

THE STRUCTURE OF NASH EQUILIBRIUM IN REPEATED GAMES WITH FINITE AUTOMATA

BY DILIP ABREU AND ARIEL RUBINSTEIN¹

We study two person infinitely repeated games in which players seek to minimize the complexity of their strategies. Players' preferences are assumed to depend both on repeated game payoffs and the complexity of the strategies they use. The model considered is that of Rubinstein (1986). Players simultaneously choose finite automata (Moore-machines) to implement their strategies. The complexity of a strategy is measured by the number of states in the automaton used to play the strategy. We analyze *Nash equilibrium* in the "machine game." Strong *necessary* conditions on the structure of equilibrium machine pairs are derived, under general assumptions about how players trade off repeated game payoffs against implementation costs. These structural results in turn place significant restrictions on equilibrium *payoffs*. We provide a complete characterization for symmetric 2×2 stage games, when repeated game payoffs are evaluated according to the limit of means, and complexity costs enter preferences lexicographically. We find that all Nash equilibrium payoffs must lie on one of the two "diagonals" of the payoff matrix, and show that "main" diagonal payoffs are always attained. Taken together our results suggest that the introduction of implementation costs results in a striking discontinuity in the Nash equilibrium set in terms of strategies, plays, and payoffs.

KEYWORDS: Strategic complexity, repeated games, finite automata, Nash equilibrium.

1. INTRODUCTION

IN THE STANDARD FORMULATION of a repeated game, players are assumed to be able to costlessly implement strategies of arbitrary complexity. We relax this assumption, pursuing a line of research initiated in Rubinstein (1986) (hereafter (Ru)). We assume instead that strategies vary in their implementation costs. As a consequence, players' strategic choices balance the twin objectives of maximizing repeated game payoffs and minimizing implementation costs.

Various features of the model presented below, such as the complexity measure we use, are rather special. Our results are therefore best regarded as suggestive, and we strongly emphasize the exploratory nature of the present paper. We provide a critical discussion of our approach, and comment on alternative formulations and possible extensions, after the detailed presentation of Section 2.

We adopt the model of (Ru): in a two player repeated game, the players are assumed to use machines (finite automata—see Hopcroft and Ullman (1979) for

¹ The authors would like to thank Thomas Marschak for helpful comments. The remarks of two anonymous referees are also gratefully acknowledged. Much of the work on this paper was done while the authors visited the Mathematical Sciences Research Institute, Berkeley, during Fall 1985. We are grateful to the Institute for its hospitality. Our stay at MSRI was supported by NSF Grants MCS-8120790 and SES-8420114. Abreu's research at Harvard was partially supported by NSF Grant SES-8509774.

an introduction) to implement their strategies. A (Moore) machine consists of a finite set of internal states (one of which is specified to be the initial state), an output function and a transition function. Given that the machine of a player is at a certain state at the t th round of the game, the output function determines the t th one-shot game action of the player as a function of the state. The transition function determines the next state as a function of the current state and of the other player's move at period t . A pair of machines induces a sequence of state pairs and a sequence of stage game action pairs, starting with the initial states of the two machines and their associated outputs.

In the "machine game" each player chooses a single machine, at the start of play. Hence a player's strategy space is the set of all finite machines. Player i 's preferences over machine pairs (M_1, M_2) depend on the repeated game payoff they yield and the number of states in M_i . The latter term incorporates the cost of implementing a repeated game strategy into a player's preferences. It embodies the assumption that this cost depends only on the number of states in the machine used to play the strategy. Note that in our usage the terms "complexity" and "implementation cost" are interchangeable: the complexity of a strategy is identified simply with the number of states in the machine used to implement the strategy.

We analyze Nash equilibrium (N.E.) in the machine game. Our main results concern the structure of equilibrium machine pairs. They provide *necessary* conditions on the form of equilibrium *strategies* and *plays*. This contrasts with previous analyses of repeated games which have typically focussed on equilibrium *payoffs*, or shown that a particular subset of strategies was *sufficient* (for example the "simple strategy profiles" of Abreu (1988)). Exceptions are (Ru), and, in the context of optimal/extremal equilibria, Abreu (1986) and Abreu, Pearce, and Stacchetti (1986).

These structural results are valid for repeated game payoffs evaluated either according to the limit of means or by discounting payoffs, and do not depend on the nature of the tradeoff between repeated game payoffs and complexity. We require only that preferences are weakly monotonic in the following sense: if (M_i, M_j) yields player i the same repeated game payoff as (M'_i, M_j) , then player i strictly prefers the pair (M_i, M_j) to (M'_i, M_j) if M_i has fewer states than M'_i . In particular our results encompass the case of lexicographic preferences between repeated game payoffs and implementation costs.

We show that in any N.E. of the machine game:

- (a) The two machines have an equal number of states, and maximize repeated game payoffs against one another.

Thus, despite complexity considerations, *in equilibrium*, players' choices are "fully" optimal. The result follows directly from a basic lemma established in Section 3: a player's machine need have no more states than his opponent's in order to maximize repeated game payoffs against the latter.

Notice that since the machines are finite, the pairs of states which are used in the play of the game must eventually repeat themselves in a *cycle*. We partition the states in a player's machine into those which are used in the cycle ("cycle states") and those which are not. A particular cycle state might conceivably appear more than once *within* the cycle, and noncycle states might be repeated in the initial periods of play before the cycle begins. However,

- (b) the states of a player's machine which appear in the cycle are all distinct. All the other states appear consecutively at the beginning of play and are never repeated.

Consequently, the periods in which noncycle states and cycle states appear, respectively, are connected. Although noncycle states appear only at the beginning of (equilibrium) play, a machine may return to these states after a deviation has occurred. Thus the initial interval of noncycle states has the potential to play the role of a punishment phase. The appearance of these states at the beginning of play may be interpreted as a "show of strength" or demonstration of punishment ability by both players which is "necessary" before the machine pair can enter the cyclical phase.

An immediate implication of (a) and (b) is that the two machines have an equal number of cycle and noncycle states respectively. A further implication is that there is a one-to-one correspondence between the states of the two machines in all periods of equilibrium play. This may be interpreted as saying that the machine of each player "knows" which state the other player's machine is in.

There is also a one-to-one correspondence, in equilibrium, between the stage-game actions of the two machines:

- (c) If in any two periods a player's machine plays the same stage-game action, this must be true of the other machine as well.

Thus whenever one machine changes its action, the other machine must change its action also. It also follows that the number of distinct G -outcomes is bounded above by the minimum of the number of one-shot strategies available to the two players.

These results apply exactly when repeated game payoffs are discounted. If the latter are evaluated according to the limit of means, analogous results hold modulo a possibly nonempty intermediate interval when cycle states are used by both machines, but, of course, not synchronized in the manner they are in the cycle.

The structural properties described above place restrictions on the set of equilibrium payoffs. These are particularly dramatic for 2×2 matrix games. In Section 5 we characterize completely the set of repeated game payoffs in N.E. of the machine game (hereafter referred to simply as N.E. payoffs) for symmetric 2×2 games with the limit of means and lexicographic preferences. In the

following version of the prisoner's dilemma, for example,

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

FIGURE 1.

only the vectors on the "cross" in Figure 2 are N.E. payoffs.²

Our results should be contrasted with those for the usual repeated game for which the Nash equilibrium set is very large in all the relevant spaces—strategies, plays, and payoffs. In particular the "folk-theorem" (see Aumann (1981), Fudenberg and Maskin (1986), and Rubinstein (1977, 1979)) applies: all individually rational payoffs are Nash equilibrium payoffs of the repeated game. While it is not surprising that complexity considerations affect equilibrium outcomes, the extent to which they do so is striking. This is underlined by the fact that the differences emerge even when implementation costs enter preferences lexicographically. Thus there is a severe discontinuity in the model when implementation costs are introduced; a "small" perturbation of the model has, because of strategic interactions, large consequences.

Notice that we ignore the complexity issues connected with *computing* optimal strategies and concentrate instead on the costs of *implementing* them. While a unified approach to both these questions would be very attractive, there are contexts in which the present formulation appears plausible. One such application is the organization of bureaucracies: sophisticated managers seek to devise simple rules of thumb which can be implemented mechanically by lower level employees operating in a strategic environment with peers in parallel hierarchies. The machine is viewed here as a set of managerial instructions in accordance with which subordinates operate. On a more individualistic and cerebral level one could think of the states in a machine as primitive representations of "states of mind"—vengefulness, conciliation, aggression, etc. Players seek to devise behavioral patterns which do not need to be constantly reassessed, and which economize on the number of states needed to operate effectively in a given strategic environment.




FIGURE 2.

² Strictly speaking only vectors on the "cross" which are rational convex combinations of the diagonal payoffs (2, 2) and (0, 0) or off-diagonal payoffs, (3, -1) and (-1, 3), respectively, are obtainable.

The paper proceeds as follows. We present the model in Section 2. Section 3 contains the results on the structure of N.E. when repeated game payoffs are discounted. In Section 4 we present the corresponding analysis for the limit of means. We start with the discounting case because it is technically simpler. The basic ideas of the proofs emerge clearly in this setting. Section 5 uses the results of Section 4 to characterize N.E. repeated game payoffs for symmetric 2×2 stage games. Section 6 is a literature review.

2. THE MODEL

Let $G = \langle S_1, S_2, u_1, u_2 \rangle$ be a two-person game in normal form. S_i is a finite set of actions for player i and $u_i: S_1 \times S_2 \rightarrow R$ is player i 's payoff function. An action pair is called an outcome.

The supergame of G consists of an infinite sequence of repetitions of G at $t = 1, 2, 3, \dots$. At period t , the players make simultaneous moves denoted by $s_i^t \in S_i$, which become common knowledge. A supergame strategy is a sequence of functions $\{\sigma_i^t\}_{t=1}^\infty$ where σ_i^t determines player i 's action at period t as a function of the previous $t - 1$ outcomes.

In the standard formulation a player has to choose a supergame strategy. In the machine game, G_m , a player chooses a (Moore) machine. A machine for player i , denoted by M_i , is a four tuple $\langle Q_i, q_i^1, \lambda_i, \mu_i \rangle$ where Q_i is a finite set, $q_i^1 \in Q_i$, $\lambda_i: Q_i \rightarrow S_i$, and $\mu_i: Q_i \times S_j \rightarrow Q_i$. The set Q_i is the set of states of M_i . The state q_i^1 is the initial state. The function λ_i is the output function, and $\lambda_i(q_i)$ is the G -action the machine plays whenever it is at state q_i . The function μ_i is the transition function. Whenever the machine is at state q_i and the other player chooses $s_j \in S_j$, then the machine's next state is $\mu_i(q_i, s_j)$.

A pair of machines M_1 and M_2 induces deterministically the sequences (s^t) and (q^t) in the following way:

$$q^1 = (q_1^1, q_2^1),$$

$$s^t = (\lambda_1(q_1^t), \lambda_2(q_2^t)),$$

$$q^{t+1} = (\mu_1(q_1^t, s_2^t), \mu_2(q_2^t, s_1^t)).$$

We will refer to the sequence (s^t) as the *action-play* of (M_1, M_2) , and to the sequence (q^t) simply as the *play* of (M_1, M_2) .

Since the machines are finite, $t_2 = \min \{m | q^{m+1} = q^n \text{ for some } n \leq m\}$, exists. Let $t_1 \leq t_2$ satisfy $q^{t_1} = q^{t_2+1}$. By the stationarity of the output and transition functions, the continuation of q^t after $t_2 + 1$ is just like after t_1 . We refer to $(q^{t_1}, \dots, q^{t_2})$ as the *cycle* of the machine pair (M_1, M_2) . The *length* of the cycle is $(t_2 - t_1 + 1)$.

$\pi_i(M_1, M_2)$ denotes the (average repeated game) payoff to player i , and $A_i(k_1, k_2)$, $k_2 \geq k_1$, the average payoff to player i between the periods k_1 and k_2

both inclusive. We deal with two cases:

(i) the limit of means:

$$\pi_i(M_1, M_2) = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T u_i(s^t) = \frac{1}{t_2 - t_1 + 1} \sum_{t=t_1}^{t_2} u_i(s^t),$$

$$A_i(k_1, k_2) = \frac{1}{k_2 - k_1 + 1} \sum_{t=k_1}^{k_2} u_i(s^t);$$

(ii) discounting:

$$\pi_i(M_1, M_2) = \frac{1 - \delta}{\delta} \sum_{t=1}^{\infty} \delta^t u_i(s^t),$$

$$A_i(k_1, k_2) = \frac{(1 - \delta)}{\delta(1 - \delta^{k_2 - k_1 + 1})} \sum_{t=k_1}^{k_2} \delta^{(t - k_1 + 1)} u_i(s^t).$$

Denote by $|M_i|$ the number of states in M_i .

Let \succ_1 be player 1's (strict) preference relation on the set of pairs (M_1, M_2) . We assume that \succ_1

(i) depends only on repeated game payoffs and the number of states:

$$\pi_i(M_1, M_2) = \pi_i(\bar{M}_1, \bar{M}_2) \quad \text{and} \quad |M_1| = |\bar{M}_1| \quad \text{implies} \\ (M_1, M_2) \sim_1 (\bar{M}_1, \bar{M}_2).$$

(ii) is increasing in the payoff:

$$\pi_i(M_1, M_2) > \pi_i(\bar{M}_1, \bar{M}_2) \quad \text{and} \quad |M_1| = |\bar{M}_1| \quad \text{implies} \\ (M_1, M_2) \succ_1 (\bar{M}_1, \bar{M}_2).$$

(iii) is decreasing in $|M_1|$:

$$\pi_i(M_1, M_2) = \pi_i(\bar{M}_1, \bar{M}_2) \quad \text{and} \quad |M_1| < |\bar{M}_1| \quad \text{implies} \\ (M_1, M_2) \succ_1 (\bar{M}_1, \bar{M}_2).$$

An analogous assumption applies to \succ_2 .

In some parts of the paper we consider lexicographic preferences, i.e.,

$$(M_1, M_2) \succ_i (\bar{M}_1, \bar{M}_2)$$

if $\pi_i(M_1, M_2) > \pi_i(\bar{M}_1, \bar{M}_2)$ or, $\pi_i(M_1, M_2) = \pi_i(\bar{M}_1, \bar{M}_2)$ and $|M_i| < |\bar{M}_i|$.

The pair of machines (M_1, M_2) is a *Nash Equilibrium* (N.E.) if there is no M or M'_2 such that

$$(M'_1, M_2) \succ_1 (M_1, M_2) \quad \text{or} \quad (M_1, M'_2) \succ_2 (M_1, M_2).$$

An *N.E. payoff* is a pair $\pi = (\pi_1, \pi_2)$ such that $\pi_i = \pi_i(M_1, M_2)$, $i = 1, 2$, for some N.E. (M_1, M_2) .

Remarks on the Model

Before proceeding to the results, we discuss various aspects of the model just described.

It is worth emphasizing that the model is embedded firmly within the standard noncooperative game-theoretic paradigm. The machine game is simply a normal form game in which a player's strategy is the choice of a machine. That is, players choose simultaneously, and once for all, machines to play the repeated game. The solution concept is just *Nash equilibrium*. The new element is in incorporating considerations of complexity explicitly. This is achieved by taking a player's strategy space to be the set of machines, and players' preferences to depend both on repeated game payoffs and the complexity costs of the player's machine. It should be noted that while not all strategies can be implemented by finite machines, the restriction to machine strategies is not in itself significant: any Nash equilibrium payoff of the usual repeated game (which is a rational convex combination of stage-game payoffs) is also a Nash equilibrium payoff of the repeated game in which players only have access to machine-implementable strategies.

The complexity measure is simply the number of states in a player's machine. It is one of many measures which have a plausible interpretation within the paper's framework. We view the choice of complexity measure to actually be part of the description of the repeated game, which ought to be determined by the economic environment being modelled.

The following complexity considerations are captured by the present measure:

(i) It is costly for a player to hold special routines to punish his opponent were the latter to deviate. Devices which may serve as potential punishments are only subscribed to if they are also employed along the equilibrium path.

(ii) Monitoring an opponent's behavior is also costly. Thus players economize on the states held to keep track of their opponent's actions.

On the other hand:

(iii) The measure neglects the desire of players to simplify their calculations during the course of play. This is reflected in the measure being independent of the specification of the output and transition functions.

In our definition, a transition function reacts only to the output of the *opponent's* machine. Thus, strictly speaking, a machine is not equivalent to a repeated game strategy since a machine does not include a specification of behavior after its own deviation. This is unobjectionable, as we assume that machines are "error free." A change in the definition of a transition function to allow it to depend on the output of *both* machines would not alter our results, since a player would not waste states to cope with his own deviations in the "error free" world we consider.

An attractive way to make behavior off the equilibrium path important, and thereby address issues of perfection, is to consider machines which make “mistakes,” for instance in the operation of the output function. However, this cannot be modelled in the by now traditional way à la Selten (1975), by letting the probability of mistakes go to zero. This is simply because if there is *any* real trade-off between repeated game payoffs and implementation costs (this excludes the limiting cases of zero discounting and lexicographic preferences) a player would not pay for extra states to protect himself against eventualities which are arbitrarily unlikely. This incidentally implies that our structural results are robust to small probabilities of error. More significantly it points to the need for a new approach to perfection in this setting. We hope to investigate this intriguing issue in future work.

Players are not allowed to choose new machines as the play of the game proceeds. One could depart from this static setup in a number of directions. If new machines could be chosen every period, and there were no change costs, then we would be back in the standard repeated game model with players choosing new one-state machines every period. If we do assume change costs, we would be in quite a different model whose implications we have not investigated. Another alternative was explored in (Ru) where dynamic considerations were incorporated directly into the solution concept rather than in the extensive form of the game. In (Ru) the cost of a state is viewed as a flow of maintenance costs. The solution concept therefore requires that a player would never want to alter his machine during the *course* of play. Hence a player will (eventually) drop a state which is not used infinitely often. In the present paper the cost of a state is best thought of as a one-shot cost. See Section 6 for a comparison of results.

We do not consider mixed strategy equilibria. This stems in part from a feeling of unease about the interpretation of mixed strategies in a number of standard economic applications. Unless one assumes that players use randomizing devices, which appears counterfactual, one needs much additional structure to provide a coherent context for mixed strategies. In our model, mixing raises additional delicate issues. It seems natural, given the motivation of the present paper, to regard mixing as a costly additional operation. Hence one could plausibly argue that players who were minimizing the complexity of their behavior would prefer to choose with probability one any particular machine which was a best response. Given these interpretational difficulties we feel that a purely technical investigation of the structure of mixed-strategy equilibria would be insufficiently motivated.

The preceding discussion addresses the question of mixing over deterministic machines. One might also consider non-deterministic machines, the output functions of which map to lotteries over actions. The difficulty noted earlier however reappears: given the other player's machine, for any “random machine” there exists a deterministic one with no more states which yields at least as high a repeated game payoff. Hence, if randomizing is costly, players will not choose to use random machines. The skeptical reader may verify that Lemma 1 im-

mediately adapts to the case of random machines, and leads directly to the earlier assertion.

3. THE STRUCTURE OF NASH EQUILIBRIUM WITH DISCOUNTING

We establish here the structural results described in the introduction for the case where repeated game payoffs are discounted. No further restrictions on γ_i are required. In particular (in contrast to (Ru)), we do not assume that preferences between repeated game payoffs and the complexity measure are lexicographic.

We start with a basic lemma which is a fundamental implication of the Markovian structure of finite automata, and of interest in its own right: for any machine M_1 there exists a machine with an equal number of states which yields player 2 at least as high a repeated game payoff against M_1 , as does any other machine. We remark that Lemma 1 and its proof continue to be valid when repeated game payoffs are evaluated according to the limit of means.

LEMMA 1: *For any finite machine M_1 there exists a machine \hat{M}_2 such that $|\hat{M}_2| = |M_1|$, and $\pi_2(M_1, \hat{M}_2) \geq \pi_2(M_1, M'_2)$ for all machines M'_2 .*

PROOF: Suppose player 2's objective is to maximize repeated game payoffs alone, ignoring complexity costs (i.e., the number of states in his machine). Then, given M_1 the choice of the optimal sequence (\hat{s}_2) of one-shot strategies is a Markovian problem, and has a stationary solution $\hat{s}_2: Q_1 \rightarrow S_2$. Consider \hat{M}_2 as defined by: $\hat{Q}_2 = Q_1$, $\hat{q}_2^1 = q_1^1$, $\hat{\lambda}_2(q) = \hat{s}_2(q)$, and $\hat{\mu}_2(q, \lambda_1(q)) = \mu_1(q, \hat{\lambda}_2(q))$. Clearly \hat{M}_2 has the required properties. *Q.E.D.*

Lemma 1 is of course true with player subscripts interchanged. Such translations of definitions, results and proofs will be left to the reader. An immediate implication is that if a machine pair is an N.E., then the machines must have an equal number of states, and maximize repeated game payoffs against one another. See (a) below.

Recall that for a machine pair (M_1, M_2) , t_1 and t_2 denote the times at which the cycle of (M_1, M_2) begins, and ends, respectively.

THEOREM 1: *Let (M_1, M_2) be an N.E. of the machine game with discounting. Then: (a) The machines have an equal number of states, and maximize repeated game payoffs against one another. In the course of equilibrium play, (b) the states of M_1 (resp. M_2) which appear in the first t_2 periods are distinct; (c) if in any two periods, M_1 (resp. M_2) plays the same stage game action, this must be true of the other machine as well.*

Before proceeding to the proof let us emphasize a few points. Observe that when preferences are not lexicographic, M_2 need not maximize repeated game payoffs against M_1 , in order to be a best response (to M_1). This is a property of

equilibrium. By (b), both machines have an equal number $((t_2 - t_1 + 1))$ of cycle states, and, since in equilibrium all states in a machine must be used, an equal number $((t_1 - 1))$ of non-cycle states also. Furthermore, there is a one-to-one correspondence between the states of player 1's machine and those of player 2's, as they appear in equilibrium play. We may interpret this result as implying that in equilibrium each machine "knows" which state the other machine is in. There will in general be an introductory phase before the cycle begins, during which the machines may be thought of as "counting" the time until the start of the cycle. Later (see Section 5) we will see that this interval may serve as a "punishment" phase. Part (c) implies that there is also a one-to-one correspondence between the stage game actions. Consequently, the number of distinct G -outcomes (or action-pairs) which appear in the equilibrium (action) play is bounded above by the minimum of the number of one-shot strategies available to the two players. That is, $|\{s^t | t \geq 1\}| \leq \min\{|S_1|, |S_2|\}$. As we demonstrate in a later section, (c) has a dramatic effect on the set of *payoffs* which can arise in equilibrium.

We turn now to the proof.

PROOF OF (a): By Lemma 1, player 2 does not need more than $|M_1|$ states to maximize repeated game payoffs. Hence in an N.E., $|M_2| \leq |M_1|$. Similarly for player 1. Therefore $|M_1| = |M_2|$. The equilibrium assumption and Lemma 1 now imply that M_1 and M_2 maximize repeated game payoffs against one another.

Q.E.D.

We next establish a useful lemma: if the states of machine M_1 which are used in any two periods k_1 and $k_2 + 1$ of equilibrium play are identical, then the average (discounted) payoff of player 2 between k_1 and k_2 must equal his average payoff from $k_2 + 1$ onwards (that is, to ∞), and therefore also his average payoff from k_1 onwards. Thus player 2 would not lose, in terms of repeated game payoffs, were he to modify his machine such that after k_2 , the sequence of G -outcomes in the interval $[k_1, k_2]$ were repeated cyclically forever after. Alternatively he could modify his machine to "skip" over the interval $[k_1, k_2]$, without loss of repeated game payoffs. Such revisions may be exploited to save states. Of course, states cannot be saved (without loss) in equilibrium. As we shall see below, this fact has important consequences for the structure of equilibrium machine pairs.

The notation $\pi_i^* \equiv \pi_i(M_1, M_2)$ and $\pi_i^*(k) \equiv A_i(k, \infty)$ is used below. The latter denotes average discounted payoffs (to player i) from period k on, along the equilibrium play of (M_1, M_2) .

LEMMA 2: Suppose $q_1^{k_1} = q_1^{k_2+1}$, $k_2 \geq k_1$. Then $A_2(k_1, k_2) = \pi_2^*(k_1) = \pi_2^*(k_2 + 1)$.

PROOF: Assume not. By definition,

$$\pi_2^*(k_1) = (1 - \delta^{k_2 - k_1 + 1})A_2(k_1, k_2) + \delta^{k_2 - k_1 + 1}\pi_2^*(k_2 + 1).$$

Either $A_2(k_1, k_2) > \pi_2^*(k_1)$, or $A_2(k_1, k_2) < \pi_2^*(k_1)$. In either case, there exists a machine \bar{M}_2 such that $\pi_2(M_1, \bar{M}_2) > \pi_2^*$, which contradicts (a). In the former case, player 2 can construct a machine \bar{M}_2 such that the action play of (M_1, \bar{M}_2) is identical to that of (M_1, M_2) until period k_2 , and thereafter repeats cyclically the sequence of G -outcomes in the interval $[k_1, k_2]$. For that, choose $\bar{Q}_2 = \{\bar{q}_2^t | t = 1, \dots, k_2\}$, $\bar{\lambda}_2(\bar{q}_2^t) = \lambda_2(q_2^t)$, and for all $s \in S_1$, $\bar{\mu}_2(\bar{q}_2^{k_2}, s) = \bar{q}_2^{k_1}$, and $\bar{\mu}_2(\bar{q}_2^t, s) = \bar{q}_2^{t+1}$, $t \neq k_2$. If, on the other hand, $A_2(k_1, k_2) < \pi_2^*(k_1)$, player 2 can obtain a higher payoff by modifying M_2 to skip over the interval (k_1, k_2) . *Q.E.D.*

The next two lemmas lead to a proof of (b). Let m_i denote the minimal t such that q_i^t is repeated. We first argue that $m_1 = m_2$.

LEMMA 3: $m_1 = m_2$.

PROOF: Suppose $m_2 > m_1$. Let $\bar{m}_1 > m_1$ satisfy $q_1^{\bar{m}_1} = q_1^{m_1}$. Player 2 can save the state $q_2^{m_1}$ by revising μ_2 as follows: $\mu_2(q_2^{m_1-1}, s_1^{m_1-1}) = q_2^{\bar{m}_1}$. Since $q_2^{m_1-1}$ and $q_2^{m_1}$ do not repeat themselves, this revision will create the play $q^1, \dots, q^{m_1-1}, q^{\bar{m}_1}, q^{\bar{m}_1+1}, \dots$. By Lemma 2, player 2's repeated game payoff is unchanged but he uses one less state, contradicting the N.E. assumption. *Q.E.D.*

We now show that $m_1 = m_2$ is the beginning of the cycle. In terms of the notation introduced earlier, we show $t_1 = m_1$. Given the definition of m_i , this establishes that all states of M_i which appear before the cycle begins, are distinct, and are never repeated in the play of (M_1, M_2) .

LEMMA 4: *The cycle begins at $t = m_1$.*

PROOF: Let \bar{m}_i be the minimal $t > m_i$ such that $q_i^t = q_i^{m_i}$. If $\bar{m}_1 = \bar{m}_2$, $q^{m_1} = q^{\bar{m}_1}$, and the lemma is true. Now suppose that $\bar{m}_1 > \bar{m}_2$. If $q_2^{\bar{m}_1} \neq q_2^t$ for all $t < \bar{m}_1$, player 2 can save $q_2^{\bar{m}_1}$ by revising μ_2 such that $\mu_2(q_2^{\bar{m}_1-1}, s_1^{\bar{m}_1-1}) = q_2^{m_1}$. Play prior to m_1 is unaffected and is followed by the cycle $(q^{m_1}, \dots, q^{\bar{m}_1-1})$. By Lemma 2, player 2's payoffs are unchanged. Alternatively suppose $q_2^{\bar{m}_1} = q_2^\ell$ for some $m_1 \leq \ell < \bar{m}_1$. If $\ell = m_1$, $q^{m_1} = q^{\bar{m}_1}$ and we are done. If not, player 1 can save $q_1^{m_1}$ by revising μ_1 as follows:

$$\mu_1(q_1^{m_1-1}, s_2^{m_1-1}) = q_1^{\bar{m}_2}, \quad \text{and}$$

$$\mu_1(q_1^{\bar{m}_1-1}, s_2^{\bar{m}_1-1}) = q_1^\ell.$$

These changes create the play: $q^1, \dots, q^{m_1-1}, q^{\bar{m}_2}, q^{\bar{m}_2+1}, \dots, q^{\bar{m}_1-1}$ followed by the perpetual repetition of $(q^\ell, q^{\ell+1}, \dots, q^{\bar{m}_1-1})$. By Lemma 2, player 1's payoffs are the same. *Q.E.D.*

Finally, we complete the proof of (b).

PROOF OF (b): From the remark preceding Lemma 4 it is clear that we need only show that q_i^1, \dots, q_i^2 , the states of M_i which appear in the cycle, are distinct. The basic construction used is similar to that in Lemma 3. Suppose not, and let m_i now denote the minimal $t \geq t_1$ such that $q_i^t = q_i^{t'}$ for some $t < t' \leq t_2$. Let \bar{m}_i be the maximal $t \leq t_2$ such that $q_i^t = q_i^{m_i}$. (Note that since both machines have an equal number of states, all of which must appear in equilibrium play, if m_1 is well defined then so is m_2 .) If $m_1 = m_2$ and $\bar{m}_1 = \bar{m}_2$ the cycle is $(q^{m_1}, \dots, q^{\bar{m}_1 - 1})$, a contradiction. If $m_2 > m_1$, or $m_1 = m_2$ and $\bar{m}_1 > \bar{m}_2$, player 2 can save $q_2^{m_1}$ by revising μ_2 as in Lemma 3. By Lemma 2, payoffs are unaffected, and we have obtained our usual contradiction. Q.E.D.

The basic idea of the proof of (c) is as follows: if player 2 uses two distinct states to play the same action, say in periods 2 and 6, it must be so as to be able to transit to two different states in periods 3 and 7 respectively. But if the corresponding actions of player 1 are distinct, a single state can be used to effect the required transition by exploiting the *switch* in player 1's actions. Thus *unsynchronized* switches of action by M_1 (relative to M_2) can be used by M_2 to "keep track" of the state of play.

PROOF OF (c): Assume not and suppose that for some t, t' $\lambda_1(q_1^t) \neq \lambda_1(q_1^{t'})$ and $\lambda_2(q_2^t) = \lambda_2(q_2^{t'})$. The states q_2^t and $q_2^{t'}$ must be distinct since otherwise by (b) $q_1^t = q_1^{t'}$, and $\lambda_1(q_1^t) = \lambda_1(q_1^{t'})$. However, although $q_2^t \neq q_2^{t'}$ the output function assigns the same action to the two states. Furthermore, in the play of the game, the transition function operates on the two states in response to two *different* actions. Player 2 can save a state by replacing the states q_2^t and $q_2^{t'}$ by the single state q^* and revising μ_2 as follows:

- (1) $\mu_2(q_2^{h-1}, s_1^{h-1}) = q^*$,
- (2) $\mu_2(q^*, s_1^h) = \begin{cases} q_2^{h+1} & \text{if } q_2^{h+1} \neq q_2^t, q_2^{t'} \\ q^* & \text{otherwise,} \end{cases}$

where $h = t, t'$. Of course, if $t = 1 < t'$, q^* is the initial state and (1) only applies to $h = t'$ (and also to $h = t_2 + 1$ if the cycle begins at $t = 1$). By (b), in the play of (M_1, M_2) q_1^h always appears with q_2^h , $h = t, t'$, and by assumption $s_1^t \neq s_1^{t'}$. Hence these revisions are consistent, and result in the same play of G -outcomes as (M_1, M_2) . Q.E.D.

4. THE STRUCTURE OF NASH EQUILIBRIUM WITH THE LIMIT OF MEANS

We now assume that repeated game payoffs are evaluated according to the limit of means. The results obtained here are basically the same as those for discounting, with one difference: between the introductory phase of states which are never repeated, and the cycle phase, there may be an intermediate interval in which cycle states are used, but not synchronized with the other machine's states as in the cycle. Let *Phase II* be the time interval during which the machines use cycle states but before the cycle begins.

THEOREM 1*: *Let (M_1, M_2) be an N.E. of the machine game with the limit of means. Then: (a) The machines have an equal number of states, and maximize repeated game payoffs, against one another. In equilibrium play, (b) the states of M_1 (resp. M_2) which appear in the cycle are distinct; all the other states appear consecutively at the beginning of play and are never repeated; (c) if in any two periods, neither of which belongs to Phase II, M_1 (resp. M_2) plays the same stage game action, this must be true of the other machine as well.*

The explanatory remarks following Theorem 1 continue to be valid *modulo Phase II*. Thus there is a one-to-one correspondence between the states of the two machines and the stage-game actions they play, in all periods of equilibrium play, except in Phase II. Now the number of distinct G outcomes in all periods other than those in Phase II is bounded above by the minimum of the number of one-shot strategies available to the two players. That is, $|\{s^h | h \notin \text{Phase II}\}| \leq \min\{|S_1|, |S_2|\}$. The one-to-one correspondence between stage-game actions in the cycle, together with the fact that for the limit of means repeated game payoffs depend only on the cycle, place strong restrictions on the set of equilibrium payoffs in the machine game. These are heavily exploited in the next section.

We show by example that Phase II may be nonempty. Consider the stage-game:

	A	B
A	1, 1	0, 0
B	0, 0	1, 1

FIGURE 3.

The “transition diagram” below denotes a machine with two states q_A and q_B , the outputs of which are A and B , respectively. The machine transits from state q_A to state q_B if the other machine’s output is A and remains in state q_A otherwise. It transits from state q_B to state q_A regardless of the other machine’s output (A or B). In the diagram, a circle (or vertice) represents a state. The output function is defined by the action (or letter) beneath the state and the transition function by directed arcs from a current state to a subsequent state in response to the action(s) indicated beside the arc.




FIGURE 4.

The machine pair (M_1, M_2) where M_1 starts at q_A and M_2 starts at q_B is an N.E. Equilibrium play is $AB, AA, BB, AA, BB, \dots$; Phase II = $\{1\}$, the cycle begins at $t_1 = 2$ and ends at $t_2 = 3$.

The proof of Theorem 1* is given in the Appendix.

5. NASH EQUILIBRIUM PAYOFFS IN REPEATED 2×2 GAMES

Theorem 1*(c) places significant restrictions on the pattern of G -outcomes generated in the course of equilibrium play. For repeated 2×2 games it implies that *all G -action pairs in the cycle must lie on one or the other diagonal of the matrix*. We provide a complete characterization of N.E. payoffs, with the limit of means, and lexicographic preferences, for symmetric 2×2 games. Away from the limit and for other preferences, the N.E. payoff set will depend in a straightforward way on the discount factor δ , and the exact trade-off between the repeated game payoff and the complexity measure. The reader is reminded that, for discounting, Theorem 1(c) implies the same diagonal restrictions on G -outcomes which can arise in equilibrium cycles as Theorem 1*(c) does for the limit of means.

Let D_i be player i 's minmax strategy, i.e., D_i solves $\min_s \max_{s_j} u_j(s_i, s_j) \equiv v_j$, and let C_i denote the other strategy of S_i . Let $MD = \{(C_1, C_2), (D_1, D_2)\}$ denote the outcomes on the main diagonal and let $AD = \{(C_1, D_2), (D_1, C_2)\}$ denote the outcomes on the auxiliary diagonal. We refer to MD and AD as the main and auxiliary diagonals respectively. For the rest of this section (C, D) should be read as (C_1, D_2) , etc. MD^u denotes the set of rational convex combinations of payoffs on the main diagonal, which are also strictly individually rational, i.e.,

$$MD^u = \{ (u_1, u_2) | u_i > v_i, \text{ and } u_i = \lambda u_i(C, C) + (1 - \lambda) u_i(D, D) \\ i = 1, 2 \text{ for some rational number } 0 \leq \lambda \leq 1 \}.$$

AD^u is defined analogously.

$NE(G)$ denotes the set of N.E. payoffs of the stage game G and $NE(G_m)$ denotes the set of N.E. (repeated game) payoffs of the machine game with the limit of means and lexicographic preferences.

We assume for convenience that no two distinct G -outcomes yield either player the same payoff. That is, if $u_i(s) = u_i(s')$, then $s = s'$.

CONCLUSION 1: Let G be a symmetric 2×2 matrix game with no equal payoffs for either player:

- (i) if $u_1(D, C) > u_1(C, C)$, then $NE(G_m) = NE(G) \cup MD^u \cup AD^u$;
- (ii) if $u_1(D, C) < u_1(C, C)$, then $NE(G_m) = NE(G) \cup MD^u$.

Thus all individually rational (rational convex combinations of) payoffs on the main diagonal are N.E. payoffs of the machine-game. The payoffs on the auxiliary diagonal are N.E. payoffs of the machine-game only if player 1 (and analogously player 2) cannot gain from deviating from (D, C) to (C, C) , a condition satisfied, for example, in the prisoners' dilemma.

Before proving the theorem we demonstrate its conclusion on several games:

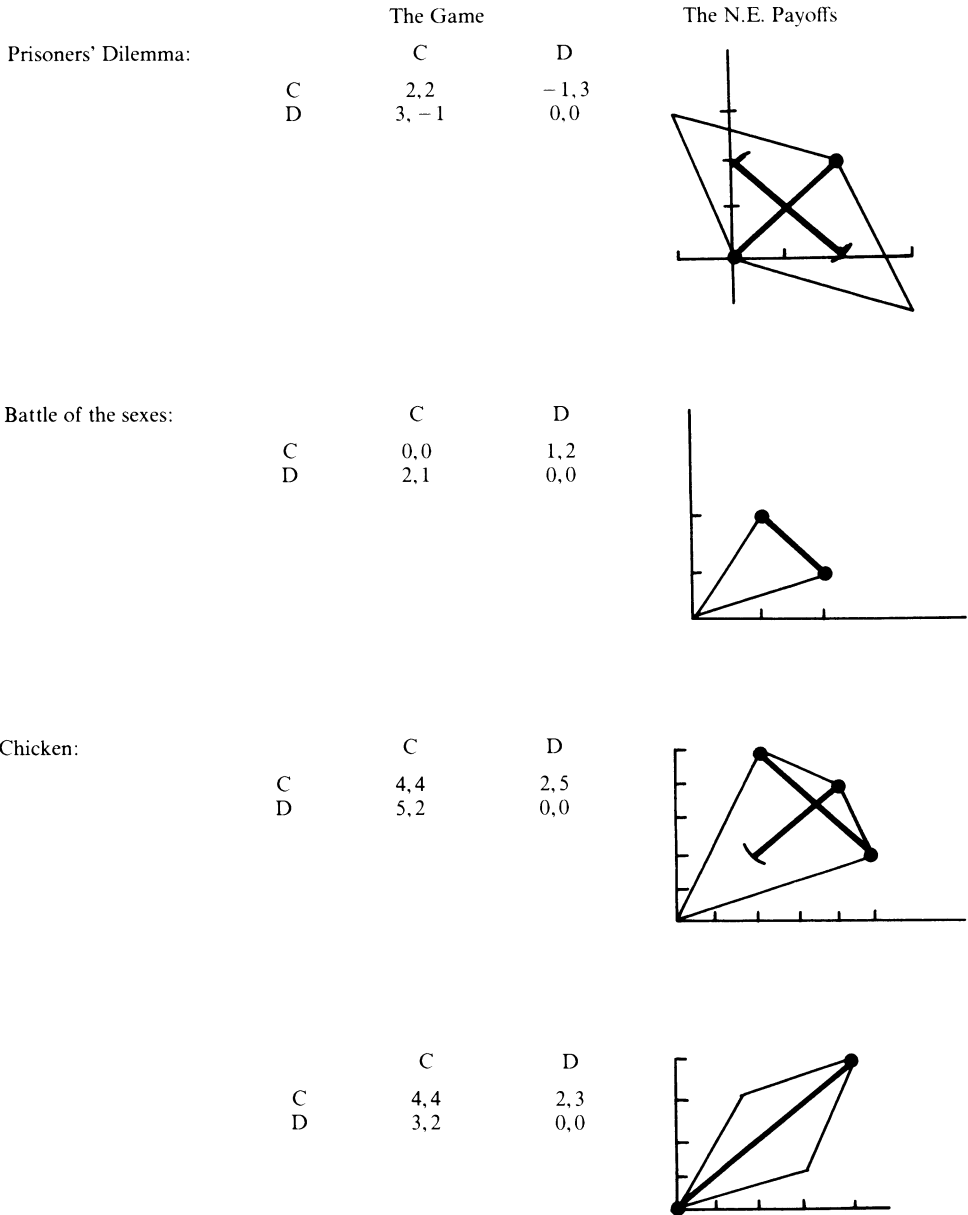


FIGURE 5.

REMARK: G_m may have no N.E. For

	C	D
C	1, -1	-1, 1
D	-1, 1	1, -1

$v_i = 1$ and no feasible payoff is individually rational.

PROOF—*Necessity*: Let (M_1, M_2) be an N.E. of G_m and (π_1^*, π_2^*) the associated payoff-vector. Clearly $\pi_i^* \geq v_i$. By Theorem 1*(c) the outcomes in the cycle of (M_1, M_2) are either only elements of MD or only elements of AD . Hence there are at most two distinct outcomes in the cycle.

We first show that $\pi_1^* = v_1$ implies $(\pi_1^*, \pi_2^*) \in NE(G)$. Obviously if $|M_1| = 1$ then $|M_2| = 1$ and $(\pi_1^*, \pi_2^*) \in NE(G)$. Assume $|M_1| > 1$. By the unequal payoff assumption and the definition of D :

(a) If there are two outcomes $(X, D), (Y, C)$ in the cycle, then $u_1(X, D) < \pi_1^* < u_1(Y, C)$. Also $u_1(Y, D) = \pi_1^* < u_1(Y, C)$. Hence player 1 can deviate to a one-state machine which plays only Y .

(b) If there is a single outcome in the cycle it must be of the form (X, D) where $u_1(X, D) = v_1$. If $X = D$ then $(D, D) \in NE(G)$. If $X = C$ and $(C, D) \notin NE(G)$, then $u_2(C, C) > u_2(C, D)$, and by symmetry $u_1(C, C) > u_1(D, C)$. Hence, by the definition of D , $u_1(C, C) > v_1$. Together with $u_1(C, D) = v_1$, this implies that player 1 can profitably deviate to a one-state machine which plays only C .

Now consider stage games such that $u_1(D, C) < u_1(C, C)$. To complete the proof we show that there are no equilibrium pairs (M_1, M_2) with cycle outcomes only on AD , and associated repeated game payoffs π_i^* which are strictly individually rational for both players. First note that since $\pi_i^* > v_i$, both (C, D) and (D, C) must appear in the cycle, and $u_1(D, C) > \pi_1^* > v_1 \geq u_1(C, D)$. The proof now proceeds in three steps.

Step 1: Phase I $\neq \emptyset$. This step is analogous to the first part of Proposition 2 of (Ru). Let $L = (t_2 - t_1 + 1)$ be the length of the cycle. Suppose by way of contradiction that Phase I is empty. Consider $k_1 \geq t_1$ such that $s_2^{k_1} = C$. Let $0 \leq \ell \leq L - 1$ satisfy $q_2^{k_1 + \ell} = \mu_2(q_2^{k_1}, C)$. Then player 1 can alter his machine to obtain the cycle of (action) outcomes $((C, C), s^{k_1 + \ell}, s^{k_1 + \ell + 1}, \dots, s^{k_1 + L - 1})$ if $\ell > 0$, and (C, C) alone if $\ell = 0$. Since $u_1(C, C) > u_1(D, C) = u_1(s^{k_1}) > \pi_1^*$, and this deviation must not increase payoffs, $\ell > 1$, and $A_1(k_1 + 1, k_1 + \ell - 1) > \pi_1^* > v_1$. Since $u_1(C, D) \leq v_1$ there exists k_2 such that $k_1 + 1 \leq k_2 \leq k_1 + \ell - 1$, $s^{k_2} = (D, C)$, and $A_1(k_1 + 1, k_2) > \pi_1^*$. Letting k_2 play the role of k_1 we can define k_3 and so on, to obtain a sequence k_1, k_2, k_3, \dots . Since M_2 has a finite number of states we must eventually find p, n such that $k_p = k_n + aL$ for some natural number a , which contradicts $A_1(k_n + 1, k_p) = \pi_1^*$.

Step 2: Phase I does not contain either (C, C) or (D, D) . This follows directly from 1*(c).

Step 3: Assume that $s^1 = (D, C)$. Define k by $q_2^k = \mu_2(q_2^1, C)$. Since $u_1(C, C)$ is the largest payoff for player 1 it is clear that $k \neq 1$. If q_2^k is a cycle state, let

$k_2 \geq t_1$ satisfy $q_2^{k_2} = q_2^k$. Otherwise let $k_2 = k$. Consider $k_1 \geq t_1$ such that $s_1^{k_1} = C$. Player 1 can save a state by dropping q_1^1 , making $q_1^{k_1}$ the initial state, and revising μ_1 as follows:

$$(1) \quad \mu_1(q_1^{k_1}, C) = q_1^{k_2},$$

and to cover the case where q_2^k is not a cycle state and $q_1^{k_1}$ is a state which appears in Phase II,

$$(2) \quad \mu_1(q_1^{t_0-1}, s_2^{t_0-1}) = q_1^n \text{ where } n \geq t_1 \text{ satisfies } q_2^n = q_2^{t_0},$$

which results in player 1's revised machine skipping over Phase II. Since $s_2^{k_1} \neq C$, the cycle is unaffected, and the proof of necessity is complete.

Sufficiency: Obviously $NE(G) \subset NE(G_m)$. We complete the proof by explicitly constructing N.E. machine pairs which yield the required payoffs.

Case A: $(\pi_1^*, \pi_2^*) \in MD^u$. Consider N_1, N_2 such that

$$\pi^* = \frac{N_1}{N_1 + N_2} u(C, C) + \frac{N_2}{N_1 + N_2} u(D, D) > v.$$

Since $u_1(D, D) \leq v_1$, $u_1(C, C) > v_1$. Let N_0 satisfy $u_1(D, C) - u_1(C, C) \leq N_0(\pi_1^* - u_1(D, D))$. Then a pair of machines of the form:




FIGURE 6.

is an N.E.

Fix M_2 as specified above, and consider deviations by Player 1. The cycles of outcomes which he can generate fall into two categories: those with an occurrence of (D, C) , and those without. The former consist of blocks of the form:

$$(1) \quad \begin{aligned} &N_1 (D, D)'s, (N_2 - x) (C, C)'s \text{ where } x \geq 1, \text{ one } (D, C), \\ &N_0 (D, D)'s \text{ and some mix (possibly empty) of } (C, D)'s \\ &\text{and } (D, D)'s. \end{aligned}$$

From the definition of N_0 it follows that the equilibrium (outcome) cycle dominates.

The latter consist of blocks of the form:

$$(2) \quad (N_0 + N_1 - x) (D, D)'s \text{ and one } (C, D),$$

which yields a payoff less than v_1 , or

$$(3) \quad N_1 (D, D)'s \text{ and } N_2 (C, C)'s, \text{ the equilibrium cycle.}$$

Thus the equilibrium cycle yields strictly higher repeated game payoffs than any other attainable cycle. It should be clear that the equilibrium cycle cannot be

attained using a machine with less than $(N_0 + N_1 + N_2)$ states. A formal proof may be constructed along the lines of Proposition 2 of (Ru).

REMARK: The analysis of (Ru) shows that Phase I is necessary for an equilibrium on the main diagonal with a cycle outcome which is not an N.E. of the stage game.

Case B: $u_1(D, C) > u_1(C, C)$ and $(\pi_1^*, \pi_2^*) \in AD^u$. Consider N_1, N_2 such that

$$\pi^* = \frac{N_1}{N_1 + N_2} u(D, C) + \frac{N_1}{N_1 + N_2} u(C, D) > v.$$

Let $N = N_1 + N_2$. Then the pair of machines below is an N.E.:




FIGURE 7.

Consider deviations by player 1 against M_2 . In any cycle when M_2 is in any one of its first N_1 states it is clearly inoptimal for player 1 to play C rather than D , since player 2's next state is unaffected and $u_1(D, C) > u_1(C, C)$. Abstracting from deviations of this sort, all attainable cycles are comprised of:

- (1) blocks of the form: $(N_2 - x) (C, D)$'s and one (D, D) , or
- (2) $N_1 (D, C)$'s followed by $n \geq 0$ blocks of type (1), and finally $N_2 (C, D)$'s.

Clearly the equilibrium cycle ($n = 0$) dominates. It is intuitively obvious that unaided by unsynchronized switches of actions, player 1 cannot play the required cycle sequence of C 's and D 's with a smaller machine than M_1 . The formal proof is almost identical to Proposition 2 of (Ru), and is omitted. Q.E.D.

REMARK: It should be clear that adding identical rows or columns will in general alter the set of equilibrium payoffs. We feel that the context determines a "natural" representation of the stage game and that a theory of complexity ought not to be insensitive to presentational changes which are regarded as "innocuous" in the standard theory. Notice that such representational changes are irrelevant if the model is reformulated to require the transition between states to depend on the player's own stage payoff rather than on the other machine's output. This

would correspond to the view that what identifies the other machine's action is its consequence rather than its label. We are grateful to Roger Myerson and Sylvain Sorin for raising and clarifying these issues.

6. RELATED LITERATURE

This paper continues the investigation begun in Rubinstein (1986). The most significant difference between (Ru) and the present paper is the solution concept used. (Ru) analyzes a much stronger solution concept: players never wish to drop states from their machines during the course of play, holding fixed the other player's machine, and given that they only pay for states which are held in the machine forever. The appropriateness of the concept depends on the interpretation of the model. In (Ru) the equilibrium concept is motivated by the view that complexity costs are a flow of maintenance costs which players can reduce by dropping states as the game proceeds. In the terminology of the present paper, this concept implies that Phase I is empty, that is, all the states in a machine must appear in the cycle. The present paper, on the other hand, adopts a different approach, using the Nash equilibrium solution concept, in a context in which players commit themselves to machines at the beginning of play.

The other differences are that (Ru) deals only with the limit of means and lexicographic preferences. For the prisoners' dilemma, the solution concept of (Ru) eliminates the payoffs on the main diagonal (except $(0,0)$) from the set of solution payoffs.

Several earlier papers on "complexity" and "bounded rationality" are referred to and commented on in (Ru). Of particular interest are Aumann (1981), who first proposed the use of finite automata in the context of repeated games, Green (1982), Radner (1978, 1980), and Smale (1980).

An alternate approach to "complexity and repeated games" has recently been proposed by Neyman (1985). He investigates a (finitely) repeated game model in which the pure strategies available to players are those which can be generated by machines which use no more than a certain number of states. He analyzes the set of mixed strategy equilibria of this game. Ben-Porath (1986), and Megiddo and Widgerson (1985) pursue this line of enquiry, the latter in the context of Turing machines.

The decisive difference, in our view, between Neyman's formulation and the present one is that in the former, the number of states is exogenously given, whereas in our work it is explicitly a choice variable. The modelling approach we pursue seems to us very economic: complexity is costly and players seek to minimize these costs. Indeed, this "cost minimization" is at the heart of our results.

The goal of Green, Neyman, Radner, and Smale is to use complexity to expand the equilibrium set, specifically to show that cooperation can emerge even with a finite horizon. We, of course, consider infinite repetitions and seek to restrict the set of equilibrium payoffs.

Other recent work includes Kalai and Stanford (1988). Of most relevance to the present paper is their demonstration that for a two-player repeated game, if a pair of strategies is a "discount robust subgame perfect equilibrium" (i.e., if the pair is a perfect equilibrium for an open interval of discount factors), then the strategies must have equal complexity. This theorem is related to our result that if (M_1, M_2) is an N.E., then $|M_1| = |M_2|$, which is an immediate corollary of Lemma 1. Their theorem is true for quite distinct reasons and the basic assumptions of our analysis are very different: in Kalai and Stanford, complexity does not enter players' preferences, nor are there restrictions (apart from finiteness) on the complexity of the strategies players may use. The reader is referred to their paper and the others cited for further details.

7. CONCLUDING REMARKS

When a decision maker formulates a rule of behavior, he confronts the following dilemma. On the one hand he wishes to choose a rule which serves his direct interests in the best possible way. On the other he attempts to make the rule as simple as possible. This paper investigates the effect of introducing the latter consideration explicitly into the description of the model. In particular we focus on whether a repeated game model without costs of implementing strategies is a good approximation of a model where such costs are "small." In our analysis a striking discontinuity emerges. It suggests a need to expand the scope of standard game theoretic analysis to include procedural aspects of decision making.

Viewed as an exercise in repeated games, our results provide strong restrictions on the structure of N.E. strategies. They also shrink dramatically the set of N.E. payoffs. It would be reckless, however, to regard these results as amounting to a new theory of repeated games: the assumptions which drive them are too special to permit such an interpretation. Nevertheless we feel they are suggestive, and serve to demonstrate that complexity considerations can be successfully incorporated into standard models.

Department of Economics, Harvard University, Cambridge, MA 02138, U.S.A.

and

Department of Economics, The Hebrew University, Jerusalem, Israel

Manuscript received September, 1986; final revision received September, 1987.

APPENDIX

The main difference in the proof of Theorem 1* relative to Theorem 1 is in part (b), the argument for which is now rather involved. The proofs of Lemma 1 and (a) are exactly as before and are omitted, and (c) requires only a minor addition. The source of the difficulty in (b) is that with the limit of means, payoffs only depend on the cycle. Hence it is only possible to establish Lemma 2 for periods k_1 and k_2 after the cycle begins. In particular, under the hypotheses of Lemma 2 it is

possible that $A_2(k_1, k_2) < \pi_2^*$, for k_1 before the cycle begins.

Note that now $\pi_i(M_1, M_2) \equiv \pi_i^* = \pi_i^*(k)$ for all k , and recall the definitions of t_1, t_2 .

LEMMA 2*: Suppose $q_1^{k_1} = q_1^{k_2+1}, k_2 \geq k_1 \geq t_1$. Then $A_2(k_1, k_2) = \pi_2^*$.

PROOF: Suppose not. Assume w.l.o.g. that $A_2(k_1, k_2) > \pi_2^*$ (otherwise there exists $k_3 > k_2 + 1$ such that $q_1^{k_2+1} = q_1^{k_3+1}$ and $A_2(k_2 + 1, k_3) > \pi_2^*$). Consider \bar{M}_2 as specified in Lemma 2. Clearly $\pi_2(M_1, \bar{M}_2) = A_2(k_1, k_2) > \pi_2^*$, which contradicts (a). Q.E.D.

Recall that m_i denotes the minimal t such that q_i^t is repeated.

LEMMA 3*: $m_1 = m_2$.

The proof is identical to Lemma 3.

Unfortunately Lemma 4 cannot be repeated; the cycle need not begin at m_1 . However, for both machines it is still the case that only cycle states appear after $m_1 = m_2$. Hence non-cycle states appear only once, and at the beginning of play, and both machines have the same number of non-cycle states.

We use the following standard convention for subscripts: if $i = 1, -i = 2$, and vice-versa. Denote by Q_i^C the states of M_i which appear in the cycle, and by Q_i^I those which do not. That is $Q_i^C = \{q_i^t | t \geq t_1\}$, and $Q_i^I = Q_i - Q_i^C$.

LEMMA 4*: $q_i^t \in Q_i^C$ for all $t \geq m_i$ and $i = 1, 2$.

PROOF: We first show that:

(i) if $q_i^{m_i} \in Q_i^C$, then $q_{-i}^t \in Q_{-i}^C$ for all $t \geq m_i, i = 1, 2$.

Let $i = 1$ and suppose $q_2^k \in Q_2^I$ for some $k \geq m_1$. Choose \bar{m}_1 such that $q_1^{\bar{m}_1} = q_1^{m_1}$ and $q_2^k \neq q_2^{\bar{m}_1}$ for all $t \geq \bar{m}_1$. Player 2 can save q_2^k by revising μ_2 as in Lemma 3. To complete the proof we argue that

(ii) $q_i^{m_i} \in Q_i^C, i = 1, 2$.

Suppose not. Then by Lemma 3* and (i), $q_i^{m_i} \notin Q_i^C, i = 1, 2$. Let \bar{m}_i be the maximal t such that $q_i^t = q_i^{m_i}$. Assume w.l.o.g. that $\bar{m}_1 > \bar{m}_2$ ($\bar{m}_1 = \bar{m}_2$ implies that the cycle starts at $t_1 = m_1 = m_2$). Player 2 can save $q_2^{m_1}$ by revising μ_2 as before. Q.E.D.

Let $t_0 = m_i$ denote the beginning of Phase II. Lemmas 3* and 4* allow us to divide equilibrium play into the following intervals:

Phase I	$1 \leq t \leq t_0 - 1$	$q_i^t \in Q_i^I$;
Phase II	$t_0 \leq t \leq t_1 - 1$	$q_i^t \in Q_i^C$ but the pair of states
		q_i^t does not repeat itself;
Phase III	$t \geq t_1$	the cycle phase.

Since all states must be used in equilibrium play, all noncycle states must appear in Phase I. That is, $Q_i^I = \{q_i^t | t \leq t_0 - 1\}$. Also, by (a), the two machines must have an equal number of cycle states.

To establish (b) it remains to show that there is no repetition of a machine's states within the cycle. If Phase II is empty the proof is exactly as in the previous section. We complete the proof below (Lemmas 5-10) assuming that Phase III is nonempty.

Let m denote any element of Phase II for which player 1's state is the same as his state at the beginning of Phase II. That is, $m \geq t_1$ and satisfies $q_1^m = q_1^{t_0}$. Let $L = |Q_1^C|$ be the number of cycle states in each of the two machines. We first show that the $(L - 1)$ states of player 1 which follow any appearance of $q_1^{t_0}$ in Phase III must be his remaining cycle states. The argument proceeds via Lemma 5.

LEMMA 5: Suppose $q_1^{k_1} = q_1^{k_2+1}, k_2 \geq k_1 \geq t_1, k_2 \geq m$. Then $L \leq k_2 - \underline{k} + 1$, where $\underline{k} = \min(m, k_1)$.

PROOF: Consider \bar{M}_2 as defined by:

$$\begin{aligned} \bar{Q}_2 &= \{ \bar{q}_2^t \mid t = 1, \dots, t_0 - 1 \text{ or } t = \underline{k}, \underline{k} + 1, \dots, k_2 \}, \\ \bar{\lambda}_2(\bar{q}_2^s) &= \lambda_2(q_2^s), \text{ and for all } s \in S_1 \\ \bar{\mu}_2(\bar{q}_2^t, s) &= \begin{cases} \bar{q}_2^m, & t = t_0 - 1, \\ \bar{q}_2^{\underline{k}}, & t = k_2, \\ \bar{q}_2^{\underline{k}+1} & \text{otherwise.} \end{cases} \end{aligned}$$

\bar{M}_2 has $|Q_2^C| + (k_2 - \underline{k} + 1)$ states and $\pi_2(M_1, \bar{M}_2) = A_2(k_1, k_2)$. Now use Lemma 2*, and note that $|M_2| = |Q_2^C| + |Q_2^C|$. Q.E.D.

LEMMA 6: *The elements of the sequence $\{q_1^m, \dots, q_1^{m+L-1}\}$ are distinct; i.e., $Q_1^C = \{q_1^m, \dots, q_1^{m+L-1}\}$.*

PROOF: Suppose not. Then there exists $\ell_1 \geq 0, \ell_2 \leq L - 2$, such that $q_1^{m+\ell_1} = q_1^{m+\ell_2+1}$. Let $k_1 = m + \ell_1, k_2 = m + \ell_2$ and check that Lemma 5 is contradicted. Q.E.D.

Consider the L -sequence defined by an appearance of $q_1^{t_0}$ in Phase III and the $(L - 1)$ states of player 1 which follow. Lemma 8 asserts that the L states of player 2 which are paired with this sequence of player 1's states, are all distinct, i.e., the corresponding sequence for player 2 contains all his cycle states. Lemma 7 is a key intermediate step.

LEMMA 7: *Suppose $q_1^{k_1} = q_1^{k_2+1}, k_2 \geq k_1 \geq t_1, k_2 \geq m$. Define $\underline{k} = \min(m, k_1)$, and $A = \{q_2^k \mid \underline{k} \leq k \leq k_2\}$. Then*

- (i) $A = Q_2^C$ if $(*) (q_2^{\underline{k}}, q_2^{\underline{k}+1}) \neq (q_2^{k_2}, q_2^{k_2+1})$ for all $\underline{k} \leq k \leq k_2 - 1$, and
- (ii) $q_2^{k_2+1} \in A$.

PROOF: (i) Suppose $(*)$ holds. Observe that it implies $s_1^k \neq s_1^{k_2}$ for all k such that $\underline{k} \leq k \leq k_2 - 1$ and $q_2^k = q_2^{k_2}$. Hence player 2 can obtain the payoff $A_2(k_1, k_2)$ by revising μ_2 as follows:

- (1) $\mu_2(q_2^{t_0-1}, s_1^{t_0-1}) = q_2^m,$
- (2) $\mu_2(q_2^{k_2}, s_1^{k_2}) = q_2^{\underline{k}}.$

Only the states $Q_2^C \cup \{q_2^k \mid \underline{k} \leq k \leq k_2\}$ now appear in the play of the game. To complete the proof note that by Lemma 2*, $A_2(k_1, k_2) = \pi_2^*$.

(ii) Suppose not. Then $(*)$ is satisfied and the contradiction with (i) is immediate. Q.E.D.

LEMMA 8: $Q_2^C = \{q_2^m, \dots, q_2^{m+L-1}\}$.

PROOF: Suppose there exists $q \in Q_2^C$ such that $q \notin \{q_2^m, \dots, q_2^{m+L-1}\}$. Let k_2 be the minimal $t \geq m$ such that $q = q_2^t$. By Lemma 6 there exists $m \leq k_1 \leq m + L - 1$ such that $q_1^{k_1} = q_1^{k_2+1}$. Lemma 7(ii) now yields a contradiction. Q.E.D.

The next lemma uses Lemma 7 with the player subscripts interchanged. It also explicitly uses the assumption that Phase II is nonempty. In particular the latter implies $q^t \neq q^{t_0}$ for all $t \geq t_0 + 1$.

LEMMA 9: $q_1^{m+L} = q_1^m$.

PROOF: Assume not. Let $k_2 = m + L - 1$ and h_2 be the minimal $h \geq m + L$ such that $q_1^{h+1} = q_1^m$. By Lemma 8 and nonempty Phase II there exists n such that $m + 1 \leq n \leq m + L - 1$, and $q_2^n = q_2^{t_0}$. By Lemma 8 (applied to player 1) $h_2 + 1 \leq n + L - 1$. Hence by Lemma 6 (applied to player 2) $q_2^{k_2+1} \neq q_2^{h_2+1}$. Thus by Lemma 8 there exist $k_1, h_1 \in \{m, \dots, m + L - 1\}, k_1 \neq h_1$, such that $q_2^{k_1} = q_2^{h_1+1}$ and $q_2^{h_1} = q_2^{h_2+1}$. Either $k_1 > m$, or $h_1 > m$. Suppose that $k_1 > m$. Then $q_1^{k_1} \notin \{q_1^k \mid \min(n, k_1) \leq k \leq k_2\}$. By Lemma 6 player 1's states in the interval $\min(n, k_1) \leq k \leq k_2$ are distinct. Hence $(*)$ is

satisfied, and Lemma 7 applied to k_1, k_2 leads to a contradiction. An analogous contradiction results if $h_1 > m$, and we consider h_1, h_2 . Q.E.D.

LEMMA 10: $q_1^{m+L+\ell} = q_1^{m+\ell}, \ell = 0, 1, \dots, L-1$.

PROOF: Let ℓ_2 be the minimal $0 \leq \ell \leq L-1$ such that $q_1^{m+L+\ell} \neq q_1^{m+\ell}$. By Lemma 6 there exists $0 \leq \ell_1 \leq L-1$ such that $q_1^{m+\ell_1} = q_1^{m+L+\ell_2}$. Observe that the preceding lemmas apply to any $k \geq t_1$ such that $q_1^k = q_1^{t_0}$. By Lemma 9 $q_1^{m+L} = q_1^{t_0}$. Hence by Lemma 6 $\ell_2 < \ell_1$. Define $k_1 = m + \ell_1, k_2 = m + L + \ell_2 - 1$, and check that Lemma 5 is contradicted. Q.E.D.

This completes the proof of (b).

Finally, we need to make only a minor addition to the proof of (c) provided in Section 4. The machine has now also to be altered to skip over Phase II when the latter is nonempty. If Phase I is empty, set the initial state to be q_2^m if $q_2^m \neq q_2^t, q_2^{t'}$ and to q^* otherwise. If Phase I is *nonempty* revise μ_2 (further) as follows:

$$\mu_2(q_2^h, s_1^{t_0-1}) = \begin{cases} q_2^m & \text{if } q_2^m \neq q_2^t, q_2^{t'}, \\ q^* & \text{otherwise,} \end{cases}$$

$$\text{where } q_h^2 = \begin{cases} q_2^{t_0-1} & \text{if } q_2^{t_0-1} \neq q_2^t, q_2^{t'}, \\ q^* & \text{otherwise.} \end{cases}$$

These revisions result in the same play of G -outcomes *in the cycle*, and thus the same repeated game payoffs.

REFERENCES

- ABREU, D. (1988): "On the Theory of Infinitely Repeated Games with Discounting," *Econometrica*, 56, 383-396.
- (1986): "Extremal Equilibria of Oligopolistic Supergames," *Journal of Economic Theory*, 39, 191-225.
- ABREU, D., D. PEARCE, AND E. STACCHETTI (1986): "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," mimeo, Harvard University.
- AUMANN, R. J. (1981): "Survey of Repeated Games" in *Essays in Game Theory and Mathematical Economics in Honor of Oskar Morgenstern*. Mannheim: Bibliographisches Institut, 11-42.
- BEN-PORATH, E. (1986): "Repeated Games with Bounded Complexity," mimeo, Graduate School of Business, Stanford University.
- FUDENBERG, D., AND E. MASKIN (1986): "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econometrica*, 54, 533-554.
- GREEN, E. (1982): "Internal Costs and Equilibrium: The Case of Repeated Prisoners' Dilemma," mimeo.
- HOPCROFT, J. E., AND J. D. ULLMAN (1979): *Introduction to Automata Theory, Language and Computation*. New York: Addison-Wesley.
- KALAI, E., AND W. STANFORD (1988): "Finite Rationality and Interpersonal Complexity in Repeated Games," *Econometrica*, 56, 397-410.
- MEGIDDO, N., AND A. WIDGERSON (1986): "On Play by Means of Computing Machines," in *Reasoning About Knowledge*, ed. by Joseph Y. Halpern. Los Altos: Kaufmann Publishers, 259-274.
- NEYMAN, A. (1985): "Bounded Complexity Justifies Cooperation in the Finitely Repeated Prisoners' Dilemma," *Economics Letters*, 19, 227-229.
- RADNER, R. (1978): "Can Bounded Rationality Resolve the Prisoner's Dilemma?" mimeo.
- (1980): "Collusive Behavior in Noncooperative Epsilon-equilibria of Oligopolies with Long but Finite Lives," *Journal of Economic Theory*, 22, 136-154.
- RUBINSTEIN, A. (1977): "Equilibrium in Supergames," Research Memorandum #25, Center for Research in Mathematical Economics and Game Theory, Hebrew University.
- (1979): "Equilibrium in Supergames with the Overtaking Criterion," *Journal of Economic Theory*, 21, 1-9.
- (1986): "Finite Automata Play the Repeated Prisoner's Dilemma," *Journal of Economic Theory*, 39, 83-96.
- SELTEN, R. (1975): "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," *International Journal of Game Theory*, 4, 25-55.
- SMALE, S. (1980): "The Prisoner's Dilemma and Dynamical Systems Associated to Non-cooperative Games," *Econometrica*, 48, 1617-1634.