prices may not belong to the competitive interval. This is true even as $\delta \to 1$. Let $p_{k\ell}^{(n)}$ denote the limiting value, as $\delta \to 1$, of the price at which k and ℓ trade when there are n agents left in the market, k is the proposer and ℓ is the responder. Solving the equations above, we see that the equilibrium payoffs in the limit as $\delta \to 1$ are

$$\begin{aligned}v_1^* &= w_1^* = \frac{1}{2}v + \frac{1}{4}w,\\ v_2^* &= w_2^* = \frac{1}{2}v - \frac{1}{4}w.\end{aligned}$$

Note that

$$\begin{aligned}v_1^* + w_1^* &= v + \frac{1}{2}w < v + w,\\ v_2^* + w_2^* &= v - \frac{1}{2}w > v - w,\\ v_2^* + w_1^* &= v_1^* + w_2^* = v.\end{aligned}$$

When there are four agents in the market, $(i,j) = (2,2)$ cannot trade, $(i,j) = (1,1)$ trade at prices depending on the identity of the proposer, and $(i,j) = (1,2),(2,1)$ trade at unique price, independently of the identity of the proposer. When there are only two agents left, they trade at a price that splits the surplus equally, independently of the identity of the proposer. Thus, we have the following corollary.

Corollary 7 *In the equilibrium described in Proposition 6, the equilibrium prices have the following limiting values as $\delta \to 1$:*

$$p_{k\ell}^{(4)} = \begin{cases} \frac{1}{2}v + \frac{1}{4}w & \text{if } \ell \text{ is either seller 1 or buyer 2,} \\ \frac{1}{2}v + \frac{3}{4}w & \text{if } \ell \text{ is either seller 2 or buyer 1,} \end{cases}$$

$$p_{k\ell}^{(2)} = \begin{cases} \frac{1}{2}(v+w) & \text{if } (k,\ell) = (1,1),(2,2), \\ \frac{1}{2}v & \text{if } \{k,\ell\} = \{i,j\}, (i,j) = (2,1), \\ \frac{1}{2}v + w & \text{if } \{k,\ell\} = \{i,j\}, (i,j) = (1,2). \end{cases}$$

Note that the example can be generalized to allow for asymmetry and a larger number of buyers and sellers: the restrictive assumptions used here are for illustrative purposes only.

5 Monotonicity

The example in Section 4.1 shows that the Markov property is not enough, by itself, to establish perfect competition. Here we introduce an additional property, which we call monotonicity. First, we note that in any Markov equilibrium, payoffs are monotonically non-increasing in an expected value sense. If f^* is a Markov equilibrium and N is a set of remaining players, let $U_k(N, f^*)$ denote the equilibrium payoff in the subgame defined by N, for any $k \in N$.

Lemma 8 *Let f^* be a Markov equilibrium and N a set of remaining players observed with positive probability on the equilibrium path. For any $k \in N$, agent k's payoff $U_i(N, f^*)$ must be at least as great as the expected payoff from the continuation game.*

Proof. By stationarity, buyer i will never trade in the subgame with remaining agents N at a price that gives him less than $U_i(N, f^*)$, since his continuation payoff will be $U_i(N, f^*)$ if he rejects. Since $U_i(N, f^*)$ is a weighted average of what he gets if he trades first and what he gets in the subgames after the first trade has occurred, the expected continuation payoff must be less than or equal to $U_i(N, f^*)$. ■

The monotonicity assumption is just a strengthening of this property. It requires that payoffs be non-increasing with probability one along the equilibrium path and not just in an expected-value sense.

Definition 9 *A Markov equilibrium f^* is* monotonic *if, along the equilibrium path, payoffs are monotonically non-increasing with probability one. Formally, a Markov equilibrium f^* is monotonic if, for any subgames N and N', with $N \subset N'$, that are reached with positive probability and any $k \in N$, $U_k(N, f^*) \leq U_k(N', f^*)$.*

From Lemma 8 and the definition, it is clear that monotonicity is automatically satisfied whenever there is no uncertainty about equilibrium payoffs. For example, if f^* is a Markov equilibrium and the matching process is deterministic, then in each subgame the payoffs in the continuation game are known with certainty and Lemma 8 implies that f^* is monotonic.

More importantly, if f^* is *competitive* in the sense that

$$\begin{aligned} U_i(N, f^*) &= \max\{v_i - p, 0\}, i = 1, ..., n, \\ U_j(N, f^*) &= \max\{p - w_j, 0\}, j = 1, ..., n. \end{aligned}$$

for some price p and every subgame N reached with positive probability on the equilibrium path, f^* is clearly monotonic.

So far, we have been deliberately vague about the definition of equilibrium used in the sequel. The game analyzed here is a game of perfect information, so subgame perfect equilibrium would be an appropriate solution concept. However, it turns out that the full power of subgame perfection is not needed. Something much closer to Nash equilibrium will suffice for the purpose of characterizing equilibrium outcomes and that is what we use. This weakening of the equilibrium concept turns out to be helpful in guaranteeing existence of equilibrium with complexity costs, as we discuss in Section 7.

Intuitively, we consider a Nash equilibrium refined by the requirement that, after any history that occurs with positive probability along the equilibrium path, agents respond optimally *to every possible proposal* and not just those that occur with positive probability in equilibrium. Let f^* be a Markov Nash equilibrium and let E denote the equilibrium path, that is, the set of finite histories that occur with positive probability in this equilibrium. We say that f^* is *response-perfect* if, for every $(N, \langle \ell, k \rangle) \in H_k^r \cap E$ and $p \in \mathbf{R}_+$,

$$f_k^*(N, \langle \ell, k \rangle, p) = \begin{cases} \textit{reject} & \text{if } k = i \text{ and } v_i - p < U_i(N, f^*) \\ \textit{accept} & \text{if } k = i \text{ and } v_i - p > U_i(N, f^*) \\ \textit{reject} & \text{if } k = j \text{ and } p - w_j < U_j(N, f^*) \\ \textit{accept} & \text{if } k = j \text{ and } p - w_j > U_j(N, f^*). \end{cases}$$

Response-perfect Markov Nash equilibrium is still quite a weak equilibrium concept because we only consider the response to a single deviation from the equilibrium path, rather than arbitrary finite numbers of deviations in the case of SPE. It is easy to see why response perfection, rather than full subgame perfection, is sufficient for our purposes. If an agent deviates, his deviation either results in trade, in which case he is out of the game, or it does not result in trade, in which case the market is the same at the next date. Markov strategies do not "remember" deviations from previous periods. Thus, a deviation that does not result in trade has no effect on the future play of the game. In determining his optimal strategy, an agent only needs to consider what happens along the equilibrium path or the response to a single deviation from the equilibrium path.

Let f^* be a response-perfect, Markov Nash equilibrium. We have seen that f^* is monotonic if it is competitive. The following series of results establishes the converse: if f^* is monotonic, then it must also be competitive.

Theorem 10 *Let f^* be a response-perfect, monotonic Nash equilibrium in Markov strategies for the market game Γ. Then exchange is efficient and there exists a price p in the competitive interval such that the equilibrium payoffs satisfy*

$$\begin{aligned} U_i(f^*) &= \max\{v_i - p, 0\}, i = 1, ..., n, \\ U_j(f^*) &= \max\{p - w_j, 0\}, j = 1, ..., n. \end{aligned}$$

Proof. Let f^* be a fixed but arbitrary response-perfect and monotonic Nash equilibrium in Markov strategies. The proof is by (backward) induction on the number of agents remaining in the game.

Starting the induction.

To start the induction, let N_0 be a minimal set of agents (possibly the empty set) observed with positive probability on the equilibrium path, that is, no proper subset of N_0 is observed along the equilibrium path. (Note that N_0 need not be unique). We claim that $U_i(N_0, f^*) = U_j(N_0, f^*) = 0$, for every i and j in N_0. The proof is by contradiction. If $N_0 = \varnothing$ or $v_i \leq w_j$ for every i and j in N_0 the result is obviously true, so we can assume without essential loss of generality that there exists a pair (i_0, j_0) in N_0 such that $v_{i_0} > w_{j_0}$. Because the strategies are Markovian, the equilibrium payoffs $U_{i_0}(N_0, f^*)$ and $U_{j_0}(N_0, f^*)$ are stationary. Buyer i_0 can guarantee a payoff of $v_{i_0} - w_{j_0} - U_{j_0}(N, f^*) - \varepsilon$, for any $\varepsilon > 0$, by waiting until he is proposer and offering a price $p = w_{j_0} + U_{j_0}(N_0, f^*) + \varepsilon$ which seller j_0 must accept. By similar reasoning, seller j_0 can guarantee a payoff of $v_{i_0} - U_{i_0}(N_0, f^*) - \varepsilon - w_{j_0}$ for any $\varepsilon > 0$. Since the choice of ε is arbitrary, this proves that $U_{i_0}(N_0, f^*) + U_{j_0}(N_0, f^*) \geq v_{i_0} - w_{j_0} > 0$. This implies that at least one of the agents trades with positive probability in the subgame defined by N_0, contradicting the assumption that N_0 is a minimal set occurring along the equilibrium path. This contradiction proves that $U_i(N_0, f^*) = U_j(N_0, f^*) = 0$, for every i and j in N_0.

The induction step
Suppose that, for some positive integer q, the following condition is satisfied for any set of agents N observed along the equilibrium path and satisfying $|N| \leq 2q$.
INDUCTION HYPOTHESIS: In the subgame beginning when N is first observed, exchange is efficient and there exists a price p in the competitive interval of the market with remaining agents N such that the payoffs are given by

$$\begin{aligned} U_i(N, f^*) &= \max\{v_i - p, 0\}, \\ U_j(N, f^*) &= \max\{p - w_j, 0\}, \end{aligned}$$

for any $i, j \in N$. (If the competitive interval is empty, we interpret the equations as implying $U_i(N, f^*) = U_j(N, f^*) = 0$ for all $i, j \in N$).

Now consider a set of agents N, observed with positive probability along the equilibrium path, such that $|N| = 2(q+1)$ and consider what happens in the subgame that begins when N is first observed. Since the equilibrium is fixed in the sequel, we suppress the reference to f^* and denote the equilibrium payoffs in this subgame by $U_i(N)$ and $U_j(N)$ for $i, j \in N$.

Step 1. Suppose that buyer i' trades first with probability 1. Then the Induction Hypothesis is satisfied.

If buyer i' trades first with probability 1, then buyer i' controls the possibility of trade. This means that

$$U_{i'}(N) \geq \max_{j \in J \cap N} \left\{ v_{i'} - w_j - U_j(N) \right\}.$$

For any $i, j \in N, i \neq i'$,

$$U_i(N) + U_j(N) \geq v_i - w_j.$$

Otherwise, i and j would trade with positive probability. The total surplus (gains from trade) when the set of agents is N is denoted by $S(N)$ and defined by

$$S(N) = \sum_{i \in I'} v_i - \sum_{j \in J'} w_j$$

where $I' \subseteq I \cap N$ and $J' \subseteq J \cap N$ are the maximal sets $|I'| = |J'|$ such that $w_j < v_i$ for all $i \in I'$ and $j \in J'$. (There is no essential loss of generality in ignoring i and j such that $w_j = v_i$). Now, individual rationality implies that $U_i(N) \geq 0$ and $U_j(N) \geq 0$ so

$$\sum_{i \in I'} U_i(N) + \sum_{j \in J'} U_j(N) \leq \sum_{i \in I \cap N} U_i(N) + \sum_{j \in J \cap N} U_j(N)$$

and feasibility implies that

$$\begin{aligned} \sum_{i \in I \cap N} U_i(N) + \sum_{j \in J \cap N} U_j(N) &\leq S(N) \\ &= \sum_{i \in I'} U_i(N) + \sum_{j \in J'} U_j(N). \end{aligned}$$

It follows immediately that $U_i(N) = 0 = U_j(N)$ for extramarginal agents $i \in N \cap (I \backslash I')$ and $j \in N \cap (J \backslash J')$ and $U_i(N) + U_j(N) = v_i - w_j$ for inframarginal agents $i \in I'$ and $j \in J'$. Then putting

$$p = v_i - U_i(N) = w_j + U_j(N), \forall i \in I', \forall j \in J',$$

establishes the claim.

Step 2. If seller j trades first with probability 1, the Induction Hypothesis is satisfied.
The argument is exactly similar to Step 1.

Step 3. Suppose that two distinct pairs have a positive probability of trading first. Then the conditions of the Induction Hypothesis are satisfied.

Fix any $i, j \in N$. By assumption there exist $i' \neq i$ and $j' \neq j$ such that (i', j') have a positive probability of trading first. But as soon as the pair (i', j') trades, the Induction Hypothesis implies that the payoffs of $i, j \in N \backslash \{i', j'\}$ in the continuation game are determined by a price p' such that

$$\begin{aligned} U_i(N \backslash \{i', j'\}, f^*) &= \max\{v_i - p', 0\}, \\ U_j(N \backslash \{i', j'\}, f^*) &= \max\{p' - w_j, 0\}. \end{aligned}$$

Since $i, j \in N \backslash \{i', j'\}$, monotonicity implies that

$$w_j + U_j(N) \geq p' \geq v_i - U_i(N).$$

Thus, for any $i, j \in N$,

$$U_i(N) + U_j(N) \geq v_i - w_j.$$

The rest of the argument follows as in Step 1. This establishes the claim. ■

Theorem 10 establishes that the equilibrium payoffs correspond to a competitive equilibrium, but it does not explicitly state that all trade occurs at a common price p. However, this is an easy corollary of the Theorem together with monotonicity.

Corollary 11 *Let f^* be a response-perfect, monotonic Nash equilibrium in Markov strategies for the market game Γ. Then all trade occurs at a common price p in every subgame observed along the equilibrium path.*

Proof. In the course of proving Theorem 10 we showed that, for any set of agents N observed on the equilibrium path, there exists a price p such that

$$\begin{aligned} U_i(N, f^*) &= \max\{v_i - p, 0\}, \forall i \in I \cap N \\ U_j(N, f^*) &= \max\{p - w_j, 0\}, \forall j \in J \cap N. \end{aligned}$$

Monotonicity implies that, for any $N' \subset N$ observed along the equilibrium path,

$$U_k(N, f^*) \geq U_k(N', f^*), \forall k \in N.$$

This immediately implies that

$$\begin{aligned} U_i(N', f^*) &= \max\{v_i - p, 0\}, \forall i \in I \cap N' \\ U_j(N', f^*) &= \max\{p - w_j, 0\}, \forall j \in J \cap N'. \end{aligned}$$

In other words, trade occurs at the same price p in the subgame defined by N'. By induction, all trade occurs at a common price p. ■

6 Global complexity

Monotonicity is a restriction on equilibrium continuation payoffs. As such, it is in need of some motivation. One way of motivating monotonicity is to show that it follows from lexicographic minimization of complexity costs. As we mentioned earlier, GS apply a concept of complexity that is 'local' in the sense that the partial ordering of strategies according to complexity is defined in terms of individual subgames. Here we use a concept of complexity that is 'global' in the sense that the partial ordering of strategies involves comparisons across subgames.

To measure complexity, we define a partial ordering $\prec^r$ on F_k as follows:

Definition 12 *For any strategies $f_k, f'_k \in F_k$, we say that f_k is (globally) less complex than f'_k, written $f_k \prec^r f'_k$, if and only if, for some partial information set $d \in \{\langle k, \ell \rangle, (\langle \ell, k \rangle, p)\}$,*

$$\begin{array}{ll} f'_k(N,d) = f'_k(N',d), & \forall N, N' \\ f_k(N,d) \neq f_k(N',d), & \text{for some } N \text{ and } N' \end{array}$$

and

$$f_k(N,d') = f'_k(N,d'), \forall N, \forall d' \neq d.$$

This notion of complexity requires two strategies to be identical everywhere except on information sets of the form (N, d) for some particular partial history d. On this information set, f'_k prescribes the same action for all subgames defined by a set of remaining agents N but f_k prescribes different actions for subgames defined by two sets of remaining agents N and N'. This is the notion of global complexity applied to the set of all strategies in GS. It is 'global' in the sense that it compares the complexity of two strategies for a given partial information set d across all subgames N. It is also stronger than is needed for the result that follows. For Theorem 15, the following notion of complexity, which is clearly implied by global complexity, is sufficient. For any information sets $h, h' \in H_k$, we write $h \geq h'$ if

$$h = (N,d), h' = (N',d), N' \subset N.$$

Definition 13 *For any strategies $f_k, f'_k \in F_k$, we say that f_k is less complex than f'_k, written $f_k \prec^r f'_k$, if and only if, for some information set $h \in H_k$,*

$$\begin{array}{ll} f'_k(h') = f'_k(h''), & \forall h', h'' \geq h \\ f_k(h') \neq f_k(h''), & \exists h', h'' \geq h, \\ f'_k(h') = f_k(h') & \forall h' \not\geq h. \end{array}$$

This weaker definition restricts the comparisons of the two strategies to subgames in which the set of remaining players is $N' \subseteq N$, for some fixed N, whereas the previous definition admits all sets N'.

Definition 14 *A* Nash equilibrium with complexity costs and perfect responses *(NECPR) is a Nash equilibrium with perfect responses f^* such that, for every agent k, there does not exist a strategy f_k that is a best response to f^*_{-k} and is less complex than f^*_k (in the sense of Definition 12 or Definition 13).*

Theorem 15 *Let f^* be a Markov equilibrium with perfect responses for the market game Γ. Then (a) if f^* is monotonic there exists a NECPR f^{**} with the same payoffs and (b) if f^* is a NECPR it is monotonic.*

Proof. To prove (a) it is sufficient to note that by Theorem 10 f^* is competitive with equilibrium price p^*, say. Then we can define the NECPR f^{**} as follows: for any $(h, \langle k, \ell \rangle) \in H^p_k$ put

$$f^{**}_k(h, \langle k, \ell \rangle) = \begin{cases} \min\{p^*, v_k\} & \text{if } k \in I \\ \max\{p^*, w_k\} & \text{if } k \in J \end{cases}$$

and for any $(h, \langle k, \ell \rangle) \in H^r_k$ and any price p put

$$f^{**}_k(h, \langle \ell, k \rangle, p) = \begin{cases} \textit{accept} & p \geq \max\{p^*, w_k\}, k \in J \\ \textit{accept} & p \leq \min\{p^*, v_k\}, k \in I \\ \textit{reject} & p < \max\{p^*, w_k\}, k \in J \\ \textit{reject} & p > \min\{p^*, v_k\}, k \in I \end{cases}$$

It is clear that f^{**} so defined is a NECPR in Markov strategies and has the same payoffs as f^*.

To prove (b) we assume that f^* is a NECPR and prove that it is competitive. The proof is the identical to the proof of Theorem 10 up to Step 3, because this is the only part of the proof where an appeal is made to monotonicity.

Step 3. Suppose that two distinct pairs have a positive probability of trading first. Then the conditions of the Induction Hypothesis are satisfied.

Fix any $i, j \in N$. Then by assumption there exist $i' \neq i$ and $j' \neq j$ such that (i', j') have a positive probability of trading first. But as soon as the pair (i', j') trades, the Induction Hypothesis implies that the payoffs of $i, j \in N \backslash \{i', j'\}$ in the continuation game are determined by a price p' such that

$$\begin{aligned} U_i(N \backslash \{i', j'\}, f^*) &= \max\{v_i - p', 0\}, \\ U_j(N \backslash \{i', j'\}, f^*) &= \max\{p' - w_j, 0\}. \end{aligned}$$

Now since f^* satisfies perfect responses, agent i must accept any price below $v_i - U_i(N)$ and agent j must accept any price above $w_j + U_j(N)$ in the subgame where N is the set of agents remaining. Since f^* is a NECPR, in all subgames, including the subgame where $N \backslash \{i', j'\}$ is the set of agents remaining, agent i must accept any price $p < v_i - U_i(N)$ offered by an agent $\ell \in N \backslash \{i', j'\}$ that does not occur on the equilibrium path. Otherwise, we could define less complex strategy $f_i \prec^r f_i^*$ by putting

$$f_i(N', (\langle \ell, i \rangle, p)) = \begin{cases} \textit{accept}, & \forall N', \\ f_i^*(N', d) & \forall d \neq (\langle \ell, i \rangle, p), \end{cases}$$

and note that f_i is a best response because it differs from f_i^* only off the equilibrium path, contradicting the definition of NECPR. By a similar argument, agent j must accept any price above $w_j + U_j(N)$ that does not occur on the equilibrium path. In any interval there is a dense set of prices that do not occur on the equilibrium path, so the preceding claim is consistent with equilibrium in the subgame with agents $N \backslash \{i', j'\}$ only if

$$w_j + U_j(N) \geq p' \geq v_i - U_i(N).$$

Thus, for any pair $i, j \in N$,

$$U_i(N) + U_j(N) \geq v_i - w_j.$$

The rest of the argument is identical to Step 1 in the proof of Theorem 10.

This establishes that the equilibrium payoffs correspond to a competitive equilibrium, but it does not guarantee that all trade occurs at the same competitive price. To show this, we merely note that, in each subgame reached along the equilibrium path, each agent i remaining in the game will accept any price below $v_i - U_i(N)$ and each agent j remaining in the game will accept any price above $w_j + U_j(N)$, where $U_i(N)$ and $U_j(N)$ are their equilibrium payoffs. Clearly, this suffices to maintain a uniform price and implies that f^* is monotonic. ■

7 Discussion

7.1 Existence of SPEC in Markov strategies

In place of the "natural" solution concept, subgame perfect equilibrium, we used the weaker concept of Nash equilibrium with perfect responses. It is interesting that a weaker notion of equilibrium is sufficient to characterize the competitive outcomes, but the choice of solution concept is to some extent driven by necessity. Adding the refinement of complexity costs to subgame perfect equilibrium may threaten existence, as the following example shows.

Proposition 16 *Consider a market consisting of four agents with valuations satisfying $w_1 < v_2 < w_2 < v_1$ and the matching process described in Section 3. Then if f^* is a NECPR in Markov strategies, it cannot also be a SPE.*

Proof. We argue by contradiction. Suppose that f^* is a NECPR in Markov strategies and a SPE. By Theorem 15, any NECPR f^* in Markov strategies will result in w_1 and v_1 trading at a price $p^* \in [v_2, w_2]$. Next, we claim that if f^* is also a SPE, then in the (off-equilibrium) subgame with remaining agents $N = \{v_1, w_2\}$, trade will occur with probability one at a price $p_{12} \in [w_2, v_1]$. To see this, note that Markov strategies imply that agents have stationary payoffs v_1^* and w_2^* as long as the set of agents is N. In a SPE, v_1 must accept any price $p < v_1 - v_1^*$ and w_2 must accept any price $p > w_2 + w_2^*$. It follows that $v_1 - w_2 = v_1^* + w_2^*$ and hence that trade occurs with probability one at $p_{12} = v_1 - v_1^* = w_2 + w_2^*$. Individual rationality implies that $p_{12} \in [w_2, v_1]$.

A similar argument shows that in the subgame with remaining players $N = \{w_1, v_2\}$, trade occurs with probability one at a price $p_{21} \in [w_1, v_2]$.

Then global complexity implies that $p_{12} = p_{21} = p^*$. For example, if $p_{12} > p^*$, then in the subgame with $N = \{w_2, v_1\}$ the buyer v_1 should accept any price $p^* < p < p_{12}$. Then a complexity argument shows that a price $p^* < p < p_{12}$ that is not offered along the equilibrium path should be accepted by v_1 everywhere. But this is clearly not consistent with equilibrium. A similar argument rules out $p_{21} < p^*$. Since $p_{12} = p_{21}$ is impossible, this contradiction shows that f^* cannot be a SPE. ■

7.2 Exogenous, simultaneous moves

The bulk of this paper is devoted to games with sequential matching, i.e., games in which exactly one pair of agents is allowed to bargain and trade at a time. In this section we briefly discusse the class of games with simultaneous matching and show that new sets of non-competitive MPE arise in these models too for reaons that are quite different from the ones discussed in Section 4.

Assume that $m = 1$ and $n = 2$ and that

$$w_1 < v_2 < w_2 < v_1.$$

Matching is random and every buyer-seller pair has equal probability and each member of the pair has an equal probability of being chosen as the proposer.

Consider the following strategies: when the pair (i, j) such that $v_i > w_j$ is formed, the proposer offers to trade at the price $p_{ij} \in [w_j, v_i]$; the responder accepts any price that is at least as good as p_{ij} and rejects any other offer. We claim that these strategies constitute a MPE for the simultaneous matching game. The strategies are clearly stationary: they depend only on the matched pair of agents. The strategies clearly form a subgame perfect equilibrium of the subgames in which only a single buyer and seller are left. Furthermore, when the pair (i, j) forms, they are either the only agents left in the market or they expect the other two agents, who are currently matched, to trade immediately and leave the market. Thus, if i and j do not trade in the current period, they expect to trade at the price p_{ij} in the continuation game. Then clearly it is optimal for the agents to follow the specified strategy in the current period.

This shows that the exogenous simultaneous move game has a continuum of stationary (Markov) non-competitive equilibria.

Random matching is not necessary for this example. The same construction works with deterministic matching as long as agents are matched simultaneously. In this case, the structure of the game is not stationary, but the strategies are stationary since they depend only on the current match. The definition of a MPE should allow strategies to depend on time because the matching rule depends on time, but a stationary strategy does not have to depend on time.

We also note that the strategies in these examples are stationary as a function of N so global complexity will not eliminate these equilibria either. The equilibria are also monotonic. All this suggests that the equilibria are quite robust in this framework. However, it is important to note that there is something contrived about exogenous simultaneous matching: it requires perfect coordination of timing in order to sustain the non-competitive equilibria. It is also possible that noisy exogenous matching processes or endogenous simultaneous matching processes will destroy such equilibria. For these reasons we think the results that are due to the exogenous and the simultaneous nature of the matching model are not appealing. In GS we investigate endogenous simultaneous matching and show that complexity type reasoning can be effective in selecting a competitive outcome in such a set-up.

References

Ausubel, Lawrence and Raymond Deneckere (1989). "Reputation in Bargaining and Durable Goods Monopoly," *Econometrica* **57**, 511-31.

— (1992). "Durable Goods Monopoly with Incomplete Information," *Review of Economic Studies* **59**, 795-812.

Binmore, Ken and Maria Herrero (1988a). "Matching and Bargaining in Dynamic Markets," *Review of Economic Studies* **55**, 17-31.

— (1988b). "Security Equilibrium," *Review of Economic Studies* **55**, 33-48.

Chatterjee, Kalyan and Hamid Sabourian (2000). "Multiperson Bargaining and Strategic Complexity," *Econometrica* **68**, 1491-1509.

Dagan, Nir, Roberto Serrano and Oscar Volij (to appear). "Bargaining, Coalitions and Competition," *Economic Theory*.

Diamond, Peter (1971). "A Model of Price Adjustment," *Journal of Economic Theory* **3**, 156-68.

Gale, Douglas (1986a). "A Simple Characterization of Bargaining Equilibrium in a Large Market Without the Assumption of Dispersed Characteristics," University of Pennsylvania Center for Analytic Research in Economics and the Social Sciences (CARESS) Working Paper: 86-05, pages 19.

— (1986b). "Bargaining and Competition Part I: Characterization," *Econometrica* **54**, pages 785-806.

— (1986c). "Bargaining and Competition Part II: Existence," *Econometrica* **54**, pages 807-818.

— (1987). "Limit Theorems for Markets with Sequential Bargaining," *Journal of Economic Theory* **43** 20-54.

— (2000). *Foundations of General Equilibrium: Dynamic Matching and Bargaining Games.* Cambridge: Cambridge University Press

Gale, Douglas and Hamid Sabourian (2002a). "Competition and Complexity, Part I: Sequential Matching," in preparation.

— and — (2002b). "Competition and Complexity, Part II: Simultaneous Endogenous Matching," in preparation.

Gul, Faruk, Hugo Sonnenschein, and Robert Wilson (1986). "Foundations of Dynamic Monopoly and the Coase Conjecture," *Journal of Economic Theory* **39**, 155-90.

Herrero, M. (1985). *A strategic bargaining approach to market institutions.* Ph.D. thesis: University of London.

Masso, Jordi and Robert Rosenthal (1989). "More on the 'Anti-Folk Theorem'," *Journal of Mathematical Economics* **18**, 281-90.

McLennan, Andrew and Hugo Sonnenschein (1991). "Sequential Bargaining as a Noncooperative Foundation for Walrasian Equilibrium," *Econometrica* **59**, 1395-1424.

Osborne, Martin and Ariel Rubinstein (1990). *Bargaining and Markets* San Diego; London; Sydney and Toronto: Harcourt Brace Jovanovich, Academic Press.

Rubinstein, Ariel (1982). "Perfect Equilibrium in a Bargaining Model," *Econometrica* **50**, 97-109.

Rubinstein, Ariel and Asher Wolinsky (1985). "Equilibrium in a Market with Sequential Bargaining," *Econometrica* **53**, 1133-50.

Rubinstein, Ariel and Asher Wolinsky (1990). "Decentralized Trading, Strategic Behaviour and the Walrasian Outcome," *Review of Economic Studies* **57**, 63-78.

Sabourian, Hamid (2001). "Bargaining and Markets: Complexity and the Walrasian Outcome," Cambridge University, unpublished

Shaked, Avner (1994). "Opting Out: Bazaars versus 'Hi Tech' Markets," *Investigaciones Economicas* **18**, 421-32.