

Discussion Paper No. 783

BOUNDED RATIONALITY AND STRATEGIC COMPLEXITY
IN REPEATED GAMES*

by

Ehud Kalai**

October 1987

Revised June 1988

* The author wishes to thank Abraham Neyman and Ariel Rubinstein for introducing him to this fascinating area, and Robert Aumann, Dov Samet, and William Stanford for many helpful conversations.

** Department of Managerial Economics and Decision Sciences, J. L. Kellogg Graduate School of Management, Northwestern University, Evanston, Illinois 60208.

Partial financial support for the preparation of this paper was made available by NSF Economics Grant No. SES-8720342.

I. Introduction

Three important areas in modern decision theory are: bounded rationality, artificial decision making, and management information systems. It is easy to illustrate fundamental questions within these areas. For example:

- What are the possible outcomes of strategic games if players are restricted to (or choose to) use "simple" strategies?
- What are the limitations or implications of delegating competitive decisions to machines?
- What information system (size and structure) should a player maintain when playing a strategic game?
- What are the effects of complexity costs on the outcomes of the game?

Recently, within the context of repeated games, meaningful answers to questions of this type were obtained. The purpose of this survey is to report some of these new methodologies and results.

Interest in bounded rationality and strategic complexity is not new. A few prominent examples from economics are: Simon [1972], Varian [1975], Futia [1977], Stelen [1978], Radner [1980], Smale [1980], Rosenthal [1982], Kreps-Milgrom-Roberts-Wilson [1982], Mount-Reiter [1983], Abreu [1984], and Lewis [1985]. However, this survey will restrict itself to recent results within the area of repeated games and with complexity measures that use notions of automata (sometimes referred to as Moore [1956] machines; more information about automata can be found in Hopcroft-Ullman [1979]).

Two recent interesting philosophical papers on related topics are: Binmore [1987] and Megiddo [1986]. Also there have been some recent studies

on bounded rationality and strategic complexity that are not included in this survey due to time and length limitation. Some of the important ones include: Aumann-Sorin [1985], Megiddo-Wigderson [1986], Gilboa [1986], Lipman-Srivastava [1987], Lehrer [1987], Kalai-Samet-Stanford [1986], Stanford [1987], Lipman [1987], and Aumann [1988].

This survey contains all the mathematical definitions needed for a mathematically able reader and does not assume prior knowledge of game theory or the concept of automata. We refer the reader who is interested in more details about repeated games and solution concepts to the excellent surveys of Aumann [1981], Mertens [1987], Sorin [1988], and Van-Damme [1983].

We turn now to a brief description of the results contained in this survey and their historical development.

The proposal to apply the model of an automaton to describe a player in a repeated game comes from Aumann's [1981] survey of repeated games. He specifically suggests this notion as a way to distinguish between simple and complicated strategies based on the number of states of automata describing them.

A second useful idea described in Aumann's survey is the one of studying a specific complexity issue by studying a modified version of the game specially designed for this purpose. Aumann reports on a result of Aumann-Cave-Kurz which deals with the infinitely repeated prisoners' dilemma game (this game and its strategies are defined later). Rather than considering the usual version of this game, they consider a modified, restricted version of it. In the restricted version the players can only use strategies that depend on the last action combination of every history

of plays, i.e., bounded recall strategies, and moreover only on their opponent's last action, i.e., reactive strategies. In this modified version of the repeated prisoners' dilemma game, they show that the famous tit-for-tat strategy is sequentially dominant. Thus by considering a variation of the game, the Aumann-Cave-Kurz approach gives us a game theoretic justification for the intuitively appealing tit-for-tat strategy. (Later studies of reactive strategies include Stanford [1986a,b] and Kalai-Samet-Stanford [1985]; later studies of finite recall strategies can be found in Kalai-Stanford [1988] and Lehrer [1987].)

Following Aumann, three path breaking papers were written by Ben-Porath [1986], Neyman [1985], and Rubinstein [1986]. These papers used the automaton notion to construct special games suitable for analyzing special complexity issues. Ben-Porath studied the advantage of having a bigger automaton in playing repeated zero sum games. He restricted his two players to use finite automata of different fixed sizes and he showed that asymptotically there is no gain (when compared with the unrestricted game) from having a bigger automaton unless it is exponentially bigger. In the exponentially bigger case the fullest possible gain is realized by the bigger player. Similar issues were later studied by Gilboa-Samet [1987]. It turns out that if the game is not zero sum then being smaller can actually turn out to be advantageous.

Neyman studied the finitely repeated prisoners' dilemma game. Cooperative play in such a game contradicts the existing solutions of game theory. Yet it is commonly observed even when the players are game theorists (see, for example, Axelrod [1980]). Neyman restricted his players to use automata of given fixed sizes in order to play the game. It turns

out that in the restricted game cooperative play is a game theoretic equilibrium. Thus the Neyman model gives us a way of resolving the cooperation paradox by considering strategies of limited complexity. Neyman's results were later elaborated upon by Zemel [1986] who modified the game further by introducing into it meaningless communication between the automata of the players during the play of the game. It turns out that the meaningless communication helps cooperation even further since it "uses up" a significant portion of the capacity of automata.

Rubinstein [1986] studied the effect of complexity costs on the outcome of the game. His approach, as the approach in the later Abreu-Rubinstein [1986] paper, is to modify the infinitely repeated game as follows. The players are restricted to use automata of any finite size in order to play the game. However, their final payoffs decrease as they use automata of bigger sizes. Thus they create a tension in a player between high overall utility and increasing complexity. In this modified version of the game the equilibrium outcomes have a nice simple structure and the set of equilibrium payoffs is dramatically reduced. This is even the case as the complexity costs approach zero and thus their model points out a fundamental discontinuity regarding complexity costs.

Kalai-Stanford [1988] introduced a modified version of the automaton which they applied to general (unmodified) infinitely repeated games. With their version it turns out that every strategy of the infinitely repeated game can be fully and uniquely described by a minimal (in the number of states) automaton. Thus this minimal number of states yields a general measure of complexity for repeated game strategies. It turns out that the complexity of a strategy equals precisely the number of distinct strategies

it induces during the play of the game (these concepts are defined in the sequel) and thus we obtain a natural game theoretic interpretation of the complexity measure which does not depend on the arbitrary choice of a machine type (say, automata versus a Turing machine). It also turns out that the complexity of a strategy is the minimal size of the information system needed by a player using the strategy.

Within the context of infinitely repeated games the following question arises. What equilibrium payoffs of the game can be obtained by strategies of finite bounded complexity. Kalai and Stanford [1988] showed that for the case of repeated games with discounting this is not a serious issue because every equilibrium payoff can be approximated by an equilibrium using bounded finite complexity strategies.

This approximation of equilibrium payoffs by ones using finite complexity strategies was extended by Ben-Porath and Peleg [1987] to two limit cases. The case of infinitely repeated games with the average payoff criterion, and the case of low discounting. In these cases the characterizations of all equilibrium payoffs is given by the well-known folk theorems (see Aumann-Shapley [1976], Rubinstein [1977], and Fudenberg-Maskin [1986]). Ben-Porath and Peleg showed that all the equilibrium payoffs described by the folk theorems can be approximated by equilibria using finite complexity strategies.

For robust equilibria of generic infinitely repeated games with discounting, interpersonal complexity bounds exist. Kalai and Stanford [1988] showed that at such equilibria the complexity of the strategy used by any one player never exceed the product of the complexity of his opponents. In particular, two players playing a game must use equal complexity

strategies.

II. What are Strategic Games?

A strategic game G is described by a triple $G = (N, S, u)$ with the following structure and interpretations. $N = \{1, 2, \dots, n\}$ is the set of players. It is assumed that n is a positive integer. Every player $i \in N$ has a set of strategies S_i . The set of strategy combinations is $S = \times_{i \in N} S_i$. The utility function $u = (u_1, u_2, \dots, u_n)$ with each $u_i: S \rightarrow \mathbb{R}$. u_i is called the utility (sometimes payoff) function of player i . The game G is played as follows: simultaneously and independently every player chooses a strategy $s_i \in S_i$. The resulting strategy combination $s = (s_1, s_2, \dots, s_n)$ determines a payoff $u_i(s)$ for every player i . A Nash equilibrium of a game G is a strategy combination $s^* \in S$ satisfying

$$u_i(s_1^*, \dots, s_{i-1}^*, s_i, s_{i+1}^*, \dots, s_n^*) - u_i(s^*) \leq 0$$

for every $i \in N$ and every $s_i \in S_i$.

In some of the applications described later it is useful to look at a slight relaxation of the Nash condition. For a given $\epsilon \geq 0$ we say that the strategy combination s^* is an ϵ -equilibrium if it satisfies the inequalities in the definition of a Nash equilibrium but with the right side zero being replaced by ϵ .

Often we are interested in situations where players are allowed to randomly choose their strategies. When a player does so we say that he is using a mixed strategy. We formally define the mixed strategies extension of a finite game $G = (N, S, u)$ to be the game $\Delta G = (N, \Delta S, u)$ described by:

$$\Delta S = \times_{i \in N} \Delta(S_i) \text{ with}$$

$$\Delta(S_i) = \{\delta_i: S_i \rightarrow \mathbb{R}: \delta_i(s_i) \geq 0 \text{ and } \sum_{s_i \in S_i} \delta_i(s_i) = 1\}.$$

The elements of $\Delta(S_i)$ are the mixed strategies of player i and elements of ΔS are mixed strategy combinations. $u = (u_1, u_2, \dots, u_n): \Delta S \rightarrow \mathbb{R}^n$. Here every $u_i(\delta)$ is defined to be the expected utility of player i when the players choose their s_i 's randomly according to the distributions δ_i , i.e.,

$$u_i(\delta) = \sum_{s \in S} u_i(s) \prod_{j \in N} \delta_j(s_j).$$

(We will be abusing notations throughout by letting u stand for different utility functions defined on different domains.)

Example II.1: The prisoner's dilemma game, $P = (\{1,2\}, A, u)$, is described as follows:

$$A = A_1 \times A_2 \text{ with } A_i = \{c, d\} \text{ for } i = 1, 2.$$

We refer to A_i as the set of actions of player i with c and d denoting a cooperating and a defecting action, respectively. A is then called the set of action combinations.

The utility functions are described by a bimatrix:

		Player 2's Actions	
		c	d
Player 1's Actions	c	3,3	0,4
	d	4,0	1,1

It is easy to see that (d,d) is the only Nash equilibrium of this game in pure or mixed strategies.

A game Z is called 0-sum if for every strategy combination $s \in S$ $\sum_{i \in N} u_i(s) = 0$. Obviously Z is 0-sum if and only if its extension to mixed strategies is also 0-sum. In a 2-person 0-sum game there is a unique Nash equilibrium payoff (in pure or mixed strategies), i.e., there is a real number v with $u(\delta^*) = (v, -v)$ for every Nash equilibrium δ^* . The unique number v is called the value of the game. This observation will be useful later when we wish to compare the payoffs of different games. Because of the uniqueness of the equilibrium payoff we will be able to say whether a player is better off in one 2-person 0-sum game than in another such game.

III. What Are Repeated Games?

III.1 Types of Repeated Games

We begin with a given strategic game $G = (N, A, u)$. We are interested in several strategic games that describe different versions of repeated play of G . To avoid confusion we will refer to G as the stage game and to the others as the repeated games. It is also convenient to refer to the strategies of the stage game as actions and use the word "strategy" for a

choice of a rule of behavior in the repeated game.

We will deal with finite and infinite repeated play of the stage game G and will use the notation G^T and G^∞ , respectively, to denote them. The game G^T is played as follows. In the first stage the players play the stage game G , creating an action combination a^1 , and getting paid $u(a^1)$. Every player is then informed of the actions of all the players. With this information available, the players proceed to play G again creating a second action combination a^2 and being paid $u(a^2)$. This process repeats itself T times. In G^∞ the same process goes on forever.

To completely specify the games G^T and G^∞ one has to describe how the players evaluate their utility for receiving (finite or infinite) streams of payoffs. We will be using two common methods. Under the first method the players use a given common discount parameter α and evaluate a stream of payoffs according to its discounted value. When this is the case we denote the resulting games by $G^{T,\alpha}$ and $G^{\infty,\alpha}$. Under the second method the players use their long run average payoff. When this is the case we denote the games by \bar{G}^T and \bar{G}^∞ .

Before turning to the formal description of these games we need to introduce some additional terminology. By a history resulting from the repeated play of G we mean a string of action combinations of any finite length, i.e., an element of $\times_{t=1}^{\bar{q}} A$ for some $\bar{q} \geq 0$. We use $\bar{q}(h)$ to denote the length of a history h . For two histories h and \bar{h} of lengths \bar{q} and $\bar{\bar{q}}$ we write $h\bar{h}$ to denote the concatenation of h with \bar{h} , i.e., the history of length $\bar{q} + \bar{\bar{q}}$ whose first \bar{q} elements are the action combinations of h followed by the $\bar{\bar{q}}$ action combinations of \bar{h} . It is also convenient to be able to speak of the empty history which denotes the fact that "nothing

happened yet." We use the letter e to denote it, define its length $l(e) = 0$, and define $eh = he = h$ for every history h .

We let H^T denote the set of all histories of length strictly less than T and H^∞ denote the set of all histories of any finite length.

III.2 The Formal Description of \bar{G}^T

The strategic game \bar{G}^T is described by the triple (N, F, \bar{u}) having the following structure and interpretations.

N is the set of players of the underlying stage game G . The set of strategy combinations $F = \times_{i \in N} F_i$ with the sets F_i denoting as usual the sets of individual strategies. Each F_i here consists of all the functions $f_i: H^T \rightarrow A_i$. Thus a strategy is a prescription of how to act after every possible history of action combinations. Notice that the domain of f_i contains histories that are not consistent with its own earlier prescriptions. This is important for two reasons. It clearly prescribes actions even if mistakes were made in carrying out f_i 's earlier prescriptions--situations that arise frequently in large economic games. Also it allows for a careful analysis of the self consistency of the rationality of playing a strategy (see, for example, Selten [1975]).

To define $\bar{u} = (\bar{u}_1, \bar{u}_2, \dots, \bar{u}_n)$ we first define the play resulting from a strategy combination $f \in F$. Inductively we define $f^1 = f(e)$ and for $t = 2, 3, \dots, T$ $f^t = f(f^1, f^2, \dots, f^{t-1})$. We now define $\bar{u}(f) = (1/T) \sum_{t=1}^T u(f^t)$. Thus every player evaluates his utility of the strategy combination f according to his average payoff.

III.3 The Formal Description \bar{G}^∞

In playing G infinitely many times, infinite streams of payoffs are created. Defining the average payoff of such streams is more difficult. It is convenient to use the notion of a Banach limit to evaluate such averages.

We let B denote the set of all bounded countable infinite sequences of real numbers and we assume that each of the players $i \in N$ has a utility function $\bar{u}_i: B \rightarrow \mathbb{R}$ satisfying the following conditions.

1. \bar{u}_i is linear;

2. For every $b = (b^t)_{t=1}^\infty \in B$

$$\liminf_{T \rightarrow \infty} (1/T) \sum_{t=1}^T b^t \leq \bar{u}_i(b) \leq \limsup_{T \rightarrow \infty} (1/T) \sum_{t=1}^T b^t, \text{ and}$$

3. $\bar{u}_i((b^t)_{t=1}^\infty) = \bar{u}_i((b^t)_{t=2}^\infty)$.

The existence of such functions \bar{u}_i is guaranteed by the Hahn-Banach theorem. There could be many functions \bar{u}_i satisfying the above conditions. However we will assume from here on that every player has chosen one such \bar{u}_i to evaluate streams of payoffs for all infinitely repeated games.

Now \bar{G}^∞ is defined by a triple (N, F, \bar{u}) described as follows. The set of strategy combinations $F = \times_{i \in N} F_i$ with every F_i containing all the functions $f_i: H^\infty \rightarrow A_i$. Just as in the definition of \bar{G}^T we define the play path (now infinite) resulting from a strategy combination f by $f^1 = f(e)$ and for $t = 2, 3, \dots$ $f^t = f(f^1, f^2, \dots, f^{t-1})$. The utility vector associated with every strategy combination $f \in F$ is defined by $\bar{u}(f) = (\bar{u}_1(f), \dots, \bar{u}_n(f))$ with $\bar{u}_i(f) = \bar{u}_i(f^1, f^2, \dots)$.

III.4 The Formal Description of $G^{\infty, \alpha}$

The infinitely repeated strategic game $G^{\infty, \alpha}$ is defined by the triple (N, F, u^{α}) . The set of strategy combinations F , is defined as in the game \bar{G}^{∞} . Also for every strategy combination $f \in F$ the play path f^1, f^2, \dots is defined as before. However, the utility functions u_i^{α} are defined for a given discount parameter α , $0 < \alpha < 1$, by $u_i^{\alpha}(f) = \sum_{t=1}^{\infty} \alpha^{t-1} u_i(f^t)$.

III.5 Examples of Strategies in the Repeated Prisoners' Dilemma Game

The strategies described here are often used in discussions of the repeated prisoners' dilemma game. They apply to the finitely and infinitely repeated games with both discounting and the average payoff criterion. We describe them as player one strategies. The symmetric definitions apply to player two.

The constant defecting strategy of player one, d_1 , is defined by $d_1(h) = d_1$ for all histories h .

The cooperate and then follow the tit-for-tat rule, $c\text{-tft}_1$, is defined as follows:

$$\begin{aligned} c\text{-tft}_1(e) &= c_1, \text{ and for every history } h \text{ of length } \ell > 0 \\ c\text{-tft}_1(h) &= h_2^{\ell}, \text{ i.e., the last action taken by player 2.} \end{aligned}$$

The grim trigger strategy, gtrg_1 , is defined by $\text{gtrg}_1(e) = c_1$, and for every history h of length $\ell > 0$ $\text{gtrg}_1(h) = c_1$ if and only if

$$h = \left(\begin{pmatrix} c_1 \\ c_2 \end{pmatrix}, \begin{pmatrix} c_1 \\ c_2 \end{pmatrix}, \dots, \begin{pmatrix} c_1 \\ c_2 \end{pmatrix} \right).$$

Other strategies that will be defined and used later are the n -period trigger strategies. For example, the two-period trigger strategy, 2trg_1 , is

the one in which player one continues cooperation until he detects a defection. When (and if) he does he plays d_1 for two periods, ignoring what player two does, and then starts the two-period trigger strategy again. It is easier to formally define these strategies recursively, through the notion of automaton. This will be done later.

III.6 Equilibria of Repeated Games

For all three types of games, \bar{G}^T , \bar{G}^∞ and $G^{\infty, \alpha}$, the concepts of Nash equilibrium and ϵ -equilibrium are defined as for general strategic form games. However the more sophisticated notion of subgame perfect equilibrium is often applied to these games.

Consider first the game \bar{G}^T , a history h of length $\ell < T$ and a strategy combination $f = (f_1, f_2, \dots, f_n)$. We denote by $f_i|h$ the strategy of player i induced by f_i and h in the game $\bar{G}^{T-\ell}$. Formally we define this strategy by $(f_i|h)(\bar{h}) = f_i(h\bar{h})$. Since $h\bar{h}$ is a history of \bar{G}^T if and only if $\ell(\bar{h}) \leq T - \ell$, $f_i|h$ is indeed a strategy of $\bar{G}^{T-\ell}$. Also, $f|h$ is defined to be the strategy combination induced by f and h , i.e., $(f_1|h, f_2|h, \dots, f_n|h)$. We say that the strategy combination f of \bar{G}^T is a subgame perfect (ϵ -)equilibrium if for every length ℓ history h of \bar{G}^T , $f|h$ is an (ϵ -)equilibrium of $\bar{G}^{T-\ell}$. In other words, at a subgame perfect equilibrium, regardless of the past, the players play an equilibrium strategy in the remaining game.

The analogous concepts are defined for \bar{G}^∞ and $G^{\infty, \alpha}$. A strategy combination f of \bar{G}^∞ (resp. $G^{\infty, \alpha}$) is a subgame perfect (ϵ -)equilibrium if for every history h , $f|h$ is an (ϵ -)equilibrium of \bar{G}^∞ (resp. $G^{\infty, \alpha}$). Here the induced strategies $f_i|h$ are defined by $(f_i|h)(\bar{h}) = f_i(h\bar{h})$ for every history

\bar{h} of any finite length.

The following are some examples of equilibria of the repeated prisoners' dilemma game P . The pair of constant defect strategies $d = (d_1, d_2)$ is a subgame perfect equilibrium for all games of the form \bar{P}^T , \bar{P}^∞ , and $P^{\infty, \alpha}$. This is the case since $d|h = d$ for every history h and d is a Nash equilibrium for all of these games. For games of the type \bar{P}^T one can actually check that d is the only Nash equilibrium.

The pair $c\text{-tft} = (c\text{-tft}_1, c\text{-tft}_2)$ is a Nash equilibrium of \bar{G}^∞ and of $G^{\infty, \alpha}$ for sufficiently large α . However, it is not a subgame perfect equilibrium of these games. Consider, for example, a history h which ends with the action combination $\begin{pmatrix} c \\ d \end{pmatrix}$. $c\text{-tft}|h = (d\text{-tft}, c\text{-tft})$ with $d\text{-tft}$ being the strategy "defect and then follow the tit-for-tat rule." It is easy to verify that $(d\text{-tft}, c\text{-tft})$ is not a Nash equilibrium of the infinitely repeated game as is required by subgame perfection since $(c\text{-tft}, c\text{-tft})$ will yield player one a higher payoff.

A pair of grim trigger strategies $gtrg = (gtrg_1, gtrg_2)$ is a subgame perfect equilibrium of \bar{P}^∞ and of $P^{\infty, \alpha}$ with sufficiently large α . This is the case since $gtrg|h = gtrg$ or $gtrg|h = d$ and both $gtrg$ and d are Nash equilibria of these games.

IV. Automata and Complexity Measures in Repeated Games

IV.1 What Types of Automata

A repeated game automaton of a player is a system that is in one of a given set of states at any time. At every such state it prescribes an action for the player and then receives as an input the action combination

of the other players. According to a fixed transition law, which depends on the current state and on the input action combination, it then transits to a new state. In this new state the same operations are repeated. It is clear that with a specification of an initial state, such an automaton can "play" the game for a player, i.e., "decide" what to do after every history of action combinations.

Two types of automata will be used later. In one type, the input to the automaton at any given state is the action combination of all other players, excluding the action of the player using the automaton. In the second type the input to the automaton is the action combination of all the players, including the player using the automaton. We refer to the former type as an exact automaton and to the latter as a full automaton. Exact automata are useful for implementing strategies in situations where the action taken by a player is always the action prescribed by the automaton. Full automata are useful for situations where the automaton's prescribed action may differ from the action actually taken.

IV.2 Formal Description of Exact Automata

For our purposes it suffices to describe exact automata for two person games. The extension to n-person games is straightforward. We will describe such automata for player one. The symmetric definitions applies to player two.

Such an automaton is described by a triple $((M, m^0), B, T)$ with the following interpretations. M is the set of states of the automaton with $m^0 \in M$ being the initial state. The behavior function $B: M \rightarrow A_1$ prescribes an action to player one at every state of the automaton. The transition

function $T: M \times A_2 \rightarrow M$ transits the automaton to a new state from an old one as a function of the action of player two.

The following table lists examples of strategies in the repeated prisoners' dilemma game and exact automata corresponding to the given strategies.

Strategies of Player One	Corresponding Exact Automata			
	<u>States</u>	<u>Behavior Function</u>	<u>Transaction Function actions of player 2</u>	
			c_2	d_2
d_1 --The constant defect strategy	$m^0 = D$	d_1	D	D
c -tft ₁ Cooperate and then tit-for-tat	$m^0 = C$ D	c_1 d_1	C C	D D
g trg ₁ The grim trigger strategy	$m^0 = C$ D	c_1 d_1	C D	D D
2 trg ₁ Two phase trigger	$m^0 = C$ P_1 P_2	c_1 d_1 d_1	C P_2 C	P_1 P_2 C

IV.3 Formal Description of Full Automata

A full automaton is described by a triple $((M, m^0), B, T)$ with the same structure and interpretation as an exact automaton. The only difference is that here $T: M \times A \rightarrow A_1$ (rather than $M \times A_2 \rightarrow A_1$). In other words, the transition function depends on the player's own last period action and not just on his opponents' last period actions. The following are examples of player one's full automata in the repeated prisoners' dilemma game.

<u>Strategies of Player One</u>	<u>Corresponding Full Automata</u>					
	<u>States</u>	<u>Behavior Function</u>	<u>Transition Function</u>			
			<u>Action Combination</u>			
			$\begin{pmatrix} c_1 \\ c_2 \end{pmatrix}$	$\begin{pmatrix} c_1 \\ d_2 \end{pmatrix}$	$\begin{pmatrix} d_1 \\ c_2 \end{pmatrix}$	$\begin{pmatrix} d_1 \\ d_2 \end{pmatrix}$
Grim trigger on opponent's defections	$m^0 = C$	c_1	C	D	C	D
	D	d_1	D	D	D	D
gtrg ₁ Grim ¹ trigger (on both)	$m^0 = C$	c_1	C	D	D	D
	D	d_1	D	D	D	D

Notice that while the full automata of the two strategies above are different, exact automata that will describe their exact (no mistakes) play will be the same. We emphasize that the full automaton gives prescriptions to the player of how to play after every history of action combinations. Thus a full automaton describes a complete strategy for a game, while an exact automaton prescribes actions only for histories which are consistent

with its earlier prescriptions.

IV.4 Complexity and the Structure of Strategies

In this section, following Kalai-Stanford [1988], we define the notion of (full) strategic complexity for strategies of infinitely repeated games through the notion of a full automaton. We then give an alternative game theoretic characterization of this measure, one that does not make use of the notion of automata but rather makes use of the internal structure of such strategies. In the following section, strategic complexity is characterized again, this time through the size and structure of the information system needed in order to implement a strategy.

The complexity of a strategy is defined to be the size (number of states) of the smallest full automaton prescribing it. It is easy to verify that this notion is well defined for every strategy provided that we allow for infinite complexities.

In the cases of infinitely repeated games, \bar{G}^∞ and $G^{\infty, \alpha}$, a full automaton turns out to be an alternative way to fully describe a strategy. Moreover, there is a natural one-to-one correspondence between the set of player's strategies of a game and the set of his minimal full automata for the game.

Consider a strategy f_i of G^∞ (recall that \bar{G}^∞ and $G^{\infty, \alpha}$ have the same strategy sets and thus we use G^∞ to denote any game of these two types). We define the set of all strategies induced by f_i after all histories of the game by $f_i|H = \{f_i|h: h \text{ is a history of } G^\infty\}$. We now define the canonical full automaton associated with f_i , A_{f_i} , by the triple $((M, m^0), B, T)$ described as follows. $M = f_i|H$, $m^0 = f_i|e$, $B(f_i|h) = f_i(h)$, and $T(f_i|h, a) = f_i|ha$.

It is easy to verify that B and T are well defined--that is, if $f_i|h = f_i|\bar{h}$ then $f_i(h) = f_i(\bar{h})$ and $f_i|ha = f_i|\bar{h}a$ for every action combination a.

Consider for example the two phase trigger strategy of player one, $2trg_1$, in the infinitely repeated prisoners' dilemma game. This strategy induces itself, for example after the empty history and after completely cooperative histories. It also induces the strategy $dd-2trg_1$ in which player one defects twice and then returns to using the $2trg_1$ strategy. For example, this strategy will be induced after the history $((\begin{smallmatrix} c_1 \\ c_2 \end{smallmatrix}), (\begin{smallmatrix} c_1 \\ c_2 \end{smallmatrix}), (\begin{smallmatrix} c_1 \\ d_2 \end{smallmatrix}))$. Similarly the strategy defect once and then go back to the $2trg_1$ strategy, $d-2trg_1$, is induced after the history $((\begin{smallmatrix} c_1 \\ c_2 \end{smallmatrix}), (\begin{smallmatrix} c_1 \\ c_2 \end{smallmatrix}), (\begin{smallmatrix} c_1 \\ d_2 \end{smallmatrix}), (\begin{smallmatrix} d_1 \\ d_2 \end{smallmatrix}))$. One can easily verify that no other strategies are induced by $2trg_1$. Now consider the following minimal automaton corresponding to $2trg_1$.

States	Behavior	<u>Transition Function</u>			
		<u>Action Combinations</u>			
		$\begin{pmatrix} c_1 \\ c_2 \end{pmatrix}$	$\begin{pmatrix} c_1 \\ d_2 \end{pmatrix}$	$\begin{pmatrix} d_1 \\ c_2 \end{pmatrix}$	$\begin{pmatrix} d_1 \\ d_2 \end{pmatrix}$
C	c_1	C	P_1	P_1	P_1
P_1	d_1	P_2	P_2	P_2	P_2
P_2	d_1	C	C	C	C

Observe that if the automaton started at its other states, rather than C, it will prescribe the induced strategies described by the following table

<u>Starting State</u>	<u>Induced Strategy</u>
C	2trg ₁
P ₁	dd-2trg ₁
P ₂	d-2trg ₁

We observe a one-to-one correspondence between the states of the automaton and induced strategies of the given strategy. This structural relationship between a strategy and a minimal full automaton prescribing it is completely general, as we describe now.

We say that a full automaton $A = ((M, m^0), B, T)$ prescribing a strategy f_i is reducible if there is a partition $L = \{L^1, L^2, \dots, L^t\}$ of M with the following two properties:

1. L is not trivial in the sense that at least one of the L^j 's is not a singleton.
2. If $m, \bar{m} \in L^j$ then:
 - a. $B(m) = B(\bar{m})$, and
 - b. for every $a \in A$ $T(m, a)$ and $T(\bar{m}, a)$ belong to the same L^k .

Clearly, if an automaton is reducible in this formal sense then one can define a **smaller** (reduced) full automaton (with the states being elements of the partition L and the initial state being the element L^j containing m_0) prescribing the same strategy.

We say that two player i automata, A and \bar{A} , of a given game are isomorphic if there is a one-to-one onto correspondence from the set of states of A to \bar{A} , $\pi: M \rightarrow \bar{M}$, which preserves behavior and transitions. This

means that for every $m \in M$ $B(m) = \bar{B}(\pi(m))$, and for every $a \in A$ $\pi(T(m,a)) = \bar{T}(\pi(m),a)$.

Theorem 4.1: Every irreducible automata prescribing a strategy f_i is isomorphic to the canonical automaton A_{f_i} .

It is easy to verify that every automaton prescribing a strategy f_i must have at least as many states as the number of strategies induced by f_i . We thus obtain the following corollary.

Corollary 4.1: The full complexity of a strategy f_i equals the number of distinct strategies it induces, i.e., $|(f_i|H)|$.

IV.5 Complexity and Information Systems

As a strategy is being implemented during a play of a given game, one has to keep track of the essential parts of past histories. It turns out that a little past information is needed in order to implement a simple strategy. As the complexity of the strategy increases, the amount of information needed for its implementation increases in a one-to-one fashion.

Consider, for example, implementing the constant defect strategy in the infinitely repeated prisoners' dilemma game (strategic complexity = 1). Essentially, no past information is ever needed in implementing it. The c-tft₁ strategy (complexity = 2), on the other hand, does need some information. Here one must be able to distinguish between two classes of histories. These are the ones that end with player two cooperating and the ones that end with player two defecting. Thus a player implementing this strategy does not need to keep track of the whole past history for future

references but only of which of the two possible types a history is. Therefore, information systems capable of being in two distinct states suffice.

In general, given a game G^∞ and a strategy of player i , f_i , we say that the histories h and \bar{h} are equivalent relative to f_i , $h \sim h_i$, if for every history string β , $f_i(h\beta) = f_i(\bar{h}\beta)$. In other words, h and \bar{h} are equivalent if no matter what actions follow h or \bar{h} player i will act the same. Thus, one could replace h by \bar{h} without ever affecting player i 's future actions. This means that an information system of player i does not need the ability to distinguish between equivalent histories; it only needs to keep track in which equivalent class of histories it is at every stage of the game. Given the strategy f_i we let H/\sim denote the quotient set of H relative to \sim .

Consider, for example, the two-phase trigger strategy of player one discussed in the previous section. It is easy to see that H/\sim here contains exactly three equivalent classes of histories: $[e]$ --the set of histories equivalent to the empty history, $[(\begin{smallmatrix} c_1 \\ d_2 \end{smallmatrix})]$ --the set of histories equivalent to $(\begin{smallmatrix} c_1 \\ d_2 \end{smallmatrix})$, and third, $[(\begin{smallmatrix} c_1 & d_1 \\ d_2 & d_2 \end{smallmatrix})]$ --the set of histories equivalent to $(\begin{smallmatrix} c_1 & d_1 \\ d_2 & d_2 \end{smallmatrix})$. Actually, one can extend the one-to-one correspondence between starting states and induced strategies discussed in the previous section, to include also the one-to-one correspondence with equivalence classes of histories. For example, for the $2trg_2$ strategy this correspondence table is the following:

Equivalence Class of Histories	Starting State Following the History	Induced Strategy
[e]	C	2trg ₁
$\left[\begin{pmatrix} c_1 \\ d_2 \end{pmatrix} \right]$	P ₁	dd-2trg ₁
$\left[\left(\begin{pmatrix} c_1 \\ d_2 \end{pmatrix} \begin{pmatrix} d_1 \\ d_2 \end{pmatrix} \right) \right]$	P ₂	d-2trg ₁

It is clear that in order to implement a strategy the minimal number of history classifications needed is precisely $|H/\sim|$. Also, from the previous discussion, the following statement is quite apparent.

Theorem 4.2: The complexity of a strategy f_i equals the number of equivalence classes of histories it induces $|H/\sim|$.

Recently, a refinement of the equivalence relation \sim was proposed by Stanford [1987]. According to it, for a given strategy f_i , two histories h and \bar{h} are equivalent, $h \ S \ \bar{h}$, if for every two histories α and β , $f_i(\alpha h \beta) = f_i(\alpha \bar{h} \beta)$. In other words, two histories are equivalent if no matter what happens prior to them and what happens following them the same behavior is induced. Clearly S is a stronger equivalence relation than \sim . If two histories h and \bar{h} are S -equivalent, then one can substitute them for each other even when augmenting them to earlier existing histories. The

number $|H/S|$ represents the size of the information system needed for implementing the strategy f_i while keeping track of the history strings that could start at any point of time, not necessarily the beginning of the game (for example, when summarizing annual activities before augmenting them to earlier activities). This is not the case for the equivalence relation \sim . Also, the set H/S can be shown to have an elegant semi-group and even group structure under some restrictions on the underlying strategy.

V. Complexity Issues Analyzed Through Specially Constructed Games

Several issues regarding the effects strategic complexity has on outcomes of games have been analyzed recently. Typically the approach in these studies is to modify the definition of the game by building in some special structure and restrictions designed to address the specific complexity issue. Then by comparing the equilibria of the modified game with the equilibria of the original game, the effect of strategic complexity is observed. We will report here on studies of three such general issues.

In the finitely repeated prisoners' dilemma game it is well known that the only Nash equilibrium is the one in which both players always defect. This is considered somewhat of a paradox, since players do not seem to follow this very convincing noncooperative game theoretic solution. In modified versions of this game the players' strategies are restricted to ones that can be played by finite automata. If the automata are restricted to be of sufficiently small size then cooperative behavior emerges as a Nash equilibrium of this modified game. Surprisingly, this result persists even when the bound on the size of the automata is quite generous. If in addition to playing the game we engage the players in some high vocabulary

meaningless back-and-forth communication, then even with very large automata at their disposal, cooperative behavior becomes a Nash equilibrium.

A second direction of research studies the effect of having the players use automata of different sizes. It is shown that in an infinitely repeated two person zero sum game, if the automaton player two can use is significantly greater than the automaton available to player one, then the value of the game will shift completely in player two's favor. However, for this phenomenon to occur, player two's automaton must be exponentially bigger than player one's automaton; otherwise player two has no gain at all. In a nonzero sum game, having a bigger automaton can actually turn out to be disadvantageous.

The third modification discussed here goes further than just restricting the game to be played by finite automata of given sizes. It takes into consideration also the cost of having a bigger automaton when playing an infinitely repeated game. When the payoff functions of the players are modified to take this cost into consideration, the resulting equilibria have a simple interesting structure, and the set of equilibria is dramatically reduced.

V.1 Cooperation in Finitely Repeated Prisoners' Dilemma Games Played Through Bounded Automata

We first summarize results of Neyman [1985]. We let P denote the two-person prisoners' dilemma game and, as before, \bar{P}^T denote the game consisting of T repetitions of P with the average payoff criterion. Now we consider a new two person strategic one shot game, \bar{P}_{m_1, m_2}^T , described as follows. The strategies of each player are choices of exact automata for playing \bar{P}^T . In

making these choices each player i is restricted to choose an automaton of size not exceeding the given positive integer m_i . Given a pair of such automata (A_1, A_2) , the utility of each player i , $u_i(A_1, A_2)$, is defined to be the utility he receives in \bar{P}^T when the prescriptions of the automata A_1 and A_2 are followed.

The first theorem states that if the players are restricted to choose automata that are too small to count the number of stages of the repeated game then both players choosing "a cooperating automaton" is a Nash equilibrium. (Recall that the only Nash equilibrium of the unrestricted repeated game has both players defecting throughout.) One may therefore think of "bounded rationality," or bounded ability to handle strategic complexities, as a way to resolve the prisoners' dilemma paradox.

Theorem 5.1: If $2 \leq m_1, m_2 \leq T - 1$ then there is a Nash equilibrium pair of automata of \bar{G}_{m_1, m_2}^T that prescribe cooperation throughout \bar{G}^T .

It is surprising that even if the players can choose large automata, then they can get arbitrarily close to the cooperative payoffs provided that they are allowed to randomize in their choices of automata. Recall that under our conventions, for a strategic form game G , ΔG denotes its extension to mixed strategies. Also the cooperative payoffs in the one shot prisoners' dilemma game are (3,3).

Theorem 5.2: For every positive integer K there is an integer T_0 such that for all $T \geq T_0$ if $T^{1/K} \leq m_1, m_2 \leq T^K$, then there is a Nash equilibrium payoff of $\Delta \bar{P}_{m_1, m_2}^T$ exceeding $3 - 1/K$ for both players.

The next result of Zemel [1985] gives an alternative way of explaining

cooperation in the finitely repeated prisoners' dilemma game. It goes further in modifying the game by incorporating artificial communications (small talk) into the stage actions of the players. The communications of the players do not directly affect the players' payoffs. However, they increase the complexity of executing a strategy. Consequently, they prevent a player from carrying on complex computations which may be required for noncooperative strategies. It is mathematically easier to prove that cooperative behavior is induced by a Nash equilibrium of the modified game with communications. Nevertheless this modified game is attractive on intuitive grounds and its methodology can be applied to some games in which the Neyman modification is not applicable.

Starting with the two person prisoners' dilemma game P , and a given finite set of messages L , we define a modified strategic game, PL , which stands for the prisoners' dilemma game with the language L attached to it. PL is described by the triple $PL = (\{1,2\}, AL, u)$ with $AL = AL_1 \times AL_2$ and $AL_i = \{c,d\} \times L$. Thus an action in PL consists of choosing an action in the original prisoners' dilemma game and a message from the language L . The utilities of both players are defined to depend only on the c,d combinations as in the unmodified prisoners' dilemma game without regard to the choices of messages. More precisely

$$u_i((x_1, l_1), (x_2, l_2)) = u_i(x_1, x_2)$$

with the right side being the utility function u_i of the prisoners' dilemma game. However, the messages will play an important role in the repeated game since they may artificially increase complexities of strategies.

We define \overline{PL}^T to be the T repetitions of PL with the average payoff criterion. $\overline{PL}_{m_1, m_2}^T$ is defined as before to be the one shot game in which the players choose automata of sizes not exceeding m_1 and m_2 , respectively, and use the prescriptions of these automata in the game \overline{PL}^T in order to determine their payoffs. The game $\Delta\overline{PL}^T$ is the one in which the players can choose these automata randomly.

Theorem 5.3: If $m_1, m_2 \geq 3$, $T \geq 5$ and the size of the message set $|L| \geq \max\{m_1, m_2\} - T + 3$, then there is a mixed strategy Nash equilibrium of $\Delta\overline{PL}_{m_1, m_2}^T$ which prescribes cooperation throughout \overline{PL}^T .

The equilibrium constructed for the proof of the above theorem has the players choose randomly a message from the large set of possible messages. They send each other the randomly selected messages in the first iteration of the game. On subsequent iterations they have to repeat to their opponent the messages they received earlier. Failing to repeat correctly a message causes the other player to trigger into defection. Thus it pays a player to keep track of the message received and to repeat it. However, keeping track of the messages received "uses up" many states of his automaton, preventing the player from counting the number of stages left in the game. This brings us back to a situation similar to that of Theorem 5.1 in which cooperation is possible.

Neyman's construction for the proof of Theorem 5.2 is similar. There, however, the irrelevant communication has to be done through the play of the game by defecting in an "agreed upon" sequence of stages. For that reason his proof is more difficult, and one actually cannot have perfect cooperation.

V.2 The Advantage of Having a Big Automaton

We first summarize results of Ben-Porath [1986]. We start with a two person zero sum game Z and its infinitely repeated version \bar{Z}^∞ with the average payoff criterion. Next we consider the strategic game $\bar{Z}_{m_1, m_2}^\infty$. This is the one shot game in which the players choose automata of sizes not exceeding m_1 and m_2 , respectively, for playing the game \bar{Z}^∞ . The payoffs of the players are then determined according to their payoffs in the game \bar{Z}^∞ when the prescriptions of the chosen automata are followed.

It is important to observe that:

1. $\bar{Z}_{m_1, m_2}^\infty$ is still a finite two person zero sum game and thus has a value (in mixed strategy).
2. Every player can guarantee his pure strategy Z -maxmin value with an automaton of size one.

More specifically, the pure-strategy Z -maxmin value of a player is the largest payoff he can secure for himself by choosing a pure action in Z and assuming that his opponent will move after him in order to minimize the original player's payoff. For example, the pure-strategy Z -maxmin payoff of player one is $\max_{a_1 \in A_1} \min_{a_2 \in A_2} u_1(a_1, a_2)$ with A_i being the sets of pure actions in Z and u_1 being the utility of player one. Let a_1^* be a maxmin action of player one, i.e., an action at which the maximum is attained. By playing a_1^* repeatedly in \bar{Z}^∞ player one can guarantee for himself this maxmin value. Thus even with a one state automaton, player one can guarantee himself his pure-strategy Z -maxmin value against any automaton of player two. The next theorem states that if player two is allowed a sufficiently large automaton, then he can indeed push player one all the way down to this

minimal security level.

Theorem 5.4: (On being much bigger): For every given positive integer m_1 there is a positive integer m_2 and an automaton A_2 of size m_2 such that for every automaton A_1 of size m_1 , $u_1(A_1, A_2) \leq$ pure strategy Z-maxmin value of player 1.

If, on the other hand, both players choose large automata of "roughly the same size," then the value of the automata game remains the same as that of the underlying stage game. It turns out roughly the same size here has a very liberal meaning.

Let Q be a function from the set of positive integers to itself. We say that $Q(n)$ is asymptotically a little bigger than n if:

1. $Q(n) \geq n$ for all n , and
2. $\ln(Q(n))/n \rightarrow 0$ as $n \rightarrow \infty$.

For example, a polynomial function of n is asymptotically a little bigger than n .

Theorem 5.5: (On being a little bigger): $\text{Value}(\bar{z}_{m_1, Q(m_1)}^\infty) \rightarrow \text{Value}(Z)$ as $m_1 \rightarrow \infty$ if $Q(m_1)$ is asymptotically a little bigger than m_1 .

Thus, asymptotically, one cannot benefit from having a bigger automaton at a zero sum game unless the automaton is much bigger.

Next we summarize results of Gilboa-Samet [1987]. They study the game in which only one of the players, say player one, is restricted to choose a finite automaton while the second player is unrestricted. It turns out that in nonzero sum games being unrestricted may work against a player.

We start with a two person game G and consider the infinitely repeated

game \bar{G}^∞ with the average payoff criterion. We now consider a game $\bar{G}_{F,\infty}^\infty$. In this game player one is restricted to choose a finite automaton of any size while player two can choose any strategy. (One can restrict player two to choices of Turing machines and obtain similar results.) We let V_{II} denote the pure strategies minmax value of player two in the game G , i.e.,

$V_{II} = \min_{a_1 \in A_1} \max_{a_2 \in A_2} u_2(a_1, a_2)$. Notice that V_{II} is the utility that player two can secure for himself in the game G provided that he chooses his strategy after observing player one's choice. Indeed, it is the case that for every player the minmax value of $G \geq$ maxmin value of G .

Theorem 5.6: (The advantage of being unbounded): Player two has a strategy in $\bar{G}_{F,\infty}^\infty$ yielding him a payoff of at least V_{II} against any automaton of player one.


This theorem is very similar to Ben-Porath's Theorem 5.4 since in a zero sum game the minmax value of player two equals the negative of player one's maxmin value. Thus player two "guaranteeing himself his minmax value" is the same as "being able to push player one down to one's maxmin value."

Now we restrict player one further by requiring that he uses automata that are "reversible." This means that for every state that player one's automaton can enter during the play of the game player two has a sequence of actions that will lead one's automaton back to its initial state. Observe that when player one uses such an automaton player two is able to "experiment" and study one's automaton without ever causing himself irreversible damage. (Strategies of reversible automata are sometimes called forgiving. See Axelrod [1980].) We let $\bar{G}_{RF,\infty}^\infty$ denote the strategic game restricting player one to use reversible finite automata.

Theorem 5.7: Player two has a dominant strategy in $\bar{G}_{RF,\infty}^\infty$.

Here a strategy f_2 of player two is said to be dominant, if for every other strategy \bar{f}_2 of player two and every automaton A_1 of player one,
 $u_2(f_2, A_1) \geq u_2(\bar{f}_2, A_1)$.

To illustrate situations advantageous to player one we need some more notation. We let C be the convex hull of the set of feasible payoffs in the game G , $\{(u_1(a), u_2(a)) : a \in A\}$. We let R be the region of feasible payoffs yielding player two more than V_{II} , $R = \{u \in C : u_2 > V_{II}\}$. Now we let W_I be the maximal payoff that player one can receive in the region R ,
 $W_I = \sup\{u_1 : (u_1, u_2) \in R \text{ for some } u_2\}$.



Theorem 5.8 (the advantages of being bounded): Suppose the region $R \neq \emptyset$ and consider the game $\bar{G}_{RF,\infty}^\infty$. For every $\epsilon > 0$ the set of ϵ -equilibria in which player two uses a dominant strategy is not empty. Moreover, at every such equilibrium the payoff to player one is within ϵ of W_I (player two's payoff is at least V_{II}).

V.3 The Reduction of the Set of Equilibria Due to Complexity Costs

Here we report on results of Abreu-Rubinstein [1986] which modify the original results of Rubinstein [1986].

We start with a finite two person stage game G and consider an infinitely repeated game G^∞ with either the average payoff criterion (\bar{G}^∞) or under discounting ($G^{\infty, \alpha}$). Next we restrict both players to use finite automata in playing G^∞ and denote the resulting game by $G_{F,F}^\infty$. However, we assume that in $G_{F,F}^\infty$ every player has a linear preference relation over the possible pairs of automata chosen. Such a preference relation depends on two parameters, the utility of the player (average or discounted) resulting from the recommended play of the automata, and the number of states in the player's own automaton. More specifically, a player has monotonically strictly increasing preference for higher utility (for a fixed size of his own automaton) and a monotonic strictly decreasing preferences in the number of states in his automaton (for a fixed utility level). Formally the following three assumptions are made about the preference relation of player i . For two pairs of automata $A = (A_1, A_2)$ and $\bar{A} = (\bar{A}_1, \bar{A}_2)$ for the game $G_{F,F}^\infty$, with number of states m_i and \bar{m}_i , respectively,

1. If $u_i(A) = u_i(\bar{A})$ and $m_i = \bar{m}_i$ then player i is indifferent between A and \bar{A} , $A \sim_i \bar{A}$;
2. If $u_i(A) > u_i(\bar{A})$ and $m_i = \bar{m}_i$ then player i strictly prefers A to \bar{A} , $A >_i \bar{A}$; and
3. If $u_i(A) = u_i(\bar{A})$ and $m_i < \bar{m}_i$ then again $A >_i \bar{A}$.

Now we define the Nash equilibrium of the game $G_{F,F}^\infty$ but relative to the preference relations $>_i$. Thus, a pair of automata (A_1, A_2) is a Nash

equilibrium of $G_{F,F}^{\infty}$ if for every automaton \bar{A}_1 of player one:

$$(A_1, A_2) \succeq_1 (\bar{A}_1, A_2)$$

with the symmetric condition holding for player two.

Theorem 5.9: (The structure of equilibrium with costly states): Let (A_1, A_2) be a Nash equilibrium of $G_{F,F}^{\infty}$ and let a^1, a^2, a^3, \dots be the play sequence of action combinations prescribed by these automata.

1. The sequence can always be decomposed into three phases, I, II, and III, as follows:

Phase III starts from the first $a^t \dots$ at which the sequence becomes cyclic (recall that the A_i 's are finite so that a a^t must exist).

Phase I is the longest initial part of the form $a^1, a^2, \dots, a^{\lambda}$ in which none of the players use any of the same states that he uses during the cycle-phase III. These may be thought of as entering states.

Phase II is $a^{\lambda+1}, \dots, a^{t-1}$. It turns out that in this phase both players use only cycle states but the cyclical order is not established yet. (Phase II is empty in the cases where G^{∞} is a repeated game with discounting.)

2. The total number of states, as well as the number of cycle states, are the same for both players.

3. In phases I and III the actions of the two players are synchronized in the sense that a player changes his action from one period to another if and only if his opponent does.

The above structure theorem facilitates the computation of the set of

equilibrium payoffs for many games. It turns out that the reduction of the equilibrium payoff set as we move from an unrestricted repeated game to the one with costly states can be dramatic. To illustrate this we will concentrate on the repeated prisoners' dilemma game with the average payoff criterion. In particular we consider the lexicographic preferences with players considering their utility first and size of their automaton second. More specifically consider two pairs of automata $A = (A_1, A_2)$ and $\bar{A} = (\bar{A}_1, \bar{A}_2)$, $A >_i \bar{A}$ if and only if:

1. $\bar{u}_i(A) > \bar{u}_i(\bar{A})$, or
2. $\bar{u}_i(A) = \bar{u}_i(\bar{A})$ and $m_i < \bar{m}_i$

with m_i and \bar{m}_i representing the number of states of the automaton A and \bar{A}_i , respectively. (Recall that $\bar{u}_i(A)$ is the average payoff to player i of the stream of payoffs created by the pair A .) In Figure 2 we illustrate the reduction of the set of equilibrium payoffs. The shaded area describes the set of all equilibrium payoffs in the unrestricted game. The two line segments within the region represent the set of equilibrium payoffs of the game with costly states.




Figure 2

VI. Complexity in General Infinitely Repeated Games

In this section we concentrate on general (unmodified) infinitely repeated games. Some of the results described here deal with the sufficiency of equilibria using finite complexity strategies. It turns out that the set of equilibria consisting of finite complexity strategies approximates well the set of all equilibrium payoffs. Thus there is no discontinuity between what can be achieved by real unrestricted players and what can be achieved by players using finite automata. This may seem surprising in view of some of the striking results reported in Section V. One has to keep in mind, though, that there the game was restricted (for example, to be played through finite fixed size automata). This restriction was common knowledge to all the players and the set equilibria of the restricted game turned out to be substantially different from the unrestricted game. The approach here is very different. It asks: In an unrestricted game, what can be achieved by "simple" strategies? This means that these simple strategies have to be optimal also when compared with non-simple strategies.

Another result in this section describes interplayer complexity bounds at discount robust subgame perfect equilibria of repeated generic games. One interesting implication of this bound is that in two player games, the players must use equal complexity strategies at equilibria.

The notion of an automaton used throughout Section VI is the one of a full automaton.

VI.1. The Case of Discounting: Finite Complexity Equilibria Suffice and

Interpersonal Complexity Bounds

We first report results due to Kalai-Stanford [1988]. They address two different issues. The first as regards the set of equilibrium payoffs that are attained by equilibria of finite full complexity. It turns out that the restriction to finite complexity is not serious, and one can uniformly approximate all equilibrium payoffs of a given game by ones of finite bounded full complexity.

The second issue regards the interplayer complexity relations at equilibria. It is shown that generically at equilibria no player's strategic complexity exceeds the product of the complexities of his opponents. In particular, in a two-person game both players must use equal complexity strategies.

We start with an n -person stage game $G = (N, A, u)$ and assume that the utility functions of all the players in G are bounded. Next we consider $G^{\infty, \alpha}$, the infinitely repeated game G with the discounting criterion for some fixed discount parameter α . A strategy combination $f = (f_1, f_2, \dots, f_n)$ is of finite complexity if each f_i is of finite complexity.

Corollary 6.1:

$\{\text{subgame-perfect equilibrium payoffs of } G^{\infty, \alpha}\} \subseteq \bigcap_{\epsilon > 0} \text{Closure}\{\text{finite-complexity subgame perfect } \epsilon\text{-equilibrium payoffs of } G^{\infty, \alpha}\}.$

This corollary follows from the following finite approximation theorem.

Theorem 6.1: Given the game $G^{\infty, \alpha}$ and an $\epsilon > 0$ there is an integer M with the following properties. For every subgame perfect equilibrium f^* there is

a subgame perfect ϵ -equilibrium g^* with $|u_i(f^*) - u_i(g^*)| \leq \epsilon$ and the complexity $(g_i^*) \leq M$ for $i = 1, 2, \dots, n$.

One can actually construct a "universal" approximating n -vector of automata that can simultaneously approximate all the subgame perfect equilibria of the game by varying only the automata's initial states. Given $G^{\omega, \alpha}$ and the $\epsilon > 0$, there is an integer M and a vector of minimal full automata $(A_i)_{i \in N}$ with the following properties:

- 1) The complexity $(A_i) \leq M$ for $i = 1, 2, \dots, n$
- 2) For every subgame perfect equilibrium f^* we can choose initial states $(m_i)_{i \in N}$ for the automata $(A_i)_{i \in N}$, respectively, so that when started at these states the automata will play a subgame perfect ϵ -equilibrium whose utilities are within ϵ of the utilities of f^* .

Turning now to the interpersonal complexity bounds, we say that the stage game G is generic if for every $a, \bar{a} \in A$, if $a \neq \bar{a}$ then for every player i $u_i(a) \neq u_i(\bar{a})$. For example, the prisoner's dilemma game is generic. A subgame perfect equilibrium f^* of $G^{\omega, \alpha}$ is discount robust if for some $\epsilon > 0$ for all $\beta \in (\alpha - \epsilon, \alpha + \epsilon)$ f^* is a subgame perfect equilibrium of $G^{\omega, \beta}$.

Theorem 6.2: Let G be a generic game and let f^* be a discount robust subgame perfect equilibrium of $G^{\omega, \alpha}$ then for every $i \in N$ complexity $(f_i^*) \leq \prod_{j \neq i} \text{complexity}(f_j^*)$.

Corollary 6.2: under the assumptions of Theorem 6.3 if $n = 2$, then

$$\text{complexity } (f_1^*) = \text{complexity } (f_2^*)$$

VI.2 The Folk Theorems with Finite Complexities

The results reported here are due to Ben-Porath and Peleg [1987]. They provide the analogue of the finite approximation theorem (Theorem 6.1) when the discount parameter approaches 1 and for games with the average payoff criterion. In these cases the approximation turns out to be stronger and can be done by using full subgame perfect equilibria as opposed to ϵ -equilibria. In these limit cases (low discount or average payoffs) the set of equilibrium payoffs has the well known characterization given by the folk-theorems (see Aumann-Shapley [1976], Rubinstein [1977] and Fudenberg-Maskin [1986]). Thus one obtains finite complexity approximation to the set of equilibrium payoffs described by the folk-theorems.

We start with a stage game $G = (N, A, u)$ with compact metric spaces A_i 's and continuous utility functions u_i 's. The pure strategy minmax value of player i , v_i , is defined by

$$v_i = \min_{a \in A} \max_{\bar{a}_i \in A_i} u_i(a_1, \dots, a_{i-1}, \bar{a}_i, a_{i+1}, \dots, a_n).$$

In other words, v_i is the utility level that player i can secure for himself if he chooses his strategy after observing the choices of his opponents. A vector $x \in \mathbb{R}^n$ is strictly individually rational if $x_i > v_i$ for $i = 1, 2, \dots, n$. A vector $x \in \mathbb{R}^n$ is a rational convex combination of pure payoffs of G if $x = \sum_{j=1}^k r_j x_j$ with entry r_j being a nonnegative rational number, $\sum_{j=1}^k r_j = 1$, and every $x_j = u(a)$ for some $a \in A$.

We first describe the folk theorem type of characterization for the finite complexity subgame perfect equilibrium payoffs of \bar{G}^∞ , the infinitely

repeated version of G with the average payoff criterion.

Theorem 6.3: Suppose x is a rational convex combination of pure payoffs of G and is strictly individually rational. There is a finite full complexity subgame perfect equilibrium of \bar{G}^∞ , f^* , with $\bar{u}(f^*) = x$.

Now we turn to the case of low discounting. We say that x is an interior rational convex combination of pure payoffs of G if it is a rational convex combination of pure payoffs of G and it is in the interior of the convex hull of the pure payoffs of G .

Theorem 6.4: Suppose x is an interior rational convex combination of pure payoffs of G which is strictly individually rational. There is a discount parameter α_0 ($0 < \alpha_0 < 1$) and finite complexity strategy combination f with $\lim_{\alpha \rightarrow 1} u_\alpha(f) = x$, and f being a subgame perfect equilibrium of $G^{\infty, \alpha}$ for all $\alpha \in (\alpha_0, 1)$.

References

- Abreu, D. [1984], "Infinitely Repeated Games with Discounting: A General Theory," manuscript.
- Abreu, D. and A. Rubinstein [1986], "The Structure of Nash Equilibrium in Repeated Games with Finite Automata," manuscript.
- Aumann, R. J. and L. S. Shapley [1976], "Long Term Competition: A Game Theoretic Analysis," mimeo.
- Aumann, R. J. [1981], "Survey of Repeated Games," in Essays in Game Theory and Mathematical Economics in Honor of Oskar Morgenstern, Bibliographisches Institut Mannheim/Wien/Zurich.
- Aumann, R. J. and S. Sorin [1985], "Bounded Rationality Forces Cooperation," manuscript.
- Aumann, R. J. [1988], "Irrationality in Game Theory," Hebrew University, mimeo.
- Axelrod, R. [1980], "Effective Choice in the Prisoner's Dilemma," Journal of Conflict Resolution, 24, 3-25.
- Ben Porath, E. [1986], "Repeated Games with Bounded Complexity," manuscript.
- Ben-Porath, E. and B. Peleg [1987], "On the Folk Theorem and Finite Automata," The Hebrew University discussion paper.
- Binmore, K. G. [1987], "Remodeled Rational Players," discussion paper, London School of Economics.
- Friedman, J. [1987], "A Non-cooperative Equilibrium for Supergames," Review of Economic Studies, 38, 1-12.
- Fudenberg, D. and E. Maskin [1986], "The Folk Theorem in Repeated Games with Discounting and with Incomplete Information," Econometrica, 54, 533-

554.

- Futia, C. [1977], "The Complexity of Economic Decision Rules," Journal of Mathematical Economics, 4, 289-299.
- Gilboa, I. [1986], "The Complexity of Computing Best Response Automaton in Repeated Games," Tel-Aviv University Discussion Paper.
- Gilboa, I. and D. Samet [1987], "Bounded Versus Unbounded Rationality: The Strength of Weakness," Tel-Aviv University discussion paper.
- Hopcroft, J. and J. Ullman [1979], Introduction to Automata Theory, Languages, and Computation. Reading, Massachusetts: Addison-Wesley.
- Kalai, E., D. Samet and W. Stanford [1985], "A Note on Reactive Equilibria in the Discounted Prisoners' Dilemma and Associated Games," forthcoming in International Journal of Game Theory.
- Kalai, E., D. Samet and W. Stanford [1986], "Complexity in Stochastic Games," Northwestern University discussion paper.
- Kalai, E. and W. Stanford [1988], "Finite Rationality and Interpersonal Complexity in Repeated Games," Econometrica, 56, No. 2, 397-410.
- Kreps, D., P. Milgrom, J. Roberts and R. Wilson [1982], "Rational Cooperation in the Finitely-Repeated Prisoner's Dilemma," Journal of Economic Theory, 27, 245-252.
- Lehrer, E. [1987], "Repeated Games with Stationary Bounded Recall Strategies," Hebrew University discussion paper.
- Lewis, A. [1985], "On Effectively Computable Realizations of Choice Functions," Mathematical Social Science, 10, 43-80.
- Lipman, B. and S. Srivastava [1987], "Informational Requirements and Strategic Complexity in Repeated Games," Carnegie-Mellon University discussion paper.

- Lipman, B. [1987], "Procedural Rationality in Decisions and Games," discussion paper, Carnegie-Mellon University.
- Megiddo, N. [1986], "Remarks on Bounded Rationality," I.B.M. manuscript, San Jose.
- Megiddo, N. and A. Wigderson [1986], "On Play by Means of Computing Machines," in Theoretical Aspects of Reasoning and Knowledge, J. Halpern, ed., pp. 259-274, Los Altos, Calif., Morgan Kaufman.
- Mertens, J. [1987], "Repeated Games," Proceedings of the International Congress of Mathematicians, Berkeley, 1528-1577.
- Moore, E. [1956], "Gedanken Experiments on Sequential Machines," in Automata Studies. Princeton, New Jersey: Princeton University Press.
- Mount, K. and S. Reiter [1983], "Computational Complexity of Resource Allocation," manuscript.
- Neyman, A. [1985], "Bounded Complexity Justifies Cooperation in the Finitely Repeated Prisoners' Dilemma," Economics Letters, 19, 227-229.
- Radner, R. [1980], "Collusive Behavior in Noncooperative Epsilon-equilibria of Oligopolies with Long but Finite Lives," Journal of Economic Theory, 22, 136-154.
- Rosenthal, R. [1982], "Games of Perfect Information, Predatory Pricing, and the Chain Store Paradox," Journal of Economic Theory, 25, 92-100.
- Rubinstein, A. [1977], "Equilibrium in Supergames," Center for Research in Mathematical Economics and Game Theory, The Hebrew University, R.M. 25.
- Rubinstein, A. [1986], "Finite Automata Play the Repeated Prisoners' Dilemma," Journal of Economic Theory, 39, 83-96.
- Rubinstein, A. [1987], "Is There a 'Finite Automata' Resolution to the Finitely Repeated Games Paradoxes?", London School of Economics

Discussion Paper.

- Selten, R. [1975], "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," International Journal of Game Theory, 4, 25-55.
- Selten, R. [1978], "The Chain Store Paradox," Theory and Decisions, 9, 127-159.
- Simon, H. [1972], "Theories of Bounded Rationality," in Decision and Organization. Edited by McGuire and R. Radner, Amsterdam: North-Holland.
- Smale, S. [1980], "The Prisoner's Dilemma and Dynamical Systems Associated to Non-cooperative Games," Econometrica, 48, 1617-1634.
- Sorin, S. [1988], "Supergames," this volume.
- Stanford, W. [1986a], "Subgame Perfect Reaction Function Equilibria in Discounted Duopoly Supergames are Trivial," Journal of Economic Theory, 39, 226-232.
- Stanford, W. [1986b], "On Continuous Reaction Function Equilibria in Duopoly Supergames with Mean Payoffs," Journal of Economic Theory, 39, 233-250.
- Stanford, W. [1987], "Some Simple Equilibria in Repeated Games," University of Illinois at Chicago discussion paper.
- Van-Damme, E. [1983], Refinements of the Nash Equilibrium Concept, Springer-Verlag.
- Varian, H. [1975], "Complexity of Social Decisions," manuscript.
- Zemel, E. [1985], "On Communication, Bounded Complexity and Cooperation," Northwestern University discussion paper.