

Game theory and strategic complexity

Kalyan Chatterjee Hamid Sabourian

May 16, 2008

Author affiliations: Chatterjee, Department of Economics, The Pennsylvania State University, University Park, Pa. USA. Sabourian, Faculty of Economics, University of Cambridge, Cambridge, UK.

Acknowledgements: We wish to thank an anonymous referee and Jihong Lee for valuable comments that improved the exposition of this chapter. We would also like to thank St. John's College, Cambridge and the Pennsylvania State University for funding Dr Chatterjee's stay in Cambridge at the time this chapter was written.

Article Outline

Glossary

1. Definition
2. Introduction
3. Games, automata and equilibrium concepts
4. Complexity considerations in repeated games
5. Complexity and bargaining
6. Complexity, market games and the competitive equilibrium
7. Discussion and future directions

Bibliography

Glossary:

Game theory, a formal model of interaction, usually in human behaviour. *Repeated games*, a series of identical interactions of this kind. *Strategy*, a complete specification of how a player will play the game. *Strategic complexity*, a measure of how complex a strategy is to implement. *Equilibrium*, a solution concept for games in which each player optimises given his correct prediction of others' behaviour. *Equilibrium path*, the outcome in terms of the play of the game if every player uses his equilibrium strategy. *Continuation game*, a description of how the play will proceed in a dynamic game once some part of the game has already occurred. *Automata*, a formal definition of a strategy that captures its complexity.

1 Definition

The subject of this chapter is at the intersection of economics and computer science and deals with the use of measures of complexity obtained from the study of finite automata to help select among multiple equilibria and other outcomes appearing in game-theoretic models of bargaining, markets and repeated interactions. The importance of the topic lies in the ability of concepts that employ bounds on available resources to generate more refined predictions of individual behaviour in markets.

2 Introduction

This chapter is concerned with the concept of strategic complexity and its use in game theory. There are many different meanings associated with the word “complexity”, as the variety of topics discussed in this volume makes clear. In this paper, we shall adopt a somewhat narrow view, confining ourselves to notions that measure, in some way, constraints on the ability of economic agents to behave with full rationality in their interactions with other agents in dynamic environments. This will be made more precise a little later. (A more general discussion is available in Rubinstein [45].)

Why is it important to study the effect of such constraints on economic decision-making? The first reason could be to increase the realism of the assumptions of economic models; it is evident from introspection and from observing others that we do not have infinite memory and cannot condition our future actions on the entire corpus of what we once knew; or, for that matter, unlimited computational power. However, only considering the assumptions of a model would not be considered enough if the increased realism were not to expand our ability to explain or to predict. The second reason therefore is that studying the effects of complexity on human decision-making might help us either to make our predictions more precise (by selecting among equilibria) or to generate explanations for behaviour that is frequently observed, but incompatible with equilibrium in models that have stronger assumptions about the abilities of agents.

A strategy in a game is an entire plan of how to play the game at every possible history/contingency/eventuality at which the player has to make a move. The particular aspect of complexity that we shall focus on is on the complexity of strategy as a function of the history. One representation

of the players' strategies in games is often in terms of (finite) automata. The finiteness need not always be assumed; it can be derived. The ideas of complexity, though often most conveniently represented this way, can also be discussed without referring to finite automata at all but purely in how a strategy depends on the past history of the game.

The number of states in the automaton can be used as a measure of complexity. This may be a natural measure of complexity in a stationary repetitive environment such as repeated games. We shall discuss this measure of complexity as well as other aspects of the complexity of a strategy that are particularly relevant in non-stationary frameworks.

Note the players are not themselves considered automata in this paper and in the literature it surveys. Also, we do not place restrictions on the ability of players to compute strategies (see Papadimitriou (1992)), only on the strategies that they can implement. The paper is also not intended as a comprehensive survey of the literature on complexity of implementation in games. The main focus of the paper is inevitably on the works that we have been personally associated with.

The remaining part of this paper is organised as follows: In the next section, we discuss strategies in a game, their representation as finite automata and the basic equilibrium concepts to be used in the paper. Section 4 will consider the use of complexity notions in repeated games. Section 5 will focus on extensive form bargaining and the effect of complexity considerations in selecting equilibria. Section 6 will extend the analysis of bargaining to markets in which several agents bargain and considers the recent literature that justifies competitive outcomes in market environments by appealing to the aversion of agents to complexity. Section 7 concludes with some thoughts on future research. This paper draws on an earlier survey paper (Chatterjee (2002)) for some of the material in Sections 3, 4 and 5.

3 Games, automata and equilibrium concepts

As mentioned in the introduction, this paper will be concerned with *dynamic games*. Though the theory of games has diffused from economics and mathematics to several other fields in the last few decades, we include an introduction to the basic concepts to keep this paper as self-contained as possible. A game is a formal model of interaction between individual agents. The basic components of a game are: (i) Players or agents, whose choices

will, in general, have consequences for each other. We assume a finite set of players, denoted by N . We shall also use N sometimes to represent the cardinality of this set. (ii) A specification of the “rules of the game” or the structure of interaction, described by the sequence of possible events in the game, the order in which the players move, what they can choose at each move and what they know about previous moves. This is usually modelled as a tree and is called the “extensive form” of the game (and will not be formalised here, though the formalisation is standard and found in all the texts on the subject). (iii) Payoffs for each player associated with every path through the tree from the root. It is easier to describe this as a finite tree and ascribe payoffs to the end nodes z . Let $u_i(z)$ be the real-valued payoff to Player i associated with end node z . The payoffs are usually assumed to satisfy conditions that are sufficient to guarantee that the utility of a probability distribution on a subset of the set of end nodes is the expectation of the utility of the individual end nodes. However, different strands of work on bounded rationality dispense with this assumption. The description above presupposes a tree of finite depth, whilst many of the applications deal with infinite horizon games. However, the definitions are easily modified by associating payoffs with a *play* of the game and defining a node as a set of plays. We shall not pursue this further here.

In the standard model of a game, players are assumed to have all orders of knowledge about the preceding description. Work on bounded rationality also has considered relaxing this assumption.

A strategy is a complete plan of action for playing a game, describing the course of action to be adopted in every possible contingency (or every information set of the player concerned). The plan has to be detailed enough so that it can be played by an agent, even if the principal is not himself or herself in town, and the agent could well be a computer, which is programmed to follow the strategy. Without any loss of generality, a strategy can be represented by an automaton (see below for illustration and Osborne and Rubinstein (1994) for a formal treatment in the context of repeated games). Often such a machine description is more convenient in terms of accounting for a complexity of a machine. For example the works that are based on the use of finite automata or Turing machines to represent *strategies* for playing a game impose a natural bound on the set of allowable strategies.

For the types of problem that we shall consider here, it is best to think of a *multistage game with observable actions*, to use the terminology of Fudenberg and Tirole (1991). The game has some temporal structure; let us call each

unit of time a period or a stage. In each period, the players choose actions simultaneously and independently. (The actions could include the dummy action.) All the actions taken in a stage are observed and the players then choose actions again. An example is a repeated normal form game, such as the famous Prisoners' Dilemma being repeated infinitely or finitely often. In each stage, players choose whether to *cooperate* or *defect*. The choices are revealed, payoffs received and the choices repeated again and so on. (The reader will recall that in the Prisoners' Dilemma played once, *Defect* is better than *Cooperate* for each player, no matter what the other player does, but both players choosing *Defect* is strictly worse for each than both choosing *Cooperate*.) What a strategy for a given player would do would be to specify the choice in a given stage as a function of the history of the game up to that stage (for every stage). A finite automaton represents a particular strategy in the following way: It partitions all possible histories in the game (at which the player concerned has to move) using a finite number of elements. Each of these elements is a *state* of the machine. Given a state, the automaton prescribes an action (for example, Cooperate after all histories in which the other party has cooperated). It also specifies how the state of the machine will change as a result of the action taken by the other player. The state-to-action mapping is called the output mapping and the rule that prescribes the state in the next period as a function of today's state and the action of one's opponent in this period is called the transition mapping. The automaton also needs to prescribe what to do in the first stage, when there is no history of past actions to rely on. Thus for example, the famous 'tit-for-tat' strategy in the repeated Prisoners' Dilemma can be represented by the following automaton.

1. Play Cooperate in the first stage. The initial state is denoted as q^1 and in this state the action prescribed is cooperate.
2. As long as the other player cooperates, stay in state q^1 .
3. If the other player defects in a state, go to state q^2 . The action specified in q^2 is Defect.
4. Stay in q^2 as long as the other player defects. If the other player cooperates in a stage, go to q^1 .

Denoting the output mapping by $\lambda(\cdot)$, we get $\lambda(q^1) = C$ and $\lambda(q^2) = D$. The transition mapping, $\mu(\cdot, \cdot)$ is as follows: $\mu(q^1, C) = q^1$, $\mu(q^1, D) = q^2$, $\mu(q^2, C) = q^1$, $\mu(q^2, D) = q^2$. Here, of course, C and D denote cooperate and defect respectively. The machine described above has two states and is an instance of a Moore machine in computer science terminology.

The use of a Moore machine to represent a strategy rules out strategies in which histories are arbitrarily finitely partitioned or arbitrarily complex. In fact, the number of states in the machine is a popular measure of the complexity of the machine and the strategy it represents.

Another kind of finite automaton used in the literature is a Mealy machine. The main difference between this and the Moore machine is that now the output is a function both of the state and of an input, unlike the Moore machine where it is only a function of the state. One can always transform a Mealy machine to a Moore machine by making transitions depend on the input and having state transitions after every input. The Mealy machine representation is more convenient for the extensive form game we shall consider in section 5. We shall briefly address why in that section.

The aim of using the machine framework to describe strategies is to take into account explicitly the cost of complexity of strategies. There is the belief for instance that short-term memory (see Miller (1956)) is capable of keeping seven things in mind at any given time and if five of them are occupied by how to play the Prisoners' Dilemma there might be less left over for other important activities.

The standard equilibrium concept in game theory is the concept of Nash equilibrium. This requires each player to choose a best strategy (in terms of payoff) given his or her conjectures about other players' strategies and, of course, in equilibrium the conjectures must be correct. Thus, a Nash equilibrium is a profile of strategies, one for each player, such that every player is choosing a best response strategy given the Nash equilibrium strategies of the other players. In dynamic games Nash equilibrium strategies may not be credible (sequentially rational). In multi-stage games, to ensure credibility, the concept of Nash equilibrium is refined by requiring the strategy of each player to be a best response to the strategies of the others at every well-defined history (subgame) within the game. This notion of equilibrium was introduced by Selten ([48]) and is called subgame perfect equilibrium. The difference between this concept and that of Nash, which it refines, is that players must specify strategies that are best responses to each other even at nodes in the game tree that would never be reached if the prescribed equilibrium were being played. The Nash concept does not require this. The notion of histories *off the equilibrium path* therefore refers to those that do not occur if every player follows his or her equilibrium strategy. Another useful concept to mention here is that of *payoff in the continuation game*. This refers to the expected payoff from the prescribed strategies in the part

of the game remaining to be played after some moves have already taken place. The restriction of the prescribed strategies to the continuation game are referred to here as *continuation strategies*.

Rubinstein (1986), Abreu and Rubinstein (1988) and others have modified the standard equilibrium concepts to account for complexity costs. This approach is somewhat different from that adopted, for example, by Neyman (1985), who restricted strategies to those of bounded complexity. We shall next present the Abreu-Rubinstein definition of *Nash equilibrium with complexity* (often referred to as NEC in the rest of the paper).

The basic idea is a very simple extension of Nash equilibrium. Complexity enters the utility function lexicographically. A player first calculates his or her best response to the conjectured strategies of the other players. If there are alternative best responses, the player chooses the less complex one. Thus a Nash equilibrium with complexity has two aspects. First, the strategies chosen by any player must be a best response given his or her conjectures about other players' strategies and, of course, in equilibrium the conjectures must be correct. Second, there must not exist an alternative strategy for a player such that his or her payoff is the same as in the candidate equilibrium strategy, given what other players do, but the alternative strategy is less complex.

In Abreu and Rubinstein ([1]), the measure of complexity is the number of states in the Moore machine that represents the strategy. The second part of their equilibrium definition restricts the extent to which punishments can be used off the equilibrium path. For example, there is a famous strategy that, if used by all players, gives cooperation in the infinitely repeated Prisoners' Dilemma (for sufficiently high discount factors), namely the "grim" strategy. This strategy can be described by the following machine: Start with Cooperate. Play Cooperate as long as the other players all cooperate. If in the last period any player has used Defect, then switch to playing Defect for ever. (That is, never play Cooperate again, no matter what the other players do in succeeding periods.) This strategy profile (each player uses the grim strategy) gives an outcome path consisting solely of players cooperating. No one defects because from then until the end of time all the players will be punishing one another.

However, this strategy profile is not a Nash equilibrium with complexity; the grim strategy is a two-state machine in which one state (the one in which a player chooses Defect) is never used given that everyone else cooperates on the equilibrium path. Some player can do better, even if lexicographically,

by switching to a one-state machine in which he or she cooperates no matter what. Thus even the weak lexicographic requirement has some bite.

Note that the complexity restriction we are considering is on the complexity of *implementation*, not the complexity of *computation*. We know that even a Turing machine, which has potentially infinite memory, might be unable to calculate best responses to all possible strategy profiles of other players in the game (see the papers by Anderlini (1990) and Binmore (1985)).

To return to the question of defining equilibrium in the machine game, the Abreu-Rubinstein approach is described by them as “buying” states in the machine at the beginning of the game. The complexity cost is therefore a fixed cost per state used. Some recent papers have taken the fixed cost approach further by requiring NEC strategies to be credible. The idea is that players pay an initial fixed cost for the complexity (the notion of complexity in some of these papers differ from counting the states approach) of his/her strategy and then the game is played with strategies being optimal at every contingency as in standard game theory. Chatterjee and Sabourian (2000a,b) model this by considering Nash equilibrium with complexity costs in (bargaining) games in which machines/strategies can make errors/trembles in output/action. The introduction of errors ensures that the equilibrium strategies are optimal after every history. As the error goes to zero, we are left with subgame perfect equilibria of the underlying game. Chatterjee and Sabourian (2000b), Sabourian (2004), Gale and Sabourian (2005) and Lee and Sabourian (2007) take a more direct method of introducing credibility into the equilibrium concept with complexity costs by restricting NEC strategies to be subgame perfect equilibrium in the underlying game with no complexity costs. We refer to such an equilibria by perfect equilibrium with complexity costs (PEC).

In contrast to the fixed cost interpretation of complexity cost, Rubinstein in his 1986 paper considers a different approach, namely the choice of “renting” states in the machine for every period the game is played. Formally, the Rubinstein notion of *semi-perfect equilibrium* requires the strategy chosen to have the minimal number of states necessary to play the game at every node on the (candidate) equilibrium outcome path. A state could therefore be dropped if it is not going to be used on the candidate equilibrium path after some period. Thus, to be in the equilibrium machine, it is not sufficient that a state be used on the path, it has to be used in every possible future. Rubinstein called this notion of equilibrium *semi-perfect*, because the complexity of a strategy could be changed in one direction (it could be de-

creased) after every period. If states could be added as well as deleted every period, we would have yet another definition of equilibrium with complexity, *machine subgame perfect equilibrium*. (See Neme and Quintas (1995).) In contrast, both the NEC and PEC concepts we use here entail a *single choice* of automaton or strategy by players at the beginning of the game.

In all these models, complexity analysis has been facilitated by considering the “machine games”. Each player chooses among machines and the complexity of a machine is taken to be the number of states of the machine. In fact, the counting-the-number-of-states measure of complexity has an equivalent measure stated in terms of the underlying strategies that the machine could implement. Kalai and Stanford (1988) define complexity of a strategy by the number of *continuation strategies* that the strategy induces at different periods/histories of the game, and establishes that such a measure is equal to the number of the states of the smallest machine that implements the strategy. Thus, one could equivalently describe any result either in terms of underlying strategies and the cardinality of the set of continuation strategies that they induce or in terms of machines and the number of states in them. The same applies to other measures of complexity discussed in this paper; they can be defined either in terms of the machine specification or in terms of the underlying strategy. In the rest of this paper, to simplify the exposition we shall at times go from one exposition to the other without further explanation.

With this preamble on the concepts of equilibrium used in this literature, we turn to a discussion of a specific game in the next section, the infinitely repeated Prisoners’ Dilemma. We will discuss mainly the approach of Abreu and Rubinstein in this section but contrast it with the literature following from Neyman. We also note that the suggestion for using finite automata in games of this kind came originally from Aumann (1981).

4 Complexity considerations in repeated games

4.1 Endogeneous Complexity

In this subsection we shall first concentrate on the Prisoners’ Dilemma and discuss the work of Abreu and Rubinstein, which was introduced briefly in the last section. For concreteness, consider the following Prisoners’ Dilemma payoffs:

	C ₂	D ₂	
C ₁	3,3	-1,4	.
D ₁	4,-1	0,0	

This is the “stage game”; each of the two players chooses an action in each stage, their actions are revealed at the end of the stage and then the next stage begins. The game is repeated infinitely often and future payoffs are discounted with a common discount factor δ .

The solution concept to be used was introduced in the last section; *NEC* or Nash equilibrium with complexity. Note that here complexity is *endogenous*. A player has a preference for less complex strategies. This preference comes into play lexicographically, that is for any strategies or machines that give the same payoff against the opponent’s equilibrium strategy, a player will choose the one with lowest complexity. Thus the cost of complexity is infinitesimal. One could also consider positive but small costs of more complex strategies, but results will then depend on how large the cost of additional complexity is compared to the additional payoff obtained with a more complex strategy.

We saw in the last section that the “grim trigger” strategy, which is a two-state automaton, is not a NEC. The reason is that if Player 2 uses such a strategy, Player 1 can be better off by deviating to a one-state strategy in which she always cooperates. (This will give the same payoff with a less complex strategy.) One-state strategies where both players cooperate clearly do not constitute NEC (deviating and choosing a one-state machine that always plays *D* is strictly better for a player). However, if both players use a one-state machine that always generates an action of *D*, this is a NEC.

The question obviously arises if the cooperative outcome in each stage can be sustained as a NEC and the preceding discussion makes clear that the answer is no. Punishments have to be used on the equilibrium path, but we can get arbitrarily close to the cooperative outcome for a high enough discount factor. For example consider the following two-state machine:

$$Q = \{q^1, q^2\}; \lambda(q^1) = D, \lambda(q^2) = C,$$

$$\mu(q^1, C) = q^1, \mu(q^1, D) = q^2, \mu(q^2, D) = q^1, \mu(q^2, C) = q^2$$

. Here both players play the same strategy, which starts out playing *D*. If both players do as they are supposed to, each plays *C* in the next period and thereafter, so the sequence of actions is $(D, D), (C, C), (C, C) \dots$ If either player plays *C* in the first period, the other player keeps playing *D* in the

next period. The transition rule prescribes that if one plays C and one's opponent plays D , one goes back to playing D , so the sequence with the deviation will be $(D,C),(D,D),(C,C),(C,C),\dots$

Suppose both players use this machine. First, we check it is a Nash equilibrium in payoffs. We only need to check what happens when a player plays C . If Player 2 deviates and plays D , she will get an immediate payoff of 4 followed by payoffs of 0, 3, 3... if she thereafter sticks to her strategy for a total payoff of $4+\delta^2\frac{3}{1-\delta}$ as opposed to $\frac{3}{1-\delta}$ if she had not deviated. The net gain from deviation is $1-3\delta$, which is negative for $\delta > \frac{1}{3}$. One can check that more complicated deviations are also worse. The second part of the definition needs to be checked as well, so we need to ensure that a player cannot do as well in terms of payoff by moving to a less complex strategy, namely a one-state machine. A one-state machine that always plays C will get the worst possible payoff, since the other machine will keep playing D against it. A one-state machine that plays D will get a payoff of 4 in periods 2,4,...or a total payoff of $\frac{4\delta}{1-\delta^2}$ as against $\frac{3\delta}{1-\delta}$. The second is strictly greater for $\delta > \frac{1}{3}$.

This machine gives a payoff close to 3 per stage for δ close to 1. As $\delta \rightarrow 1$, the payoff of each player goes to 3, the cooperative outcome.

The paper by Abreu and Rubinstein obtains a basic result on the characterisation of payoffs obtained as NEC in the infinitely repeated Prisoners' Dilemma. We recall that the "Folk Theorem" for repeated games tells us that all outcome paths that give a payoff per stage strictly greater for each player than the minmax payoff for that player in the stage game can be sustained by Nash equilibrium strategies. Using endogenous complexity, one can obtain a refinement; now only payoffs on a so-called "cross" are sustainable as NEC. This result is obtained from two observations. First, in any NEC of a two-player game, the number of states in the players' machines must be equal. This follows from the following intuitive reasoning (we refer readers to the original paper for the proofs). Suppose we fix the machine used by one of the players (say Player 1), so that to the other player it becomes part of the "environment". For Player 2 to calculate a best response or an optimal strategy to Player 1's given machine, it is clearly not necessary to partition past histories more finely than the other player has done in obtaining her strategy; therefore the number of states in Player 2's machine need not (and therefore will not, if there are complexity costs) exceed the number in Player 1's machine in equilibrium. The same holds true in the other direction, so

the number of states must be equal. (This does not hold for more than two players.) Another way of interpreting this result is that it restates the result from Markov decision processes on the existence of an optimal “stationary” policy (that is depending only on the states of the environment, which are here the same as the states of the other player’s machine). See also Piccione (1992)

Thus there is a one-to-one correspondence between the states of the two machines. (Since the number of states is finite and the game is infinitely repeated, the machine must visit at least one of the states infinitely often for each player.) One can strengthen this further to establish a one-to-one correspondence between *actions*. Suppose Player 1’s machine has $a_t^1 = a_s^1$, where these denote the actions taken at two distinct periods and states by Player 1, with $a_t^2 \neq a_s^2$ for Player 2. Since the states in t and s are distinct for Player 1 and the actions taken are the same, the transitions must be different following the two distinct states. But then Player 1 does not need two distinct states, he can drop one and condition the transition after, say, s on the different action used by Player 2. (Recall the transition is a function of the state and the opponent’s action.) But then Player 1 would be able to obtain the same payoff with a less complex machine; so the original one could not have been a NEC machine.

Therefore the actions played must be some combination of (C, C) and (D, D) (the correspondence is between the two C s and the two D s) or some combination of (C, D) and (D, C) . (By combination, we mean combination over time. For example, $\{C, C\}$ is played, say, 10 times for every 3 plays of (D, D) .) In the payoff space, sustainable payoffs are either on the line joining $(3,3)$ and $(0,0)$ or on the line joining the payoffs on the other diagonal; hence the evocative name chosen to describe the result—the cross of the two diagonals.

While this is certainly a selection of equilibrium outcomes, it does not go as far as we would wish. We would hope that some equilibrium selection argument might deliver us the co-operative outcome $(3,3)$ uniquely (even in the limit as $\delta \rightarrow 1$), instead of the actual result obtained. There is work that does this, but it uses evolutionary arguments (see Binmore and Samuelson (1992)). An alternative learning argument is used by Maenner (2008). In his model, a player tries to infer what machine is being used by his opponent and chooses the simplest automaton that is consistent with the observed pattern of play as his model of his opponent. A player then chooses a best response to this inference. It turns out complexity is not sufficient to pin

down an inference and one must use optimistic or pessimistic rules to select among the simplest inferences. One of these gives only (D, D) repeated, whilst the other reproduces the Abreu-Rubinstein NEC results. Piccione and Rubinstein (1993) show that the NEC profile of 2-player repeated *extensive form games* is unique if the stage game is one of perfect information. This unique equilibrium involves all players playing their one-shot myopic non-cooperative actions at every stage. This is a strong selection result and involves stage game strategies not being observable (only the path of play is) as well as the result on the equilibrium numbers of states being equal in the two players' machines.

In repeated games with more than two players or with more than two actions at each stage the multiplicity problem may be more acute than just not being able to select uniquely a "cooperative outcome". In some such games complexity by itself may not have any bite and the Folk Theorem may survive even when the players care for the complexity of their strategies. (See Bloise [12] who shows robust examples of two-player repeated games with three actions at each stage such that every individually rational payoff can be sustained as a NEC if players are sufficiently patient.)

4.2 Exogenous complexity

We now consider the different approach taken by Neyman (1985, 1997), Ben Porath(1986, 1993), Zemel (1989) and others. We shall confine ourselves to the papers by Neyman and Zemel on the Prisoners' Dilemma, without discussing the more general results these authors and others have obtained.

Neyman's approach treats complexity as exogenous. Let Player i be restricted to use strategies/automata with the number of states not to exceed m_i . He also considers finitely repeated games, unlike the infinitely repeated games we have discussed up to this point. With the stage game being the Prisoners' Dilemma and the number of repetitions being T (for convenience, this includes the first time the game is played). We can write the game being considered as $G^T(m_1, m_2)$. Note that without the complexity restrictions, the finitely repeated Prisoners' Dilemma has a unique Nash equilibrium outcome path (and a unique subgame perfect equilibrium)- (D, D) in all stages. Thus sustaining cooperation in this setting is obtaining non-equilibrium behaviour, though one that is frequently observed in real life. This approach therefore is an example of bounded rationality being used to explain observed behaviour that is not predicted in equilibrium.

If the complexity restrictions are severe, it turns out that (C, C) in each period is an equilibrium. For this, we need $2 \leq m_1, m_2 \leq T - 1$. To see this consider the grim trigger strategy mentioned earlier-representable as a two-state automaton- and let $T = 3$. Here $\lambda(q^1) = C; \lambda(q^2) = D; \mu(q^1, C) = q^1; \mu(q^1, D) = q^2; \mu(q^2, C \text{ or } D) = q^2$. If each player uses this strategy, (C, C) will be observed. Such a pair of strategies is clearly not a Nash equilibrium-given Player 1's strategy, Player 2 can do better by playing D in stage 3. But if Player 2 defects in the second stage, by choosing a two-state machine where $\mu(q^1, C) = D$, he will gain 1 in the second stage and lose 3 in the third stage as compared to the machine listed above, so he is worse off. But defecting in stage 3 requires an automaton with three states-two states in which C is played and one in which D is played. The transitions in state q^1 will be similar but, if q^2 is the second cooperative state, the transition from q^2 to the defect state will take place no matter whether the other player plays C or D . However, automata with three states violate the constraint that the number of states be no more than 2, so the profitable deviation is out of reach.

Whilst this is easy to see, it is not clear what happens when the complexity is high. Neyman shows the following result: *For any integer k , there exists a T_0 , such that for $T \geq T_0$ and $T^{\frac{1}{k}} \leq m_1, m_2 \leq T^k$, there is a mixed strategy equilibrium of $G^T(m_1, m_2)$ in which the expected average payoff to each player is at least $3 - \frac{1}{k}$.*

The basic idea is that rather than playing (C, C) at each stage, players are required to play a complex sequence of C and D and keeping track of this sequence uses up a sufficient number of states in the automaton so that profitable deviations again hit the constraint on the number of states. But since D cannot be avoided on the equilibrium path, only something close to (C, C) each period can be obtained rather than (C, C) all the time.

Zemel's paper adds a clever little twist to this argument by introducing communication. In his game, there are two actions each player chooses at each stage, either C or D as before and a message to be communicated. The message does not directly affect payoffs as the choice of C or D does. The communication requirements are now made sufficiently stringent, and deviation from them is considered a deviation, so that once again the states "left over" to count up to N are inadequate in number and (C, C) can once again be played in each stage/period. This is an interesting explanation of the rigid "scripts" that many have observed to be followed, for example, in negotiations.

Neyman (1997) surveys his own work and that of Ben Porath (1986, 1993). He also generalises his earlier work on the finitely-repeated Prisoners' Dilemma to show how small the complexity bounds would have to be in order to obtain outcomes outside the set of (unconstrained) equilibrium payoffs in the finitely-repeated, normal-form game (just as (C, C) is not part of an unconstrained equilibrium outcome path in the Prisoners' Dilemma). Essentially, if the complexity permitted grows exponentially or faster with the number of repetitions, the equilibrium payoff sets of the constrained and the unconstrained games will coincide. For sub-exponential growth, a version of the folk-theorem is proved for two-person games. The first result says: *For every game G in strategic form and with m_i being the bound on the complexity of i 's strategy and T the number of times the game G is played, there exists a constant c such that if $m_i \geq \exp(cT)$, then $E(G^T) = E(G^T(m_1, m_2))$ where $E(\cdot)$ is the set of equilibrium payoffs in the game concerned.* The second result, which generalises the Prisoners' Dilemma result already stated, considers a sequence of triples $(m_1(n), m_2(n), T(n))$ for a two-player strategic form game, with $m_2 \geq m_1$ and shows that the lim inf of the set of equilibrium payoffs of the automata game as $n \rightarrow \infty$ includes essentially the strictly individually rational payoffs of the stage game if $m_1(n) \rightarrow \infty$ and $\frac{\log m_1(n)}{T(n)} \rightarrow 0$ as $n \rightarrow \infty$. Thus a version of the folk theorem holds provided the complexity of the players' machines does not grow too fast with the number of repetitions.

5 Complexity and bargaining

5.1 Complexity and the unanimity game

The well-known alternating offers bargaining model of Rubinstein has two players alternating in making proposals and responding to proposals. Each period or unit of time consists of one proposal and one response. If the response is "reject", the player who rejects makes the next proposal but in the following period. Since there is discounting with discount factor δ per period, a rejection has a cost. The unanimity game we consider is a multiperson generalisation of this bargaining game, with n players arranged in a fixed order, say 1,2,3... n . Player 1 makes a proposal on how to divide a pie of size unity among the n people; players 2,3,... n respond sequentially, either accepting or rejecting. If everyone accepts, the game ends. If someone rejects, Player 2 now gets to make a proposal but in the next period. The

responses to Player 2's proposal are made sequentially by Players 3,4,5.... n ,1. If Player i gets a share x_i in an eventual agreement at time t , his payoff is $\delta^{t-1}x_i$.

Avner Shaked had shown in 1986 that the unanimity game had the disturbing feature that all individually rational (that is non-negative payoffs for each player) outcomes could be supported as subgame perfect equilibria. Thus the sharp result of Rubinstein (1982), who found a unique subgame perfect equilibrium in the two-play stood in complete contrast with the multiplicity of subgame perfect equilibria in the multiplayer game.

Shaked's proof had involved complex changes in expectations of the players if a deviation from the candidate equilibrium were to be observed. For example, in the three-player game with common discount factor δ , the three extreme points $(1, 0, 0)$, $(0, 1, 0)$, $(0, 0, 1)$ sustain one another in the following way. Suppose Player 1 is to propose $(0, 1, 0)$, which is not a very sensible offer for him or her to propose, since it gives everything to the second player. If Player 1 deviates and proposes, say, $((1 - \delta)/2, \delta, (1 - \delta)/2)$, then it might be reasoned that Player 2 would have no incentive to reject because in any case he or she can't get more than 1 in the following period and Player 3 would surely prefer a positive payoff to 0. However, there is a counter-argument. In the subgame following Player 1's deviation, Player 3's expectations have been raised so that he (and everyone else, including Player 1) now expect the outcome to be $(0, 0, 1)$, instead of the earlier expected outcome. For sufficiently high discount factor, Player 3 would reject Player 1's insufficiently generous offer. Thus Player 1 would have no incentive to deviate. Player 1 is thus in a bind; if he offers Player 2 less than δ and offers Player 3 more in the deviation, the expectation that the outcome next period will be $(0, 1, 0)$ remains unchanged, so now Player 2 rejects his offer. So no deviation is profitable, because each deviation generates an expectation of future outcomes, an expectation that is confirmed in equilibrium. (This is what equilibrium means.) Summarising, $(0, 1, 0)$ is sustained as follows: Player 1 offers $(0, 1, 0)$, Player 2 accepts any offer of at least 1 and Player 3 any offer of at least 0. If one of them rejects Player 1's offer, the next player in order offers $(0, 1, 0)$ and the others accept. If any proposer, say Player 1, deviates from the offer $(0, 1, 0)$ to (x_1, x_2, x_3) the player with the lower of $\{x_2, x_3\}$ rejects. Suppose it is Player i who rejects. In the following period, the offer made gives 1 to Player i and 0 to the others, and this is accepted.

Various attempts were made to get around the continuum of equilibria problem in bargaining games with more than two players; most of them

involved changing the game. (See [15], [16] for a discussion of this literature.) An alternative to changing the game might be to introduce a cost for this additional complexity, in the belief that players who value simplicity will end up choosing simple, that is history independent, strategies. This seems to be a promising approach because it is clear from Shaked's construction that the large number of equilibria results from the players choosing strategies that are history-dependent. In fact, if the strategies are restricted to those that are history-independent (also referred to as *stationary* or *Markov*) then it can be shown (see Herrero 1985) that the subgame perfect equilibrium is unique and induces equal division of the pie as $\delta \rightarrow 1$.

The two papers ([15], [16]) in fact seek to address the issue of complex strategies with players having a preference for simplicity, just as in Abreu and Rubinstein. However, now we have a game of more than two players, and a single extensive form game rather than a repeated game as in Abreu-Rubinstein. It was natural that the framework had to be broadened somewhat to take this into account.

For each of n players playing the unanimity game, we define a machine or an implementation of the strategy as follows.

A *stage* of the game is defined to be n periods, such that if a stage were to be completed, each player would play each role at most once. A *role* could be as proposer or $n - 1^{th}$ responder or $n - 2^{th}$ responderupto first responder (the last role would occur in the period before the player concerned had to make another proposal). An *outcome* of a stage is defined as a sequence of offers and responses, for example $e = (x, A, A, R; y, R; z, A, R; b, A, A, A)$ in a four-player game where the (x, y, z, b) are proposals made in the four periods and (A, R) refer to accept and reject respectively. From the point of view of the first player to propose (for convenience, let's call him Player 1), he makes an offer x , which is accepted by Players 2 and 3 but rejected by Player 4. Now it is Player 2's turn to offer, but this offer, y , is rejected by the first responder Player 3. Player 1 gets to play as second responder in the next period, where he rejects Player 3's proposal. In the last period of this stage, a proposal b is made by Player 4 and everyone accepts (including Player 1 as first responder). Any partial history within a stage is denoted by s , For example, when Player 2 makes an offer, he does so after a partial history $s = (x, A, A, R)$. Let the set of possible outcomes of a stage be denoted by E and the set of possible partial histories by S . Let Q_i denote the set of states used in the i^{th} player's machine M_i . The output mapping is given by $\lambda_i : S \times Q_i \rightarrow \Lambda$, where Λ is the set of possible actions (that is the set of

possible proposals, plus accept or reject). The transition between states now takes place at the end of each stage, so the transition mapping is given as $\mu_i : E \times Q_i \rightarrow Q_i$. As before, in the Abreu-Rubinstein setup, there is an initial state $q_{initial,i}$ specified for each player. There is also a termination state F , which is supposed to indicate agreement. Once in the termination state, players will play the null action and make transitions to this state.

Note that our formulation of a strategy naturally uses a Mealy machine. The output mapping $\lambda_i(\cdot, \cdot)$ has two arguments, the state of the machine and the *input* s , which lists the outcomes of previous moves within the stage. The transitions take place at the end of the stage. The benefit of using this formulation is that the continuation game is the same at the beginning of each stage. In Chatterjee and Sabourian (2000a), we investigate the effects of modifying this formulation, including studying the effects of having a sub-machine to play each role. The different formulations can all implement the same strategies, but the complexities in terms of various measures could differ. We refer the reader to that paper for details, but emphasise that in the general unanimity game, the results from other formulations are similar to the one developed here, though they could differ for special cases, like three-player games.

We now consider a machine game, where players first choose machines and then the machines play the unanimity game in analogy with Abreu-Rubinstein. Using the same lexicographic utility, with complexity coming after bargaining payoffs, what do we find for Nash equilibria of the machine game?

As it turns out, the addition of complexity costs in this setting has some bite but not much. In particular, any division of the pie can be sustained in some Nash equilibrium of the machine game. Perpetual disagreement can, in fact, be sustained by a stationary machine, that is one that makes the same offers and responses each time, irrespective of past history. Nor can we prove, for general n – *player* games that the equilibrium machines will be one-state. (A three-player counter-example exists in [15]; it does not appear to be possible to generate in games that lasted less than thirty periods.) For two-player games, the result that machines must be one-state in equilibrium can be shown neatly ([15]); another illustration that in this particular area, there is a substantial increase of analytical difficulty in going from two to three players.

One reason why complexity does not appear important here is that the definition of complexity used is too restrictive. Counting the number of states

is fine, so long as we don't consider how complex a response might be for partial histories *within* a stage. The next attempt at a solution is based on this observation.

We devise the following definition of complexity: Given the machine and the states, if a machine made the same response to different partial stage histories in different states and another machine made different responses, then the second one was more complex (given that the machines were identical in all other respects). We refer to this notion as *response complexity*. (In [16] the concept of response complexity is in fact stated in terms of the underlying strategy rather than in terms of machines.) It captures the intuition that counting states is not enough; two machines could have the same number of states, for example because each generated the same number of distinct offers, but the complexity of responses in one machine could be much lower than that in the other. Note that this notion would only arise in extensive-form games. In normal form games, counting states could be an adequate measure of complexity. Nor is this notion of complexity derivable from notions of transition complexity, due to Banks and Sundaram, for example, which also apply in normal-form games.

The main result of Chatterjee and Sabourian (2000b) is that this new aspect of complexity enables us to limit the amount of delay that can occur in equilibrium and hence to infer that only one-state machines are equilibrium machines.

The formal proofs using two different approaches are available in Chatterjee and Sabourian (2000a,b). We mention the basic intuition behind these results. Suppose, in the three player game, there is an agreement in period 4 (this is in the second stage). Why doesn't this agreement take place in period 1 instead? It must be because if the same offer and responses are seen in period 1 some player will reject the offer. But of course, he or she does not have to do so because the required offer never happens. But a strategy that accepts the offer in period 4 and rejects it *off the equilibrium path in period 1 must be more complex, by our definition*, than one that always accepts it whenever it might happen, on or off the expected path. Repeated application of this argument by backwards induction gives the result. (The details are more complicated but are in the papers cited above.) Note that this uses the definition that two machines might have the same number of states and yet one could be simpler than the other. It is interesting, as mentioned earlier, that for two players one can obtain an analogous result without invoking the response simplicity criterion, but from three players on this criterion is

essential.

The above result (equilibrium machines have one state each and there are no delays beyond the first stage) is still not enough to refine the set of equilibria to a single allocation. In order to do this, we consider machines that can make errors/trembles in output. As the error goes to zero, we are left with perfect equilibria of our game. With one-state machines, the only subgame perfect equilibria are the ones that give equal division of the pie as $\delta \rightarrow 1$. Thus a combination of two techniques, one essentially recognising that players can make mistakes and the other that players prefer simpler strategies if the payoffs are the same as those given by a more complex strategy, resolves the problem of multiplicity of equilibria in the multiperson bargaining game.

As we mentioned before, the introduction of errors ensures that the equilibrium strategies are credible at every history. We could also take the more direct (and easier) way of obtaining the uniqueness result with complexity costs by considering NEC strategies that are subgame perfect in the underlying game (PEC) (as done in [16]). Then since an history-independent subgame perfect equilibrium of the game is unique and any NEC automaton profile has one state and hence is history-independent, it follows immediately that any PEC is unique and induces equal division as $\delta \rightarrow 1$.

5.2 Complexity and repeated negotiations

In addition to standard repeated games or standard bargaining games, multiplicity of equilibria often appear in dynamic repeated interactions, where a repeated game is superimposed on an alternating offers bargaining game. For instance, consider two firms, in an ongoing vertical relationship, negotiating the terms of a merger. Such situations have been analyzed in several “negotiation models” by Busch and Wen (1995), Fernandez and Glazer (1991) and Haller and Holden (1990). These models can be interpreted as combining the features of both repeated and alternating-offers bargaining games. In each period, one of the two players first makes an offer on how to divide the total available periodic (flow) surplus; if the offer is accepted, the game ends with the players obtaining the corresponding payoffs in the current and every period thereafter. If the offer is rejected, they play some normal form game to determine their flow payoffs for that period and then the game moves on to the next period in which the same play continues with the players’ bargaining roles reversed. One can think of the normal form game played in the

event of a rejection as a “threat game” in which a player takes actions that could punish the other player by reducing his total payoffs.

If the bargaining had not existed, the game would be a standard repeated normal form game. Introducing bargaining and the prospect of permanent exit, the negotiation model still admits a large number of equilibria, like standard repeated games. Some of these equilibria involve delay in agreement (even perpetual disagreement) and inefficiency, while some are efficient.

Lee and Sabourian (2007) apply complexity considerations to this model. As in Abreu and Rubinstein (1988) and others, the players choose among automata and the equilibrium notion is that of NEC and PEC. One important difference however is that in this paper the authors do not assume the automata to be *finite*. Also, the paper introduces a new machine specification that formally distinguishes between the two *roles* - proposer and responder - played by each player in a given period.

Complexity considerations select only *efficient* equilibria in the negotiation model players are sufficiently patient. First, it is shown that if an agreement occurs in some finite period as a NEC outcome then it must occur within the first two periods of the game. This is because if a NEC induces an agreement beyond the first two periods then one of the players must be able to drop the last period’s state of his machine without affecting the outcome of the game. Second, given sufficiently patient players, every PEC in the negotiation model that induces perpetual disagreement is at least *long-run* almost efficient; that is, the game must reach a finite date at which the continuation game then on is almost efficient.

Thus, these results take the study of complexity in repeated games a step further from the previous literature in which complexity or bargaining alone has produced only limited selection results. While, as we discussed above, many inefficient equilibria survive complexity refinement, Lee and Sabourian (2007) demonstrate that complexity and bargaining in tandem ensure efficiency in repeated interactions. Complexity considerations also allow Lee and Sabourian to highlight the role of transaction costs in the negotiation game. Transaction costs take the form of paying a cost to enter the bargaining stage of the negotiation game. In contrast to the efficiency result in the negotiation game with complexity costs, Lee and Sabourian also show that introducing transaction costs into the negotiation game dramatically alters the selection result from efficiency to inefficiency. In particular, they show that, for any discount factor and any transaction cost, every PEC in the costly negotiation game induces perpetual disagreement if the stage game

normal form (after any disagreement) has a unique Nash equilibrium.

6 Complexity, market games and the competitive equilibrium

There has been a long tradition in economics of trying to provide a theory of how a competitive market with many buyers and sellers operates. The concept of competitive (Walrasian) equilibrium (see Debreu (1959)) is a simple description of such markets. In such an equilibrium each trader chooses rationally the amount he wants to trade taking the prices as given, and the prices are set (or adjust) to ensure that total demanded is equal to the total supplied. The important feature of the set-up is that agents assume that they cannot influence (set) the prices and this is often justified by appealing to the idea that each individual agent is small relative to the market.

There are conceptual as well as technical problems associated with such a justification. First, if no agent can influence the prices then who sets them? Second, even in a large but finite market a change in the behaviour of a single individual agent may affect the decisions of some others, which in turn might influence the behaviour of some other agents and so on and so forth; thus the market as a whole may end up being affected by the decision of a single individual.

Game theoretic analyses of markets have tried to address these issues (e.g. see [21] and [47]). This has turned out to be a difficult task because the strategic analyses of markets, in contrast to the simple and elegant model of competitive equilibrium, tends to be complex and intractable. In particular, dynamic market games have many equilibria, in which a variety of different kinds of behaviour are sustained by threats and counter-threats.

More than 60 years ago Hayek (1945) noted the competitive markets are simple mechanisms in which economic agents only need to know their own endowments, preferences and technologies and the vector of prices at which trade takes place. In such environments, economic agents maximizing utility subject to constraints make efficient choices in equilibrium. Below we report some recent work, which suggests that the converse might also be true: *“If rational agents have, at least at the margin, an aversion to complex behaviour, then their maximizing behavior will result in simple behavioral rules and thereby in a perfectly competitive equilibrium”* (Gale and Sabourian

(2005)).

6.1 Homogeneous markets

In a seminal paper, Rubinstein and Wolinsky (1990), henceforth RW, considered a market for a single indivisible good in which a finite number of homogeneous buyers and homogeneous sellers are matched in pairs and bargain over the terms of trade. In their set-up, each seller has one unit of an indivisible good and each buyer wants to buy at most one unit of the good. Each seller's valuation of the good is 0 and each buyer's valuation is 1. Time is divided into discrete periods and at each date, buyers and sellers are matched randomly in pairs and one member of the pair is randomly chosen to be the proposer and the other the responder. In any such match the proposer offers a price $p \in [0, 1]$ and the responder accepts or rejects the offer. If the offer is accepted the two agents trade at the agreed price p and the game ends with the seller receiving a payoff p and the buyer in the trade obtaining a payoff $1 - p$. If the offer is rejected the pair return to the market and the process continues. RW further assume that there is no discounting to capture the idea that there is no friction (cost to waiting) in the market.

Assuming that the number of buyers and sellers is not the same, RW showed that this dynamic matching and bargaining game has, in addition to a perfectly competitive outcome, a large set of other subgame perfect equilibrium outcomes, a result reminiscent of the Folk Theorem for repeated games. To see the intuition for this, consider the case in which there is one seller s and many buyers. Since there are more buyers than sellers the price of 1, at which the seller receives all the surplus, is the unique competitive equilibrium; furthermore, since there are no frictions $p = 1$ seems to be the most plausible price. RW's precise result, however, establishes that for any price $p^* \in [0, 1]$ and any buyer b^* there is a subgame perfect equilibrium that results in s and b^* trading at p^* . The idea behind the result is to construct an equilibrium strategy profile such that buyer b^* is identified as the intended recipient of the good at a price p^* . This means that the strategies are such that (i) when s meets b^* , whichever is chosen as the proposer offers price p^* and the responder accepts, (ii) when s is the proposer in a match with some buyer $b \neq b^*$, s offers the good at a price of $p = 1$ and b rejects and (iii) when a buyer $b \neq b^*$ is the proposer he offers to buy the good at a price of $p = 0$ and s rejects. These strategies produce the required outcome. Furthermore, the equilibrium strategies make use of the following punishment strategies

to deter deviations. If the seller s deviates by proposing to a buyer b a price $p \neq p^*$, b rejects this offer and the play continues with b becoming the intended recipient of the item at a price of zero. Thus, after rejection by b strategies are the same as those given earlier with the price zero in place of p^* and buyer b in place of buyer b^* . Similarly, if a buyer b deviates by offering a price $p \neq p^*$ then the seller rejects, another buyer $b' \neq b$ is chosen to be the intended recipient and the price at which the unit is traded changes to 1. Further deviations from these punishment strategies can be treated in an exactly similar way.

The strong impression left by RW is that indeterminacy of equilibrium is a robust feature of dynamic market games and, in particular, there is no reason to expect the outcome to be perfectly competitive. However, the strategies required to support the family of equilibria in RW are quite complex. In particular, when a proposer deviates, the strategies are tailor-made so that the responder is rewarded for rejecting the deviating proposal. This requires coordinating on a large amount of information so that at every information set the players know (and agree) what constitutes a deviation.

In fact, RW show that if the amount of information available to the agents is strictly limited so that the agents do not recall the history of past play then the only equilibrium outcome is the competitive one. This suggests that the competitive outcome may result if agents use simple strategies. Furthermore, the equilibrium strategies used described in RW to support non-competitive outcomes are particularly unattractive because they require all players, including those buyers who do not end up trading, to follow complex non-stationary strategies in order to support a non-competitive outcome. But buyers who do not trade and receive zero payoff on the equilibrium path could always obtain at least zero by following a less complex strategy than the ones specified in RW's construction. Thus, RW's construction of non-competitive equilibria is not robust if players prefer, at least at the margin, a simpler strategy to a more complex one.

Following the above observation, Sabourian (2003), henceforth S, addresses the role of complexity (simplicity) in sustaining a multiplicity of non-competitive equilibria in RW's model. The concept of complexity in S is similar to that in Chatterjee and Sabourian (2000b). It is defined by a partial ordering on the set of individual strategies (or automata) that very informally satisfies the following: if two strategies are otherwise identical except that in some role the second strategy uses more information than that available in the current period of bargaining and the first uses only the infor-

mation available in the current period, then the second strategy is said to be more complex than the first. S also introduces complexity costs lexicographically into the RW game and shows that any PEC is history-independent and induces the competitive outcome in the sense that all trades take place at the unique competitive price of 1.

Informally, S's conclusions *in the case of a single seller s and many buyers* follows from the following three steps. First, since trading at the competitive price of 1 is the worst outcome for a buyer and the best outcome for the seller, by appealing to complexity type reasoning it can be shown that in any NEC a trader's response to a price offer of 1 is always history-independent and thus he either always rejects 1 or always accepts 1. For example, if in the case of a buyer this were not the case, then since accepting 1 is a worst possible outcome, he could economise on complexity and obtain at least the same payoff by adopting another strategy that is otherwise the same as the equilibrium strategy except that it always rejects 1.

Second, in any non-competitive NEC in which s receives a payoff of less than 1, there cannot be an agreement at a price of 1 between s and a buyer at *any* history. For example, if at some history, a buyer is offered $p = 1$ and he accepts then by the first step the buyer should accept $p = 1$ whenever it is offered; but this is a contradiction because it means that the seller can guarantee himself an equilibrium payoff of one by waiting until he has a chance to make a proposal to this buyer.

Third, in any non-competitive PEC the continuation payoffs of all buyers are positive at every history. This follows immediately from the previous step because if there is no trade at $p = 1$ at *any* history it follows that each buyer can always obtain a positive payoff by offering the seller more than he can obtain in any subgame.

Finally, because of competition between the buyers (there is one seller and many buyers), in any subgame perfect equilibrium there must be a buyer with a zero continuation payoff after some history. To illustrate the basic intuition for this claim, let m be the worst continuation payoff for s at any history and suppose that there exists a subgame at which s is the proposer in a match with a buyer \bar{b} and the continuation payoff of s at this subgame is m . Then if at this subgame s proposes $m + \epsilon$ ($\epsilon > 0$), \bar{b} must reject (otherwise s can get more than m). Since the total surplus is 1, \bar{b} must obtain at least $1 - m - \epsilon$ in the continuation game in order to reject s 's offer and s gets at least m , this implies that the continuation payoff of all $b \neq \bar{b}$ after \bar{b} 's rejection is less than ϵ . The result follows by making ϵ arbitrarily small (and by appealing to

the finiteness of f).

But the last two claims contradict each other unless the equilibrium is competitive. This establishes the result for the case in which there is one seller and many buyers. The case of a market with more than one seller is established by induction on the number of sellers.

The matching technology in the above model is random. RW also consider another market game with the matching is endogenous: at each date each seller (the short side of the market) chooses his trading partner. Here, they show that non-competitive outcomes and multiplicity of equilibria survive even when the players discount the future. By strengthening the notion of complexity S also shows that in the endogenous matching model of RW the competitive outcome is the only equilibrium if complexity considerations are present.

These results suggest perfectly competitive behavior may result if agents have, at least at the margin, preferences for simple strategies. Unfortunately, both RW and S have too simple a market set-up; for example, it is assumed that the buyers are all identical, similarly for the sellers and each agent trades at most one unit of the good. Do the conclusions extend to richer models of trade?

6.2 Heterogeneous markets

There are good reasons to think that it may be too difficult (or even impossible) to establish a similar set of conclusions as in S in a richer framework. For example, consider a heterogeneous market for a single indivisible good, where buyers (and sellers) have a range of valuations of the good and each buyer wants at most one unit of the good and each seller has one unit of the good for sale. In this case the analysis of S will not suffice. First, in the homogeneous market of RW, except for the special case where the number of buyers is equal to the number of sellers, the competitive equilibrium price is either 0 or 1 and all of the surplus goes to one side of the market. S's selection result crucially uses this property of the competitive equilibrium. By contrast, in a heterogeneous market, in general there will be agents receiving positive payoffs on both sides of the market in a competitive equilibrium. Therefore, one cannot justify the competitive outcome simply by focusing on extreme outcomes in which there is no surplus for one party from trade. Second, in a homogeneous market individually rational trade is by definition efficient. This may not be the case in a heterogeneous market (an inefficient

trade between inframarginal and an extramarginal agent can be individually rational). Third, in a homogeneous market, the set of competitive prices remains constant, independently of the set of agents remaining in the market. In the heterogeneous market, this need not be so and in some cases, the new competitive interval may not even intersect the old one. The change in the competitive interval of prices as the result of trade exacerbates the problems associated with using an induction hypothesis because here future prices may be conditioned on past trades even if prices are restricted to be competitive ones.

Despite these difficulties associated with a market with a heterogeneous set of buyers and sellers, Gale and Sabourian (2005), henceforth GS, show that the conclusions of S can be extended to the case of a heterogeneous market in which each agent trades at most one unit of the good. GS, however, focus on *deterministic* sequential matching models in which one pair of agents are matched at each date and they leave the market if they reach an agreement. In particular, they start by considering exogenous matching processes in which the identities of the proposer and responder at each date are an exogenous and deterministic function of the set of agents remaining in the market and the date. The main result of the paper is that a PEC is always competitive in such a heterogeneous market, thus supporting the view that competitive equilibrium may arise in a finite market where complex behavior is costly.

The notion of complexity in GS is similar to that in S (and Chatterjee and Sabourian (2000b)). However, in the GS set-up with heterogeneous buyers and sellers the set of remaining agents changes depending who has traded and left the market and who is remaining, and this affects the market conditions. (In the homogeneous case, only the *number* of remaining agents matters.) Therefore, the definition of complexity in GS is with reference to a given set of remaining agents. GS also discuss an alternative notion of complexity that is independent of the set of remaining agents; such a definition may be too strong and may result in an equilibrium set being empty.

To show their result, GS first establish two very useful restrictions on the strategies that form a NEC (similar to the no delay result in Chatterjee and Sabourian (2000b)). First, they show that if along the equilibrium path a pair of agents k and ℓ trade at a price p with k as the proposer and ℓ as the responder then k and ℓ always trade at p , irrespective of the previous history, whenever the two agents are matched in the same way with the same remaining set of agents. To show this consider first the case of the

responder ℓ . Then it must be that at every history with the same remaining set of agents ℓ always accepts p by k . Otherwise, ℓ could economise on complexity by choosing another strategy that is otherwise identical to his equilibrium strategy except that it always accepts p from k without sacrificing any payoff: such a change of behaviour is clearly more simple than sometimes accepting and sometimes rejecting the offer and moreover, it results in either agent k proposing p and ℓ accepting, so the payoff to agent ℓ is the same as from the equilibrium strategy, or agent k not offering p , in which case the change in the strategy is not observed and the play of the game is unaffected by the deviation. Furthermore, it must also be that at every history with the same remaining set of agents agent k proposes p in any match with ℓ . Otherwise, k could economise on complexity by choosing another strategy that is otherwise identical to his equilibrium strategy except that it always proposes p to ℓ without sacrificing any payoff on the equilibrium path: such a change of behaviour is clearly more simple and moreover k 's payoff is not affected because either agent k and ℓ are matched and k proposes p and ℓ by the previous argument accepts, so the payoff to agent k is the same as from the equilibrium strategy, or agent k and ℓ are not matched with k as the proposer, in which case the change in the strategy is not observed and the play of the game is unaffected by the deviation.

GS's shows a second restriction, again with the same remaining set of agents, namely that in any NEC for any pair of agents k and ℓ , player ℓ 's response to k 's (on or off-the-equilibrium path) offer is always the same. Otherwise, it follows that ℓ sometimes accepts an offer p by k and sometimes rejects (with the same remaining set of agents). Then by the first restriction it must be that if such an offer is made by k to ℓ on the equilibrium path it is rejected. But then ℓ can could economise on complexity by always rejecting p by k without sacrificing any payoff on the equilibrium path: such a change of behaviour is clearly more simple and furthermore ℓ 's payoff is not affected because such a behaviour is the same as what the equilibrium strategy prescribes on the equilibrium path.

By appealing to the above two properties of NEC and to the competitive nature of the market GS establish, using a complicated induction argument, that every PEC induces a competitive outcome in which each trade occurs at the same competitive price.

The matching model we have described so far is deterministic and exogenous. The selection result of GS however extends to richer deterministic matching models. In particular, GS also consider a semi-endogenous sequen-

tial matching model in which the choice of partners is endogenous but the identity of the proposer at any date is exogenous. Their results extends to this variation, with an endogenous choice of responders. A more radical departure change would be to consider the case where at any date any agent can choose his partner and make a proposal. Such a totally endogenous model of trade generates new conceptual problems. In a recent working paper Gale and Sabourian (2008) consider a continuous time version of such a matching model and show that complexity considerations allows one to select a competitive outcome in the case of totally endogenous matching. Since the selection result holds for all the different matching models we can conclude that complexity considerations inducing a competitive outcome seem to be a robust result in deterministic matching and bargaining market games with heterogeneous agents.

Random matching is commonly used in economic models because of its tractability. The basic framework of GS, however, does not extend to such a framework if either the buyers or the sellers are not identical. This is for two different reasons. First, in general in any random framework there is more than one outcome path that can occur in equilibrium with a positive probability; as a result introducing complexity lexicographically may not be enough to induce agents to behave in a simple way (they will have to be complex enough to play optimally along all paths that occur with a positive probability). Second, in Gale and Sabourian (2006) it is shown that subgame perfect equilibria in Markov strategies are not necessarily perfectly competitive for the random matching model with heterogeneous agents. Since the definition of complexity in GS is such that Markov strategies are the least complex ones, it follows that with random matching the complexity definition used in GS is not sufficient to select a competitive outcome.

6.3 Complexity and off-the-equilibrium path play

The concept of the PEC (or NEC) used in S, GS and elsewhere was defined to be such that for each player the strategy/automaton has minimal complexity amongst all strategies/automata that are best responses to the equilibrium strategies/automata of others. Although, these concepts are very mild in the treatment of complexity, it should be noted that there are other ways of introducing complexity into the equilibrium concept. One extension of the above set-up is to treat complexity as a (small) positive fixed cost of choosing a more complex strategy and define a Nash (subgame perfect) equilibrium

with a fixed positive complexity costs accordingly. All the selection results based on lexicographic complexity in the papers we discuss in this survey also hold for positive small complexity costs. This is not surprising because with positive costs complexity has at least as much bite as in the lexicographic case; there is at least as much refinement of the equilibrium concept with the former as with the latter. In particular, in the case of a NEC (or a PEC), *in considering complexity*, players ignore any consideration of payoffs off the equilibrium path and the trade-off is between the equilibrium payoffs of two strategies and the complexity of the two. As a result these concepts put more weight on complexity costs than on being “prepared” for off-the-equilibrium-path moves. Therefore, although complexity costs are insignificant, they take priority over optimal behavior after deviations. (See [15] for a discussion.)

A different approach would be to assume that complexity is a less significant criterion than the off-the-equilibrium payoffs. In the extreme case, one would require agents to choose minimally complex strategies among the set of strategies that are best responses on and off the equilibrium path (see Kalai and Neme (1992)).

An alternative way of illustrating the differences between the different approaches is by introducing two kinds of vanishingly small perturbations into the underlying game. One perturbation is to impose a small but positive cost of choosing a more complex strategy. Another perturbation is to introduce a small but positive probability of making an error (off-the-equilibrium-path move). Since a PEC requires each agents to choose a minimally complex strategy within the set of best responses, it follows that the limit points of Nash equilibria of the above perturbed game correspond to the concept of PEC if we first let the probability of making an off-the-equilibrium-path move go to zero and then let the cost of choosing a more complex strategy go to zero (this is what Chatterjee and Sabourian (2000b) do). On the other hand, in terms of the above limiting arguments, if we let the cost of choosing a more complex strategy go to zero and then let the probability of making an off-the-equilibrium-path move go to zero then any limit corresponds to the equilibrium definition in Kalai and Neme (1992) where agents choose minimally complex strategies among the set of strategies that are best responses on and off the equilibrium path.

Most of the results reported in this paper on refinement and endogenous complexity (for example Abreu-Rubinstein (1988), Chatterjee and Sabourian (2000b), Gale and Sabourian (2005) and Lee and Sabourian (2007) hold only for the concept of NEC and its variations and thus depend crucially on

assuming that complexity costs are more important than off-the-equilibrium payoffs. This is because these results always appeal to an argument that involves economising on complexity if the complexity is not used off the equilibrium path. Therefore, they may be a good predictor of what may happen only if complexity costs are more significant than the perturbations that induce off-the-equilibrium-path behaviour. The one exception is the selection result in S [47]. Here, although the result we have reported is stated for NEC and its variations, it turns out that the selection of competitive equilibrium does not in fact depend on the relative importance of complexity costs and off-the-equilibrium path payoffs. It remains true even for the case where the strategies are required to be least complex amongst those that are best responses at every information set. This is because in S's analysis complexity is only used to show that every agent's response to the price offer of 1 is always the same irrespective of the past history of play. This conclusion holds irrespective of the relative importance of complexity costs and off-the-equilibrium payoff because trading at the price of 1 is the best outcome that any seller can achieve at any information set (including those off-the-equilibrium) and a worst outcome for any buyer. Therefore, irrespective of the order, the strategy of sometimes accepting a price of 1 and sometimes rejecting cannot be an equilibrium for a buyer (similar arguments applies for a seller) because the buyer can economise on complexity by always rejecting the offer without sacrificing any payoff off or on-the-equilibrium path (accepting $p = 1$ is a worse possible outcome).

7 Discussion and future directions

The use of finite automata as a model of players in a game has been criticised as being inadequate, especially because as the number of states becomes lower it becomes more and more difficult for the small automaton to do routine calculations, let alone the best response calculations necessary for game-theoretic equilibria. Some of the papers we have explored address other aspects of complexity that arise from the concrete nature of the games under consideration. Alternative models of complexity are also suggested, such as computational complexity and communication complexity.

While our work and the earlier work on which it builds focuses on equilibrium, an alternative approach might seek to see whether simplicity evolves in some reasonable learning model. Maenner (2008) has undertaken such an

investigation with the infinitely repeated Prisoners' Dilemma (studied in the equilibrium context by Abreu and Rubinstein). Maenner provides an argument for "learning to be simple". On the other hand, there are arguments for increasing complexity in competitive games ([42]). It is an open question, therefore, whether simplicity could arise endogenously through learning, though it seems to be a feature of most human preferences and aesthetics (see [11]).

The broader research programme of explicitly considering complexity in economic settings might be a very fruitful one. Auction mechanisms are designed with an eye towards how complex they are – simplicity is a desideratum. The complexity of contracting has given rise to a whole literature on incomplete contracts, where some models postulate a fixed cost per contingency described in the contract. All this is apart from the popular literature on complexity, which seeks to understand complex, adaptive systems from biology. The use of formal complexity measures such as those considered in this survey and the research we describe might throw some light on whether incompleteness of contracts, or simplicity of mechanisms, is an assumption or a result (of explicitly considering choice of level of complexity).

References

- [1] Abreu, Dilip and Ariel Rubinstein (1988): "The Structure of Nash Equilibria in Repeated Games with Finite Automata," *Econometrica*, **56**, pp 1259-1282.
- [2] Anderlini, Luca (1990): "Some notes on Church's Thesis and the Theory of Games", *Theory and Decision*, 29, 19-52
- [3] Anderlini, Luca and Hamid Sabourian (1995): "Cooperation and Effective Computability", *Econometrica*, 63, 1337-1369
- [4] Aumann, R.J. (1981), "Survey of repeated games" in *Essays in Game Theory and Mathematical Economics in Honor of Oskar Morgenstern*, 11-42, Mannheim/Vienna/Zurich Bibliographisches Institute
- [5] Banks, J. and R. Sundaram (1990): "Repeated games, finite automata and complexity", *Games and Economic Behaviour*, 2, 97-117

- [6] Ben Porath, E. (1986), "Repeated games with bounded complexity", mimeo Stanford University.
- [7] Ben Porath, E. (1993), "Repeated games with finite automata", *Journal of Economic Theory*, **59**, 17-32.
- [8] Binmore, Kenneth G. (1987): "Modelling rational players I", *Economics and Philosophy*, 3 (1987), 179 - 214.
- [9] Binmore, K. and L. Samuelson, (1992), "Evolutionary stability in repeated games played by finite automata", *Journal of Economic Theory*, **57**, 278-305.
- [10] Binmore, K., M. Piccione, and L. Samuelson, (1998), "Evolutionary stability in alternating-offers bargaining games", *Journal of Economic Theory*, 80 (1998), 257-291.
- [11] Birkhoff, George David (1933), *Aesthetic Measure*, Cambridge, Harvard University Press.
- [12] Bloise, G., *Strategic Complexity and Equilibrium in Repeated Games*, unpublished doctoral dissertation, University of Cambridge, 1998.
- [13] Busch, L-A., and Q. Wen, (1995), "Perfect equilibria in a negotiation model", *Econometrica*, 63 (1995), 545-565.
- [14] Chatterjee, Kalyan (2002), "Complexity of strategies and multiplicity of Nash equilibria", *Group Decision and Negotiation*, **11**, 3, May, 223-230
- [15] Chatterjee, Kalyan and Hamid Sabourian (2000a): "N-person bargaining and strategic complexity," mimeo, University of Cambridge and the Pennsylvania State University.
- [16] Chatterjee, Kalyan and Hamid Sabourian (2000b), "Multiperson bargaining and strategic complexity", *Econometrica*, 68, 1491-1509.
- [17] Debreu, Gerard, (1959), *Theory of Value*, Yale University Press, New Haven and London.
- [18] Fernandez, R., and J. Glazer, (1991) "Striking for a bargain between two completely informed agents", *American Economic Review*, 81 (1991), 240-252.

- [19] Fudenberg, Drew and Eric Maskin (1990), “Evolution and repeated games”, mimeo Harvard and Princeton.
- [20] Fudenberg, Drew and Jean Tirole (1991): *Game Theory*, MIT Press, Cambridge, Mass.
- [21] Gale, Douglas, (2000) *Strategic Foundations of General Equilibrium: Dynamic Matching and Bargaining Games*, Cambridge University Press, Cambridge and New York.
- [22] Gale, Douglas, and H. Sabourian, (2005) “Complexity and competition”, *Econometrica*, **73** , 739-770.
- [23] Gale, Douglas and Hamid Sabourian (2006), “Markov equilibria in dynamic matching and bargaining games”, *Games and Economic Behavior*, **54**, 2, February, 336-352.
- [24] Gale, Douglas and Hamid Sabourian (2008), “Complexity and competition II; endogenous matching”, mimeo New York University and University of Cambridge.
- [25] Haller, H., and S. Holden,(1990) “A letter to the Editor on wage bargaining”, *Journal of Economic Theory*, **52** , 232-236.
- [26] Hayek, F. (1945) “The use of knowledge in society,” *American Economic Review* **35**, 519-530.
- [27] Herrero, Maria (1985), “A Strategic Theory of Market Institutions”, unpublished doctoral dissertation, London School of Economics.
- [28] E. Kalai and A. Neme, (1992), “The strength of a little perfection”, *Int. J. Game Theory* **20**, 335-355.
- [29] Kalai, E., and W. Stanford, (1988), “Finite rationality and interpersonal complexity in repeated games”, *Econometrica*, **56** , 397-410.
- [30] Klemperer, Paul (Ed.) (2000), *The Economic Theory of Auctions*, Edward Elgar.
- [31] Lee, J., and H. Sabourian,(2007), “Coase Theorem, complexity and transaction costs”, *Journal of Economic Theory*, **135**, 214-235.

- [32] Maenner, Eliot (2008), “Adaptation and complexity in repeated games”, *Games and Economic Behavior*, **63**, 1, May, 166-187.
- [33] Miller, George A. (1956), “The magical number seven plus or minus two: Some limits on our capacity to process information”, *The Psychological Review*, **63**, 81-97.
- [34] Neme, Alejandro and Luis Quintas (1995), “Subgame perfect equilibrium of repeated games with implementation cost” *Journal of Economic Theory*, *66*, pp 599-608.
- [35] Neyman, Abraham, (1985), “Bounded complexity justifies cooperation in the finitely-repeated Prisoners’ Dilemma”, *Economics Letters*, **19**, 227-229.
- [36] Neyman, Abraham, (1997), “Cooperation, Repetition and Automata” in
Cooperation: Game-Theoretic Approaches, NATO ASI Series F, Vol. 155, S. Hart and A. Mas-Colell (eds.), Springer-Verlag (1997), pp. 233–255
- [37] Osborne, Martin and Ariel Rubinstein (1990), *Bargaining and Markets*, Academic Press.
- [38] Osborne, Martin and Ariel Rubinstein (1994), *A Course in Game Theory*, MIT Press, Cambridge, Massachusetts.
- [39] Papadimitriou, Christos H. (1992), “On games with a bounded number of states,” *Games and Econ. Behav.* **4** (1992), 122-131.
- [40] Piccione, M., “Finite automata equilibria with discounting”, *Journal of Economic Theory*, *56* (1992), 180-193.
- [41] Piccione, M., and Ariel Rubinstein, “Finite automata play a repeated extensive game”, *Journal of Economic Theory*, *61* (1993), 160-168.
- [42] Robson, Arthur, (2003), “The evolution of rationality and the Red Queen”, *Journal of Economic Theory*, **111**, 1-22.
- [43] Rubinstein, Ariel (1982): “Perfect equilibrium in a bargaining model”, *Econometrica*, **50**, pp 97-109 .

- [44] Rubinstein, Ariel (1986): “Finite automata play the repeated Prisoners’ Dilemma”, *Journal of Economic Theory*, **39**,1, June, pp. 83-96.
- [45] Rubinstein, Ariel (1998), *Modeling Bounded Rationality*, MIT Press, Cambridge, Mass.
- [46] Rubinstein, Ariel and Asher Wolinsky,(1990), “Decentralized trading, strategic behaviour and the Walrasian outcome”, *Rev. Econ. Stud.* **57** (1990), 63-78.
- [47] Sabourian, H., (2003), “Bargaining and markets: complexity and the competitive outcome”, *Journal of Economic Theory*, **116** , 189-228.
- [48] Selten, Reinhard (1965), "Spieltheoretische Behandlung eines Oligopolmodells mit Nachfragenträgheit," *Zeitschrift für die gesamte Staatswissenschaft*,12, 201-324.
- [49] Shaked, Avner (1986): “A three-Person unanimity game”, talk given at the Los Angeles national meetings of the Institute of Management Sciences and the Operations Research Society of America.
- [50] Zemel, Eitan (1989), “Small talk and cooperation: A note on bounded rationality”, *Journal of Economic Theory*, **49**, 1-9.