

Discussion Paper No. 679

FINITE RATIONALITY AND
INTERPERSONAL COMPLEXITY
IN REPEATED GAMES*

by

Ehud Kalai**

and

William Stanford**

April 1986

*The authors wish to thank Robert Aumann and Eitan Zemel for helpful observations and suggestions.

**Department of Managerial Economics and Decision Sciences, J. L. Kellogg Graduate School of Management, Northwestern University, Evanston, Illinois 60201.

This research was partly supported by NSF Economics Grant No. 84-09798.

FINITE RATIONALITY AND INTERPERSONAL COMPLEXITY IN REPEATED GAMES

by

Ehud Kalai and William Stanford

Abstract

Every subgame perfect equilibrium of a repeated game with discounting can be approximated by one of finite complexity. Generically at equilibria, interpersonal complexity bounds exist and at equilibria of two person games the players must use strategies of equal complexity and equal memory. The measure of the complexity of a strategy, introduced in this paper, is defined to be the number of distinct strategies induced by it in the various subgames. It also equals the smallest number of states of an automaton which can fully implement the strategy. New computational methods for the subgame perfect equilibrium payoffs of the game are available by the use of complexity properties at equilibria. Further related results are also proved.

1. Introduction

In this paper we introduce a measure of strategic complexity for a strategy of a repeated game. Having such a measure is important for several reasons:

1. It gives us a measure of the difficulty involved in conceptualizing or describing a player's strategy. This is important because some strategies are so difficult to conceptualize that it is unlikely that players would ever use them. Moreover, strategies which are not describable in finite time could not be used in games where a player has several agents working for him--for example, a firm whose strategy has to be carried out by its employees.

2. It measures the possible difficulty that a player may face in checking the rationality (for example Nash, or (subgame) perfect equilibrium) of a strategy combination. As we will point out later, this task may vary from being trivial to impossible depending upon the complexities of the underlying strategies.

3. It gives us better understanding of the structure of strategies and equilibria. This contrasts with previous analyses which have concentrated on characterizing equilibrium payoffs. (Exceptions are Rubinstein [1985], Abreu [1985], and Stanford [1985].) This new understanding aids us in many ways--for example, it enables us to generate new methods for computing and approximating equilibria and equilibrium payoffs.

In order to illustrate these points and our results we consider the two person repeated prisoners' dilemma game with the labels *c* for "cooperate" and *d* for "defect" describing the stage game actions available to each of the two players. Let *C* and *D* denote the constant cooperation and defection strategies, respectively. Thus, for example under *D*, a player defects in

every stage of the game regardless of past history. These are the simplest strategies that one may think of in this repeated game. On the other hand, the tit-for-tat rule (choose for this period the action that your opponent chose in the previous period) induces two strategies which are specified by the choice of an initial action. We denote them by c-tft and d-tft. While these strategies are more complicated than the constant strategies they are still quite simple.

To illustrate an example of a highly complicated strategy we do the following. We consider the binary expansion of the irrational number π (or any irrational number will do), and label the periods in the repeated game as 0-periods or 1-periods according to the corresponding entries in the binary expansion (this assures us that there is no regularity in the labeling of periods). A player playing the π strategy will defect on every 0-period. On 1-periods the player will play the "trigger two period punishment" strategy, which means the following. After observing a defection by anyone (including himself) on a 1-period the player will defect for the following next two 1-periods (he will not trigger, however, on the punishment phase itself). Otherwise he will continue cooperating.

It seems that the constant and the tit-for-tat strategies are much more likely to be used by players because of their simplicity. Our measure will assign complexity 1 to the constant strategies, complexity 2 to the tft strategies and infinite complexity (aleph-naught) to the π strategy. Constant strategies are assigned complexity one because they induce only one mode of behavior (strategy) regardless of the histories of plays that might have led to the current stage. The cooperate and then tit-for-tat strategy, c-tft, is somewhat more involved. Depending on the history of plays that led to the current stage the player strategy from the current stage on may be c-tft or

d-tft (depending on his opponent's last period action). Thus, the c-tft strategy induces two different strategies and for this reason its complexity is defined to be 2. It is easy to see that the π strategy induces infinitely many different strategies as we consider all the possible histories that may be played, and thus its complexity is infinite. We define the complexity of a strategy to be the number of different strategies that it may induce after the various histories of plays.

The task of checking the rationality (Nash equilibrium, perfect equilibrium, etc.) of a strategy combination becomes more difficult as the players use more complex strategies. To check the constant strategy pair (D,D) for perfection a player has to solve at most one optimization problem. This is the case because under this pair only one situation ($1 = \text{complexity}(D) \times \text{complexity}(D)$) occurs repeatedly. To check for perfection of the pair of strategies (c-tft, c-tft) no more than four optimization problems will have to be solved. Because four ($4 = \text{complexity}(c-tft) \times \text{complexity}(c-tft)$) is the largest number of different induced strategy pairs that may arise after all the possible histories of play.

A natural interpretation of the complexity measure is through the computer science notion of an automaton (finite or infinite). With this interpretation we think of a player as having a set of states of mind. Each state of mind completely summarizes the past for the player, and the action of the player in every stage of the game is completely determined by his current state. There is one unique state, called the initial state, with which he starts the game. As the game progresses the player transits from one state to another depending upon the action combination taken by all the players (including his own) in the previous stage. After this transition, his new state determines his next period action and so on. It is important to note,

and will be expanded upon later, that the transition from a state to the next is defined also for action combinations in which the player himself took an action which was not prescribed by his automaton. In other words, the transition rule from a state to a state can cope with mistakes. With this notion of an automaton we argue that every strategy in a repeated game can be fully described by an automaton. Moreover, we argue that the measure of complexity of a strategy discussed above equals the number of states of the smallest automaton that describes it. Thus one may think of the complexity of a strategy as the minimal number of states of mind that the player must possess in order to consider the strategy.

The results can be divided naturally into two categories. The first concerns the possibility of low complexity approximations for subgame perfect equilibria. More explicitly, suppose we are given a repeated game and a vector of discounted payoffs for the players which is realizable in perfect equilibrium. Then we can find a finite complexity strategy vector for the repeated game which is an approximate perfect equilibrium and yielding a vector of discounted payoffs very close to the payoffs in the initial equilibrium. Moreover, the simultaneous approach to full subgame perfection and to the initial equilibrium payoffs can be made as close as we desire, retaining the finite complexity character of the approximating strategy vector. Thus we find that imposing the condition of finite complexity on players' decision procedures is unrestrictive in the sense made precise above. This result appears to be well within the spirit of any natural theory of practical computations since any such theory would be based upon the idea of finite but potentially unbounded computing power.

The second category of results has to do with the structure of perfect equilibria in discounted repeated games. By way of introduction, note that

perfection requires the players to find themselves in an equilibrium state under all possible contingencies. However, in the standard repeated games model, contingencies are merely histories. Thus, in some sense, there must be agreement among the players concerning the important aspects of a given history and what should be done by all players in the light of that agreement. This would seem to argue in favor of necessary complexity relationships among the strategies. Our second category of results helps to make this precise. For an interesting class of stage games (including the Prisoner's Dilemma) and a reasonable class of perfect equilibria, we demonstrate that certain complexity relationships must indeed hold. In particular, the complexity of each player's strategy must be bounded above by the product of the complexities of strategies adopted by all other players. In the two player case, this reduces to equal complexity. Actually, all of this is driven by a stronger result which says that equilibrium is completely determined by the strategy choice of any group of $N - 1$ players. Thus two-player perfect equilibria are prescribed by the strategy choice of either player, and we may as well refer to an equilibrium as a strategy instead of as a strategy pair.

A second result in the second category focuses on another interesting class of two-player games and related ideas. In this class of stage games, players may well employ strategies with different complexities in perfect equilibria of the repeated game. However, it will prove possible to modify the strategy of one or both players in a manner which simultaneously equates complexities and preserves the payoff structure of the equilibrium in all contingencies, and hence the equilibrium itself. Thus, in this case, little is lost by restricting the search for perfect equilibria to strategy pairs of equal complexity. Results of this kind are then used to outline some

computational aspects of subgame perfect equilibria in repeated games with finite action stage games.

The importance of incorporating bounded rationality into economics has been argued by Simon since the 1950s (see Simon [1972]). Notions of bounded rationality and computational complexity have been studied by many mathematicians and economists. See, for example Mount-Reiter [1983], Smale [1980], and Futia [1977]. The notion of an automaton in economics has been used by Varian [1975] and Lewis [1985]. A notion of bounded rationality which is especially important for this paper is the notion of ϵ -equilibrium used by Radner [1980] in his analysis of repeated oligopoly games.

The idea of using finite automata in order to separate out a "simple" class of repeated game strategies was proposed by Aumann [1981] in his well-known survey of repeated games. In this survey he proposed the automaton notion with states describing the states of minds of the players. He then proposed that simple strategies are those that use a finite number of states. In two recent path-breaking papers, Neyman [1985] and Rubinstein [1985] used Aumann's automaton notion in order to model two different problems of bounded rationality within the context of repeated games.

Neyman studies repeated games in which the strategies of the players are restricted. His restrictions consist of requiring that a player's strategy be describable by an Aumann type of automaton of a given fixed size. It could also be thought of as the game being played among the programmers who have different size machines at their disposal (the Axelrod [1980] type of experiments are a very good illustration of this type of game). With this model he is able to carry a meaningful analysis of the effects of computing ability on the final outcome of the repeated game.

Rubinstein uses Aumann type automata in an infinitely repeated prisoners'

dilemma game. However, he incorporates it straight into the solution concept. He modifies the notion of Nash equilibrium to still have the players maximize their utility but to also minimize the number of states they use at the given equilibrium. As he illustrates, this results in a significant reduction in possible payoffs of the kind found in the "folk-theorem." Extensions of these results and other interesting recent studies of bounded rationality can be found in Ben-Porath [1986], Megiddo-Wigderson [1985], Zemel [1986], Abreu-Rubinstein [1986], and Aumann-Sorin [1985].

Our approach is different from the ones cited in several important respects. Unlike them we work with standard unmodified notions of repeated games and strategies. It is within this context that we define the complexity measure which allows us to sort strategies according to their complexity. The classical Kuhn's notion of a strategy is defined for every history, whether it contradicts its own earlier prescribed moves or not [see Owen [1982]]. For this reason we choose an automaton notion which is richer than the one proposed by Aumann. The key difference here is that our automaton has as its input the player's own action in addition to the actions of his opponents, while in Aumann, Neyman, and Rubinstein automata, only the actions of the opponents are input. This difference enables our automata to read every history of past plays and not just the self consistent histories. Viewing the automata as programs to implement strategies, the Aumann automata may be thought of as "exact implementations" of plays of strategies and have no prescriptions of how to react to self-made mistakes. Our automata are programmed also with reactions to self-made mistakes, and in this sense can be thought of as the complexity of full (even if sometimes wasteful) implementations. We could follow the other route and define the complexity of a strategy to be the number of states required for its "exact

implementation." This would result in a smaller measure of complexity and thus our finite approximation result will hold for it as well. We believe also that the interpersonal complexity relations will hold for this smaller measure. However, we feel that if it is the strategic complexity, and not the exact implementation complexity, that we want to measure, then the one we use is the appropriate one.

2. Notation and Definitions

By an N-person stage game we mean a pair (A, u) . $A = \prod_{i=1}^N A_i$ describes the action combinations available to the N players. $u = (u_1, u_2, \dots, u_N)$ is a vector of utility functions where $u_i: A \rightarrow \mathbb{R}$ represents player i's preferences. We assume u_i is bounded, i.e., there is an $M_i \in \mathbb{R}$ with $|u_i(a)| < M_i$ for every $a \in A$, and $i = 1, 2, \dots, N$.

For a discount parameter α , $0 < \alpha < 1$, we define the infinitely repeated game $G^\infty(A, \alpha)$ as follows. The set of histories of length 0 is a singleton set denoted by H^0 . Its single element will be denoted by e . For a positive interger n the set of histories of length n is defined by $H^n = A^n$. the set of all histories is defined by $H = \bigcup_{n=0}^{\infty} H^n$.

For every $h \in H$, define $h^r \in A$ to be the projection of h onto its r-th coordinate. For every $h \in H$ we let $\ell(h)$ denote the length of h. Thus, for $h \in H^n$, we have $\ell(h) = n$. For two positive length histories h and $\bar{h} \in H$ we define the concatenation of h and \bar{h} , in that order, to be the history $(h \cdot \bar{h})$ of length $\ell(h) + \ell(\bar{h})$:

$$(h \cdot \bar{h})^r = h^r \text{ if } 1 \leq r \leq \ell(h),$$

and

$$(h \cdot \bar{h})^r = \bar{h}^{r-\ell(h)} \text{ if } \ell(h) < r \leq \ell(h) + \ell(\bar{h}).$$

We also make the convention that $e \cdot h = h \cdot e = h$ for every $h \in H$.

An individual strategy in the repeated game is a function $f_i: H \rightarrow A_i$. We let F_i denote the set of all such individual strategies for player i and we let $F = \prod_{i=1}^N F_i$ denote the set of strategy vectors. $f \in F$ is then a function $f: H \rightarrow A$ defined by $f(h) = (f_1(h), f_2(h), \dots, f_N(h))$. Define $(\bar{f}_{-i}, \bar{f}_i) = (f_1, \dots, f_{i-1}, \bar{f}_i, f_{i+1}, \dots, f_N)$ for $i = 1, 2, \dots, N$.

Given $f \in F$, define a sequence of functions $f^n: H \rightarrow H$ by:

$$\begin{aligned} f^0(h) &= h \\ f^1(h) &= h \cdot f(h) \\ &\vdots \\ f^n(h) &= f^{n-1}(h) \cdot f(f^{n-1}(h)). \end{aligned}$$

Now extend the utility functions from A to F by defining

$$u_i^\alpha(f) = \sum_{r=1}^{\infty} \alpha^{r-1} u_i(f(f^{r-1}(e))).$$

Finally, the repeated game derived from (A, u) with discount parameter α , $G^\infty(A, u, \alpha)$ is defined to be the pair (F, u^α) .

For a history $h \in H$ and an integer $0 \leq m \leq \ell(h)$, the m -stage end tail of h is denoted by $E^m(h) \in H$:

$$(E^m(h))^i = h^{\ell(h)-m+1} \quad \text{for } i = 1, 2, \dots, m$$

and the m -stage beginning of tail h is denoted by $B^m(h) \in H$:

$$(B^m(h))^i = h^i \text{ for } i = 1, 2, \dots, m.$$

Given an individual strategy $f_i \in F_i$ and a history h we denote the individual strategy induced by f_i at h by $f_i|_h$. $f_i|_h$ is defined pointwise on H :

$$(f_i|_h)(\bar{h}) = f_i(h \cdot \bar{h}), \text{ for every } \bar{h} \in H.$$

Clearly $f_i|_h \in F_i$, $f_i|_e = f_i$ and for every history h and integer m with $0 < m < l(h)$,

$$f_i(h) = (f_i|_{B^m(h)})(E^{l(h)-m}(h)).$$

We will use $(f|_h)$ to denote $(f_1|_h, f_2|_h, \dots, f_N|_h)$ for every $f \in F$ and $h \in H$. We let $F_i(f_i) = \{f_i|_h : h \in H\}$ and for $f \in F$, we let $F(f) = \{(f|_h) : h \in H\}$. Clearly, $F_i(f_i) \subset F_i$, $F(f) \subset F$, and $F(f) \subset \prod_{i=1}^N F_i(f_i)$. Typically, these three containments will be proper and $F(f)$ need not be the product of subsets of the $F_i(f_i)$.

A strategy vector $f \in F$ is a Nash equilibrium of $G^\alpha(A, u, \alpha)$ if

$$u_i^\alpha(f) > u_i^\alpha(f_{-i}, \hat{f}_i)$$

for all $\hat{f}_i \in F_i$ and $i = 1, 2, \dots, N$. A strategy vector $f \in F$ is a subgame perfect equilibrium of $G^\alpha(A, u, \alpha)$ if every $\bar{f} \in F(f)$ is a Nash equilibrium (see Selten [1975]). $f \in F$ is a discount robust subgame perfect equilibrium of $G^\alpha(A, u, \alpha)$ if f is a subgame perfect equilibrium of $G^\beta(A, u, \beta)$ for every β in some neighborhood of α . Since this condition on equilibrium is extensively

used in this paper we abbreviate it by saying that f is a DRSP equilibrium.

Two strategy vectors, f and g in F , are payoff equivalent if $u^\alpha(f) = u^\alpha(g)$. They are perfectly payoff equivalent if $u^\alpha(f|h) = u^\alpha(g|h)$ for every $h \in H$.

3. Complexity and Automaton Implementations

In this section, we give the complexity definition for repeated game strategies and establish the simple but basic relationship between strategies and computing machines called automata. There are at least three reasons for doing the latter. First, the equivalence gives substance to the complexity notion which can then be viewed as a measure of computing power inherent in the strategy. Second, in the case of finite complexity strategies--or equivalently, finite automata--we are restricting attention to a class of objects whose properties and capabilities have been studied extensively and independently of game theoretic considerations. See, for example, Hopcroft and Ullman [1979] for a state-of-the-art introduction and references. Finally, the equivalence between strategies and automata makes available an interpretation and way of thinking about strategies which has materially aided us in our research. It also makes available a notation scheme which we will find convenient to use at several points in the paper.

We now proceed to the complexity definition for strategies and a discussion of automata.

Given a strategy $f_i \in F_i$ we define the complexity of f_i , $\text{comp}(f_i)$, to be the cardinality of the set $F_i(f_i) = \{f_i|h: h \in H\}$. Thus the complexity of a strategy is measured by asking how many strategies it induces.

As we have seen, constant strategies are assigned complexity one; while the tit-for-tat strategy is judged to be of complexity two. Also, again in the repeated Prisoner's Dilemma, consider an n -period "trigger" strategy.

This strategy plays cooperatively until either player chooses his noncooperative action and then plays noncooperatively for n periods before returning to cooperation. The complexity of such a strategy is judged to be $n + 1$, corresponding to a cooperative "state," and n punishment states across which the player must "transition" as a punishment cycle is played out.

That the above reference to "states" and "transitions" is not an aimless convenience is brought home in the following discussion.

By an automaton implementation for player i , we mean a triple $I_i = ((S_i, s_i^0), T_i, B_i)$ where:

S_i is a set of states;

s_i^0 is an initial state;

$T_i: S_i \times A \rightarrow S_i$ is a transition function; and

$B_i: S_i \rightarrow A_i$ is a behavior function.

Given an implementation I_i for player i , it induces a strategy f_i for player i as follows:

For a history h with $l(h) = m$ define inductively

$$s_i^0(h) = s_i^0, \quad s_i^r(h) = T_i(s_i^{r-1}, h^r) \text{ for } r = 1, 2, \dots, m$$

and then

$$f_i(h) = B_i(s_i^m(h)).$$

The cardinality of an implementation I_i is defined to be the cardinality of

its S_i , $\text{Card}(S_i)$.

In the case of finite S_i , the above definition corresponds to a class of finite automata called Moore [1956] machines. Moore machines are distinguished among finite automata essentially because the behavior function B_i is allowed to take on more than two values. Again, see Hopcroft and Ullman [1979] for details and further references.

The fact that the transition function of a player's automaton depends on the player's own move ($T_i: S_i \times A \rightarrow S_i$ rather than the more restrictive one, $T_i: S_i \times A_{-i} \rightarrow S_i$) is crucial. It represents one of the significant differences between the Aumann [1981], Neyman [1985], Rubinstein [1985] papers and ours. When one restricts oneself to the second type of automata, many strategies cannot be modeled as such. For example, the grim trigger strategy in the repeated prisoners' dilemma game (trigger defection on anybody's defection including one's own) could not be represented by a restricted automaton. Thus, in order to obtain a measure of complexity for every strategy we need the larger class of automata. We then have the following lemma.

Lemma 3.1: Every strategy $f_i \in F_i$ is implementable by some automaton.

Proof: We let $I_i = ((H,e), T_i, B_i)$ be defined by $T_i(h,a) = h \cdot a$ and $B_i(h) = f_i(h)$. \square

Theorem 3.1 below is the main result of this section. In light of this result, it is reasonable to think of a repeated game strategy in the following terms: in each period, the player finds himself in some state of mind, which gives rise to action as a function of state. The player then observes the actions of all players, including his own. This datum, together with his original state of mind determines a new state which in turn governs his action

choice in the next period. This process continues through the course of the game. As we will see, states of mind correspond to induced strategies, action as a function of state is determined as the first period action of the corresponding induced strategy, and transition from state to state is governed by the transition law for induced strategies.

Theorem 3.1: For every $f_i \in F_i$ $\text{comp}(f_i)$ is the cardinality of the smallest automaton implementing f_i .

Proof: Let $I_i = ((S_i, s_i^0), T_i, B_i)$ be an implementation of f_i . We first show that $\text{Card}(S_i) \geq \text{Card}(F_i(f_i)) = \text{comp}(f_i)$. We assume without loss of generality that for every $s_i \in S_i$ there is a history $h \in H$ with $s_i^r(h) = s_i$ for some $r > 0$. Now define $C_i: S_i \rightarrow F_i(f_i)$ by letting $C_i(s_i)$ be the strategy induced by $\bar{I}_i = ((S_i, s_i), T_i, B_i)$. It suffices to show that C_i is onto. But for every $\bar{f}_i \in F_i(f_i)$ we have $\bar{f}_i = f_i|_h$ for some $h \in H$ and hence $C_i(s_i) = \bar{f}_i$ if we let $s_i = s_i^{\lambda(h)}(h)$.

To see that the $\text{comp}(f_i)$ equals the cardinality of some implementation of f_i define

$$I_i = ((F_i(f_i), f_i), T_i, B_i) \text{ with}$$

$$T_i(\bar{f}_i, a) = \bar{f}_i|_a \text{ and } B_i(\bar{f}_i) = \bar{f}_i(e).$$

It is easy to see that I_i implements f_i and obviously it has the right cardinality. \square

In the computer science literature, this result is known as the Myhill-Nerode theorem,¹ which is proved there for the finite cardinality case. In the computer science vocabulary, the theorem says that there is an essentially

unique minimum state finite automaton for every "regular set" of symbol strings. See Hopcroft and Ullman [1979], Myhill [1957], and Nerode [1958].

To illustrate examples of automaton implementations we consider first the example where player I plays cooperatively and then tit-for-tat, c-tft, in the repeated prisoners' dilemma with the stage actions being labeled C_i and N_i describing the "cooperative" and "noncooperative" choices of the two players.

States	<u>Transition Function</u>				Behavior Function
	<u>Action Combinations</u>				
	(C_1, C_2)	(C_1, N_2)	(N_1, C_2)	(N_1, N_2)	
$s_1^0 = \bar{C}_1$	\bar{C}_1	\bar{N}_1	\bar{C}_1	\bar{N}_1	C_1
\bar{N}_1	\bar{C}_1	\bar{N}_1	\bar{C}_1	\bar{N}_1	N_1

On the other hand, the following player I automaton describes the trigger 2 phase punishment strategy. This is related to the trigger strategies introduced by Friedman [1971].

States	<u>Transition Function</u>				Behavior Function
	<u>Action Combinations</u>				
	(C_1, C_2)	(C_1, N_2)	(N_1, C_2)	(N_1, N_2)	
\bar{C}_1	\bar{C}_1	\bar{P}_1	\bar{P}_1	\bar{P}_1	C_1
\bar{P}_1	\bar{P}_2	\bar{P}_2	\bar{P}_2	\bar{P}_2	N_1
\bar{P}_2	\bar{C}_1	\bar{C}_1	\bar{C}_1	\bar{C}_1	N_1

¹We are grateful to Eitan Zemel for pointing out the relationship of our Theorem 3.1 to the Myhill-Nerode theorem.

4. Finite Complexity Approximations

In this section, we consider finite complexity approximations to (potentially) infinitely complex perfect equilibria. We contrast this with the approach of Fudenberg and Levine [1983] in which perfect equilibria of infinite horizon games are characterized as the limits of perfect approximate equilibria of finite horizon truncations of the original game.

Before proceeding to the results, and in order to bring some of the relevant questions more sharply into focus, we will first discuss an example. This example is based on a family of infinitely complex perfect equilibria in the discounted Prisoner's Dilemma.

Example 4.1: Our stage game is represented by the following bimatrix game:

		Player 2	
		N_2	C_2
Player 1	N_1	0,0	a_1, b_2
	C_1	a_2, b_1	1,1

To make this a Prisoner's Dilemma, we require $a_1, b_1 > 1$ and $a_2, b_2 < 0$. C_i is associated with "cooperation" and N_i with "noncooperation." Thus $A_i = \{C_i, N_i\}$, and $u_i: A \rightarrow R$ is the payoff to player i as given in the above matrix. Given α , we then form $G^\alpha(A, \alpha)$.

In the following, we construct an infinitely complex perfect equilibrium for $G^\alpha(A, \alpha)$. Consider an irrational number x between zero and one. x has a

binary expansion $\langle b_n \rangle_{n=1}^{\infty}$ which is an infinite vector of 0's and 1's. Thus $x = \sum_{n=1}^{\infty} (b_n/2^n)$ and this expansion is unique. Moreover, $\langle b_n \rangle$ is aperiodic, which means precisely that

$$\langle b_{n+m} \rangle_{n=1}^{\infty} \neq \langle b_{n+l} \rangle_{n=1}^{\infty}$$

whenever $m, l > 0$ and $m \neq l$. Thus infinite ending tails of $\langle b_n \rangle$ which are complements of different finite initial tails must in fact be different. We construct a pair of strategies based on x by first modifying $\langle b_n \rangle$. Define a new infinite binary vector $\langle c_n \rangle$ by: $c_{2i} = 1$ for $i = 1, 2, \dots$, and $c_{2i-1} = b_i$ for $i = 1, 2, \dots$. Thus $\langle c_n \rangle$ is just $\langle b_n \rangle$ with 1's interspersed between the components. Clearly, $\langle c_n \rangle$ is also aperiodic. Define the strategy f_i for player i as follows: along the equilibrium path, play cooperatively in period n if and only if $c_n = 1$. In the event of defection from this rule by either player, transition to an absorbing state for which the behavior rule specifies the noncooperative action. One might label this the "grim" trigger strategy. This strategy has (countably) infinite complexity because infinitely many outcome paths are induced as play proceeds along the equilibrium path. If both players adopt this strategy, it is also easy to verify that they are in perfect equilibrium if α exceeds some critical lower bound. (This is the reason for interspersing 1's between the components of $\langle b_n \rangle$. In the absence of this device, it is conceivable that arbitrarily large gaps between 1's might reasonably tempt defection for any pre-specified value of α).

The point of this example is simply that infinitely complex perfect equilibria exist, even in discounted repeated games based on finite bimatrix stage games. The existence question for finite complexity approximations is

easily answered in this family of equilibria. One needs only to replace a sufficiently distant infinite tail of $\langle c_n \rangle$ with an infinite vector of 1's and leave the specifications of the f_i as is. This results in a perfect equilibrium of finite complexity in which the payoffs are close to those in the equilibrium based on $\langle c_n \rangle$. What makes this work is that the worst punishment payoff for each player can be realized simultaneously in finite complexity perfect equilibrium, which is a consequence of the structural characteristics of the Prisoner's Dilemma as a stage game. The rest of this section is concerned with extending this kind of result to repeated games based on much more general stage games.

In the rest of this section, we assume only that u_i is bounded, i.e., there is an $M_i \in \mathbb{R}$ with $|u_i(a)| \leq M_i$ for every $a \in A$ and $i = 1, 2, \dots, N$. The main result is given in Theorem 4.1. It shows that perfect equilibria in discounted games can be approached (with regard to payoffs) by finite complexity approximately perfect equilibria of the repeated game.

Before proceeding, we need to make precise the notion of approximately perfect equilibrium. A strategy vector $f \in F$ is a (Nash) ϵ -equilibrium of $G^\omega(A, u, \alpha)$ if

$$u_i^\alpha(f) \geq u_i^\alpha(f_{-i}, \hat{f}_i) - \epsilon$$

for all $\hat{f}_i \in F_i$ and $i = 1, 2, \dots, N$. A strategy vector $f \in F$ is a subgame perfect ϵ -equilibrium of $G^\omega(A, u, \alpha)$ if every $\bar{f} \in F(f)$ is an ϵ -equilibrium of $G^\omega(A, u, \alpha)$.

Our first lemma, which is auxiliary in nature, is adopted from a result of Harris [1984]. According to this lemma, in discounted games, it is easy to check whether a given strategy vector is a perfect ϵ -equilibrium; we need only

consider one-shot deviations by the players. This of course contrasts with the case of evaluation relations based on a zero interest rate where it must be verified that arbitrary sequences of deviations are not overly profitable in any subgame.

Lemma 4.1: Let $f \in F$ and suppose that no single period deviation against f is more than $(1 - \alpha)\epsilon$ profitable for any player in any subgame. Then f is a subgame perfect ϵ -equilibrium.

Proof: Fix $h \in H$ and suppose player i considers employing the strategy g_i instead of f_i . For each $k \geq 0$, define $g_{i,k}$ (a strategy for player i) by

$$g_{i,k}(\bar{h}) = \begin{cases} g_i(\bar{h}) & \text{if } \lambda(\bar{h}) < k \\ f_i(\bar{h}) & \text{if } \lambda(\bar{h}) \geq k. \end{cases}$$

Then if $\lambda(h) < k$,

$$\begin{aligned} u_i^\alpha(f_{-i}, g_{i,k} | h) &\leq u_i^\alpha(f_{-i}, g_{i,k-1} | h) + \alpha^{k-\lambda(h)-1} (1 - \alpha)\epsilon \\ &\leq u_i^\alpha(f_{-i}, g_{i,k-2} | h) + (\alpha^{k-\lambda(h)-2} + \alpha^{k-\lambda(h)-1}) (1 - \alpha)\epsilon \\ &\quad \cdot \\ &\quad \cdot \\ &\leq u_i^\alpha(f_{-i}, g_{i,\lambda(h)} | h) + \left(\sum_{r=1}^{k-\lambda(h)} \alpha^{r-1} \right) (1 - \alpha)\epsilon \\ &= u_i^\alpha(f | h) + \left(\sum_{r=1}^{k-\lambda(h)} \alpha^{r-1} \right) (1 - \alpha)\epsilon. \end{aligned}$$

Taking limits as $k \rightarrow \infty$ on both sides of the above inequality yields the result. \square

We can now state the main result of this section.

Theorem 4.1: Consider the game $G^\alpha(A, u, \alpha)$ and let $\varepsilon > 0$. There exists a positive integer W such that for every subgame perfect equilibrium f there is a strategy combination $g \in F$ satisfying:

1. $\text{comp}(g_i) \leq W$ for $i = 1, 2, \dots, N$,
2. $|u_i^\alpha(f) - u_i^\alpha(g)| < \varepsilon$ for $i = 1, 2, \dots, N$, and
3. g is a subgame perfect ε -equilibrium of $G^\alpha(A, u, \alpha)$.

To prove Theorem 4.1 we first define W and then g and show, in a series of lemmas, that it satisfies the properties of the theorem.

To begin, we partition the set of subgame perfect equilibria, SPE, into a set of equivalence classes as follows: first enclose $u^\alpha(\text{SPE})$, the set of discounted payoff vectors to subgame perfect equilibria, by an N -cube K with sides of length $\max_i \{2M_i / (1 - \alpha)\}$. Next, partition \mathbb{R}^N into disjoint half-open cubes with sides of length $(1 - \alpha)^2 \varepsilon / 2$. This partition of \mathbb{R}^N clearly induces a finite disjoint partition of K , $P(K)$, by way of intersection. We define W to be the cardinality of this partition. Two members of SPE are now declared to be equivalent if and only if they yield vectors of discounted payoffs which occupy the same member of $P(K)$. Now choose one representative member of every equivalence class in the partition of SPE. Define $C: \text{SPE} \rightarrow \text{SPE}$ to be this selection rule, and let $C_i: \text{SPE} \rightarrow F_i$ be the projection operator for C , $i = 1, 2, \dots, N$. Then $C(\text{SPE}) = \{C(\bar{f}): \bar{f} \in \text{SPE}\}$ is a finite set with at most W elements.

Given a subgame perfect equilibrium $f \in \text{SPE}$, we are now in a position to define g , the approximating strategy vector for f : let g_i be the strategy

implemented by the automaton

$$I_i = ((C(F(f)), C(f)), T_i, B_i)$$

defined by

$$T_i(\bar{f}, a) = C(\bar{f}|a) \text{ and } B_i(\bar{f}) = \bar{f}_i(e).$$

First, it is clear that $\text{comp}(g_i) \leq W$ for $i = 1, 2, \dots, N$ because each implementing automaton uses the same finite state space (and, we note, the same transition function) and as we have seen in the proof of Theorem 3.1, the complexity of a strategy is never larger than the cardinality of any implementing automaton. Thus part one of Theorem 4.1 is proved. For the second part, we have:

Lemma 4.2: $|u_i^\alpha(f) - u_i^\alpha(g)| < (1 - \alpha)\varepsilon$ for $i = 1, 2, \dots, N$.

Proof: First, define inductively a sequence of strategy vectors g_k by:

$$\begin{aligned} g_1 &= C(f) \\ g_2 &= C(g_1|g_1(e)) \\ &\vdots \\ g_k &= C(g_{k-1}|g_{k-1}(e)), \end{aligned}$$

and recall the convention that

$$\begin{aligned} g^0(h) &= h \\ g^1(h) &= h \cdot g(h) \\ &\vdots \\ g^n(h) &= g^{n-1}(h) \cdot g(g^{n-1}(h)). \end{aligned}$$

It is straightforward to verify that $g_k(e) = g(g^{k-1}(e))$. Thus we can find n so large that

$$\left| \sum_{r=1}^{n-1} \alpha^{r-1} u_i(g_r(e)) + \alpha^{n-1} u_i^\alpha(g_n) - u_i^\alpha(g) \right| < (1 - \alpha)\epsilon/2.$$

For such n , we have:

$$\begin{aligned} & \left| u_i^\alpha(f) - u_i^\alpha(g) \right| < \\ & \left| u_i^\alpha(g) - u_i^\alpha(g_1) \right| + \left| u_i^\alpha(g_1) - u_i(g_1(e)) - \alpha u_i^\alpha(g_2) \right| + \\ & \sum_{k=3}^n \left| \sum_{r=1}^{k-2} \alpha^{r-1} u_i(g_r(e)) + \alpha^{k-2} u_i^\alpha(g_{k-1}) - \sum_{r=1}^{k-1} \alpha^{r-1} u_i(g_r(e)) - \alpha^{k-1} u_i^\alpha(g_k) \right| + \\ & \left| \sum_{r=1}^{n-1} \alpha^{r-1} u_i(g_r(e)) + \alpha^{n-1} u_i^\alpha(g_n) - u_i^\alpha(g) \right| < \\ & \left(\sum_{r=1}^n \alpha^{r-1} \right) (1 - \alpha)^2 \epsilon/2 + (1 - \alpha)\epsilon/2 < (1 - \alpha)\epsilon. \quad \square \end{aligned}$$

The next lemma finally establishes Theorem 4.1.

Lemma 4.3: g is a subgame perfect ϵ -equilibrium of $G^\infty(A, u, \alpha)$.

Proof: The idea is to show that no single period deviation against g is more than $(1 - \alpha)\epsilon$ profitable for any player in any subgame. The result then follows by Lemma 4.1.

Fix $h \in H$. Without loss of generality, we assume that player i deviates once from g_i immediately following h , and then all players conform with g in all subsequent periods.

Now, there is a strategy vector $f_{\lambda(h)} \in F(f)$ with which $g|h$ has a natural association. $f_{\lambda(h)}$ is defined inductively as follows:

$$\begin{aligned} f_0 &= C(f) \\ f_1 &= C(f_0|h^1) \\ &\vdots \\ &\vdots \\ &\vdots \\ f_{\lambda(h)} &= C(f_{\lambda(h)-1}|h^{\lambda(h)}), \end{aligned}$$

and we have $f_{\lambda(h)}(e) = (g|h)(e)$. Thus any deviation by player i from $(g_i|h)(e)$ yields the same immediate payoff to player i as it would against $f_{\lambda(h)}(e)$. Let $d = ((g_{-i}|h)(e), d_i) \in A$ represent the vector of actions indicating i 's deviation. Then given that all players conform, post-deviation (discounted) payoffs to player i from $g|h \cdot d$ are within $\alpha(1 - \alpha)\epsilon$ of post-deviation payoffs to i from $f_{\lambda(h)}|d$. (See the proof of Lemma 4.2.) Now observe that $f_{\lambda(h)} \in F(f)$ means that d followed by conformity is not profitable for player i against $f_{\lambda(h)}$, since $f_{\lambda(h)}$ is perfect. \square

At this point, let us briefly consider some observations concerning the structure of Theorem 4.1 and its proof.

First, note that no particular relationship is implied concerning the complexities of the approximating strategies g_i . Of course, each is finite and bounded above by the cardinality of $C(F(f))$; but the structure of g could be quite complicated, with complexities of all feasible orders represented. We will return to this issue in Section 5, where sufficient conditions for complexity relationships are considered. These conditions are stated primarily in terms of the stage game payoff structure which, apart from the boundedness condition, has been left arbitrary in this section.

Note also that given α , ε , and $u^\alpha(F(f))$, an upper bound on $\text{Card}(C(F(f)))$ is easily calculated. In fact, the way g is constructed gives rise to the reasonable presumption of increasing complexity for g with decreasing ε and increasing α . This simply says that the more closely we wish to credibly approximate a complex equilibrium, the more computing power we need.

Finally, we will informally outline a way in which the results of Theorem 4.1 might be sharpened. We have in mind a sufficient condition on the perfect equilibrium f which guarantees the approximating vector g can be taken as a full perfect equilibrium. Thus we can dispense with one of the two approximating aspects of g under this condition, with the rest of the theorem remaining as is. The condition on f might be termed uniform strict perfection. By this, we mean that one-shot deviations by any player are always strictly unprofitable, and in fact uniformly so across players, contingencies, and deviations. Thus there will exist a $\delta > 0$ such that under all conditions, one-shot deviations by any player followed by conformity yield the deviating player a discounted payoff bounded above by the induced equilibrium payoff decremented by δ . By the result of Harris [1984], if $f \in F$ is uniformly strictly perfect, then f is a perfect equilibrium. The equilibria of Example 4.1 will satisfy this condition if α is large enough. Referring the reader to the proof of Lemma 4.3 for justification, we will simply assert that if f and g are as in Theorem 4.1., and f is uniformly strictly perfect, then g (as constructed) will be a perfect equilibrium if ε is small enough. Under this condition on f and in the case of finite-action games ($\text{Card}(A)$ is finite), we also assert that g (as constructed) will be a DRSP equilibrium if ε is small enough.

5. Complexity and the Structure of Perfect Equilibria

As we remarked in the introduction, perfection would seem to carry with

it at least the presumption of complexity relationships between players' strategies. As some examples will make clear, this is not universally so, at least with respect to our current definition of complexity. In particular, when we move away from discounting to zero interest rate evaluation relations, dramatically large differences between players' strategy complexities can be found in perfect equilibria. We will discuss this and some other examples later in this section.

Our purpose here is to give some simple sufficient conditions on discounted games for patterns to begin to emerge. We begin by singling out a particular class of stage games and note that finite matrix games fall generically into this class. We say that a stage game (A,u) has individually responsive payoffs if for every $i = 1,2,\dots,N$ and every $a_{-i} \in A_{-i}$, $u_i(a_{-i},a_i) \neq u_i(a_{-i},\bar{a}_i)$ whenever $a_i \neq \bar{a}_i$. Thus individually responsive payoffs have the property that player i 's payoffs are one-to-one on his action set given an action combination of the other $N - 1$ players. Recall that $f \in F$ is a discount robust subgame perfect (DRSP) equilibrium of $G^\omega(A,u,\alpha)$ if f is a subgame perfect equilibrium of $G^\omega(A,u,\beta)$ for every β in some neighborhood of α . We consider only DRSP equilibria in this section. Given a strategy of player i , $f_i \in F_i$, we say the memory of f_i is m , $\text{Mem}(f_i) = m$ if m is the smallest integer satisfying the property: for every $h \in H$ with $\ell(h) \geq m$, $f_i|_h = f_i|_{E^m(h)}$. If such an m does not exist, we say that $\text{Mem}(f_i) = \infty$.

Our first theorem shows that DRSP equilibria of $G^\omega(A,u,\alpha)$ are completely determined by the strategy choice of any group of $N - 1$ players if the stage game has individually responsive payoffs. This gives rise to some interesting corollaries on complexity relationships and the memory of strategies.

Theorem 5.1: Suppose (A,u) has individually responsive payoffs and f is a DRSP equilibrium of $G^\omega(A,u,\alpha)$. Then for all i ,

$$\{\hat{f}_i: (f_{-i}, \hat{f}_i) \text{ is a DRSP equilibrium of } G^\infty(A, u, \alpha)\}$$

is a singleton set.

Proof: Suppose $\hat{f} = (f_{-i}, \hat{f}_i)$ is a DRSP equilibrium of $G^\infty(A, u, \alpha)$. We will show that $\hat{f}_i \equiv f_i$ on H . Fix $h \in H$. By the best reply property of perfection applied to player i , we must have:

$$u_i^\beta(f|h) = u_i^\beta(\hat{f}|h) \text{ for all } \beta \text{ in an open neighborhood of } \alpha.$$

Thus we have two power series which converge to the same function on some open neighborhood of α . A well-known result from real analysis then yields term-by-term equality of stage game payoffs to player i along the two induced equilibrium paths. In particular, we have

$$u_i((f|h)(e)) = u_i((\hat{f}|h)(e)).$$

However, the action combination of players other than i is the same in these two cases. The fact that payoffs are individually response gives

$$f_i(h) = \hat{f}_i(h). \quad \square$$

Our first corollary applies to the general case of $N > 2$ players:

Corollary 5.1: Under the hypothesis of Theorem 5.1:

1. $\text{comp}(f_i) \leq \text{Card}(\prod_{j \neq i} F_j(f_j))$ for $i = 1, 2, \dots, N$.

In particular, if all complexities are finite:

$$\text{comp}(f_i) \leq \prod_{j \neq i} \text{comp}(f_j), \text{ for } i = 1, 2, \dots, N.$$

$$2. \text{ Mem}(f_i) \leq \max_{j \neq i} \{\text{Mem}(f_j)\}, \text{ for } i = 1, 2, \dots, N.$$

In particular, if we assume strategies are ordered by memory size:

$$\text{Mem}(f_1) \leq \text{Mem}(f_2) \leq \dots \leq \text{Mem}(f_N), \text{ then}$$

$$\text{Mem}(f_{N-1}) = \text{Mem}(f_N).$$

Proof: For the first part, since DRSP equilibria of $G^\infty(A, u, \alpha)$ induce DRSP equilibria of $G^\infty(A, u, \alpha)$ after all histories, Theorem 5.1 provides an onto map $M: F_{-i}(f_{-i}) \rightarrow F_i(f_i)$. But clearly,

$$\text{Card}(F_{-i}(f_{-i})) \leq \text{Card}\left(\prod_{j \neq i} F_j(f_j)\right).$$

For the second part, if for some i , $\text{Mem}(f_i) > \max_{j \neq i} \{\text{Mem}(f_j)\}$, we could find a history h with $\ell(h) = \max_{j \neq i} \{\text{Mem}(f_j)\}$ and two histories \hat{h} and \bar{h} such that $f_i | \hat{h} \cdot h \neq f_i | \bar{h} \cdot h$. But by Theorem 5.1, and the fact that DRSP equilibria induce DRSP equilibria, we get a contradiction when we note that $f_{-i} | \hat{h} \cdot h = f_{-i} | \bar{h} \cdot h$. \square

Especially in the finite case, we would expect the complexity bounds given in the first part of the corollary to be conservative, since typically we should have:

$$\text{comp}(f_i) \leq \text{Card}(F_{-i}(f_{-i})) < \text{Card}\left(\prod_{j \neq i} F_j(f_j)\right).$$

In fact, if all complexities are finite and $\text{comp}(f_i) > 1$ for at least three players, then the structure of the inequality system itself guarantees that

$$\text{comp}(f_i) = \prod_{j \neq i} \text{comp}(f_j)$$

for at most one value of i . Otherwise we would have

$$(\text{comp}(f_i))(\text{comp}(f_k)) = \left(\prod_{j \neq i} \text{comp}(f_j) \right) \left(\prod_{j \neq k} \text{comp}(f_j) \right),$$

which yields

$$1 = \left(\prod_{j \neq i, k} \text{comp}(f_j) \right)^2.$$

We might paraphrase the result of Theorem 5.1 and its first corollary in terms of computing machines. They say the maximum amount of computing power (size of automaton or equivalently, complexity of strategy) which a player can profitably bring to bear in equilibria of repeated conflict and cooperation situations is strictly limited by the capabilities of the other players. We think this makes precise a reasonable and intuitive truth: in equilibrium, such interactions must be governed by the capacities of the least capable participants.

The case of two players deserves to be singled out for special attention. Theorem 5.1 says directly that if $N = 2$, equilibria are completely determined by the strategy choice of either player. Thus all information carried by the equilibrium is encoded in each strategy. A weaker condition is expressed in the following definition. If $N = 2$, we say that a pair of strategies $f = (f_1, f_2) \in F$ is conjugate if there is a one-to-one and onto correspondence $M: F_1(f_1) \rightarrow F_2(f_2)$ satisfying

$$f_2|_h = M(f_1|h) \text{ for every } h \in H.$$

It is obvious that a conjugate pair of strategies must be of equal complexity and of equal memory.

Corollary 5.2: With $N = 2$ and under the hypotheses of Theorem 5.1, $f = (f_1, f_2) \in F$ is a conjugate pair. Hence, $\text{comp}(f_1) = \text{comp}(f_2)$ and $\text{Mem}(f_1) = \text{Mem}(f_2)$.

Proof: Again by Theorem 5.1 and the fact that DRSP equilibria induce DRSP equilibria, we have $f_1|h = f_1|\bar{h}$ if and only if $f_2|h = f_2|\bar{h}$ for all $h, \bar{h} \in H$. \square

Again, in the language of machines, the minimal automata implementing the two strategies may be taken to use the same state space, the same initial state, and the same transition rule, with differences in the strategies showing up only in the behavior functions. Thus, the dynamics of the two strategies are identical. This contrasts with the case of N players because, at least with respect to our present knowledge, the minimal implementing automata will have cardinalities subject only to the bounds given in Corollary 5.1.

As an extreme example of how these results can fail, consider a duopoly game with the limit of the means evaluation relation. In this case, we give up any form of discounting, and payoffs do not have the individual responsiveness property. Stanford [1984] considers linear reaction function equilibria in this context. A linear reaction function strategy specifies a production quantity for the player in the first period, and a decision rule which selects a production level for the player in period $t + 1$ as a linear function of the other player's production level in period t . From our present

viewpoint, an interesting result is that perfect equilibria of this form exist with the property that exactly one of the players uses a constant reaction function, and thus a strategy of complexity one. The other player uses a nontrivial reaction function which, because the stage game strategy sets are continua, means a strategy of uncountable complexity. This results in collusive outcomes along the equilibrium path. We would guess that this possibility is driven more by the zero interest rate evaluation relation than by the lack of individual responsiveness in the payoffs. In any case, perfection itself is insufficient to guarantee relationships like the ones we find here.

Another (less extreme) example of failure comes from the repeated Prisoners' Dilemma game. In this case, we have individually responsive payoffs and return to discounting as the evaluation relation. Kalai, Samet, and Stanford [1985] have shown that tit-for-tat versus tit-for-tat can be a perfect equilibrium in this framework for only one value of the discount parameter α . To be precise, recall the notation of Example 4.1. The result is that given α , tit-for-tat versus tit-for-tat will be a perfect equilibrium of $G^\infty(A, u, \alpha)$ if and only if $a_1 = b_1 = 1/(1 - \alpha)$ and $a_2 = b_2 = -\alpha/(1 - \alpha)$. Now, regarding a_i and b_i as fixed at these levels, we find a perfect equilibrium of a very knife-edge character with respect to discounting. In other words, we have a perfect equilibrium of $G^\infty(A, u, \alpha)$ which is not a DRSP equilibrium of $G^\infty(A, u, \alpha)$. It is also easy to see that tit-for-tat versus tit-for-tat is not a conjugate pair. Thus Corollary 5.2 fails even though we have a perfect equilibrium in a discounted game based on a stage game with individually responsive payoffs. This means some form of discount robustness of perfection is necessary for the results. As a final remark we observe though, that the strategies involved have the same complexities.

Finally, if we drop the requirement of individually responsive payoffs in the stage game, it is easy to find examples of DRSP equilibria in which the players employ strategies with arbitrarily large differences in complexity. We have in mind replicating the first row in the payoff matrix for the Prisoners' Dilemma, giving player one an extra action in the stage game:

		Player 2	
		N_2	C_2
	N_1	$0,0$	a_1, b_2
	C_1	a_2, b_1	$1,1$
	\bar{N}_1	$0,0$	a_1, b_2

Thus, a DRSP equilibrium of $G^\infty(A,u,\alpha)$ could be constructed where player 1 alternates between N_1 and \bar{N}_1 in an aperiodic manner, while player 2 uses the constant N_2 strategy. This equilibrium has player 1 using an infinitely complex repeated game strategy and player 2 with a strategy of complexity one. This last is an example of a two-player stage game drawn from a general class with which we will be concerned in the remainder of this section.

We say the two-player stage game (A,u) has jointly varying payoffs if $u_1(a) = u_1(\bar{a})$ if and only if $u_2(a) = u_2(\bar{a})$ for every $a, \bar{a} \in A = A_1 \times A_2$. Note that the Prisoners' Dilemma as well as its row-augmented version are both examples of stage games with jointly varying payoffs. We might suppose that giving player 1 an additional inconsequential action in the stage game as above gives rise to potential complexity differences only as an artifact.

That this is so follows from our next result. Before considering this result, recall that two strategy pairs $f, g \in F$ are perfectly payoff equivalent if $u^\alpha(f|h) = u^\alpha(g|h)$ for all $h \in H$.


Theorem 5.2: For $N = 2$, let (A, u) be a stage game with jointly varying payoffs. Suppose $f = (f_1, f_2) \in F$ is a DRSP equilibrium of $G^\alpha(A, u, \alpha)$. Then there exists $g = (g_1, g_2) \in F$ satisfying:

1. g is a DRSP equilibrium of $G^\alpha(A, u, \alpha)$;
2. f and g are perfectly payoff equivalent;
3. g is a conjugate strategy pair;
4. $\text{comp}(g_1) = \text{comp}(g_2) \leq \min\{\text{comp}(f_1), \text{comp}(f_2)\}$; and
5. $\text{Mem}(g_1) = \text{Mem}(g_2) \leq \min\{\text{Mem}(f_1), \text{Mem}(f_2)\}$.

Moreover, perfect payoff equivalence is uniform in the sense of holding for all β in some open neighborhood of α .

Theorem 5.2 says that any search for DRSP equilibria in games of this kind may as well be restricted to conjugate strategy pairs satisfying the conclusions of the theorem.

We should remark that stage games with jointly varying payoffs can arise naturally when we reduce certain extensive form games to their normal forms. For example:



Thus, even though finite matrix games fall generically into the class of games with individually responsive payoffs, other natural structures such as jointly varying payoff games deserve investigation as well.²

The proof of Theorem 5.2 is deferred to later in this section. It uses two propositions which may be of independent interest, since they do not rely on the assumption of jointly varying payoffs.

Proposition 5.1: Compositions of perfectly payoff equivalent strategy pairs are perfectly payoff equivalent. Let f and g be strategy pairs of $G^\alpha(A, u, \alpha)$. Suppose g satisfies the following condition: for every $h \in H$, $g(h) = \bar{f}^h(e)$ for some strategy pair \bar{f}^h which is perfectly payoff equivalent to $f|h$. Then g is perfectly payoff equivalent to f .

Proof: Define a sequence of strategy pairs g_k for $k = 0, 1, 2, \dots$, by

$$g_k(h) = \begin{cases} g(h) & \text{if } \lambda(h) < k \\ f(h) & \text{if } \lambda(h) \geq k. \end{cases}$$

It is clear (by the discounting) that for all $h \in H$, $\lim_{k \rightarrow \infty} u^\alpha(g_k|h) = u^\alpha(g|h)$. Thus, it suffices to show that for $k = 0, 1, 2, \dots$, g_k is perfectly payoff equivalent to f . This is clearly true for $k = 0$. Now consider $k > 0$ and $h \in H$. We want to show that $u^\alpha(g_k|h) = u^\alpha(f|h)$. This is proved by backwards induction on $\lambda(h)$. Clearly, if $\lambda(h) \geq k$, equality holds by the definition of g_k . So assume equality holds for all histories with $k \geq \lambda(h) \geq K$ and suppose b is a history with $\lambda(b) = K - 1$. Then

$$u^\alpha(g_k|b) = u(g_k(b)) + \alpha u^\alpha(g_k|b \cdot g_k(b))$$

²We are grateful to Robert Aumann for pointing this out.

$$\begin{aligned}
 &= u(\bar{f}^b(e)) + \alpha u^\alpha(f|b \cdot g_k(b)) \text{ (by induction hypothesis)} \\
 &= u(\bar{f}^b(e)) + \alpha u^\alpha(\bar{f}^b|\bar{f}^b(e)) \text{ (by the definition of } \bar{f}^b \text{ and its equivalence to } f|b) \\
 &= u^\alpha(\bar{f}^b) \\
 &= u^\alpha(f|b) \text{ (by the equivalence of } \bar{f}^b \text{ and } f|b). \quad \square
 \end{aligned}$$

Proposition 5.2: Compositions of perfectly payoff equivalent subgame perfect equilibria are perfectly payoff equivalent subgame perfect equilibria. Let f and g be strategy pairs of $G^\infty(A, u, \alpha)$. Suppose g satisfies the following condition: for every $h \in H$, $g(h) = \bar{f}^h(e)$ for some subgame perfect equilibrium \bar{f}^h which is perfectly payoff equivalent to $f|h$. Then g is a subgame perfect equilibrium which is perfectly payoff equivalent to f .

Proof: By Proposition 5.1, g is perfectly payoff equivalent to f . If g is not subgame perfect, by Lemma 4.1 with $\varepsilon = 0$, we may assume there is a history h and an action $a_1 \in A_1$ such that

$$u_1(a_1, g_2(h)) + \alpha u_1^\alpha(g|h \cdot (a_1, g_2(h))) > u_1^\alpha(g|h).$$

Now,

$$\begin{aligned}
 &u_1(a_1, g_2(h)) + \alpha u_1^\alpha(g|h \cdot (a_1, g_2(h))) \\
 &= u_1(a_1, g_2(h)) + \alpha u_1^\alpha(f|h \cdot (a_1, g_2(h))) \text{ (by Proposition 5.1)} \\
 &= u_1(a_1, \bar{f}_2^h(e)) + \alpha u_1^\alpha(\bar{f}^h|(a_1, \bar{f}_2^h(e))) \text{ (by perfect payoff equivalence of } \bar{f}^h \text{ and } f|h)
 \end{aligned}$$

$$\begin{aligned}
 &< u_1^\alpha(\bar{f}^h) \text{ (by perfection of } \bar{f}^h) \\
 &= u_1(\bar{f}^h(e)) + \omega_1^\alpha(\bar{f}^h | \bar{f}^h(e)) \\
 &= u_1(g(h)) + \omega_1^\alpha(f|h \cdot g(h)) \text{ (by perfect payoff equivalence of } \bar{f}^h \text{ and } f|h) \\
 &= u_1(g(h)) + \omega_1^\alpha(g|h \cdot g(h)) \text{ (by perfect payoff equivalence of } f \text{ and } g) \\
 &= u_1^\alpha(g|h), \text{ which is a contradiction. } \quad \square
 \end{aligned}$$

Proof of Theorem 5.2: For $\bar{g}, g \in F$, we write $\bar{g} \sim g$ if and only if $u^\beta(\bar{g}|\hat{h}) = u^\beta(g|\hat{h})$ for every $\hat{h} \in H$ and every β in some open neighborhood of α . Clearly, \sim is an equivalence relation. Thus $F(f)$ can be partitioned into \sim equivalence classes. From every equivalence class, we choose one representative strategy pair (using the axiom of choice if $F(f)$ is large). Define $C: F(f) \rightarrow F(f)$ to be this selection rule, and let $C_i: F(f) \rightarrow F_i(f_i)$ be the projection operator for C . Thus, if $g, \bar{g} \in F(f)$ then $g \sim \bar{g}$ if and only if $C(g) = C(\bar{g})$. Also $g \sim C(g)$.

Define the strategy pair g by

$$g(h) = (C(f|h))(e), \text{ for all } h \in H.$$

Then by Proposition 5.2., g is a DRSP equilibrium of $G^\alpha(A, \alpha)$ which is perfectly payoff equivalent to f . Moreover, perfect payoff equivalence is uniform in the sense of holding for all β in some open neighborhood of α .

We next claim that $g = (g_1, g_2)$ is a conjugate pair of strategies. If we

can show $g_1|h = g_1|\bar{h}$ implies $g_2|h = g_2|\bar{h}$ for all $h, \bar{h} \in H$, then by symmetry, the claim will be proved. So suppose $g_1|h = g_1|\bar{h}$. Since g is a DRSP equilibrium of $G^\infty(A, u, \alpha)$,

$$u_2^\beta(g|h \cdot \hat{h}) = u_2^\beta(g|\bar{h} \cdot \hat{h}) \text{ for all } \hat{h} \in H \text{ and all } \beta.$$

Again, we have two power series converging to the same function on an open neighborhood of α , and so there is term-by-term equality of stage game payoffs to player 2. Since stage game payoffs are jointly varying, this means

$$u_1^\beta(g|h \cdot \hat{h}) = u_1^\beta(g|\bar{h} \cdot \hat{h}) \text{ for all } \hat{h} \in H \text{ and all } \beta.$$

Together, these facts imply $g|h \sim g|\bar{h}$. Since g is uniformly perfectly payoff equivalent to f , we know then that $f|h \sim f|\bar{h}$. Thus

$$C(f|h \cdot \hat{h}) = C(f|\bar{h} \cdot \hat{h}) \text{ for all } \hat{h} \in H,$$

and so

$$g(h \cdot \hat{h}) = g(\bar{h} \cdot \hat{h}) \text{ for all } \hat{h} \in H.$$

In particular, this means

$$(g_2|h)(\hat{h}) = (g_2|\bar{h})(\hat{h}) \text{ for all } \hat{h} \in H,$$

or

$$g_2|h = g_2|\bar{h}.$$

Next we note that $\text{Card}(C(F(f))) < \text{Card}(F_i(f_i))$. This follows from an argument similar to the one above. In particular, we can show that $f|h \not\sim f|\bar{h}$

implies $f_i|_h \neq f_i|\bar{h}$ for $i = 1, 2$. Thus the projection from $C(F(f))$ into $F_i(f_i)$ is one-to-one. Now observe that g_i can be implemented by the automaton $I_i = ((C(F(f)), C(f)), T_i, B_i)$ defined by $T_i(\bar{f}, a) = C(\bar{f}|a)$ and $B_i(\bar{f}) = \bar{f}_i(e)$. Again, by the proof of Theorem 3.1, we have

$$\text{comp}(g_i) \leq \text{Card}(I_i) = \text{Card}(C(F(f))) \leq \text{comp}(f_i).$$

Equality of complexity follows from conjugacy, and the fourth result of the theorem follows from $\text{comp}(g_1) = \text{comp}(g_2)$.

Finally, we want to show that

$$\text{Mem}(g_i) \leq \min\{\text{Mem}(f_1), \text{Mem}(f_2)\}, \text{ for } i = 1, 2.$$

Without loss of generality, assume

$$m = \text{Mem}(f_1) \leq \text{Mem}(f_2), \text{ with } m < \infty.$$

Let h be a history with $\ell(h) > m$. Since $\text{Mem}(f_1) = m$, we have

$f_1|_h = f_1|_{E^m(h)}$. This implies (as before) that $f|_h \sim f|_{E^m(h)}$. Thus

$C(f|_h) = C(f|_{E^m(h)})$, which gives $g(h) = g(E^m(h))$, and we see that g_1 and g_2

are both of memory at most m . $\text{Mem}(g_1) = \text{Mem}(g_2)$ follows from conjugacy. \square

6. Finite Complexity and the Computation of Perfect Equilibria

This section contains a brief discussion of some computational aspects of perfect equilibria. We will describe a strategy tuple (or equilibrium) as being of finite complexity if the strategies of all players are of finite complexity. With modifications, the following remarks will apply to all finite complexity equilibria in repeated games based on finite action stage

games.

As a concrete example, consider the Prisoners' Dilemma as a stage game and form $G^\infty(A,u,\alpha)$ for given α . If we specify complexity levels for the strategies of both players and are interested in the corresponding set of finite complexity perfect equilibria, it should be an easy task to write a computer program listing all of these as output. Our aim here is to give an outline for such a program.

Restricting attention to DRSP equilibrium of $G^\infty(A,u,\alpha)$ simplifies things since the results of section 5 show that we need consider only conjugate strategy pairs. As we have seen, this means the minimal implementing automata may be taken to use the same state space, the same initial state, and the same transition rule. Thus if the fixed (common) complexity level is denoted by c , we start with a finite set S such that $\text{Card}(S) = c$, and distinguish one element of S as the initial state. Denote this element by s^0 . Recalling the notation of Example 4.1, we have $A_i = \{N_i, C_i\}$ for $i = 1, 2$, and $A = A_1 \times A_2$. We now program our computer to generate all possible transition rules $T: S \times A \rightarrow S$. (Actually, we need not consider all such transition rules. Since we are dealing with minimal implementing automata we can, for example, restrict attention to those T which are onto the set $S - \{s^0\}$. This is because states which cannot be reached from s^0 via T and suitable histories are superfluous.) This set of transition rules is finite. The next step is to generate the sets of all possible behavior functions $B_i: S \rightarrow A_i$ for $i = 1, 2$. These two sets of behavior functions are both finite. We proceed by forming all possible triples (T, B_1, B_2) . Since, under a given transition rule and behavior functions, states correspond precisely to induced strategies, in this way we can generate a finite set of payoff pairs which can be indexed over states. Thus $u^\alpha(s)$ represents discounted payoffs to the induced strategy

pair corresponding to the state s . These payoff pairs are easily computed given (T, B_1, B_2) because the finiteness of S guarantees that starting from any state, the sequence of states induced by the transition rule and behavior functions must eventually cycle repeatedly through some finite set of states. Our computer program can be written to detect this condition. Hence induced per-period payoffs must also cycle repeatedly through a finite set of values.

Now observe that the one-shot deviation condition of Lemma 4.1 (with $\epsilon = 0$) is both necessary and sufficient for perfection. Since the Prisoners' Dilemma is a finite action game, in each state there are finitely many deviations to consider. Thus for every state s , we need check only finitely many inequalities of the form

$$u_i(d) + \alpha u_i^\alpha(T(s, d)) \leq u_i^\alpha(s),$$

where d represents a pair of actions with $d_i \neq B_i(s)$ for precisely one value of i . In summary, there will be finitely many conditions to check in verifying that (T, B_1, B_2) does or does not correspond to a perfect equilibrium $G^\omega(A, \alpha)$. This means we have a finite algorithm for generating all perfect equilibria of $G^\omega(A, \alpha)$ which consist of conjugate strategy pairs of given complexity. Since S is finite and (A, u) is a finite action game, a sufficient condition for one of these to be a DRSP equilibrium is that all of the inequalities shown above be satisfied as strict inequalities.

7. Concluding Remarks

Broadly, the goal of this paper was to study the structure of infinitely repeated games with discounting. The development of our understanding of such games has been slow, largely because their equilibria can be extremely

complex. Our approach was to attack this feature of the problem directly by proposing and examining a precise complexity concept for strategies in repeated games. The relevance of the concept was tested in studying some of its relationships to subgame perfect equilibria.

With the foregoing results in hand, we are willing to conclude that measuring the complexity (or simplicity) of a strategy by the number of strategies it induces gives rise to a useful classification scheme particularly in the case of finite complexity. With respect to perfection, the observations of section 6 show that finite complexity equilibria are indeed simple in the relevant sense: in a finite number of steps, we can decide the perfection question regarding proposed strategy tuples, at least for finite action stage games. Also the simplest strategies, the ones with complexity one, capture the idea of no collusion. Because one can easily see that every subgame perfect equilibrium in which the players use strategies of complexity one must consist of repeated use of the equilibria of the one shot game.

Further, we have seen that finite complexity strategies correspond precisely to the simplest computing machines: the finite automata. If we adopt the views that "rationality" is a finitely describable phenomenon, then the automaton interpretation lends real substance to the approximation results of section 4. Adopting this viewpoint, particularly in infinite action games leads naturally to a focus on approximate equilibria implemented by finite automata. We would paraphrase the approximation result by saying that under the restriction of "finite rationality," all equilibria remain relevant. None can be entirely ruled out on the basis of complexity considerations. This is particularly true in light of the constructive nature of Theorem 4.1's proof. The structure of the equilibrium is partially echoed in the structure

of the approximation. Of course, assigning costs to increasing complexity, an idea which can be found in Rubinstein [1985], might yield very different results and deserves investigation.

We are less hopeful regarding the appropriateness of the complexity definition when strategies are judged to be of infinite complexity. It strains credulity to claim that a linear reaction function in a duopoly game gives rise to an infinitely (in fact, uncountably) complex strategy. Such a possibility comes, of course, from considering stage games with continua as strategy sets. To summarize the foregoing, we think the complexity concept is arguably the "right" one for finite complexity strategies. It is impossible to implement such a strategy with a strictly smaller automaton and adding superfluous states accomplishes nothing. On the other hand, in the infinite complexity case, we doubt that a strategy can be found which will be assigned a complexity which is too low by our definition relative to any other reasonable measure. Thus, if "errors" occur in our classification scheme, they should uniformly consist of assigning complexities which are too high.

The results on structure in section 5 capture at least some of the essence of subgame perfection as a condition on equilibria. The fact that equilibrium is completely determined by the strategy choice of any group of $N - 1$ players sharply emphasizes the strategic interdependence among players. The corollary complexity relationships have the natural interpretation that players must face worthy opponents, at least in the aggregate, in order to realize their full strategic potential. This is true both with respect to memory and with respect to the complexity of their strategies.

With regard to extensions to mixed or behavior strategies of the infinitely repeated game we see no difficulty in extending our definitions and

the approximation result. The complexity of a behavior strategy will still be defined as the number of distinct (behavior) strategies it induces in the various subgames. The automaton notion would then have to be modified to become an automaton with a randomizing device. This means that the behavior function would be allowed to randomize over the set of actions for every given state. We see no difficulty in repeating the proof of Theorem 4.1 for this case. We doubt, however, that the interpersonal complexity bounds will carry over to the mixed strategy case.

References

1. D. Abreu, Extremal equilibria of oligopolistic supergames, manuscript, 1985.
2. D. Abreu and A. Rubinstein, The structure of Nash equilibrium in repeated games with finite automata, manuscript, 1986.
3. R. Aumann, Survey of repeated games, in "Essays in Game Theory and Mathematical Economics in Honor of Oskar Morgenstern," Bibliographisches Institut Mannheim/Wien/Zurich, 1981.
4. R. Aumann and S. Sorin, Bounded rationality forces cooperation, manuscript, 1985.
5. R. Axelrod, Effective choice in the prisoner's dilemma, Journal of Conflict Resolution, 24 (1980), 3-25.
6. E. Ben Porath, Repeated games with bounded complexity, manuscript, 1986.
7. J. Friedman, A non-cooperative equilibrium for supergames, Rev. Econ. Stud. 38 (1971), 1-12.
8. D. Fudenberg and D. Levine, Subgame-perfect equilibria of finite and infinite horizon games, J. Econ. Theory 31 (1983), 251-268.
9. C. Futia, The complexity of economic decision rules, J. Math. Econ. 4 (1977), 289-299.
10. C. Harris, A note on the existence of optimal simple penal codes, manuscript, Nuffield College, Oxford, 1984.
11. J. Hopcroft and J. Ullman, "Introduction to Automata Theory, Languages, and Computation," Addison-Wesley, Reading, Massachusetts, 1979.
12. E. Kalai, D. Samet, and W. Stanford, A note on reactive equilibria in the discounted prisoners' dilemma and associated games, manuscript, 1985.
13. A. Lewis, On effectively computable realizations of choice functions, Math. Soc. Sci. 10 (1985), 43-80.

14. N. Megiddo and A. Wigderson, On play by means of computing machines, manuscript, 1985.
15. E. Moore, Gedanken experiments on sequential machines, in "Automata Studies," Princeton University Press, Princeton, New Jersey, 1956.
16. K. Mount and S. Reiter, Computational complexity of resource allocation, manuscript, 1983.
17. J. Myhill, Finite automata and the representation of events, WADD TR-57-624, Wright Patterson AFB, Ohio, 1957.
18. A. Nerode, Linear automaton transformations, Proc. AMS, 9 (1958), 541-544.
19. A. Neyman, Bounded complexity justifies cooperation in the finitely repeated prisoners' dilemma, manuscript, 1985.
20. G. Owen, "Game Theory," Academic Press, Orlando, Florida, 1982.
21. R. Radner, Collusive behavior in noncooperative epsilon-equilibria of oligopolies with long but finite lives, J. Econ. Theory 22 (1980), 136-154.
22. A. Rubinstein, Finite automata play--the repeated prisoners' dilemma, Research Report No. 149, Department of Economics, The Hebrew University, Jerusalem, Israel, 1985.
23. R. Selten, Reexamination of the perfectness concept for equilibrium points in extensive games, Int. J. Game Theory 4 (1975), 25-55.
24. H. Simon, Theories of bounded rationality, in "Decision and Organization," ed. by C. McGuire and R. Radner, North-Holland, Amsterdam, 1972.
25. S. Smale, The prisoner's dilemma and dynamical systems associated to non-cooperative games, Econometrica 48 (1980), 1617-1634.
26. W. Stanford, On continuous reaction function equilibria in duopoly supergames with mean payoffs, manuscript, University of Illinois at Chicago, 1984.
27. W. Stanford, Subgame perfect reaction function equilibria in discounted duopoly supergames are trivial, manuscript, 1985.
28. H. Varian, Complexity of social decisions, manuscript, 1975.

29. E. Zemel, On communication, bounded complexity and cooperation, manuscript, 1986.