Bargaining and markets: complexity and the Walrasian outcome*

Hamid Sabourian,
Faculty of Economics and Politics,
Unoversity of Cambridge,
Sidgwick Site, Sidgwick Avenue,
Cambridge CB3 9DD, UK
Email: hamid.sabourian@econ.cam.ac.uk

This version: April 2001

Abstract

Rubinstein and Wolinsky (1990b) consider a simple decentralized market in which agents either meet randomly or choose their partners volunatarily and bargain over the terms on which they are willing to trade. Intuition suggests that if there are no transaction costs, the outcome of this matching and bargaining game should be the unique competitive equilibrium. This does not happen. In fact, Rubinstein and Wolinsky show that any price can be sustained as a sequential equilibrium of this game. In this paper, I consider Rubinstein and Wolinsky's model and show that if the complexity costs of implementing strategies enter players' preferences, together with the standard payoff in the game, then the only equilibrium that survives is the unique competitive outcome. This is done both for the random matching and for the voluntary matching models. Thus, the paper demonstrates that complexity costs might have a role in providing a justification for the competitive outcome.

JEL Classification: C72, C78, D5

Keywords: Bargaining, matching, complexity, automata, bounded rationality, competitive outcome, Walrasian equilibrium.

*The first version of this paper was written while visiting NYU Economics Department and the Cowles Foundation Yale University. I would like to thank both institutions for their hospitality. I am also grateful to Kalyan Chatterjee for reading the first draft of this paper. I would also like to thank the seminar participants at Caltech, Cambridge, LSE, NYU, Northwestern, Penn State, Stanford and Warwick Game Theory Conference for helpful comments.

1. Introduction

In a competitive market agents take prices parametrically. This is usually justified by saying that agents are 'negligible'. In a dynamic game-theoretic context this can sometimes be formalized by assuming that there is a continuum of anonymous agents. The equilibria of these models can be shown to coincide with their competitive equilibria under some regularity conditions. On the other hand, the equilibria of dynamic games with a (large but) finite number of players can be shown to be radically different from those in a model with a continuum of players (see Gale (2000)) While this paradox may seem narrow, the issue has broad economic significance. The rationale for the continuum economy is that it is useful idealization for an economy with a large but finite number of agents. Clearly this idealization is of limited value if the equilibria of finite economies are radically different from those of the continuum case.

The reason for this paradox in dynamic games with a finite number of players is that players can choose history-dependent strategies. The possibility of conditioning behaviour on histories induces different expectations for future play depending on the history preceding the play. This allows one to construct a large number of (history-dependent) equilibria in which a single agent has a large effect. The best example of this is the Folk Theorem of the repeated game, which holds for an arbitrary number of players.[1] In these dynamic games a player has to consider the possible reaction of others. As a result, these equilibria will depart from the competitive outcome even in a frictionless market with a large but finite number of players. Thus, even in environments in which the competitive outcome might appear as the natural outcome (e.g. the case of one seller of an indivisible good who faces two buyers who bid for the unit), one can show that in general non-competitive outcomes might emerge as equilibria if the environment is modelled as a dynamic game.

Considering explicitly bounds on or costs of computation and memory (bounded rationality) is one research strategy for dealing with the large number of equilibria in dynamic games. (For example, see Kalai (1990) for a survey of the literature on modelling players as automata in 2-player repeated games.) Such bounds and or costs induce natural restrictions on the way strategies depend on past history.

In this paper, I investigate the effect of introducing complexity costs in the dynamic matching and bargaining games. In particular, I will show that complexity considerations (some elements of 'bounded rationality') can provide a game-theoretic foundation for the competitive behaviour in decentralized markets with a finite number of agents.

[1]With a continuum of players, the Folk Theorem remains valid if the game is non-anonymous. With a finite number of players, Folk Theorem type results can survive even with anonymity. One needs noise (in strong form) and some form of anonymity to eliminate history-dependent equilibria in repeated games with a large but finite number of players (see Green (1980), Sabourian (1989), Levine and Pesendorfer (1995), Gale (2000) and Al-Najjar and Smordinsky (1999)).

Also, I will show that in these models the introduction of complexity costs into players' preferences ensures that in equilibrium players choose history-independent (sometimes referred to as stationary or Markov) strategies.

There is already a large literature on dynamic matching models with explicit non-cooperative bargaining. (See for example Rubinstein and Wolinsky (1985), Binmore and Herrero (1988a,b), Gale (1986a,b and 1987), Mclennan and Sonnenschein (1991); also see the texts by Osborne and Rubinstein (1990a) and by Gale (2000)). By assuming a continuum of agents and/or by restricting the strategy sets to the stationary ones, such models have been used to provide a game-theoretic foundation for the competitive equilibrium .

One of the few papers that deals with bargaining and matching with a finite number of players and with an unrestricted set of strategies, is that of Rubinstein and Wolinsky (1990b) - henceforth referred to as RW. This paper considers a simple decentralized market in which agents either meet randomly or choose their partners voluntarily and bargain over the terms on which they are willing to trade. Intuition suggests that if there are no transaction costs, the outcome of bargaining should be the competitive equilibrium. This intuition turns out to valid if players are restricted to choosing history-independent strategies. However, RW demonstrate that if there are no restrictions on the set of strategies and players can condition their behaviour on past history of plays, then a continuum of non-competitive sequential equilibria emerges. With random matching, this result is only established for the case of zero discounting. However, RW argue convincingly that with random matching zero discounting is the appropriate way to model frictionless markets; otherwise staying with one's partner could be very costly. Thus, they argue that in a model with discounting agents should be able to choose their partners. In such a model, with voluntary matching, they also demonstrate the continuum result even for the case in which players discount the future.

In sections 2 and **??** below, I will consider RW's models and discuss how their predictions differ from that of the competitive behaviour. I shall then show that in RW's models if players attach some arbitrarily small weight to the complexity of their strategies then the only perfect Bayesian equilibrium outcome that survives is the competitive one, and that these equilibrium strategies are unique and stationary. This will be done both for the random matching (Section 2) and for the voluntary matching (Section **??**) models of RW. The equilibrium concept I use in sections 2 and **??** is such that, *in considering complexity*, agents ignore any consideration of payoffs off-the-equilibrium path. Section 4 extends the selection results to weaker equilibrium concepts in which complexity is less important than the off-the-equilibrium payoff. Section 5 contains some concluding remarks. Most of the proofs are in Appendices A1, A2 and B.

In the literature on dynamic games, the strategy of concentrating on the history-independent/stationary/Markov equilibria is very common. Very little justification is usually provided for this approach. Sometimes these Markov equilibria are proposed as "focal" points. Intuitively, one might argue that a player concerned with the cost of implementing complex strategies would choose a stationary strategy, where behaviour in each period is independent of payoff-irrelevant past history.[2] This paper, in addition to providing a justification for the competitive outcomes, attempts to formalise this intuition, in the context of dynamic matching and bargaining.

In this paper, complexity costs are introduced with the standard payoff into the players' preference ordering as in Rubinstein (1986), Abreu and Rubinstein (1988), Piccione and Rubinstein (1993) and others. In these papers, players are modelled as finite-state automata involved in a two-player repeated game. Complexity is measured by the (arbitrarily small) cost of maintaining an additional "machine" state.

Here, I will also focus on the complexity of implementation[3] rather than on computational complexity (see Papadimitriou, 1992). But, because of the asymmetric nature of bargaining, my notion of complexity of strategies is somewhat different from that in the above literature. Informally, the measure of complexity adopted in the random matching model has the following property: if two strategies are otherwise identical except that in some instance the first strategy uses more information than that available in the current period of bargaining and the second uses only the information available in the current period, then the first strategy is more complex than the second. This notion of complexity is a very weak measure of the complexity of response rules within a period. Thus, I shall refer to it as response-complexity. Chatterjee and Sabourian (1999, 2000) also use a similar complexity criterion to justify stationary equilibrium in alternating n-player bargaining games. This notion of complexity neither implies nor is implied by the 'counting states' notion of complexity. In the voluntary matching model of section 4, I shall use both the counting states measure together with the response-complexity to select uniquely the competitive outcome.

RW's dynamic matching and bargaining game is rather special. Other games might give different results. The point, here, is not that there is a right way of modelling competitive behaviour but to give an example of what it takes, in terms of the primitives of the model, to obtain the competitive outcome. In particular, this paper demonstrates that complexity costs might have a role in providing a justification for a competitive equilibrium.

Gale (2000) also discusses how the introduction of 'bounded rationality' can provide a justification for the competitive equilibrium in RW's model. He obtains his results by either putting a bound on the complexity of the *strategy profiles* or by introducing noise in the implementation of the strategies. Our motivation is similar to that of Gale; however the approach taken in this paper is somewhat different from his.

[2] Osborne and Rubinstein (1994) have, however, provided arguments against such intuition. They argue that if equilibrium strategies are thought of as equilibrium in beliefs then it is not clear why players should believe that other players follow the same actions after histories which have involved highly non-stationary past plays.

[3] Other examples involving complexity of implementation include Baskar and Vega-Redonda (2000) Binmore et al (1988) Banks and Sundrum (1990), Neme and Quintas (1995).

2. Random matching model

RW has a model of a market with B identical buyers and S identical sellers. Let $\mathcal{B}$ denote the set of buyers and $\mathcal{S}$ denote the set of sellers. Each seller has one unit of an indivisible good. Each buyer wants to buy at most one unit of the good. The valuations of the buyers and the sellers for one unit of the good are one and zero respectively. Throughout this paper, I assume that

$$B > S.$$

Time is discrete and each player has a discount factor $\delta \in [0, 1]$. Thus, if a seller sells one unit of the commodity to a buyer at a price p in any period $t = 0, 1, 2..$, the payoff of the seller and that of the buyer are given by $\delta^t p$ and $\delta^t(1 - p)$, respectively.

At each period $t \geq 1$, the agents remaining in the market are randomly matched in pairs of one seller and one buyer (all possible matches are equally likely). One member of each matched pair is then randomly chosen (with probability 1/2) to propose a price $p \in [0, 1]$. Then the other agent accepts (A) or rejects (R) the offer. I shall denote such a match between a seller s and a buyer b with s as the proposer by the ordered pair (s, b) and a match between s and b with b as the proposer by the ordered pair (b, s). If a proposal is accepted by the responder, the parties implement it and leave the market. Rejection dissolves the match, in which case the agents proceed to the next matching stage. Any unmatched buyers are forced to remain inactive for the period.

The game is such that Nature effectively chooses the matching and the choice of proposer and responder. Thus, at any period a typical choice of Nature consists of a one-to-one function $q : \mathcal{S} \to \mathcal{B}$ assigning to each seller s a buyer $q(s)$ together with the identity of the proposer and the responder when s is matched with $q(s)$.

At each period t each agent has perfect information about all the past events of the game, including all the past play in matches in which he did not participate. In addition, at each period t, each proposer in a match knows the identity of his opponent in that match and each responder knows both the identity of the proposer in his match and the actual price on offer. However, when the agents choose their actions in any period they do not know the identity of the other matches in that period and what actions are being simultaneously chosen by other agents.

2.1. RW's equilibrium characterisation with random matching for the case of $\delta = 1$.

The competitive model corresponds to the case in which the frictions and the transaction costs in the market are negligible. In RW's model this corresponds to the case in which the players do not discount the future. The main equilibrium characterisation result of RW corresponds to this case.

Theorem 2.1. *(See RW) If $\delta = 1$ then for any price $\overline{p} \in [0, 1]$ and for every one-to-one function $q : \mathcal{S} \to \mathcal{B}$ there exists a perfect Bayesian equilibrium in which each seller s sells one unit of the good to buyer $q(s)$ for a price $\overline{p}$.*[4][5]

Thus, there is a continuum of prices that can be sustained as a perfect Bayesian equilibria. On the other hand, in a competitive equilibrium all goods are sold at the unique price of 1 because $B > S$. Therefore, the competitive outcome is not the unique perfect Bayesian equilibria of the matching and bargaining game with a finite number of agents (irrespective of the numbers B and S).

For the case of a single seller s, the intuition for the proof of Theorem 2.1 is as follows. Construct an equilibrium strategy profile such that at the beginning of the game a distinguished buyer $\overline{b}$ has the 'privilege' to buy the good of s at $\overline{p}$. Having such a 'privilege' has the following meaning: (i) whenever s meets the distinguished buyer $\overline{b}$ an agreement at $\overline{p}$ is reached and (ii) whenever s meets a buyer $b \neq \overline{b}$, no agreement is reached.

Clearly, the outcome of such a relationship is that the seller eventually sells the good to $\overline{b}$ at $\overline{p}$. Moreover, to deter deviations the equilibrium strategy profile specifies the following reward and punishment scheme. If at any stage player i proposes a different price to some player j from that specified by the equilibrium strategy profile then the responder j rejects and subsequently j receives the entire surplus. If j is a buyer this is achieved by the continuation strategy profile giving j the 'privilege' to buy the good at the price of 0 and if j is the seller then this is obtained by the continuation strategy profile giving some buyer the 'privilege' to buy the good at the price of 1.

The above reward and punishment scheme clearly deters any deviation. It does not pay any proposer to deviate given that the responder rejects and receives the entire surplus . Moreover, it is optimal for any responder j to reject any proposed deviation because he obtains the entire surplus following the rejection.

Notice that the above strategies are quite complicated and the behaviour of each agent at any period depends on the history play up to that period - there are potentially an indefinite number of potential deviations and for each deviation the above strategy profile specifies a tailor-made response in order to deter the deviation. Thus the agents

[4]Here, a perfect Bayesian equilibrium refers to a profile of strategies such that, at every information set, each player's strategy maximizes his expected continuation payoff given the strategies of the others, where expectation is with respect the choice of Nature. Of course, in this randon matching game, with symmetric uncertainty over the choice of nature, perfect Bayesian equilibrium is equivalent to sequential equilibrium; however, I shall use the former concept because it is easier to define. (RW, however, state their Theorems in terms of the latter concept.)

[5]Note that when the number of the sellers exceeds one, the above game is one of imperfect information (there is more than one match per period) and the appropriate equilibrium concept is perfect Bayesian equilibrium (or sequential equilibrium). When there is only one seller, the game has perfect information and the set of perfect Bayesian equilibrium coincides with the set of subgame perfect equilibrium.

need a large amount of information to implement the above strategy profile.[6]

If, on the other hand, one assumes that at any period the agents cannot condition their behaviour on the previous history of plays, then it is not possible to implement the above strategies. In particular, if a proposer deviates then it is not possible to reward the responder for rejecting the deviation; this is simply because the initial deviation by the proposer is not remembered. RW go further and establish the following.

Theorem 2.2. *(See RW) If at any period t each player can condition its behaviour only on the set of players that are present in the market at t and on t itself then the unique perfect Bayesian equilibrium price is the competitive price of 1.*

2.2. Complexity and equilibrium

Before introducing the notion of complexity used in this paper, I need some further notation.

An outcome of a match at any period is described by an ordered four-tuple (i, j, p, l) where $i \in \mathcal{B} \cup \mathcal{S}$ is the proposer in this match, $j \in \mathcal{B} \cup \mathcal{S}$ is the responder, $p \in [0, 1]$ is the proposal by i and $l \in \{A, R\}$ is the response by j. I also denote a history of outcomes in a period of the game by e. Thus e consists of outcomes of at most S different matches, one for each seller; it describes everything that happens at a period of the game. For example e could be $\{q(s), m(s), r(s), p(s), l(s)\}_{s\in S}$; namely that each seller s was matched with buyer $q(s)$ and the proposer, the responder, the price offer and the response in this match involving s were $m(s) \in \{s, q(s)\}$, $r(s) \in \{s, q(s)\}$, $p(s) \in [0, 1]$ and $l(s) \in \{A, R\}$, respectively. Let E be the set of such outcomes.

The history of outcome at any time t is denoted by $e^t \in E$ and the history of outcomes of the game up to and including period t consists of a sequence of outcomes $h^t = (e^1, ..., e^t)$. I shall denote the set of such t-period history of outcomes by H^t. Also, let $H^\infty = \cup_{t=0}^{\infty} H^t$ be the set of all possible finite histories of periods. (H^0 is assumed to be the null set).

At each date t, in addition to the history of the outcomes h^t of the preceding periods, players also receive information about the preceding moves by Nature and/or other players during the current period. I also need notation to describe these partial descriptions of outcomes (partial history) a player receives within a bargaining period. I shall denote such a partial history by d and the set of such partial histories by D. Thus $d \in D$ is either the ordered pair (i, j) describing the match between player i and j with i as the proposer, or the ordered triplet (i, j, p) describing the match between players i and j followed by a price offer p by i. If $d = (i, j)$ the bargaining is just beginning and an offer has yet to be made by i to j, and if $d = (i, j, p)$ it is player $j's$ turn to respond to an offer price of p by player i. Also, I shall denote the set of

[6] To implement the above strategies the agents need to observe past outcomes of all matches including those in which they did not participate. However, RW also show that Theorem 2.1 is still valid in the case in which the players only observe their own personal history.

information about the preceding moves by Nature and/or other players during the any period that player i receives (the sets of partial histories for i) by D_i. Thus

$$D_i \equiv \{d \in D \mid \text{it is } i's \text{ turn to play after } d\}$$

Let

$$C = [0,1] \cup A \cup R \cup \varnothing,$$

where $\varnothing$ is the null set. Denote the set of choices available to player i, given a partial description $d \epsilon D_i$, by $C_i(d)$. Thus

$$C_i(d) = \begin{cases} [0,1] & \text{if } d \text{ is such that } i \text{ is the proposer} \\ \{A,\ R\} & \text{if } d \text{ is such that } i \text{ is the responder to some offer } p. \end{cases}$$

Now a strategy for player i is described by a function $f_i : H^\infty \times D_i \longrightarrow C$, such that for any $(h,d) \in H^\infty \times D_i$, $f_i(h,d) \in C_i(d)$ for any $h \in H^t$ such that i has not left the market and $f_i(h,d) = \varnothing$, otherwise.

I shall denote the set of strategies for player i by F_i. Also, f_{-i} refers to the strategy of all players other than that of player i. And finally, for any strategy f_i and for any history $h \in H^\infty$, I shall define the strategy induced by f_i after h by $\langle f_i \mid h\rangle \in F_i$. Thus, for any $h \in H^\infty$

$$\langle f_i \mid h\rangle(h',d) = f_i(h,h',d) \text{ for all } h' \in H^\infty \text{ and } d \in D_i$$

Similarly, the strategy induced by f_i after any (h,d) is denoted by $\langle f_i \mid h,d\rangle$.

2.2.1. Automata and Complexity

The definitions of complexity used in this paper will be defined directly in terms of strategies used by the players in the game. However, in this section, I shall first discuss complexity in terms of automata that implement strategies in order to facilitate comparisons with the existing literature on complexity in games.

Any strategy in the game can be implemented by an automaton (machine) consisting of a set of states (not necessarily finite), an initial state, a terminal state, an output function describing the output of the machine as a function of its current state (and its current input) and a transition function determining the next state of the machine as a function of its current state and current input (the outcome in the current period).

In the literature on automata in repeated one-shot games, there is a natural specification of a machine. Here, I am dealing with a repeated extensive form game. Moreover, since each player has to play a different role (of a proposer and a responder) the extensive form bargaining game in each period has a certain degree of asymmetry built in. As a result, one can specify a machine to implement a particular strategy in several different ways. For example, one could assume

(i) the states of the machine do not change during each period of the game and transitions from a state to another state in the same player's machine take place at the end of a period;

(ii) each state of the machine for player i would specify an action for every $d \in D_i$ - the partial history of the period.

A referee (called "Master of the Game" by Piccione and Rubinstein, (1993)) would activate each player's machine when needed.

The formal definition of such a specification would be the following.

Definition 1. *A machine M_i is a five-tuple $(Q_i, q_i^1, T, \lambda_i, \mu_i)$, where*

Q_i is a set of states;

$q_i^1 \in Q_i$ is a distinguished initial state;

$T \notin Q_i$ is a distinguished terminal state such that if the machine enters this state it shuts off;

$\lambda_i : Q_i \times D_i \to C$, describes the output function of the machine given the state of the machine and given the partial history that has occurred during the current period of the game before i's move, such that $\lambda_i(q_i, d) \in C_i(d)$, $\forall q_i \in Q_i$ and $\forall d \in D_i$;

$\mu_i : Q_i \times E \to Q_i \cup T$ is the transition function, specifying the state of the machine in the next period of the game as a function of the current state and the realised history of the current period.[7]

Remark 1. *If we denote the set of strategies for i in any period of the game by*

$$\mathcal{G}_i \equiv \{g : D_i \to C \,|g(d) \in C_i(d) \; \forall d \in D_i\}$$

then the output function λ_i in Definition 1 can be thought of as a mapping $\widetilde{\lambda}_i : Q_i \to \mathcal{G}_i$ where $\widetilde{\lambda}_i(q_i)(d) = \lambda_i(q_i, d)$. Thus each q_i specifies a mapping $\widetilde{\lambda}_i(q_i) \in \mathcal{G}_i$ from the information set within a period to the set of choices.

The fact that the game is identical at the beginning of each period (though behaviour could be different depending on past histories as encapsulated in the state) provides the basic rationale for using this specification of a machine. Thus, the nature of the output and transition functions remain the same in each period. There are other possible specifications: for example the state of the machine changes before a player has to move or each player has different sub-automaton to play the different roles of a proposer and a responder. However, these alternative specifications do not have the "game-stationarity" features that the above specification does.

Next I need to define the strategy that is implemented by a given machine. Before addressing this, with some abuse of notation, denote the state of machine

[7]Henceforth, I shall not always explicitly refer to the terminal state T. I am assuming that if an offer is accepted then the machine of each participant to this agreement enters state T and shuts off. Thus $\mu_k(q, i, j, p, A) = T$ for any player $k = i, j$, any state $q \in Q_k$, and any price p. Also, I shall simply refer to the members of the set Q_i as the states of the machine.

$M_i = (Q_i, q_i^1, T, \lambda_i, \mu_i\ )$ after any history $h = (e^1, ..., e^t)$ by $\mu_i(q_i^1, h)$. Thus $\mu_i(q_i^1, h)$ can be defined iteratively by

$$\mu_i(q_i^1, e^1, ..., e^\tau) = \mu_i(\mu_i(q_i^1, e^1, ..., e^{\tau-1}), e^\tau)) \text{ for any } 1 < \tau \leq t \tag{2.1}$$

Definition 2. *Machine* $M_i = (Q_i, q_i^1, T, \lambda_i, \mu_i\ )$ *implements strategy* $f_i \in F_i$ *if*

$$f_i(h, d) = \lambda_i(\mu_i(q_i^1, h), d) \text{ for all } h \in H^\infty \text{ and for all } d \in D_i$$

The complexity of a machine (or of a strategy) can be measured in many different ways. In the literature on repeated games played by automata the number of states of the machine is often used as a measure of complexity. Henceforth, I shall refer to this measure of complexity by state-complexity (or simply by s-complexity). This is because the set of states of the machine can be regarded as a partition of possible histories. (See footnote 7 below and Kalai and Stanford 1989)

Definition 3. *(State-complexity) A machine* $M_i = (Q_i, q_i^1, T, \lambda_i, \mu_i\ )$ *is more s-complex than another machine* $M_i' = (Q_i', q_i^{1'}, T, \lambda_i', \mu_i'\ )$*, denoted by* $M_i \succ^s M_i'$*, if* $|Q_i| > |Q_i'|$*, where, for any set* W*,* $|W|$ *refers to the cardinality of the set* W*.*

Now it can be shown for any strategy $f_i \in F_i$ the size (number of states) of the smallest (in terms of the number of states) machine that implements f_i is equal to the number of strategies $|\{\langle f_i \mid h\rangle \, |h \in H^\infty\}|$ induced by f_i after different histories. (See Kalai and Stanford (1989).) Thus, I could also define s-complexity in terms of the underlying strategies in the game as follows.

Definition 4. *S-complexity of any strategy* f_i *is defined to be the number of induced strategies* $\{\langle f_i \mid h\rangle \, |h \in H^\infty\}$ *after different histories.*[8] *Thus a strategy* f_i *is more s-complex than* f_i'*, denoted by* $f_i \succ^s f_i'$*, if*

$$|\{\langle f_i \mid h\rangle \, |h \in H^\infty\}| > |\{\langle f_i' \mid h\rangle \, |h \in H^\infty\}|$$

Next I define finite machines and finite strategies.

Definition 5. *An automaton* M_i *is finite if it has a finite number of states. A strategy* f_i *is finite if it can be implemented by a finite machine or equivalently if it has a finite number of induced strategies* $\{\langle f_i \mid h\rangle \, |h \in H^\infty\}$*. A profile of machines (strategies) is finite if each of its components is finite.*

[8] Any f_i defines a partition (call it s-partition) on H^∞ given by an equivalence relation

$$h \sim^{f_i} h' \text{ if and only if } f_i \mid \langle h\rangle = f_i \mid \langle h'\rangle$$

S-complexity simply reflects the size of these partitions.

Counting the number of states does not fully measure the complexity of the machine (strategy) during a period, specifically the complexity of different choices following the same partial history. This is because the states of a machine do not change during a period of the game. More formally, s-complexity is a measure of the complexity of the domain of $\lambda_i(.,d)$ for each d (the cardinality of the domain of $\widetilde{\lambda}_i(.)$) but it does not capture the complexity of the range of the mapping $\lambda_i(.,d)$. The following examples illustrate the point further.

Example 1. *There are two machines M_i and M_i'. Both machines have two states q_i^1 and q_i^2. Both are in state q_i^1 in the odd periods and in state q_i^2 in the even periods (thus they have the same transition functions). Also as a proposer, in state q_i^l ($l = 1, 2$), both machines offer price p_i^l to any player j. As a responder, M_i' always rejects all offers. Machine M_i, on the other hand, responds differently to the same proposal by any player j (by conditioning on the two states q_i^1 and q_i^2). In particular, for any offer p by j, M_i rejects p in the odd periods and accepts p in the even periods.*

Example 2. *There are two machines M_i and M_i'. Both machines have two states q_i^A and q_i^R. Both are in state q_i^A in the odd periods and in state q_i^R in the even periods (thus they have the same transition functions). Also as a responder, in state q_i^A both machines accept any price offer and in state q_i^R both machines reject any price offer. As a proposer, M_i' always offers a price p. Machine M_i, on the other hand, makes different proposal to any player j (by conditioning on the two states q_i^A and q_i^R). In particular, M_i proposes p in the odd periods and p' in the even periods.*

According to s-complexity M_i and M_i' are of equal complexity in both examples, despite the fact that in the first example the strategy that machine M_i implements has the additional complexity of different responses to the same offer and in the second example the strategy that machine M_i implements has the additional complexity of making different proposals in different periods. This is not a desirable property.

A plausible (and minimal) way of capturing the complexity of strategy *during a period* - complexity of different behaviour after the same partial history - is to assume that the complexity criterion satisfies the following two conditions.

(i) If two machines (and therefore two strategies) M_i and $M_i^{'}$ are otherwise identical except that as a responder to some price offer p by some player j, M_i' always responds the same way (always accepts or always rejects) whereas M_i sometimes accepts and sometimes rejects the offer p by j, then M_i should be considered as being more complex than M_i'.

(ii) If M_i makes at least two different proposals p and p' to some player j depending on the history of actions before the current period and if M_i' is otherwise identical to M_i except that as a proposer to player j it drops the offer p' in favour of p (after all histories at which M_i proposes p' to j M_i' proposes p) , then M_i should be considered as being more complex than M_i'.

I call such notion of complexity response-complexity (r-complexity). A similar definition can be found in Chatterjee and Sabourian (1999, 2000). The formal definition of r-complexity consists of a partial order (the weakest) that captures (i) and (ii) above. Such a partial order can be defined either on the set of machines or on the set of strategies.

Definition 6. *(Response complexity) A machine* $M_i = \{Q_i, q_i^1, T, \lambda_i, \mu_i\}$ *is more r-complex than another machine* $M_i' = \left\{Q_i', q_i^{1\prime}, T, \lambda_i', \mu_i'\right\}$, *denoted by* $M_i \succ^r M_i'$, *if the machines* M_i *and* M_i' *are otherwise identical except that in response to some partial history* $d' \in D_i$ *machine* M_i' *(strategy* f_i'*) is conditioning less on history than machine* M_i. *Formally,* $M_i \succ^r M_i'$ *if* $Q_i = Q_i', q_i^1 = q_i^{1\prime}$, $\mu_i = \mu_i'$ *and there exists a partial history* $d' \in D_i$ *and a set of states* $\overline{Q}_i \subset Q_i (= Q_i')$ *such that*

$$\left.\begin{array}{ll} \lambda_i(q_i, d) = \lambda_i'(q_i, d) & \text{if } d \neq d' \text{ or if } q_i \notin \overline{Q_i} \\ \lambda_i'(q_i, d') = \lambda_i'(q_i', d') & \forall q_i, q_i' \in \overline{Q}_i, \\ \lambda_i(q_i, d') \neq \lambda_i(q_i', d') & \text{for some } q_i, q_i' \in \overline{Q}_i \\ \lambda_i(q_i, d') \neq \lambda_i(q_i', d') & \forall q_i \in Q_i/\overline{Q}_i \text{ and } \forall q_i' \in \overline{Q}_i \end{array}\right\} \quad (2.2)$$

The first three conditions in (2.2) capture precisely the idea that M_i and M_i' are everywhere identical except in response to some d', M_i' always takes the same action in all states $q_i' \in \overline{Q}_i$ whereas M_i does not. The fourth condition is imposed so that the partial order $\succ^r$ is asymmetric. The fourth condition in (2.2) guarantees that we cannot have both $M_i \succ^r M_i'$ and $M_i' \succ^r M_i$.

Since states of a machine encapsulate past history, I could also define r-complexity directly in terms of the underlying strategies in the game.

Definition 7. *A strategy* f_i *is more r-complex than* f_i', *denoted by* $f_i \succ^r f_i'$, *if there exists a partial history* $d' \in D_i$ *and a set of histories* $\overline{H} \subset H^\infty$ *such that*

$$\left.\begin{array}{ll} f_i(h, d) = f_i'(h, d) & \text{if } d \neq d' \text{ or if } h \notin \overline{H} \\ f_i'(h, d') = f_i'(h, d') & \forall h, h' \in \overline{H}, \\ f_i(h, d') \neq f_i(h', d') & \text{for some } h, h' \in \overline{H} \\ f_i(h, d') \neq f_i(h', d') & \forall h \in H^\infty/\overline{H} \text{ and } \forall h' \in \overline{H} \end{array}\right\} \quad (2.3)$$

The four conditions defined in (2.3) are anaologous to those in (2.2), defined in terms of strategies.

Remark 2. *Clearly, conditions (2.2) and (2.3) above imply respectively that*

$$\left.\begin{array}{ll} \lambda_i(q_i, d) = \lambda_i'(q_i, d) & \forall q_i \text{ and } \forall d \neq d' \\ \lambda_i'(Q_i, d') \subset \lambda_i(Q_i', d') & \end{array}\right\} \quad (2.4)$$

$$\left.\begin{array}{ll} f_i(h, d) = f_i'(h, d) & \forall h \text{ and } \forall d \neq d' \\ f_i'(H^\infty, d') \subset f_i(H^\infty, d') & \end{array}\right\} \quad (2.5)$$

where $\lambda_i(Q_i, d')$, $\lambda'_i(Q'_i, d')$,$f_i(H^\infty, d')$ and $f'_i(H^\infty, d')$ refer to the ranges of the functions $\lambda_i(., d')$, $\lambda'_i(., d')$, $f_i(., d')$ and $f'_i(., d')$ respectively. The results of this paper on equilibrium selection remain valid if in the definition of r-complexity we use the stronger conditions (2.4) and (2.5) instead of conditions (2.2) and (2.3), respectively.

The r-complexity definition is a very weak *local (partial)* concept - local in the sense that M_i and M'_i in Definition 6 (f_i and f'_i in Definition 7) are everywhere identical except in response to some d', machine M'_i (strategy f'_i) always takes the same action in all states $q'_i \in \overline{Q}_i$ (after any history $h' \in \overline{H}$) whereas machine M_i (strategy f_i) does not. In other words, $M_i \succ^r M'_i$ ($f_i \succ^r f'_i$) if the behaviour of machine M_i (strategy f_i) and machine M'_i (strategy f'_i) are everywhere identical except that, given some d', the response of M'_i to d' is simpler than that of M_i

Notice that, given the specification of automata adopted here and given that we are dealing with a repeated *extensive form* game, s-complexity and r-complexity measure the complexity of different aspects of behaviour - the number of induced strategies at the beginning of each period versus the complexity of behaviour within a period.

For the results in this section with random matching, I only need to introduce this minimal notion of r-complexity into the standard game-theoretic set-up. In section 4 with voluntary matching, I use both r-complexity and s-complexity.

2.2.2. Equilibrium with complexity costs

I could define equilibrium with complexity costs either in terms of the underlying strategies and use the complexity criteria defined over the strategy sets, as in Definitions 4 and 7 or in terms of machines and use the complexity criteria defined on the set of machines, as in Definitions 3 and 6. For ease of exposition, henceforth, I will refer to complexity costs defined over the set of strategies rather than over the set of machines, and define the equilibrium concepts also in terms of the former.

Note that any strategy profile $f = \{f_i\}_{i \in \mathcal{B} \cup \mathcal{S}}$ induces a random outcome path because of Nature's moves (the random matching and random choice of proposers). I shall denote the expected payoff to each player i if strategy profile $f = \{f_i\}_{i \in \mathcal{B} \cup \mathcal{S}}$ is chosen by $\pi_i(f)$. Since I only allow for pure strategies the expectation is with respect to the moves of Nature.[9]

Definition 8. *For any $\epsilon \geq 0$, a strategy f_i is said to be a ϵ-best response to f_{-i} if*

$$\pi_i(f_i,\, f_{-i}) + \epsilon \geq \pi_i(f'_i, f_{-i}) \text{ for all } f'_i \in F_i$$

A strategy is a best response if it is a 0-best response.

[9]Formally, one needs to define an underlying probability space and expectation is taken with respect to the appropriate probability measure.

Now, I would like define the *Nash equilibrium* for the game in which each player's preference depends both on the player's actual payoff and on the complexity of the strategy adopted. Clearly, if the complexity costs are a great deal more important than the actual payoffs then it is easy to discard highly complex (non-stationary) strategies as candidate equilibrium. To make the framework most amenable to indeterminacy type results, as well as to keep the equilibrium concepts as close as possible to the standard ones, I will consider equilibrium concepts in which complexity has arbitrarily small weight. One way of doing this is the following.

Definition 9. *A strategy profile $f = f = \{f_i\}_{i \in \mathcal{B} \cup \mathcal{S}}$ constitutes a Nash equilibrium with lexicographic $l-$complexity cost (for $l = r, s$) if for each player i the following two conditions hold*

$$f_i \text{ is a best response to } f_{-i} \tag{2.6}$$

$$\nexists\ f_i' \in F_i \text{ such that } f_i' \text{ is a best response to } f_{-i} \text{ and } f_i \succ^l f_i'. \tag{2.7}$$

Complexity costs are treated lexicographically in the above definition of equilibrium. I could also have introduced complexity as a small fixed cost of choosing a more complex strategy and defined a Nash Equilibrium with a fixed complexity cost as follows.

Definition 10. *A strategy profile $f = \{f_i\}_{i \in \mathcal{B} \cup \mathcal{S}}$ constitutes a Nash equilibrium with a (small) fixed $l-$complexity cost ($l = r, s$) $c \geq 0$, denoted by NECl(c), if for each player i the following two conditions hold*

$$f_i \text{ is a best response to } f_{-i} \tag{2.8}$$

$$\nexists\ f_i' \in F_i \text{ such that } f_i' \text{ is a c-best response to } f_{\sim i} \text{ and } f_i \succ^l f_i'. \tag{2.9}$$

Remark 3. *Notice that a NECl(c) profile refers to a Nash equilibrium with lexicograhic l-complexity cost (Definition 10 is identical to Definition 9) if and only if $c = 0$ and it refers to a Nash equilibrium with positive fixed costs if and only if $c > 0$. Clearly, positive fixed complexity costs induce at least as much economy as the lexicographic criterion and thus any NECl(c) strategy profile is also a NECl(0) profile.*

NEC$l(c)$ strategy (machine) profiles are not necessarily 'credible' for the usual reasons. To ensure credibility, I could, as in Chatterjee and Sabourian (1999, 2000), introduce noise into the system and consider the limit of extensive form trembling hand equilibrium (Nash equilibrium with independent trembles at each information set) with complexity costs as the trembles become small. The trembles will ensure that strategies are optimal (allowing for complexity) after all histories that occur with a positive probability. However, any precise characterisation of such an equilibrium may depend on *the order* in which complexity costs and trembles enter the limiting arguments (more on this later).

A more direct, and simpler, approach of introducing credibility would be to consider NEC$l(c)$ strategy profiles that are perfect Bayesian equilibria (subgame perfect equilibrium for the case of a single seller) of the game without complexity costs.

Definition 11. *A profile f constitutes a perfect Bayesian equilibrium strategy profile with a (small) fixed l-complexity cost $c \geq 0$, denoted by PECl(c), if f is both a NECl(c) strategy profile and a perfect Bayesian equilibrium of the underlying game.*

Thus, it follows from the above definition that a profile f is a PEC$l(c)$ if and only if it is a perfect Bayesian equilibrium of the underlying game and it satisfies condition (2.9).

The selection result in this section holds for the case of PEC$r(0)$ with lexicographic r-complexity costs. Trivially, since any PEC$r(c)$ strategy profile is also a PEC$r(0)$, the selection result of this section applies equally to the case of any fixed r-complexity cost $c \geq 0$. In the next section with voluntary matching, the selection result is established only for equilibria where, in addition to r-complexity, s-complexity enters players prefences as a fixed positive cost.

2.3. The selection result

Before stating the results of this section, I shall formally define stationary behaviour.

Definition 12. *A strategy f_i is stationary if and only if $f_i(h, d) = f_i(h', d)$ $\forall$ h,$h' \in H^\infty$ and $\forall$ $d \in D_i$. Also, a profile $f = \{f_i\}_{i \in \mathcal{B} \cup \mathcal{S}}$ is stationary if f_i is stationary $\forall i$.*

Thus the behaviour of such strategies at any time may depend on the information (partial history) in the current period but not on the previous history of the game before the current period.

Remark 4. *A strategy is stationary if and only if it has a minimal l-complexity ($l =$ r or s). Formally:*

$$f_i \text{ is stationary if and only if } \nexists\ f'_i \text{ such that } f_i \succ^l f'_i$$

Thus, a stationary Nash (perfect Bayesian) equilibrium strategy profile of the underlying game is a NECl(c) (PECl(c))of the game with complexity cost.

Clearly, for any $c \geq 0$, a stationary PEC$l(c)$ ($l = r, s$) strategy profile exists. Consider the following stationary profile of strategies: all players always offer 1, each seller accepts an offer if and only if the offer is 1 and buyers accept all offers. This profile induces the competitive outcome. Trivially, it also constitutes a perfect Bayesian equilibrium and is stationary (has minimal l-complexity). Therefore this profile is a PEC$l(c)$.

Notice, however, that the strategies used by RW in the proof of Theorem 2.1, to support *non-competitive* outcomes trivially cannot constitute a PEC$l(c)$ - or even a NECl(c) - for any $c \geq 0$. This is because these strategies are non-stationary. In particular, all those buyers who do not have any 'privileges' with respect to any of the goods (buyers who do not end up buying the goods on the equilibrium path constructed in the proof of Theorem 2.1) also follow non-stationary strategies. But such buyers receive zero payoff on the equilibrium path. Thus, these strategies could not be a NECl(c) because each such buyer could always obtain at least a zero payoff by following a less complex strategy than the non-stationary one specified in RW's construction (this is always possible because non-stationary strategies do not have minimal l-complexity).

The next Theorem is the main result of this section. It demonstrates that *all* credible equilibria of the game with r-complexity costs are stationary and induce the unique competitive outcome.

Theorem 2.3. *Consider any $c \geq 0$ and any PEC$r(c)$ strategy profile $f = \{f_i\}_{i \in \mathcal{B} \cup \mathcal{S}}$. If f is finite then $\pi_s(\langle f \mid h \rangle) = 1$ and $\pi_b(\langle f \mid h \rangle) = 0$ for all h, for all s and for all b, the unique induced price is the competitive price of 1 and f is stationary.*

Remark 5. *The assumption that each PEC$r(c)$ f is finite in the above Theorem is only needed for the case in which r-complexity costs enters the players' preferences lexicographically ($c = 0$). If fixed positive r-complexity cost ($c > 0$) is assumed then it can be shown that the above Theorem holds without making such an assumption (finiteness) on the set of PEC$r(c)$ strategy profiles. For ease of exposition, I will not demonstrate this more general result and refer the reader to footnote 10 and Sabourian (2001a).*

Remark 6. *If one uses s-complexity (instead of r-complexity) as the complexity criterion, then the selection result in Theorem 2.3 does not hold and RW's continuum result remains valid . In particular, it can be shown that for any price $1 \geq p > 0$, there exists a finite NECs(0) strategy profile such that each seller ends up selling his commodity to some buyer at p. To keep the paper no longer than it actually is, I shall omit this result - see Sabourian (2001a) for the proof of this result for the case of a market with a single seller.*

Next, I turn to the formal proof of Theorem 2.3. The next two Lemmas characterise some crucial properties of NEC$r(c)$ strategy profiles.[10]

Lemma 1. *For any $c \geq 0$, any NEC$r(c)$ profile of strategies $f = (\{f_i\}_{i \in \mathcal{B} \cup \mathcal{S}}$, any buyer $b \in \mathcal{B}$ and any seller $s \in \mathcal{S}$ the following holds:*

$$f_b(h, s, b, 1) = f_b(h', s, b, 1) \quad \text{for all } h \text{ and } h' \in H^\infty \tag{2.10}$$

$$f_s(h, b, s, 1) = f_s(h', b, s, 1) \quad \text{for all } h \text{ and } h' \in H^\infty \tag{2.11}$$

[10] These properties do not hold in general for Nash equilibrium profiles.

Proof. To show that condition (2.10), suppose otherwise. Then there exist a seller s, a buyer b, and histories h and h' such that

$$f_b(h, s, b, 1) = A \text{ and } f_b(h', s, b, 1) = R$$

Now consider another strategy f_b' for player b such that for all $(h'', d) \in (H^\infty \times D_b)$

$$\begin{array}{ll} f_b'(h'', d) = R & \text{if } d = (s, b, 1) \\ f_b'(h'', d) = f_b(h'', d) & \text{if } d \neq (s, b, 1) \end{array}$$

Clearly, the only difference between f_b' and f_b is that f_b' always rejects an offer of 1 by s and f_b' does not. Thus, f_b' induces at least the same payoff as f_b and moreover it is less r-complex than f_b according to Definition 7. But this contradicts f being a NEC$r(c)$. Thus condition (2.10) holds.

Using a similar reasoning as above, I now show that condition (2.11) holds. Suppose not; then for some s, for some b, for some h and for some h' the following holds

$$f_s(h, b, s, 1) = A \text{ and } f_s(h', b, s, 1) = R$$

Now consider another strategy f_s' for player s such that for all $(h'', d) \in (H^\infty \times D_s)$

$$\begin{array}{ll} f_s'(h'', d) = A & \text{if } d = (b, s, 1) \\ f_s'(h'', d) = f_s(h'', d) & \text{if } d \neq (b, s, 1) \end{array}$$

Clearly, the only difference between f_s' and f_s is that f_s' always accepts an offer of 1 by b and f_s does not. Thus, f_s' induces at least the same payoff as f_s and moreover it is less r-complex than f_s according to Definition 7. But this contradicts f being a NEC$r(c)$. Thus condition (2.11) holds. ■

Lemma 2. *For any $c \geq 0$, any NECr(c) profile of strategies $f = (\{f_i\}_{i \in \mathcal{B} \cup \mathcal{S}}$, any buyer $b \in \mathcal{B}$ and any seller $s \in \mathcal{S}$ the following holds:*

$$\text{if } f_s(h, b, s, 1) = A \text{ for some } h \text{ then } \forall\ h \text{ either } f_b(h, b, s) = 1 \text{ or } f_b(h, b, s) \neq 1 \quad (2.12)$$

Proof. Suppose not; then there exist a buyer b and a seller s such that $f_s(h, b, s, 1) = A$ for some h, $f_b(h', s, b) = p \neq 1$ for some h' and $f_s(h'', s, b) = 1$ for some h''. Since $f_s(h, b, s, p) = A$ for some h, by Lemma 1, we have

$$f_s(h, s, b, p) = A \text{ for all } h \tag{2.13}$$

Now consider a strategy f_b' such that

$$\begin{array}{ll} f_b'(h, d) = p & \text{for all } (h, d) \text{ such that } d = (b, s) \text{ and } f_b(h, d) = 1 \\ f_b'(h, d) = f_b(h, d) & \text{otherwise} \end{array}$$

Clearly, the only difference between f_b and f_b' is that in some instance f_b proposes an offer of 1 and f_b' does not. Thus,

- f'_b induces at least the same payoff as f_b (this is because by (2.13), s always accepts a price offer of 1 by b)

- f'_b is less r-complex than f_b according to Definition 7.

But this contradicts f being a NEC$r(c)$. ■

Since any PEC$r(c)$ is a NEC$r(c)$ it follows that any PEC$r(c)$ also satisfies the above properties of NEC$r(c)$ given in Lemmas 1 and 2. In fact, these properties are so critical for establishing the selection result of this section[11] that I will refer to a perfect Bayesian equilibrium strategy profile that satisfies condition (2.10) of Lemma 1 and condition (2.12) of Lemma 2 by Pr*.

Definition 13. *A strategy profile f is a Pr* if f is both a perfect Bayesian equilibrium and it satisfies the following conditions:*

$$f_b(h, s, b, 1) = f_b(h', s, b, 1) \quad \text{for all } h \text{ and } h' \in H^\infty \tag{2.17}$$

$$\left.\begin{array}{l} \text{if } f_s(h, b, s, 1) = A \text{ for some } h \text{ then for all } h \\ \text{either } f_b(h, b, s) = 1 \text{ or } f_b(h, b, s) \neq 1 \end{array}\right\} \tag{2.18}$$

Clearly, by Lemmas 1 and 2, any PEC$r(c)$ strategy profile is also a Pr*.

Next I shall state a selection result in which the sellers receive the entire surplus (the competitive outcome) after every history with the weaker solution concept of Pr*. This result is then used to establish Theorem 2.3.

Theorem 2.4. *Consider any Pr* strategy profile $f = \{f_i\}_{i \in \mathcal{B} \cup \mathcal{S}}$. If f is finite, then $\pi_s(\langle f \mid h \rangle) = 1$ for all h and for all s.*

To provide some basic intuition for this result, notice that the strategies used by RW in the proof of Theorem 2.1, to support *non-competitive* outcomes trivially cannot constitute a Pr* either. This is because these strategies require any buyer b to accept a price offer of 1 (or to propose a price of 1) after histories at which he has the 'privilege' to buy the good at a price offer of 1 and to reject a price offer of 1 (or to propose a

[11] In Remark 5, I mentioned that Theorem 2.3 remains valid for the case of positive fixed cost $c > 0$ even without assuming that the PECr(c) strategy profile is finite. This is because for the case of $c > 0$, the critical conditions (2.10), (2.11) and (2.12) described in the previous two Lemmas, can be, respectively, strengthened to show that for any $p > 1 - c$

$$f_b(h, s, b, p) = f_b(h', s, b, p) \quad \text{for all } h \text{ and } h' \in H^\infty \tag{2.14}$$

$$f_s(h, b, s, p) = f_s(h', b, s, p) \quad \text{for all } h \text{ and } h' \in H^\infty \tag{2.15}$$

$$\text{if } f_s(h, b, s, p) = A \text{ for some } h \text{ then either } f_b(h, b, s) = p \ \forall\ h \text{ or } f_b(h, b, s) \neq p \ \forall\ h \tag{2.16}$$

With these stronger condition, one does not need to make the assumption of finiteness.(See Sabourian (2001) for details.)

price other than 1) otherwise. But this contradicts condition (2.17) (condition (2.18)) in the definition of Pr*.

The actual proof of Theorem 2.4 is based on generalising this intuition. In particular, it turns out that any non-competitive Pr* profile cannot result in an agreement at a price of 1 after any history. This feature of any Pr* is sufficient to obtain the selection result.

The proof of the Theorem 2.4 for the case of a market with one seller can be found in Appendix A1. The proof for the general case with an arbitrary number of sellers is by induction on the number of sellers and it can be found in Appendix A2.[12] In the next subsection, I will provide a sketch of the proof of Theorem 2.4.

Next, I will explain how the main result of this section - Theorem 2.3 - follows from Theorem 2.4.

Proof of Theorem 2.3: Since any PEC$r(c)$ strategy profile is a Pr*, it follows from Theorem 2.4 that for any finite PECr(c) profile f, $\pi_s(\langle f \mid h \rangle) = 1$ for all h and for all s. This implies that profile f is such that after any h, with probability 1 the seller s reaches an agreement at $p = 1$ with some buyer and $\pi_b(\langle f \mid h \rangle) = 0$ for all h and for all b. Therefore, f_b is stationary for all b; otherwise, b could economize on r-complexity by playing a simpler strategy and obtain at least a zero payoff. (Non-stationary strategies are not of minimal r-complexity).

Last, I need to show that f_s is stationary for all s. Consider any s and any b. Since $\pi_s(\langle f \mid h \rangle) = 1$ and $\pi_b(\langle f \mid h \rangle) = 0$ for all h, it follows that

$$f_s(h, s, b) = 1 \text{ for any } h \tag{2.19}$$

$$f_s(h, b, s, p) = R \text{ for any } (h, b, s) \text{and } \forall\ p < 1 \tag{2.20}$$

Condition (**??**) holds because if s offers a price less than 1 to b after any h it will be accepted (rejection would result in a zero continuation payoff for b); but then s would receive a continuation payoff that is less than 1; a contradiction. Condition (**??**) follows because if b offers a price less than 1 after any h and if s accepts then s would receive a continuation payoff that is less than 1; a contradiction.

Also, by Lemma 1, we have that $f_s(h, b, s, 1) = f_s(h, b, s, 1)$ for all h and h'. Thus f_s is stationary. ■

2.4. Sketch of the steps of Theorem 2.4

Market with a single seller s

The proof of this result in this case basically consists of establishing the following three steps.

[12] It turns out that the induction argument on the number of sellers, applied to the concept of Pr* in Appendix A2, cannot be applied directly to the concept of PBEC$r(c)$. It is for this reason that the selection result is first established for the weaker concept of Pr*. (See the sketch of this induction argument in the next subsection.)

Step 1 (see Lemma **??**, **??**): If any Pr* strategy profile f is such that $\pi_s(\langle f \mid h\rangle) < 1$ for some history h (the outcome is non-competitive after h), then there does not exist a history after which the seller reaches an agreement at a price of 1 with some buyer.

This step follows from from conditions (2.17) and (2.18).

Step 2 (see Lemma **??**): If any Pr* strategy profile f is such that $\pi_s(\langle f \mid h\rangle) < 1$ for some history h (the outcome is non-competitive after h), then for any buyer b and for any history h the continuation payoff to b after the ordered triple (h, b, s) is positive.

The intuition for this step is as follows. Since there is no agreement at a price of 1 (Step 1), it follows that the continuation payoff to the seller is always less than 1 (Lemma **??**). This, together with the finiteness of the Pr* strategy profile, imply that, after any history, if a buyer has the opportunity to make an offer to the seller he can obtain a positive payoff by offering a price that is both less than 1 and more than the continuation payoff of the seller.

Step 3 (see Lemma **??**): This involves showing that for any finite perfect Bayesian equilibrium, there exists a buyer b and a history h such that b's continuation payoff after (h, b, s) is zero.

This step follows from considering histories at which the continuation payoff of the seller is minimised. At such histories, competition between buyers ensures that the continuation payoff of at least one buyer is zero.

Steps 2 and 3 contradict each other unless the Pr* strategy profile results in a continuation payoff of 1 for the seller after any history. This establishes the result.

Market with an arbitrary number of sellers

The proof for the case of more than one seller is by induction on the number of sellers S in the market. Given that the result for the case of $S = 1$, to complete the proof of Theorem 2.4 with an arbitrary number of sellers, I need to show that if any Pr* profile for a market with $S' < S$ sellers results in continuation payoffs of 1, after all histories, for each of the S' sellers, then any Pr* profile for a market with S sellers also results in continuation payoffs of 1, after all histories, for each of the S sellers. This is done by establishing the following four steps.

Step1 (Lemma **??**):If a strategy profile is a Pr* then after any history it is also a Pr* in the continuation game.

This step follows immediately from the definition of a Pr* profile. Note also that this step is not necessarily valid if we replace the concept Pr* by PECr (see footnote 11).

Step 2 (Lemmas **??** and **??**): If Theorem 2.4 holds when the number of sellers is less than S then for any Pr* profile f in a market with exactly S sellers the following holds: if $\pi_s(\langle f \mid h\rangle) < 1$ for some history h (the outcome is non-competitive after h), then f does not result in an agreement at a price of 1 in any match between a seller and a buyer after any history.

The step is proved by appealing to Step 1 and by repeating some of the arguments used for the case of $S = 1$.

Step 3 (Lemma **??**): If Theorem 2.4 holds when the number of sellers is less than S then for any Pr* profile f in a market with exactly S sellers the following holds

$$\text{if } \pi_s(\langle f \mid h \rangle) < 1 \text{ for some } s \text{ and for some } h \text{ then } \pi_s(\langle f \mid h \rangle) < 1 \text{ for all } h \qquad (2.21)$$

This step follows immediately from Step 2.

Step 4 (Lemma **??**): If Theorem 2.4 holds when the number of sellers is less than S, then for any Pr* profile with S sellers there exits a history such that the continuation payoff of some seller s is one.

Intuitively, this is because if at some point a pair of buyers and sellers, excluding s, leave the market then by Step 1 the continuation strategy is also a Pr* for the smaller market. But then, by assumption, the remaining sellers will receive a continuation payoff of 1.

But Step 4 contradicts Step 3 unless the Pr* strategy profile results in a continuation payoff of 1, after any history and for each of the S sellers. Thus if Theorem 2.4 holds when the number of the sellers is less than S it also holds when there are exactly S sellers.

3. Voluntary matching with endogenous choice of partners, discounting and complexity

The no discounting assumption is important in establishing the existence a of continuum of perfect Bayesian equilibrium prices in Theorem 2.1. This Theorem works because after any history there are special 'relationships' between buyers and sellers - after every history, for each unit of the good of a seller, a buyer has the 'privilege' to buy it at a particular price. Each deviation from the equilibrium strategies is deterred by the creation of a new relationship. With random matching, with probability one, the two sides of the new relationship will meet in a finite time. With no discounting, the length of the period it will take for the two sides to meet is unimportant. However, with discounting the cost of maintaining these relationships may be high if it takes a long time for the designated buyers and sellers to meet each other. Thus, with discounting it may not be optimal for players to play the appropriate punishments needed to support the equilibria in Theorem 2.1. Therefore, discounting eliminates a large number of equilibria. For the one seller model, RW have the following result.

Theorem 3.1. *(See RW) Suppose that $S = 1$ and $\delta \in (0, 1)$. Then there exists a unique subgame perfect equilibrium in which trade takes place at $t = 1$. Moreover, as $B \to \infty$ or as $\delta \to 1$ the unique equilibrium converges to the competitive price of 1.*

There is no known equivalent result to the above Theorem for the case of more than one seller. Nevertheless, Theorem **??** (in particular the part on convergence of the equilibrium prices to the competitive one as $\delta \to 1$) seems to throw some doubt on the multiplicity result in Theorem 2.1. However, RW argue that discounting imposes a cost on having a relationship because the formation and the termination of matches are random. But staying with one's current partner should not be costly. Thus, they consider a voluntary matching model with an endogenous choice of partner that is otherwise identical to the random matching model except that in each period the sellers (the short side of the market) choose the buyers with whom they wish to bargain. For this matching model, they demonstrate the existence of a large number of (non-competitive) equilibria even for the case in which $\delta < 1$.

Theorem 3.2. *(See RW) If $S = 1$ and the seller can choose in each period the buyer with whom he wishes to bargain then for each buyer b and any price $\frac{1}{1+\delta} \leq p \leq 1$ there exists a subgame perfect equilibrium in which b receives the good at the price equal to either p or $\frac{\delta p}{2-\delta}$, according to whether the seller or the buyer b is the proposer in their first encounter.*

Thus the indeterminacy and non-competitive outcomes are present in the model with discounting as well, irrespective of the number of buyers.[13] But the strategies needed to implement the above equilibria for any $p < 1$ turn out to be unnecessarily complex. To establish Theorem **??**, for any price $\frac{1}{1+\delta} \leq p \leq 1$, RW construct the following subgame perfect equilibrium strategy profile. The seller s always offers p and agrees to accept $\frac{\delta p}{2-\delta}$ or more. A buyer always offers $\frac{\delta p}{2-\delta}$ and accepts p or less. In the first period s picks buyer b and in the case of disagreement s continues with the same buyer only if the same buyer did not deviate. If a buyer deviates at any period from the above strategy the seller discontinues the bargaining with him and picks a new buyer.

The above strategy, clearly, results in an agreement between the seller and buyer b at price p in the first period. But then why should the seller choose a strategy that involves selecting different partners depending on the previous history of moves? Consider a simpler strategy for the seller that *always* chooses buyer b, always offers p to b and agrees to an offer if and only if the offer is $\frac{\delta p}{2-\delta}$. Clearly, if all buyers follow the above strategies, this simple strategy for the seller results in the same payoff (after all histories) as before but with less complexity (in terms of choosing the same partner irrespective of past history).

In this section, I extend the result of the previous section by showing that with complexity costs the only perfect Bayesian equilibrium of the above game with endogenous choice of partners is also the competitive outcome. However, as was mentioned before, I obtain this result

- by using s-complexity (for the sellers) in addition to r-complexity
- by assuming a positive fixed cost of s-complexity rather than s-complexity costs introduced lexicographically.

[13]RW state that Theorem **??** can be extended to the case in which there are more than one seller. However, in the case of more than one seller, one has to specify further the precise description of the matching process - see below.

Before, describing the result, notice that although with one seller the voluntary matching is well defined (at the beginning of each period the seller chooses the buyer with whom he wishes to bargain with for that period), with more than one seller there are several different ways of describing a voluntary matching model. For example, one could have at every period each seller simultaneously choosing the buyer with whom he wishes to bargain and if more than one seller choose the same buyer in a given period then the buyer is matched to one of the sellers according to some exogenous (possibly random) mechanism. Or alternatively, one could have sellers choosing their partners sequentially. In this case at every period there is only one match, with a single seller selected according to some exogenous (possibly random) mechanism amongst the remaining sellers. The chosen seller then chooses the buyer with whom he wishes to bargain at that period. In this latter approach, a unit of time could represent the time it takes for a pair of a buyer and a seller to be matched with the actual round of baragining taking place almost instantaneously. At this stage I will not describe the precise formualtion of the voluntary matching model for the case of more than one seller and instead return to this issue below.

The notation in this section is the same as in the previous section. In particular, h and d refer respectively to a finite history of outcomes and a partial history of actions within a period. The definition of strategy in this section is the same as that in the case of random matching case except that here, with an endogenous choice of partners, a seller has to choose a partner at the beginning of each period. Formally, I represent the beginning of each period at which the seller s has to choose a buyer by φ. Also, denote the buyers left in the market after any history h by $\mathcal{B}_h$. Now, I define a strategy for seller s by a function

$$f_s : H^\infty \times (D_s \cup \varphi) \to C \cup \mathcal{B}$$

such that for any $h \in H^\infty$ and for any $d \in D_s$, $f_s(h,d) \in C_s(d)$ and $f_s(h,\varphi) \in \mathcal{B}_h$ if h is such that s has not left the market and $f_s(h,d) = \varnothing$ and $f_s(h,\varphi) = \varnothing$ otherwise.[14,15]

The two definitions of complexity (both r-complexity and s-complexity) in this section with voluntary matching are also identical to those in the previous section with random matching (Definitions 4 and 7). In the latter, however, r-complexity (measuring the complexity of responses during a period) was sufficient to select uniquely the competitive outcome. In this section, we have an additional element of complexity of behaviour - the complexity of each seller's decision at the beginning of each period at which he has to choose a buyer. I need to strengthen the definition of complexity to

[14]Similarly, one can define the automaton representing a seller's strategy $M_s = \{Q_s, q_s^1, T, \lambda_s, \mu_s\}$ in the same way as before except that the output function of any seller is now defined by

$$\lambda_s : Q_s \times (D_s \cup \varphi) \to C \cup \mathcal{B}$$

where $\lambda_s(q_i, d) \in C_s(d)$ and $\lambda_s(q_s, \varphi) \in \mathcal{B}$ for any $q_s \in Q_s$ and for any $d \in D_s$.

[15]Clearly, the strategy (machine) of a buyer is defined in the same way as in the previous section.

capture the complexity of conditioning the choice of the buyer at any period on the history of the game prior to that period.

One way of capturing the complexity, for each seller, of choosing a buyer at the beginning of a period is to strengthen r-complexity definition to allow for responses to φ. Thus, in addition to r-complexity, I could require the complexity criterion to rank strategies for seller s according to the following criterion.

Definition 14. *Strategy f_s is more choice-complex than strategy f'_s, denoted by $f_s \succ^c f'_s$, if f_s and f'_s are otherwise identical except that the choice of a buyer for f'_s at the beginning of a period is independent of past history whereas $f'_s s$. choice depends on the previous history of outcomes Formally, $f_s \succ f'_s$ if there exists a set of histories $\overline{H} \subset H$ such that*

$$\left.\begin{array}{ll} f_s(h,\varphi) = f'_s(h,\varphi) & \text{if } h \notin \overline{H} \\ f'_s(h,\varphi) = f'_s(h',\varphi) & \forall h, h' \in \overline{H}, \\ f_s(h,\varphi) \neq f_s(h',\varphi) & \text{for some } h, h' \in \overline{H} \\ f_s(h,\varphi) \neq f_s(h',\varphi) & \forall h \in H/\overline{H} \text{ and } \forall h' \in \overline{H} \end{array}\right\} \tag{3.1}$$

However, it turns out that r-complexity together with c-complexity are not sufficient to select the competitive outcome. (I have a counter-example demonstrating the existence of a non-competitive outcome with this stronger complexity criterion[16] for the no discounting case.)

Another candidate for measuring the complexity of each seller's choice of partners at the beginning of a period is s-complexity. Clearly, a seller's strategy (machine) needs to have as many induced strategies (states) as the number of possible partners he chooses in the game. Putting it differently, if two strategies (machines) for seller s are otherwise identical except that the first chooses fewer partners than the second, then the second strategy (machine) must have more induced strategies (states) than the first.

I shall demonstrate below that s-complexity (a measure of the number of induced rules at the beginning of each period) for the sellers together with r-complexity (a measure of the complexity within a period) are sufficient to give us the selection result in the voluntary matching model with endogenous choice of partners. As mentioned before, this will be done for the case in which s-complexity enters the sellers' preferences as a positive fixed cost. The formal definition of equilibrium both NECr(c) and PECr(c) with endogenous choice of partners is the same as those in the previous section (see Definitions 9, 10 and 11).

Definition 15. *A strategy profile $f = \{f_i\}_{i \in \mathcal{B} \cup \mathcal{S}}$ constitutes a Nash equilibrium with $r-$complexity cost $c \geq 0$ and s-complexity cost $c' \geq 0$, denoted by NECr$(c)s(c')$, if the following three conditions hold:*

$$\text{for each player } i,\ f_i \text{ is a best response to } f_{-i} \tag{3.2}$$

[16] This stronger definition of complexity (r-complexity together with c-complexity) is used in Chatterjee and Sabourian (2000).

for each player i, $\nexists\ f_i' \in F_i$ such that f_i' is a c-best response to $f_{\sim i}$ and $f_i \succ^r f_i'$
$$(3.3)$$
$$\nexists\ f_s' \in F_s \text{ such that } f_s' \text{ is a } c'\text{-best response to } f_{\sim s} \text{ and } f_s \succ^l f_s'. \tag{3.4}$$

Definition 16. *A strategy profile $f = \{f_i\}_{i \in \mathcal{B} \cup \mathcal{S}}$ constitutes a perfect Bayesian equilibrium with $r-$complexity cost $c \geq 0$ and s-complexity cost $c' \geq 0$, denoted by PECr$(c)s(c')$, if f is both a NECr$(c)s(c')$ and a perfect Bayesian equilibrium of the underlying game.*

Remark 7. *It is clear from the above that a NECr$(c)s(c')$ (PECr$(c)s(c')$) is simply a NECr(c) (PECr(c)) that satisfies condition (**??**). Also, notice that this condition (and thus both definitions of NECr$(c)s(c')$ and PECr$(c)s(c')$) refers only to the s-complexity of the sellers' strategies.*

Next, I state the result of this section for the set of PEC$r(c)s(c')$ profiles with $c \geq 0$ and $c' > 0$.[17] This selection result, however, is stated for the no discounting case since this appears to be most amenable to indeterminacy type results.

Theorem 3.3. *Suppose $S = 1$ and the seller can choose in each period the buyer with whom he wishes to bargain (voluntary matching) and $\delta = 1$. Then consider any $c \geq 0$, any $c' > 0$ and any PECr$(c)s(c')$ strategy profile $f = \{f_i\}_{i \in \mathcal{B} \cup \mathcal{S}}$ If f is finite then $\pi_s(\langle f \mid h \rangle) = 1$ and $\pi_b(\langle f \mid h \rangle) = 0$ for all s, b and h, the unique induced price is the competitive price of 1 and f is stationary.*

Remark 8. *As in the proof of Theorem 2.3, the assumption that f is finite in Theorem **??** is only needed for the case in which r-complexity enters the players preference lexicographically ($c = 0$). If positive fixed r-complexity cost ($c > 0$) is assumed then it can be shown that Theorem **??** holds without assuming that f is finite. For ease of exposition, I will not demonstrate this result and refer the reader to Sabourian (2001a).*

The proof of the Theorem **??** can be found in Appendix B. In the rest of this section, I will provide a brief sketch of the proof of Theorem **??** and then discuss an extention of Theorem **??** to the case in which there is more than one seller.

A brief sketch of the proof of Theorem ??: Consider any profile f. For any b, let
$$z(b) = \max_h \pi_b(\langle f \mid h \rangle).$$

First, it is shown that $z(b)$ is the same for all b (this is because the seller selects a buyer at each period). Denote $z(b)$ by z and consider the set of histories
$$H(b) \equiv \{h \in H^\infty \mid \pi_b(\langle f \mid h \rangle) = z\}.$$

[17] I do not have any equivalent result for the case in which $s-$complexity enters the sellers' preferences lexicographically ($c' = 0$).

Next, define the following property. A strategy profile f is said to satisfy property α if

$$\begin{aligned}&\text{there exists } b, b' \neq b \text{ and } h_b \in H(b) \text{ such that the probability}\\&\text{that the seller chooses } b' \text{ at some point after } h_b \text{ is postive.}\end{aligned} \tag{3.5}$$

Also, define a Pr* strategy profile as in the previous section and note that any PECr(c) is a Pr* profile.

The rest of the proof is divided into two steps:

Step 1: Consider any $c' > 0$. Then every PEC$r(c)s(c\prime)$ (in fact any perfect Bayesian equilibrium that satisfies condition (**??**)), satisfies property α.

Step 2: If f is a Pr* and it satisfies property α then $\pi_s(\langle f \mid h\rangle) = 1$ for all h.

Since any PEC$r(c)s(c\prime)$ strategy profile is a Pr*, the two steps together establish that for any $c' > 0$ and any PBECr(c)s(c) profile f we have $\pi_s(\langle f \mid h\rangle) = 1$ for all h. This is sufficient to complete the proof of the Theorem.

If f satisfies property α then after any h_b defined in condition (**??**) - there is a positive probability of choosing another buyer $b' \neq b$. Therefore, after h_b, the game is effectively similar to the random matching model and thus the proof of the statement in Step 2 is similar to that of Theorem 2.3 in Appendix A.1.

To demonstrate Step 1, I assume otherwise. Then, I show that there exists another strategy f'_s such that the following two conditions hold

$$\pi_s(\langle f \mid h\rangle) - \pi_s(\langle f'_s, f_{-s} \mid h\rangle) \leq \epsilon < c \quad \forall h \tag{3.6}$$

$$f_s \succ^s f'_s \tag{3.7}$$

But then f does not satisfy condition (**??**); a contradiction.

Strategy f'_s, mentioned in the previous paragraph, is constructed such that it is otherwise identical to the equilibrium strategy f_s except that the continuation strategies $\{\langle f_s \mid h_b\rangle \,|h_b \in \cup_b H(b)\,\}$ of f_s are replaced by a single continuation strategy f''_s that always chooses buyer b, always offers $1 - z - \epsilon$ for some $0 < \epsilon < c$ and always accepts an offer if and only if the offer is not less than $1 - z - \epsilon$. This construction together with the assumption that f does not satisfy property α ensures that condition (**??**) holds.

To demonstrate condition (**??**), note that the only difference between f_s and f'_s is that the continuation strategies $\{\langle f_s \mid h_b\rangle \,|h_b \in \cup_b H(b)\,\}$ of the former are replaced by a single continuation strategy f''_s. Since, by assumption, f does not satisfy property α, it follows that the continuation strategies $\{\langle f_s \mid h_b\rangle \,|h_b \in \cup_b H(b)\,\}$ of f_s are not singleton (this is because for all b and for all $h_b \in H(b)$ the probability that the seller s chooses another buyer $b' \neq b$ at some point after h_b is zero.) But this implies that f_s is more s-complex than f'_s.

Extension of Theorem ?? with more than one seller. When $S > 1$, as was mentioned before, there are several different ways to model voluntary matching with endogenous choice of partners. It can be shown that Theorem **??** remains valid with more than one seller if one adopts the following sequential specification of the voluntary matching.

At each period there is only one match per period. (With this specification of volunatary matching, a period effectively refers to the time it takes for a match to be consumated.) Thus at any time t, one seller s is chosen according to some exogeneous mechanism; and this seller chooses the buyer b with whom he wishes to bargain at t. The bargaining between s and b is as before with the choice of the proposer being random. The mechanism for selecting the seller at each period may even be random and may also be time-dependent (we could even allow for the selection of the seller at any time to depend on the history of outcomes). The only assumption I impose on any such selection mechanism is that for each seller s, the probability that s is selected after some finite time is one.

The proof of the extension of Theorem **??** to the case of more than one seller with the above voluntary matching specification involves proving the two steps described in the sketch of Theorem **??**. The proof of Step 1 with more than one seller is identical to the one seller model. The proof of Step 2 with more than one seller is by induction on the number of sellers in a similar fashion as in the proof of Theorem 2.4. In fact, as was mentioned above, property α is similar to the random matching model after any h_b defined by condition (**??**). Therefore, given the conclusion of Theorem 2.4, it should not be surprising that for any Pr* profile f that satisfies property α, $\pi_s(\langle f \mid h\rangle) = 1$ for all h. To keep the length of the paper no longer than it is, I will not descibe the proof of the extension of Theorem **??** to the case of more than one seller with the above sequential voluntary matching specification and refer the reader to Sabourian (2001a).

I do not have any result on the extension of Theorem **??** to the case of more than one seller with simultaneous voluntary matching specification (where at each period every seller chooses the buyer with whom he wishes to bargain with simultaneously).

4. An extension of the basic selection results: Complexity and off-the-equilibrium payoff

In Section 2, a NEC$r(c)$ (PEC$r(c)$) was defined as a profile of strategies $f = (f_i)_{i \in \mathcal{B}\cup\mathcal{S}}$ such that it is both a Nash (perfect Bayesian) equilibrium and is such that for all i

$$\nexists\ f_i' \text{ s.t. } f_i' \text{ is a } c\text{-best responses to } f \text{ and } f_i \succ^r f_i'$$

Thus in the definitions of both NEC$r(c)$ and PBEC$r(c)$, equilibrium strategies are supposed to be least r-complex among all the $c-$best response strategies. Thus, for example, a strategy profile $f = (f_i)_{i \in \mathcal{B}\cup\mathcal{S}}$ is a NEC$r(0)$ (PEC$r(0)$) if it is a it is both a Nash (perfect Bayesian) equilibrium and, for all i, f_i has minimal $l-$complexity amongst all strategies for i that are best responses to f_{-i}.

Both NEC$r(c)$ and PEC$r(c)$ definitions can be criticised on the grounds that, *in considering complexity*, players ignore any consideration of payoffs off-the-equilibrium path. In other words, in the above definitions, in comparing two strategies, the trade-off is between complexity of the two strategies and the payoffs of the two strategies against the *equilibrium* strategies of the others. The trade off between complexity versus off-the-equilibrium payoffs does not arise. Therefore, although complexity costs may be small (negligible, in the lexicographic case), they take priority over optimal behaviour after deviations. Hence, a profile $f = (f_i, f_{-i})$ does not constitute a NEC$r(c)$ (PEC$r(c)$), if there exists a strategy f_i' that is a $c-$best response to f and is less $r-$complex than f_i. This is the case irrespective of whether f_i' is optimal after histories which are inconsistent with f (the candidate equilibrium).

On the other hand it may be argued that a player, in comparing two strategies that give the same payoff, may prefer the more complex strategy because it generates a higher payoff off-the-equilibrium path than does the less complex one. This would be the case if complexity were a less significant criterion than the off-the-equilibrium payoff.[18] An alternative approach to PEC$r(c)$ (NEC$r(c)$) that allows for this and gives priority to the off-the-equilibrium behaviour over the complexity costs is the following weaker equilibrium concept (Kalai and Neme (1992) use a similar notion of equilibrium).

Definition 17. *A strategy profile f is weakly perfect Bayesian equilibrium strategy profile with a fixed r-complexity cost $c \geq 0$, denoted by WPEC$r(c)$, if it is both a perfect Bayesian equilibrium and is such that for all i*

$$\nexists\ f_i' \text{ s.t. } \forall\ h \in H^{\infty} \text{ and } \forall d \in D_i$$

$$\langle f_i' \mid h, d\rangle \text{ is a c-best responses to } \langle f_{-i} \mid h, d\rangle$$

$$\text{and } f_i \succ^r f_i'$$

In the above definition, complexity costs impose a restriction only among strategies that are $c-$best responses at *every* information set. Notice that any PEC$r(c)$ strategy profile is a WPEC$r(c)$.[19] Thus, any equilibrium concept with $r-$complexity costs that is optimal at every information set (credible) must satisfy WPECr irrespective of the relative importance of complexity and off-the equilibrium behaviour.

The selection result for the random matching model holds also for the weaker definition of WPEC$r(c)$.

Theorem 4.1. *Consider any $c \geq 0$ and any WPEC$r(c)$ strategy profile $f = \{f_i\}_{i\in\mathcal{B}\cup\mathcal{S}}$. If f is finite then $\pi_s(\langle f \mid h\rangle) = 1$ and $\pi_b(\langle f \mid h\rangle) = 0$ for all h, for all s and for all b, the unique induced price is the competitive price of 1 and f is stationary.*

[18]If one models credibility by introducing noise, the analogue of this issue (relative importance of complexity and off-the-equilibrium) is the relative importance of the trembles and complexity costs. More on this later.

[19]The same is not true for NECl and WPBECl.

Thus, the emergence of the competitive equilibrium as the only equilibrium outcome of random matching game is independent of the relative importance of the complexity costs and the off-the equilibrium payoff.[20]

The proof of the above result is similar to that of Theorem 2.3. The first step is to show that any WPECr(c) profile is also a Pr* - satisfies conditions (2.17) and (2.18). The proof of this is almost identical to those of Lemmas 1 and 2. The rest of the proof Theorem **??** follows exactly the same reasoning as that found in the Theorem 2.3.

In Section **??** with voluntary matching, the selection result was obtained for the set of PECr(c)s(c') profiles; these are strategies that are PECr and satisfy the following additional condition

$$\nexists\ f'_s \text{ s.t. } f'_s \text{ is a } c' - \text{best responses to } f_{-s} \tag{4.1}$$

$$\text{and } f_s \succ^s f'_s.$$

This additional condition was introduced to capture s-complexity of the sellers' strategies. As in the case of the random matching model, one can strengthen the selection result for the voluntary matching model by using a weaker concept than PECr(c)s(c') that imposes less restriction on the relative importance of complexity and off-the equilibrium payoff.

Definition 18. *A strategy profile f is a weakly perfect Bayesian equilibrium strategy profile with r-complexity $c \geq 0$ and s-complexity $c' \geq 0$, denoted by WPECr(c)s(c'), if it is both a WPECr(c) and is such that*

$$\left.\begin{array}{c} \nexists\ f'_s \textit{ s.t. } h \in H^\infty \textit{ and } \forall b \\ \langle f'_s \mid h\rangle \textit{ is a c-best responses to } \langle f_{-s} \mid h\rangle \textit{ and} \\ \langle f'_s \mid h, b\rangle \textit{ is a c-best responses to } \langle f_{-s} \mid h, b\rangle, \\ \textit{and } f_s \succ^s f'_s \end{array}\right\} \tag{4.2}$$

where (h, b) denotes history h followed by seller s choosing b.

Theorem 4.2. *Suppose $S = 1$ and the seller can choose in each period the buyer with whom he wishes to bargain (voluntary matching) and $\delta = 1$. Then consider any $c \geq 0$, any $c' > 0$ and any WPECr(c)s(c') strategy profile $f = \{f_i\}_{i \in \mathcal{B} \cup \mathcal{S}}$ If f is finite then $\pi_s(\langle f \mid h\rangle) = 1$ and $\pi_b(\langle f \mid h\rangle) = 0$ for all s, b and h, the unique induced price is the competitive price of 1 and f is stationary.*

[20]The selection results in other dynamic models with complexity costs often depend on the relative importance of the complexity costs and the off-the-equilibrium payoff. For example, in contrast to Abreu and Rubinstein's (1988) selection results with NECs(0) in 2-player repeated games, Kalai and Neme (1992) demonstrate a Folk Theorem type result for the WPECs(0) strategies in the repeated Prisoner's Dilemma.

The proof of the above result is similar to that of Theorem **??**. The first step, as in the proof of Theorem **??**, involves showing that that any WPECr(c) profile is also a Pr*. The rest of the proof Theorem **??** follows exactly the same reasoning as that found in the Theorem **??**. I shall omit the proof for reasons of space (see Sabourian 2001a).

Notice that in the definition of WPEC$r(c)s(c\prime)$

- r-complexity costs impose a restriction only among strategies that are $c-$best responses at every information set
- s-complexity costs impose a restriction only among strategies that are $c-$best responses after *every* history h and (h, b) - rather than at every information set.

Thus, Theorem **??** demostrates that the selection result with both r-complexity together with s-complexity is valid irrespective of the relative importance of complexity and off-the equilibrium behaviour after different histories h and (h, b).

An even weaker equilibrium concept than WPEC$r(c)s(c\prime)$ would be to replace (**??**) in the definition of WPECr(c)s(c') by

$$\left.\begin{array}{c} \nexists\ f'_s \text{ s.t. } h \in H^\infty \text{ and } \forall d \in D_i \\ \langle f'_s \mid h\rangle \text{ is a c-best responses to } \langle f_{-s} \mid h\rangle \text{ and} \\ \langle f'_s \mid h, d\rangle \text{ is a c-best responses to } \langle f_{-s} \mid h, d\rangle, \\ \text{and } f_s \succ^s f'_s \end{array}\right\} \tag{4.3}$$

This implies that s-complexity costs also impose a restriction only among strategies that are $c-$best responses at *every information set.* I do not have any result for this weaker concept.

5. Concluding Remarks

Finally, I would like to conclude this paper with some remarks and conjectures on the various ways of extending and expanding its results.

5.1. Equal number of buyers and sellers

The selection result in this paper shows that those on the short side of the market (the sellers in the model presented) receive all the surplus generated by exchange in any equilibrium with appropriate complexity costs. What if the number of buyers equals the number of sellers? In this case, complexity considerations do not select among the set of possible equilibrium prices. But notice that this is consistent with the competitive outcome; when $B = S$, any price between 0 and 1 is a competitive price.

5.2. Complexity criterion and alternative machine specification

R-complexity used to obtain the competitive outcome in the random matching model is a very weak concept. In the voluntary matching model, I use r-complexity together with s-complexity. Clearly, this division between the two notions of complexity reflects the machine specification I have adopted in this paper. It is possible that with a different machine specification (e.g. states of the machines changing within a period) one may be able to establish the selection results of this paper with a different notion of complexity.

5.3. Equilibrium concept

5.3.1. Noise and dominance

As I mentioned before, PEC$r(c)$ (and WPEC$r(c)$) imposes the notion of credibility directly on the set of NECr profiles. Another way of ensuring that NEC$r(c)$ strategy profiles are credible is to allow strategies (machines) to tremble and consider the limit of Nash equilibrium with trembles and r-complexity as the trembles become small. This approach is adopted by Chatterjee and Sabourian (1999, 2000). Given that the selection result for the random matching model in this paper extends to the concept of WPEC$r(c)$ my conjecture is that this result remains valid with this alternative formulation of credibility, irrespective of *the order* in which complexity costs and trembles enter the limiting arguments. For the voluntary matching model, given the discussion at the end of the last section, the selection result may depend on the relative weight of the trembles and complexity costs.

5.3.2. Dominance

Another way of weakening the equilibrium concept in this paper is to use solution concepts based on the notion of strict dominance rather than NEC$r(0)$ or PEC$r(0)$, which are based on the idea of Nash equilibrium.[21] For example, the set of NEC$r(0)$ strategy profiles could be replaced by the set that survives iterative deletion of strictly dominated strategies with r-complexity cost (denoted by ISDr), where the later concept.is defined as follows: a strategy f_i is defined to be strictly dominated with r-complexity cost (denoted by SDr) if there exist a strategy f_i' such that for all f_{-i}

$$\begin{array}{l}\text{either } \pi_i(f_i', f_{-i}) > \pi_i(f_i, f_{-i}) \\ \text{or } \pi_i(f_i, f_{-i}) = \pi_i(f_i', f_{-i}) \text{ and } f_i \succ^r f_i'\end{array}$$

I could also replace PEC$r(0)$ profiles with something like Pearce's(1984) extensive form rationalizability together with $r-$complexity (or with strategies that survive

[21] The equilibrium concepts based on the strict dominance criterion are more attractive than those based on Nash equilibrium because it is easier to justify them in terms of either rationality arguments or in terms of evolutionary/learning stories.

iterated conditional dominance together with $r-$complexity; see Fudenberg and Tirole (1991) section 4.6 for the definition of conditional dominance). It is my conjecture that one may obtain some of the results of this paper with these weaker solution concepts.[22]

5.4. Richer models of trade

RW's model considered in this paper is very simple. It is my conjecture that the results of this paper hold if one introduces a different matching/bargaining arrangement into RW's model. A more interesting issue would be to consider complexity costs in richer models of exchange than that considered by RW. For example, one could address the issues considered in this paper with a heterogeneous set of buyers and sellers and/or in the context of a model in which trade decision is not restricted to a single unit of a good. Or one could look at an exchange economy with many goods where agents trade their endowments sequentially. (For example, Gale 1986a, 2000.) It is an open question whether complexity costs also allow one to select the competitive outcomes among the set of equilibria in these richer models of exchange.

5.5. Complexity and the properties of bargaining games

Chatterjee and Sabourian (1999,2000) and Sabourian (2001b) also use complexity costs to select (uniquely) among the large number of equilibria in an n-person complete information alternating bargaining game and in a 2-person one-sided incomplete information bargaining game, respectively. In particular, these papers try to provide a justification for stationary equilibria in these classes of dynamic games. Complexity costs, however, do not always select a unique equilibrium or provide a justification for stationary/Markov strategies in dynamic games (for example repeated games; see Abreu and Rubinstein (1988) and Bloise (1998)). This paper, together with Chatterjee and Sabourian (1999, 2000) and Sabourian (2001b) demonstrate that non-stationary equilibria of the dynamic models involving bargaining are not always robust to the introduction of complexity considerations. Bargaining games have the following two properties:

(i) the (last) responder can always leave the game (the market) by accepting an offer;

(ii) the payoffs the players receive depend on the final agreement price (and if there is discounting, on the time at which the agreement is reached) and not on the history of play up to the final agreement.

These two features give complexity considerations a role in selecting among a large number of equilibria in these classes of dynamic games.

[22] In fact, it is reasonably easy to show that the crucial properties of NECr(c) profile, described in Lemmas 1 and 2, are equally satisfied by the set of ISDr strategy profiles.

6. Appendix A1: Proof of Theorem 2.4 for the case of one single seller s

Lemma 3. *Suppose $S = 1$. Then for any Pr* strategy profile f such that $\pi_s(\langle f \mid h\rangle) < 1$ for some $h \in H^\infty$ we have $f_b(h, s, b, 1) = R$ for all b and for all h.*

Proof. Suppose not; then $f_b(h, s, b, 1) = A$ for some b and for some h. Then, by condition (2.17) in the definition of Pr*, we have

$$f_b(h, s, b, 1) = A \text{ for all } h \tag{6.1}$$

Now consider another strategy f'_s for s that always proposes 1 and rejects all offers. Since after any history, the ordered pair (s, b) occurs with probability 1 in a finite time, it follows from (**??**) that f'_s can guarantee s a payoff of 1 after all histories; but this is a contradiction. ■

Lemma 4. *Suppose $S = 1$. Consider any Pr* strategy profile f; if f does not result in a continuation payoff of 1 for the seller s after all histories then there does not exist a buyer b and a history h such that the ordered pair (b, s) reaches an agreement at a price of 1 after h. Formally, for any Pr* profile f, if $\pi_s(\langle f \mid h\rangle) < 1$ for some $h \in H^\infty$ then for all b and for all h*

$$\text{either} \qquad f_b(h, b, s) \neq 1 \qquad \text{or} \qquad f_s(h, b, s, 1) = R.$$

Proof. Suppose not; then there exists b and h such that $f_b(h, b, s) = 1$ and $f_s(h, b, s, 1) = A$. Then, by condition (2.18) in the definition of Pr*, we have

$$f_b(h, b, s) = 1 \text{ for all } h \tag{6.2}$$

But then s could always obtain, after any history h, a continuation payoff of 1 by following a strategy that, irrespective of the past, always proposes 1 and accepts an offer if and only if b offers a price of 1. This is because after h either some buyer accepts the offer of 1 by s or by the law of large numbers s will eventually be matched with b and will receive an offer of 1 from b (by (**??**)). But this contradicts $\pi_s(\langle f \mid h\rangle) < 1$ for some $h \in H^\infty$. ■

Lemma 5. *Suppose $S = 1$. Then for any Pr* strategy profile f such that $\pi_s(\langle f \mid h\rangle) < 1$ for some $h \in H^\infty$, we have*

$$\pi_s(\langle f \mid h\rangle) < 1 \text{ for all } h \in H^\infty$$

Proof. This follows from b never accepting an offer of 1 (Lemma **??**) and from the ordered pair (b, s) never reaching an agreement at a price of 1 after any history (Lemma **??**). ■

Lemma 6. *Suppose $S = 1$. Then for any Pr* profile f such that $\pi_s(\langle f \mid h\rangle) < 1$ for some $h \in H^\infty$, we have*

$$\pi_b(\langle f \mid h, b, s\rangle) > 0 \text{ for all } h \text{ and for all } b$$

Proof. Since $\pi_s(\langle f \mid h\rangle) < 1$ for some $h \in H^\infty$ we have, by Lemma **??**, $\pi_s(\langle f \mid h\rangle) < 1$ for all h. Thus K, defined by

$$K \equiv \max_{h \in H^\infty} \pi_s(\langle f \mid h\rangle),$$

is less than 1 (K is well-defined because f is finite). Since f constitutes a perfect Bayesian equilibrium, it follows that if, after any (h, b, s), buyer b offers a price $K + \epsilon$ for some ϵ such that $K + \epsilon < 1$, it will be accepted by s (otherwise s obtains at most K). Thus b can always obtain at least $1 - K - \epsilon > 0$. This establishes the result. ■

Now for any strategy profile f let

$$m_s^b(f) = \min_{h \in H^\infty} \pi_s(\langle f \mid h, s, b\rangle)$$

$$m_b^b(f) = \min_{h \in H^\infty} \pi_s(\langle f \mid h, b, s\rangle)$$

$$z(b, f) = \max_{h \in H^\infty} \pi_b(\langle f \mid h\rangle) \tag{6.3}$$

$$\overline{b}(f) \equiv \text{ arg min}_b \ (m_b^b(f) + m_s^b(f)) \tag{6.4}$$

Note that if f is finite then $z(b, f)$ and $m_i^b(f)$ are well defined for $i = b, s$.

For the rest of this section, I fix a profile f and refer to $m_s^b(f)$, $m_b^b(f)$, $z(b; f)$ and $\overline{b}(f)$ by m_s^b, m_b^b, $z(b)$ and $\overline{b}$ respectively.

Lemma 7. *Suppose $S = 1$. Then for any finite perfect Bayesian (subgame perfect) equilibrium strategy profile f we have $m_{\overline{b}}^{\overline{b}} \geq m_s^{\overline{b}}$.*

Proof. Suppose not; then $m_{\overline{b}}^{\overline{b}} < 1/2(m_s^{\overline{b}} + m_{\overline{b}}^{\overline{b}})$. Now it follows from the definition of $m_{\overline{b}}^{\overline{b}}$ that there exists a history h such that $m_{\overline{b}}^{\overline{b}} = \pi_s(\langle f \mid h, \overline{b}, s\rangle)$. Now suppose p is the offer of $\overline{b}$ after $(h, \overline{b}, s)$. Now if s rejects p after $(h, \overline{b}, s)$ he can get at least $1/2(m_s^{\overline{b}} + m_{\overline{b}}^{\overline{b}})$ in the next period. But this exceeds $m_{\overline{b}}^{\overline{b}} = \pi_s(\langle f \mid h, \overline{b}, s\rangle)$. This contradicts the supposition that f is a perfect Bayesian equilibrium. ■

Lemma 8. *Suppose $S = 1$. Then for any finite perfect Bayesian equilibrium strategy profile f there exists a buyer b and a history h such that $\pi_b(\langle f \mid h, b, s\rangle) = 0$.*

Proof. Consider any perfect Bayesian equilibrium f and let $\overline{b}$ be defined as in (**??**). First, I establish that

$$m_s^{\overline{b}} \geq 1 - z(\overline{b}) \tag{6.5}$$

To show this, suppose otherwise; then

$$m_s^{\overline{b}} < 1 - z(\overline{b}) - \epsilon \text{ for some } \epsilon > 0. \tag{6.6}$$

Now consider any history h and suppose that players s and $\overline{b}$ are matched and s makes a price offer of $(1 - z(\overline{b}) - \epsilon)$ to $\overline{b}$ after $(h, s, \overline{b})$. By the definition of $z(\overline{b})$, given in condition (**??**), this offer will be accepted by $\overline{b}$. Thus $m_s^{\overline{b}} \geq 1 - z(\overline{b}) - \epsilon$. But this contradicts condition (**??**). Therefore, condition (**??**) holds.

Now it follows from the definition of $z(\overline{b})$ that there exists a h such that $\pi_{\overline{b}}(f \mid h) = z(\overline{b})$. Therefore,

$$z(\overline{b}) \leq \tfrac{1}{B}\left\{1/2(1-m_s^{\overline{b}}) + 1/2(1-m_{\overline{b}}^{\overline{b}})\right\} + \textstyle\sum_{b\neq\overline{b}} \tfrac{1}{B}$$
$$\{1/2\left[1 - \pi_s(\langle f \mid h, s, b\rangle) - \pi_b(\langle f \mid h, s, b\rangle)\right] + 1/2\left[1 - \pi_s(\langle f \mid h, b, s\rangle) - \pi_b(\langle f \mid h, b, s\rangle)\right]\}$$

(The expression in the RHS of the above inequality gives an upper bound on $z(\overline{b})$. The term $1/2(1 - m_s^{\overline{b}}) + 1/2(1 - m_{\overline{b}}^{\overline{b}})$ bounds $\overline{b}$'s expected payoff 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 second term on the RHS of the last inequality is the weighted sum of the maximum payoff $\overline{b}$ can receive in the event that in the next period the seller meets one of the other buyers, weighted by the probability of each such event.) Therefore, from the definitions of m_s^b and m_b^b we have

$$z(\overline{b}) \leq \tfrac{1}{B}\left\{1/2(1-m_s^{\overline{b}}) + 1/2(1-m_{\overline{b}}^{\overline{b}})\right\} +$$
$$\textstyle\sum_{b\neq\overline{b}} \tfrac{1}{B}\left\{1/2(1 - m_s^b - \pi_b(\langle f \mid h, s, b\rangle)) + 1/2(1 - m_b^b - \pi_b(\langle f \mid h, b, s\rangle))\right\}$$

The last condition together with condition (**??**) imply that

$$m_s^{\overline{b}} \geq 1 - z(\overline{b}) \geq \frac{1}{2B}\left\{(m_s^{\overline{b}} + m_{\overline{b}}^{\overline{b}}) + \sum_{b\neq\overline{b}}\left[(m_s^b + m_b^b) + (\pi_b(\langle f \mid h, s, b\rangle) + \pi_b(\langle f \mid h, b, s\rangle)\right]\right\}$$

But this together with the definition of $\overline{b}$ imply that

$$m_s^{\overline{b}} \geq \frac{1}{2B}\left\{B(m_s^{\overline{b}} + m_{\overline{b}}^{\overline{b}}) + \sum_{b\neq\overline{b}}(\pi_b(\langle f \mid h, s, b\rangle) + \pi_b(\langle f \mid h, b, s\rangle))\right\}$$

Therefore, it follows from Lemma **??** that

$$m_s^{\overline{b}} \geq \frac{1}{2B}\left\{2Bm_s^{\overline{b}} + \sum_{b\neq\overline{b}}(\pi_b(\langle f \mid h, s, b\rangle) + \pi_b(\langle f \mid h, b, s\rangle))\right\}$$

Hence,

$$0 \geq \frac{1}{2B}\left\{\sum_{b\neq\bar{b}}(\pi_b(\langle f \mid h,s,b\rangle) + \pi_b(\langle f \mid h,b,s\rangle)\right\}$$

Since the continuation payoffs are always non-negative, it follows from the previous inequality that

$$\pi_b(\langle f \mid h,b,s\rangle) = 0 \text{ for all } b \neq \bar{b} \tag{6.7}$$

This completes the proof of this Lemma. ■

Now, by Lemmas **??** and **??**, we have for any Pr* profile f

$$\pi_s(\langle f \mid h\rangle) = 1 \text{ for all } h \in H^\infty.$$

7. Appendix A2: Proof of Theorem 2.4 with an arbitrary number of sellers

The proof for the case of more than one seller is by induction on the number of sellers S in the market. In Appendix A1, it was shown that the result holds for the case of $S = 1$. To complete the proof of Theorem 2.4 with an arbitrary number of sellers, I need to show that if for any $S' < S$, all Pr* profiles for a market with S' sellers results in continuation payoffs of 1, after all histories, for each of the S' sellers, then any Pr* profile in a market with S sellers also results in continuation payoffs of 1, after all histories, for each of the S sellers.

Lemma 9. *Consider any Pr* strategy profile f in a market with exactly S sellers. Let $\overline{h}$ be any history after which there are less than S sellers left in the market. Denote, respectively, the set of buyers and sellers left in the market after $\overline{h}$ by $\mathcal{B}' \subset \mathcal{B}$ and $\mathcal{S}' \subset \mathcal{S}$. Then the strategy $\langle f \mid \overline{h}\rangle$, when restricted to the remaining agents $(\mathcal{B}', \mathcal{S}')$, is also a Pr* for this smaller market.*

Proof. Let $H^\infty(\mathcal{B}', \mathcal{S}')$ be the set of finite histories when there are $(\mathcal{B}', \mathcal{S}')$ agents in the market. Next, consider any $b \in \mathcal{B}'$ and $s \in \mathcal{S}'$ remaining in the market after $\overline{h}$. Since f is a Pr* strategy profile and $\langle f_i \mid \overline{h}\rangle(h,d) = f_i(\overline{h},h,d)$ for all i, for all h and for all d, it follows, respectively, from conditions (2.17) and (2.18) in the definition of a Pr* that

$$\langle f_b \mid \overline{h}\rangle(h,s,b,1) = \langle f_b \mid \overline{h}\rangle(h',s,b,1) \quad \text{for all } h \text{ and } h' \in H^\infty(\mathcal{B}', \mathcal{S}') \tag{7.1}$$

$$\left.\begin{array}{l}\text{if } \langle f_s \mid \overline{h}\rangle(h,b,s,1) = A \text{ for some } h \in H^\infty(\mathcal{B}', \mathcal{S}') \text{ then either} \\ \langle f_b \mid \overline{h}\rangle(h,b,s) = 1 \ \forall h \in H^\infty(\mathcal{B}', \mathcal{S}') \text{ or } \langle f_b \mid \overline{h}\rangle(h,b,s) \neq 1 \ \forall h \in H^\infty(\mathcal{B}', \mathcal{S}').\end{array}\right\} \tag{7.2}$$

Therefore, since any strategy induced by a perfect Bayesian equilibrium after any history is also a perfect Bayesian equilibrium of the continuation game, it follows from conditions (**??**) and (**??**) that $\langle f \mid \overline{h}\rangle$ is a Pr*. ■

Lemma 10. *Suppose that for any positive integer $S' < S$ all Pr* strategy profiles in markets with S' sellers result in a continuation payoff of 1, after any history, for each of the S' seller. Then for any Pr* strategy profile f in a market with exactly S sellers the following holds: if $\pi_s(\langle f \mid h \rangle) < 1$ for some seller s and for some h then we have $f_b(h, s, b, 1) = R$ for all h, for all s and for all b.*

Proof. Suppose not; then $f_b(h, s, b, 1) = A$ for some h , for some s and for some b. By condition (2.17) in the definition of Pr*, this implies that

$$f_b(h, s, b, 1) = A \text{ for all } h \tag{7.3}$$

Now consider strategy f'_s for s such that for any $(h, d) \in H^\infty \times D_s$

$$f'_s(h,d) = \begin{cases} 1 & \text{if } h \text{ s.t. there are } S \text{ sellers \& } d = (s, b') \text{ for some } b' \\ A & \text{if } h \text{ is s.t. there are } S \text{ sellers \& } d = (b', s, 1) \text{ for some } b' \\ R & \text{if } h \text{ is s.t. there are } S \text{ sellers, \& } d = (b', s, p) \text{ for some } b' \text{ \& } p < 1 \\ f_s(h,d) & \text{otherwise.} \end{cases}$$

Fix any history $\overline{h}$ at which there are still S sellers in the market. Next, notice that if s chooses strategy f'_s and the other players choose the profile $(\langle f_{-s} \mid \overline{h} \rangle)$ then with probability one some seller is going to reach an agreement with a buyer in finite time. Otherwise, by the law of large numbers, seller s will eventually be matched with buyer b with s as the proposer; but then, by the definition of f'_s and (**??**), s will make an offer of 1 and b will accept; but this is a contradiction.

Next, I would like to show that the expected payoff to s if $(f'_s, \langle f_{-s} \mid \overline{h} \rangle)$ is chosen is 1. This is done in several steps.

Step 1: $\pi_s(\langle f'_s, f_{-s} \mid h \rangle) = 1$ for any history of outcomes $h = (e^1, ..., e^t)$ such that

$$\left.\begin{array}{l} \text{no agreement is reached before period } t \text{ and} \\ \text{seller } s \text{ reaches an agreement at period t with some buyer } b' \end{array}\right\} \tag{7.4}$$

Given any $h = (e^1, ..., e^t)$ that satisfies (**??**), there are still S sellers remaining in the market after $(e^1, ..., e^{t-1})$. Therefore, it follows from the definition of f'_s that any agreement by s after $(e^1, ..., e^{t-1})$ must be at the price 1. Thus, in this case $\pi_s(\langle f'_s, f_{-s} \mid h \rangle) = 1$.

Step 2: $\pi_s(\langle f'_s, f_{-s} \mid h \rangle) = 1$ for any history of outcomes $h = (e^1, ..., e^t)$ such that

$$\left.\begin{array}{l} \text{no agreement is reached before period } t, \\ \text{some seller } s' \neq s \text{ reaches an agreement at period t with some buyer } b' \\ \text{ and seller } s \text{ does not reach an agreement at period } t. \end{array}\right\} \tag{7.5}$$

Notice that after any history $h = (e^1, ..., e^t)$ that satisfies (**??**), there are less than S sellers in the market and seller s is one of the remaining sellers. Therefore, by the definition of f_s', we have

$$(\langle f_s', f_{-s} \mid h\rangle) = (\langle f_s \mid h\rangle, \langle f_{-s} \mid h\rangle) \tag{7.6}$$

Moreover, since (f_s, f_{-s}) constitutes a Pr*, it follows from Lemma **??** that $(\langle f_s, f_{-s} \mid h\rangle)$ constitutes a Pr* for the remaining agents. This, together with (**??**) imply that $(\langle f_s', f_{-s} \mid h\rangle)$ is also also a Pr* for the remaining agents. But since there are less than S sellers in the market after h and s is one of the remaining sellers in the market after h., we have, by the supposition of the induction argument, $\pi_s(\langle f_s', f_{-s} \mid h\rangle) = 1$. This completes the proof of this step.

Step 3: $\pi_s(f_s', \langle f_{-s} \mid \overline{h}\rangle) = 1$ for any history $\overline{h}$ at which there are still S sellers in the market.

Since $(f_s', \langle f_{-s} \mid \overline{h}\rangle)$ results in an agreement between some pair of agents with probability 1, it follows that any outcome path that $(f_s', \langle f_{-s} \mid \overline{h}\rangle)$ induces with a positive probability must include a finite history h satisfying either (**??**) or (**??**). But in both cases the expected continuation payoff after h to s if $(f_s', \langle f_{-s} \mid \overline{h}\rangle)$ is chosen is 1. This completes the proof of this step.

Finally, since the strategy profile f is a perfect Bayesian equilibrium, it follows from Step 3 that $\pi_s(\langle f \mid \overline{h}\rangle) = 1$ for any history $\overline{h}$ at which there are still S sellers in the market. But this, together with the induction assumption contradicts the hypothesis that $\pi_s(\langle f \mid h\rangle) < 1$ for some h. ■

Suppose that for any positive integer $S' < S$ all Pr* strategy profiles in markets with S' sellers result in a continuation payoff of 1, after any history, for each of the S' sellers. Then for any Pr* strategy profile f in a market with exactly S sellers the following holds: if $\pi_s(\langle f \mid h\rangle) < 1$ for some seller s and for some h then we have for all b and for all h

$$\text{either} \qquad f_b(h, b, s) \neq 1 \qquad \text{or} \qquad f_s(h, b, s, 1) = R.$$

Proof. Suppose not; then there exists a buyer b and a history h such that $f_b(h, b, s) = 1$ and $f_s(h, b, s, 1) = A$. By condition (2.18) in the definition of Pr*, this implies that

$$f_b(h, b, s) = 1 \text{ for all } h \tag{7.7}$$

Next consider any strategy f_s' for s defined by

$$f_s'(h, d) = \begin{cases} 1 & \text{if } h \text{ s.t. there are } S \text{ sellers } \& \ d = (s, b') \text{ for some } b' \\ A & \text{if } h \text{ is s.t. there are } S \text{ sellers } \& \ d = (b', s, 1) \text{ for some } b' \\ R & \text{if } h \text{ is s.t. there are } S \text{ sellers, } \& \ d = (b', s, p) \text{ for some } b' \ \& \ p < 1 \\ f_s(h, d) & \text{otherwise.} \end{cases}$$

Now fix any history $\overline{h}$ at which there are still S sellers in the market. Notice that if the players choose the profile $(f'_s, \langle f_{-s} \mid \overline{h}\rangle)$ then with probability one some seller is going to reach an agreement with a buyer in finite time. Otherwise, by the law of large numbers, seller s will eventually be matched with buyer b with b as the proposer; but then, by condition (**??**) and the definition of f'_s, b will make an offer of 1 and s will accept; but this is a contradiction.

Now, by exactly the same arguments as in the previous Lemma (Steps1-3), it can be shown that $\pi_s(\langle f'_s, \langle f_{-s} \mid \overline{h}\rangle) = 1$ for any history $\overline{h}$ at which there are still S sellers in the market. But since the strategy profile f is a perfect Bayesian equilibrium it follows that $\pi_s(\langle f \mid \overline{h}\rangle) = 1$ for any such $\overline{h}$. But this, together with the induction assumption, contradict the hypothesis that $\pi_s(\langle f \mid h\rangle) < 1$ for some h. ■

Lemma 11. *Suppose that for any positive integer $S' < S$ all Pr* strategy profiles in markets with less than S sellers result in a continuation payoff of* 1*, after any history, for each of the S' sellers. Then for any Pr* strategy profile f in a market with exactly S sellers the following holds: if for some seller s, $\pi_s(\langle f \mid h\rangle) < 1$ for some h then $\pi_s(\langle f \mid h\rangle) < 1$ for all h.*

Proof. This follows from no buyer accepting an offer of 1 (Lemma **??**) and from the ordered pair (b, s) never reaching an agreement at a price of 1, for all b and all s (Lemma **??**). ■

Lemma 12. *Suppose that all Pr* strategy profiles in markets with less than $S(> 1)$ sellers result in a continuation payoff of* 1*, after any history, for each seller. Then any Pr* strategy profile f in a market with exactly S sellers results in continuation payoff of 1, after each history, for each seller.*

Proof. Suppose not; then there exists a Pr* strategy profile f in a market with exactly S sellers such that for some seller s, $\pi_s(\langle f \mid h\rangle) < 1$ after some h. This implies, by Lemma (**??**), that

$$\pi_s(\langle f \mid h\rangle) < 1 \qquad \text{for all } h \tag{7.8}$$

Now, consider any history of outcomes $h = (e^1, ..., e^t)$ such that

(i) seller $s' \neq s$ and buyer b' are matched and reach an agreement at some period $t' \leq t$;

(ii) no other pair of agents reach an agreement at any period $t' \leq t$.

Therefore, after history h there are $S - 1$ sellers in the market and seller s is one of the remaining sellers. Now, since f is a Pr*, it follows from Lemma **??** and the supposition of the induction argument that seller s has a continuation payoff of 1 after h. This contradicts condition (**??**). ■

Now it follows from Lemma **??**, the proof of Theorem 2.4 for the case of a single seller (in Appendix A1) and by the induction on the number of sellers that for any Pr* strategy profile f, $\pi_s(\langle f \mid h\rangle) = 1$ for all s and all h.

8. Appendix B: Proof of Theorem 4.3 for the case of a single seller s

Here with voluntary matching, I define a Pr* strategy profile in exactly the same way as with random matching - see Definition 13. Also, Lemmas 1, 2, **??**, **??**, **??**, **??** hold for the voluntary matching model and the proofs of these Lemmas in this case with endogenous matching arrangement are similar to those with random matching. In fact, the proofs of Lemmas 1, 2, **??** and **??** are identical to those in Section **??** and Appendix A1 and I will not repeat the arguments. Here, I shall only provide the proofs for Lemmas **??** and **??** when the trading arrangement is voluntary.

Proof of Lemma ?? for the voluntary matching model: Suppose not; then $f_b(h, s, b, 1) = A$ for some b and for some h. This, together with condition (2.17) of the definition of Pr*, imply that

$$f_b(h, s, b, 1) = A \text{ for all } h \tag{8.1}$$

Now consider any strategy f_s' for s that always chooses player b, always proposes 1 and always rejects all offers. Since, after any history, with probability 1, s will have the opportunity to make a proposal to player b in finite time, it follows from (**??**) that f_s' can guarantee s a continuation payoff of 1 after any history; but this contradicts the assumption that $\pi_s(\langle f \mid h \rangle) < 1$ for some h. ■

Proof of Lemma ?? for the voluntary matching model: Suppose not; then there exists b and h such that $f_b(h, b, s) = 1$ and $f_s(h, b, s, 1) = A$. This, together with condition (2.18) of the definition of Pr*, imply that

$$f_b(h, b, s) = 1 \text{ for all } h \tag{8.2}$$

But then, after any history h, s could always obtain 1 by following a stationary strategy that irrespective of the past history always chooses b, always makes an offer of 1 and accepts an offer p if and only if $p = 1$. (This is because after h this strategy results in either b accepting the offer of 1 by s or, by condition (**??**), in b making an offer of 1 to s.) But this contradicts the assumption that $\pi_s(\langle f \mid h \rangle) < 1$ for some h. ■

Now, as in Appendix A1, for any f let

$$\begin{aligned} m_s^b(f) &= \min_{h \in H^\infty} \pi_s(\langle f \mid h, s, b \rangle) \\ m_b^b(f) &= \min_{h \in H^\infty} \pi_s(\langle f \mid h, b, s \rangle) \\ \bar{b}(f) &\equiv \arg\min_b \left(m_b^b(f) + m_s^b(f)\right) \\ z(b; f) &= \max_{h \in H^\infty} \pi_b(\langle f \mid h \rangle) \\ H(b; f) &= \{h \in H^\infty \,|\pi_b(\langle f \mid h \rangle) = z(b; f)\} \end{aligned} \tag{8.3}$$

Note that since f is finite, $z(b; f)$ and $m_i^b(f)$ are well defined for $i = b, s$, and $H(b; f)$ is not empty.

Next, not that any strategy profile f defines a probability distribution on the set of outcome paths in this game. From this, one can compute the probability of any finite history $h \in H^{\infty}$, given that the players choose a given strategy profile f. I shall denote such a probability by $\theta(h; f)$. Also, with some abuse of notation, let

$$\theta(h, d; f) = \text{probability of } (h, d) \in H^{\infty} \times D \text{ given that the players choose profile } f$$

$$\theta(h, b; f) = \text{probability of } (h, b) \text{ given that the players choose profile } f$$

where (h, b) refers to history h followed by the seller s choosing b as the partner in the next period. Finally, for any strategy profile f, I denote the probability that s chooses a buyer b for the first time after history h by $\beta(h, b; f)$. Thus,

$$\beta(h, b; f) = \sum_{h' \in \Sigma^b} \theta(h, h', b; f)$$

where

$$\Sigma^b = \left\{ h = (e^1, ..., e^t) \in H^{\infty} \,|e^{\tau} \text{ does not involve a match between } s \text{ and } b \text{ for all } \tau \leq t \right\}$$

Henceforth, I fix a strategy profile f and refer to $m_s^b(f)$, $m_b^b(f)$, $\overline{b}(f)$, $z(b; f)$, $H(b; f)$, $\theta(h, d; f)$ and $\beta(h, b; f)$ by m_s^b, m_b^b, $\overline{b}$, $z(b)$, $H(b)$, $\theta(h, d)$ and $\beta(h, b)$, respectively.

Lemma 13. *Suppose $S = 1$. Then for any finite perfect Bayesian equilibrium strategy profile f we have $z(b) = z(b')$ for all b and b'.*

Proof. Suppose not; then

$$z(b') > z(b) + \epsilon$$

for some b, for some b' and for some $\epsilon > 0$.

Consider any $h_{b'} \in H(b')$. Since $\pi_b(\langle f \mid h_{b'} \rangle) = z(b')$ it follows that

$$\pi_s(\langle f \mid h_{b'} \rangle) \leq 1 - z(b') < 1 - z(b) - \epsilon \tag{8.4}$$

Now consider a strategy f_s' for s that always chooses buyer b, rejects all offers and always makes the proposal $1 - z(b) - \epsilon$. Clearly, b always accepts the proposal $1 - z(b) - \epsilon$. Therefore $(f_s', f_{-s} \mid h_{b'})$ guarantees seller s a payoff of $1 - z(b) - \epsilon$. Since f is a perfect Bayesian equilibrium we have $\pi_s(\langle f \mid h_{b'} \rangle) \geq \pi_s(\langle f_s', f_{-s} \mid h_{b'} \rangle) \geq 1 - z(b) - \epsilon$. But this contradicts condition (**??**). ■

Since for any finite perfect Bayesian equilibrium f we have $z(b) = z(b')$ for all b and b' henceforth, I shall refer to $z(b)$ by z.

Definition 19. *A strategy profile f is said to satisfy property α if there exists a buyer b and a history $h_b \in H(b)$ such that the probability that s chooses some $b' \neq b$ at some point after h_b is postive. Formally, f satisfies property α if $\exists$ b and $h_b \in H(b)$ s.t.*

$$\beta(h_b, b') > 0 \text{ for some } b' \neq b \tag{8.5}$$

Next, for any strategy f_s define an equivalence relation $\sim^{f_s}$ as follows:

$$h \sim^{f_s} h' \text{ if and only if } \langle f_s \mid h \rangle = \langle f_s \mid h' \rangle.$$[23]

Lemma 14. *For any perfect Bayesian equilibrium f, for any b and for any $h_b \in H(b)$ we have*

$$\textit{if } f \textit{ satisfies does not satisfy property } \alpha \textit{ then } \pi_s(\langle f \mid \langle h \rangle) \leq 1 - z \textit{ for any } h \sim^{f_s} h_b \tag{8.6}$$

Proof. Fix any b and any $h_b \in H(b)$. Since f does not satisfy property α the induced strategy $\langle f_s \mid h_b \rangle$ always chooses buyer b after any history h' such that $\theta(h'; \langle f \mid h_b \rangle) > 0$. Therefore, since $\pi_b(\langle f \mid h_b \rangle) = z$, it follows that $\langle f_b \mid h_b \rangle$ and $\langle f_s \mid h_b \rangle$ result in a payoff of z for player b irrespective of the strategies adopted by the other players. Thus,

$$\pi_b(\langle f_b \mid h_b \rangle, \langle f_s \mid h_b \rangle, f'_{-b,s}) = z \text{ for all } f'_{-b,s} \tag{8.7}$$

Now fix any $h \sim^{f_s} h_b$. By the definition of the equivalent relation $\sim^{f_s}$ we have

$$\langle f_s \mid h \rangle = \langle f_s \mid h_b \rangle \tag{8.8}$$

Conditions (**??**) and (**??**) together imply

$$\left.\begin{array}{c} \pi_b(\langle f_b \mid h_b \rangle, \langle f_{-b} \mid h \rangle) = \pi_b(\langle f_b \mid h_b \rangle, \langle f_s \mid h \rangle, \langle f_{-b,s} \mid h \rangle) = \\ \pi_b(\langle f_b \mid h_b \rangle, \langle f_s \mid h_b \rangle, \langle f_{-b,s} \mid h \rangle) = z \end{array}\right\} \tag{8.9}$$

But since f is a perfect Bayesian equilibrium, it follows from (**??**) that

$$\pi_b(\langle f \mid h \rangle) \geq \pi_b(\langle f_b \mid h_b \rangle, \langle f_{-b} \mid h \rangle) = z \tag{8.10}$$

Therefore, $\pi_s(\langle f \mid \langle h \rangle) \leq 1 - z$ for any $h \sim^{f_s} h_b$. ■

Lemma 15. *Suppose $S = 1$. Then for any $c' > 0$, every PECr(c)s(c') (in fact any perfect Bayesian equilibrium that satisfies condition (**??**)) profile f satisfies property α.*

Proof. Suppose not; then for some $c' > 0$ there exists a PEC$r(c)s(c')$ (perfect Bayesian equilibrium that satisfies condition (**??**)) profile f that does not satisfy property α.

Now, fix any $\epsilon > 0$ such that

$$\epsilon < c'$$

and let

$$p' = \max\{0, 1 - z - \epsilon\}$$

[23] As mentioned in Section 2, the size of the partition on the set of histories H^∞ defined by the equivalence relation $\sim^{f_s}$ is a measure of s-complexity of strategy f_s.

Next notice that since f does not satisfy property α, $f_s(h_b) = b$ for every buyer b. Thus,

$$\langle f_s \mid h_b \rangle \neq \langle f_s \mid h_{b'} \rangle \text{ for all } b' \neq b \tag{8.11}$$

Now, define another machine f'_s for s that is otherwise identical to f_s except that the induced strategies $\{\langle f_s \mid h_b \rangle\}_{b \in \mathcal{B}}$ are replaced by a single induced strategy that always chooses a single buyer $\overline{b}$, always offers a price p' and accepts a price offer if and only if the price offer is no less than p'. Formally, let

$$\widetilde{H} = \left\{ h \in H^\infty \,\middle|\, h \sim^{f_s} h_b \text{ for some } b \right\}$$

$$\overline{H} = \left\{ h \in H^\infty \,\middle|\, h = (h^1, h^2) \text{ for some } h^1 \in \widetilde{H} \text{ and } h^2 \in H^\infty \right\}$$

Then, f'_s is defined as follows.

$$f'_s(h, \varphi) = \begin{cases} \overline{b} & \text{if } h \in \overline{H} \\ f_s(h, \varphi) & \text{otherwise} \end{cases}$$

$$f'_s(h, d) = \begin{cases} p' & \text{if } h \in \overline{H} \text{ and } d = (s, b) \text{ for some } b \\ A & \text{if } h \in \overline{H} \text{ and } d = (b, s, p) \text{ for some } b \text{ and } p \geq p' \\ R & \text{if } h \in \overline{H} \text{ and } d = (b, s, p) \text{ for some } b \text{ and } p < p' \\ f_s(h, d) & \text{otherwise} \end{cases}$$

By condition (**??**), $\widetilde{H}$ has $B > 1$ elements. Moreover, by construction

$$\langle f'_s \mid h \rangle = \langle f'_s \mid h' \rangle \text{ for all } h, h' \in \widetilde{H}$$

Thus, f_s is more s-complex than f'_s.

Now, I demonstrate a contradiction by showing that $\pi_s(\langle f_s, f_{-s} \mid h \rangle) - \pi_s(\langle f'_s, f_{-s} \mid h \rangle) < c$ for all h.

Since f does not satisfy property α, by Lemma **??**, we have

$$\text{if } h \in \widetilde{H} \text{ then } \pi_s(\langle f \mid h \rangle) \leq 1 - z \tag{8.12}$$

Now consider 2 cases.

Case 1: $1 - z = 0$. Since $\langle f'_s \mid h \rangle$ guarantees at least a payoff of zero, it follows that in this case

$$\pi_s(\langle f'_s \mid h \rangle, \langle f_{-s} \mid h \rangle) \geq 1 - z \tag{8.13}$$

Case 2: $1 - z > 0$. Since the continuation payoff of any buyer is at most z it follows that any buyer accepts any price offer less than $1 - z$. But since $1 - z > 0$ we have $p' = \max\{0, 1 - z - \epsilon\} < 1 - z$. Therefore, b always accepts p'. Thus,

$$\text{if } h \in \widetilde{H} \text{ then } \pi_s(\langle f'_s \mid h \rangle, \langle f_{-s} \mid h \rangle) = p' \geq 1 - z - \epsilon. \tag{8.14}$$

(This is because f'_s always selects $\overline{b}$, always offers $p' < 1 - z$ and always accepts an offer if the price offer is no less than p').

Thus, in both cases we have

$$\text{if } h \in \widetilde{H} \text{ then } \pi_s(\langle f_s' \mid h\rangle, \langle f_{-s} \mid h\rangle) \geq 1 - z - \epsilon. \tag{8.15}$$

But this, together with condition (**??**) imply that

$$\text{if } h \in \widetilde{H} \text{ then } \pi_s(\langle f \mid h\rangle) - \pi_s(\langle f_s' \mid h\rangle, \langle f_{-s} \mid h\rangle) \leq \epsilon < c \tag{8.16}$$

If, on the other hand , h is such that $h \notin \widetilde{H}$ then by the definition of f_s' the profiles (f_s, f_{-s}) and (f_s', f_{-s}) behave in exactly the same way at any period following such h. This, together with (**??**) imply that

$$\pi_s(\langle f_s', f_{-s} \mid h\rangle) - \pi_s(\langle f_s, f_{-s} \mid h\rangle) \leq c \text{ for all } h.$$

But this is a contradiction because $f_s^{'}$ is less complex than f_s and is a c-best response to f_{-s}. ■

Lemma 16. *Suppose $S = 1$. Then for any finite perfect Bayesian equilibrium strategy profile f with voluntary matching we have $m_{\overline{b}}^{\overline{b}} \geq m_s^{\overline{b}}$.*

Lemma **??** is a restatement of Lemma **??** for the voluntary matching model. The steps of the proofs of the two Lemmas are identical and therefore I will omit stating the proof of Lemma **??**.

Lemma 17. *Suppose $S = 1$. Then for any finite Pr* strategy profile f that satisfies Property α we have $\pi_s(\langle f \mid h\rangle) = 1$ for all h.*

Proof. Suppose not; then there exits a finite Pr* strategy profile f that satisfies Property α and $\pi_s(\langle f \mid h\rangle) < 1$ for some h. Now let

$$\overline{\epsilon} = \min_{b \in \mathcal{B},\ h \in H^\infty} \pi_b(\langle f \mid h, b, s\rangle) \tag{8.17}$$

(Since f is finite $\overline{\epsilon}$ is well defined). Then by Lemma **??**, we have

$$\overline{\epsilon} > 0.$$

Now since f satisfies property α there exists b, $h_b \in H(b)$ and $b' \neq b$ such that

$$\beta(h_b, b') > 0.$$

By the definition of $z(b)$ and Lemma **??** we have that $z = z(b) = \pi_b(\langle f \mid h_b\rangle)$. Therefore, since the seller's minimum continuation payoff is at least $1/2(m_s^{\overline{b}} + m_{\overline{b}}^{\overline{b}})$, we can write an upper bound for z as follows

$$z \leq 1 - 1/2(m_s^{\overline{b}} + m_{\overline{b}}^{\overline{b}}) - \sum_{h \in \Sigma^{b'}} \theta(h_b, h, b')[1/2(\pi_{b'}(\langle f \mid h_b, h, s, b' \rangle) + 1/2\pi_{b'}(\langle f \mid h_b, h, b', s \rangle)] \tag{8.18}$$

(The third terms on the RHS of the last inequality is simply the sum of the expected continuation payoff of $b' \neq b$ after history h_b.)

Therefore, it follows from (**??**)

$$z \leq 1 - 1/2(m_s^{\overline{b}} + m_{\overline{b}}^{\overline{b}}) - \frac{\overline{\epsilon}}{2} \sum_{h \in \Sigma^{b'}} \theta(h_b, h, b') \tag{8.19}$$

Thus it follows from (**??**) and from the definition of $\beta(h_b, b')$ that

$$z \leq 1 - 1/2(m_s^{\overline{b}} + m_{\overline{b}}^{\overline{b}}) - \frac{\overline{\epsilon}\beta(h_b, b')}{2} \tag{8.20}$$

Now, by the same argument as that which follows (**??**) in the proof of Lemma **??**, I show that

$$m_s^{\overline{b}} \geq 1 - z \tag{8.21}$$

To show this, suppose otherwise; then

$$m_s^{\overline{b}} < 1 - z - \epsilon \text{ for some } \epsilon > 0. \tag{8.22}$$

Now consider any history h and suppose that s makes a price offer of $(1 - z - \epsilon)$ to $\overline{b}$ after $(h, s, \overline{b})$. Since $z = z(\overline{b})$ is the maximum continuation payoff of $\overline{b}$, it follows that this offer will be accepted by $\overline{b}$. Thus $m_s^{\overline{b}} \geq 1 - z(\overline{b}) - \epsilon$. But this contradicts condition (**??**). Therefore, condition (**??**) holds.

But (**??**), together with condition (**??**), imply that

$$m_s^{\overline{b}} \geq 1/2(m_s^{\overline{b}} + m_{\overline{b}}^{\overline{b}}) + \frac{\overline{\epsilon}\beta(h_b, b')}{2}$$

Therefore, it follows from Lemma **??** that

$$m_s^{\overline{b}} \geq m_s^{\overline{b}} + \frac{\overline{\epsilon}\beta(h_b, b')}{2}$$

But since $\overline{\epsilon} > 0$ and $\beta(h_b, b') > 0$ this is a contradiction. Therefore, $\pi_s(\langle f \mid h \rangle) = 1$ for all h. ■

Since any PEC$r(c)s(c')$ is a PEC$r(c)$ strategy profile and any PEC$r(c)$ is a Pr*, it follows from Lemmas **??** and **??** that for any $c > 0$, any $c' > 0$ and for any finite PEC$r(c)s(c')$ strategy profile f, $\pi_s(\langle f \mid h \rangle) = 1$ for all h. But this implies that for any such f, $\pi_b(\langle f \mid h \rangle) = 0$ for all h and that the unique equilibrium price is one. Moreover, by the same arguments as that in the proof of Theorem 2.3 in Section 2, f_i is stationary for all i.

References

[1] Abreu, D. and A. Rubinstein (1988):"The Structure of Nash Equilibria in Repeated Games with Finite Automata," *Econometrica*, 56**,** pp 1259-1282.

[2] Al-Najjar N. and R. Smordinsky (1998) 'Pivotal Players and the Characterization of Influence' *Center for Mathematical Studies in Economics and Management Sciences Discussion paper 1147R, Northwestern University.*

[3] Banks, Jeffrey and Rangarajan Sundaram (1990):"Repeated Games,Finite Automata and Complexity", *Games and Economic Behavior,***2**, pp 97-117.

[4] Binmore, K.G. and M. Herrero (1988a) 'Matching and Bargaining in Dynamic Markets' *Review of Economic Studies*, 55, 17-32.

[5] Binmore, K.G. and M. Herrero (1988b) 'Security equilibrium' *Review of Economic Studies*, 55, 33-48.

[6] Binmore, K.G., M. Piccione and L. Samuelson (1998): "Evolutionary Stability in Alternating Offers Bargaining Games", *Journal of Economic Theory, 80,2, pp 257-291.*

[7] Chatterjee, K. and H. Sabourian (1999): " Strategic Complexity and n-person Bargaining," *Penn State working paper* # 5-99-1.

[8] Chatterjee, K. and H. Sabourian (2000): "Multilateral Bargaining and Strategic Complexity," forthcoming; *Econometrica.*

[9] Fudenberg D. and J. Tirole (1991) *Game Theory,* MIT press.

[10] Gale D. (1986a) 'Bargaining and Competition Part I: Characterization' *Econometrica* 54, 785-806.

[11] Gale D. (1986b) 'Bargaining and Competition Part II: Existence' *Econometrica* 54, 807-818.

[12] Gale D. (1987) 'Limits Theorems for Markets with Sequential Bargaining' *Journal of Economic Theory*, 43, 20-54.

[13] Gale D. (2000) *Strategic Foundations of General Equilibrium: Dynamic Matching and Bargaining Games,* Cambridge University Press.

[14] Green E.J. (1980) 'Non-cooperative Price Taking in Large Dynamic Markets' *Journal of Economic Theory* Vol 22.

[15] Kalai, E. (1990):"Bounded Rationality and Strategic Complexity in Repeated Games",*Game Theory and Applications* ,edited by T.Ichiishi, A.Neyman and Y.Tauman,pp 131-157.

[16] Kalai, E. and A. Neme (1992) 'The strength of a little perfection' *International Journal of game theory.*

[17] Levine D.K. and W. Pesendorfer (1995) 'When Are Agents Negligible? *American Economic Review* Vol 85 No 5.

[18] Mclennan A. and H. Sonnenschein (1991) 'Sequential Bargaining as a Non-Cooperative Foundation of Competitive Equilibrium' *Econometrica* 59,1395-1424.

[19] Osborne, M. and A. Rubinstein (1990): *Bargaining and Markets,* Academic Press.

[20] Osborne, M. and A. Rubinstein (1994): *A Course in Game Theory,* MIT Press, Cambridge, Massachusetts.

[21] Neme, Alejandro and Luis Quintas (1995): "Subgame perfect Equilibrium of Repeated Games with Implementation Cost" *Journal of Economic Theory, 66,* pp 599-608.

[22] Papadimitriou, C. (1992): "On Games with a Bounded Number of States", *Games and Economic Behavior,* 4 , pp 122-131.

[23] Pearce D. (1984) 'Rationalizable Strategic Behavior and the Problem of Perfection' *Econometrica* 52, 1029-1050.

[24] Piccione, M. and A. Rubinstein (1993): "Finite Automata Play a Repeated Extensive Form Game," *Journal of Economic Theory,* 61, pp 160-168.

[25] Rubinstein A. and A. Wolinsky (1985) 'Equilibrium in a Market with Sequential Trading' *Econometrica,* 53,1133-1150.

[26] Rubinstein A. and A. Wolinsky (1990) 'Decentralized Trading, Strategic Behaviour and the Walrasian Ou*tcome' Review of Economic Studies, 57.*

[27] Sabourian H. (1990) 'Anonymous Repeated Games with a Large Number of Players and Random Outcomes', *Journal of Economic Theory* Vol 51, No 1.

[28] Sabourian H. (2001a) 'Bargaining and Markets: Complexity and the Competitive Outcome' Mimeo, University of Cambridge.

[29] Sabourian H. (2001b) 'Complexity and One-sided Incomplete Information Bargaining' Mimeo, University of Cambridge.