

Evolutionary Stability in Repeated Games Played by Finite Automata*

KENNETH G. BINMORE

*Department of Economics, University College of London, London, WC1E 6BT, England, and
Department of Economics, University of Michigan, Ann Arbor, Michigan 48109*

AND

LARRY SAMUELSON

Department of Economics, University of Wisconsin, Madison, Wisconsin 53706

Received December 5, 1990; revised November 20, 1991

We consider a game in which "meta-players" choose finite automata to play a repeated stage game. Meta-players' utilities are lexicographic, first increasing in the (limit-of-the-means) payoffs of the repeated game and second decreasing in the number of states in their automaton. We examine the outcomes in this game which satisfy a version of evolutionary stability that has been modified to permit existence. We find that such automata must be efficient, in that they must maximize the sum of the (limit-of-the-means) payoffs from the repeated game. *Journal of Economic Literature* Classification Numbers: C70, C72. © 1992 Academic Press, Inc.

1. PREVIEW

Although the results in this paper are applicable in general, attention will often be focused on the Prisoners' Dilemma in the form given on the left of Fig. 1. The shaded region in the diagram on the right of Fig. 1 is the set of payoff pairs that the "Folk Theorem" (Aumann [3]) shows to be achievable by equilibrium play in the infinitely repeated version of the Prisoners' Dilemma with "limit-of-the-means" payoffs (which we call "profits" in this paper).

Abreu and Rubinstein [1] study a model in which each of two "metaplayers" choose a finite automaton to play the infinitely repeated Prisoners' Dilemma on their behalf. In one specification of this model, the metaplayers have lexicographic preferences. They seek to maximize their

* Financial support from the NWO (Netherlands B46-276) is gratefully acknowledged, as is the hospitality and support of the Center for Economic Research at Tilburg University.




FIG. 1. The Prisoner's Dilemma

profit; but if two automata achieve the same profit, a metaplayer prefers whichever is less complex. Nash equilibrium outcomes in this automaton selection game are the rational points lying on the cross shown on the right of Fig. 1.

Abreu and Rubinstein interpret their result in terms of decision-makers of unbounded rationality who must delegate authority to subordinates who can only reliably execute simple decision rules. The interpretation to be explored in this paper is that the metaplayers are a metaphor for an evolutionary process. That is to say, the automata represent rules-of-thumb that have evolved during past plays of the (infinitely repeated) game. If metaplayers are to be seen as a metaphor for an evolutionary process, then it is natural to replace the notion of a Nash equilibrium by an appropriate version of the idea of an evolutionarily stable strategy. This paper examines the evolutionary viability of strategies in automaton selection games of the Abreu-Rubinstein variety.

With the notion of evolutionary viability that we employ, only (2, 2) in the diagram on the right of Fig. 1 survives as a possible equilibrium outcome of the infinitely repeated Prisoners' Dilemma. The same argument, applied to a general two-player, normal-form game¹ shows that only *utilitarian* outcomes are sustainable as evolutionary viable equilibria. However, we are anxious that the paper not be seen as a stylized defense of utilitarianism valid for all societies without qualification. Some of our assumptions need to be treated with great reserve for the reasons explained in Section 4.

¹ The extensions to n -player games are straightforward, but we do not discuss them.

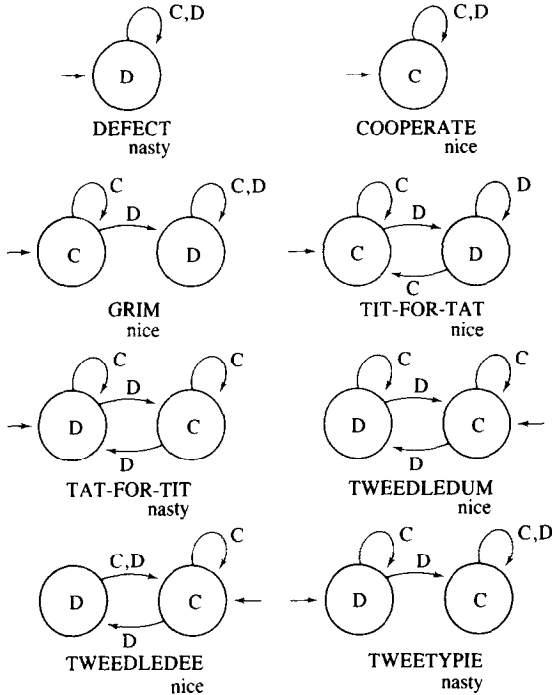


FIG. 2. Some automata

The idea behind our argument is of ancient vintage. Recent papers that use the idea include Binmore [8], Fudenberg and Maskin [13] and Robson [23]. An initially nonutilitarian population is vulnerable to invasion by mutants who recognize each other by means of what Robson [23] calls a “secret handshake.” This private signal allows the mutants to form an insider group who cooperate among themselves but treat outsiders as outsiders treat each other. As a result, insiders earn a higher average payoff than outsiders, and so the latter are displaced. Only utilitarian machines can be immune to such invasions and so they are the only possible candidates for evolutionary viability.

Axelrod and Hamilton [6] are perhaps the most notable of the contributors to what is now a large literature on the “evolution of cooperation.” Their evolutionary defense of TIT-FOR-TAT as a strategy for the infinitely repeated Prisoners’ Dilemma is much cited.

Figure 2 shows a representation of TIT-FOR-TAT as a finite automaton² along with a number of other strategies for the infinitely

² The particular type of automaton we use is called a Moore machine (Hopcraft and Ullman [14]). Appendix B discusses some of the formal issues, but the paper can be read without penetrating these mysteries.

repeated Prisoners' Dilemma.³ Each circle in a diagram of Fig. 2 indicates a possible *state* of the machine. The letter written inside the circle shows the action the machine will take in the Prisoners' Dilemma when in that state. The arrows indicate how machines make *transitions* from one state to another. For example, after TIT-FOR-TAT has just defected (*D*), it will remain in its *D* state if the opponent just played *D*, but will shift to its *C* state if the opponent just cooperated (*C*). The *initial state* of a machine is shown by an arrow with no source.

When only profits are considered, all the strategies represented by the automata of Fig. 2 are Nash equilibria against themselves except for COOPERATE and TWEETYPIC.⁴ In particular, as emphasized by Axelrod and Hamilton [6], TIT-FOR-TAT is a Nash equilibrium against itself. However, two TIT-FOR-TAT machines do *not* constitute a Nash equilibrium for the automaton selection game of Abreu and Rubinstein based on the Prisoners' Dilemma. COOPERATE is a better reply to TIT-FOR-TAT than TIT-FOR-TAT is to itself, because COOPERATE gets the same profit in the repeated Prisoners' Dilemma when playing TIT-FOR-TAT, but COOPERATE has the additional advantage of being less complex than TIT-FOR-TAT. In our evolutionary model, this implies that a population consisting entirely of TIT-FOR-TAT machines can be invaded by COOPERATE mutants.

In Axelrod's [5] terminology, a *nice* machine is one that is never the first to defect. TIT-FOR-TAT is therefore a "nice" machine. Axelrod [5] emphasizes the success of nice machines in his setting. However, in our model the same reasons that destabilize TIT-FOR-TAT also apply to any population consisting entirely of nice machines.⁵ A population therefore cannot be evolutionarily viable in our model unless it contains at least some machines that are *nasty* (i.e. not nice). To this extent, our work is at variance with Axelrod's.

The machine TAT-FOR-TIT of Fig. 2 is a *nasty* machine. It begins by defecting and treats a corresponding defection by the opponent as a secret handshake signaling a switch to its *C* state. In contrast to the case of two TIT-FOR-TAT machines, two TAT-FOR-TIT machines constitute a Nash equilibrium in the Abreu-Rubinstein automaton selection game. Moreover, two TAT-FOR-TIT machines achieve the utilitarian outcome (2, 2) and so do no worse than two TIT-FOR-TAT machines.

³ Appendix A contains a listing of *all* one and two state machines capable of playing the infinitely repeated Prisoners' Dilemma.

⁴ Appendix A contains a payoff table for *all* pairs of one and two state machines.

⁵ Any population of nice machines is vulnerable to invasion by COOPERATE mutants except a population consisting entirely of COOPERATE machines. But such a population is vulnerable to invasion by DEFECT mutants.

The fact that two TAT-FOR-TIT machines are a Nash equilibrium does not guarantee that a population of TAT-FOR-TITs is evolutionarily stable in the Abreu-Rubinstein automaton selection game. In fact, the game has *no* strategies that are evolutionarily stable according to the standard definition. This is not a problem that is special to Abreu and Rubinstein's model. Existence problem for evolutionarily stable strategies have been well-documented in a number of other variants of the infinitely repeated Prisoners' Dilemma (Boyd and Lorberbaum [9], Farrell and Ware [11], and Kim [16]). The nonexistence problem arises because the standard definition requires that *any* mutant bridgehead will be entirely eliminated. However, a mutant player whose behavior would differ from that of an existing strategy only off the equilibrium path will necessarily be indistinguishable from existing strategies on the equilibrium path, and hence cannot be driven out.

Using a modified definition of evolutionary stability that takes account of this difficulty, we show that a population consisting entirely of TAT-FOR-TIT machines is evolutionarily viable. However, the new definition makes it impossible to ignore the case of polymorphous populations in which different machines may co-exist in a symbiotic relationship. Extending our analysis to include such populations, we show in general that only utilitarian outcomes are evolutionarily viable.

The paper is organized in the following way. Sections 2 and 3 provide a discussion of the issues that motivate our work. Section 4 considers some of the modeling problems that our approach entails and explains how we deal with some of these problems and evade others. Sections 5 and 6 introduce the necessary formalities, and introduce our modified definition of an evolutionary stable strategy. Section 7 establishes that a population consisting of only one type of machine must be utilitarian if it is to be evolutionarily stable according to our modified definition. Section 8 seeks to explain why such a result needs to be generalized to the case of polymorphous populations of machines. Section 9 provides such a generalization.

2. EQUILIBRIUM SELECTION

When more than one equilibrium exists, the problem of selecting one from among them is not easy. But progress on this front is necessary if game theory is to break out of the beachhead it has established in the social sciences.

In seeking insight into the equilibrium selection problem, it is sensible to first look at simple examples. In such simple examples, the "right" equilibrium often seems glaringly obvious, and it may be easy to give lists

of plausible *ad hoc* reasons why the right equilibrium should be selected. However, such principles are notoriously unreliable when applied in general. The purpose of the enterprise is not to pluck a selection criterion from the air that happens to be intuitively satisfying in particular cases. It is to find selection criteria that are defensible *from first principles*. Our selection criteria, for example, choose the utilitarian outcome. One could give a hundred reasons drawn from political philosophy why the utilitarian outcome should be desirable, but none of these reasons are relevant to our concerns.⁶

Our approach to the equilibrium selection problem is based on studying “trembles,” as is the widely advertised “refinements of Nash equilibrium” literature. Indeed, this seems inevitable. In traditional analyses, players use their equilibrium strategies because of what would happen if they did not. It is therefore necessary to contemplate counterfactual hypotheses of the type: suppose a player who is sure to play his equilibrium strategy, doesn’t. Philosophers (e.g., D. Lewis [18]) discuss the meaning of such counterfactuals in terms of “possible worlds.” It is true that no player will deviate from his equilibrium strategy in *this* world, but there are *possible* worlds in which he would. Our task is to find the *closest* possible world to our own in which the deviation occurs. A tremble can then be thought of as being the change in our own world necessary to bring about the closest possible world in which the deviation occurs. Equilibrium selection proceeds by retaining only those equilibria that remain equilibria in the analyst’s favored closest possible world.

But the word “closest” has no *a priori* definition in this context. An equilibrium selection theory that employs this methodology must assign the word a meaning by choosing the nature of the trembles it takes as being relevant. It is natural that economists should be reluctant to abandon the paradigm of an economic agent as a “perfectly rational” optimizer. Traditionally, trembles are therefore imposed only on the rules of the game, the preference of the players, or the beliefs that the players hold. A strong bias exists in favor of the first of these. This is explicit in Selten’s [25] “trembling-hand perfect” equilibrium, which is obtained by adding appropriate chance moves to the original game. A similar attitude pervades the refinements literature, being implicit, for example, in the definitions of sequential or subgame-perfect equilibrium.⁷

⁶ Our aims are the same as the even more easily misunderstood work of Aumann and Sorin [4] or Anderlini [2] who are concerned with equilibrium selection in games of pure coordination in which one equilibrium Pareto-dominates the others. Nothing could be easier or less relevant than to solve the problem they set for themselves by inventing “collective rationality” principles like: reject any equilibrium Pareto-dominated by another.

⁷ Even the successive deletion of weakly dominated strategies requires a similar treatment of counterfactuals.

Unlike the refinements literature and the otherwise closely related work of Fudenberg and Maskin [13], we see the traditional types of tremble only as epiphenomena. Such trembles will indeed be significant in many cases, especially in evolutionary animal biology, but they are seen here as secondary in games played by humans. For this reason, we do not consider environmental trembles in this paper.⁸

The trembles this paper takes as primary are those in the thinking processes of the players. The idea is that, if an opponent plays irrationally, the explanation of first resort should be that he or she reasoned irrationally (and therefore may perhaps reason irrationally in the future). The trembles we wish to study are therefore *internal* to the players rather than *external*, as in traditional treatments.⁹ Such an approach requires modeling the thinking processes of the players explicitly.

How are the thinking processes of a player to be modeled? The avenue of investigation that we regard as most promising abandons the theory that people think deeply about their behavior when interacting with others in game-like situations. Instead, they are seen as hosts for “memes.” Dawkins [10] uses the term meme to include rules-of-thumb, social norms, conventions, or other more complex idea systems that a human being may use in translating a stimulus into a response. Evolution is seen as being responsible for a selection being made from the pool of possible memes. After evolution has operated, non-selected memes play a role in interpreting counterfactuals much like that played by trembles in traditional refinement theory. In brief, the non-selected memes serve as “explanations” for what would happen if selected memes were to deviate from equilibrium play.

3. BOUNDED RATIONALITY

The previous section explains why we seek to model the thinking processes of the players in a game explicitly. The device used for this purpose is the idea of a finite automaton as outlined in Section 1 and, more formally, in Appendix B. The exposition is often less clumsy if players are identified with the finite automaton that represents the strategy they are using. However, a more generally applicable paradigm sees the finite automaton as something like a virus that controls the strategy of a player

⁸ The common criticism of Axelrod's work, that a pair of TIT-FOR-TAT strategies is not subgame-perfect, is therefore not relevant to our concerns. Nor is Kalai and Neme's [15] similar criticism of the results of Abreu and Rubinstein [1].

⁹ With the notable exception of the “gang of four” paper [17] and related work.

who gets infected with it. The virus then spreads at a rate determined by how well the strategy performs relative to other strategies currently in use.

One virtue of working with finite automata is that they focus attention on the complexity of strategies and the possible cost of complexity. It is not, however, obvious how such costs are to be incorporated into a formal model. It is often taken for granted in the literature on “bounded rationality” that one should work with a fixed, exogenously determined upper bound on the complexity of the automata to be studied. We would prefer to characterize work in which this assumption is made, notably that of Neyman [22], as being concerned with *uniformly* bounded rationality. With such an assumption, it is easy to see that a pair of GRIM machines (Fig. 2) is a Nash equilibrium for an automaton selection game based on a 100 times repeated Prisoners’ Dilemma *provided* that machines with more than 99 states are excluded. This cooperation is achieved because a machine that plays better against GRIM than GRIM itself will necessarily cooperate in every round but the last, when it will defect. Such a machine must be able to count to 100 in order to know when the last round has arrived, but a machine with only 99 states can count only to 99.¹⁰

The difficulty with models incorporating an exogenously determined upper bound on the complexity of a machine is that cases can arise in which a highly complex machine may achieve only a relatively small payoff in equilibrium, although a much higher payoff might be achieved if the machine were replaced by another machine with just one additional state. Because of this difficulty, we follow Abreu and Rubinstein [1] in studying models in which the bound on complexity is *endogenously* determined.¹¹ Such an assumption always permits a machine to be displaced by another machine of slightly greater complexity if the resulting improvement in game payoffs is sufficiently large. (One may always “think a little harder” if the prospective benefits make it seem worthwhile.)

In measuring complexity, we simply count states. The complexity $|a|$ of a machine a is therefore how many states it has. All the machines of Figure 2 have complexity 2 except for DEFECT and COOPERATE which have complexity 1. A discussion of this complexity measure, which is arbitrary to a considerable degree, is postponed until the next section.

¹⁰ Neyman [22] shows that a clever construction of the implicit messages that the machines send each other through their behavior in the early stages of the game allows equilibrium payoffs very close to the utilitarian outcome to be achieved in a 100 times repeated Prisoners’ Dilemma, even when the uniform bound on the complexity of a machine is very much higher than 99.

¹¹ It is sometimes argued that such models are not properly models of “bounded rationality”. If this argument is thought persuasive, then one can refer to “imperfect rationality” or “costly rationality” to describe our approach.

4. MODELING ISSUES

In evolutionary game theory, a game G is seen as being played repeatedly. Each time it is played, Nature chooses its players from a population whose composition changes over time. The players do not think about how to play G . They are endowed with strategies by a process of mutation and selection that tends to eliminate strategies that are relatively less successful. In the literature, G is usually a normal-form game, but in our paper it is itself the repeated game G^∞ .

Ideally, one would evaluate the players' income streams in G^∞ using a discount factor p that represents the probability that the game will continue beyond any stage reached in G^∞ . Thus, although G^∞ would be modeled as an infinite game, it would end in finite time with probability one, so that no difficulty arises in constructing an evolutionary model in which the "infinitely repeated" game G^∞ is repeatedly played. Within this context, we are interested in cases in which p is relatively large, so that the future is relatively important. In order to construct a mathematically simple model, however, we first assign payoffs to income streams using the limit-of-the-means criterion rather than discounting. This can be viewed as focusing on the case in which p approaches unity while glossing over the problems that arise in going to the limit. Appendix C examines these limiting problems in the course of extending the results to the case of discounting. Whether discounting or limit-of-the-means payoffs are used, it is essential to our analysis that players be *patient*. This ensures that the mutants we construct can deliver their secret handshakes without suffering undue payoff losses.

The strategies that players use in G^∞ are modeled as finite automata. Our basic intuition is that interesting evolutionary processes will tend to select against complex machines unless their complexity generates real gains compared with less complex machines. We model this by introducing costs of complexity into players' utility functions.

The crudeness of such an approach is compounded by our measuring complexity by simply counting states. It is important to our analysis that states be costly, as this ensures that existing machines will not contain states that are not used in equilibrium but could be used to inflict harsh punishments on machines which offer secret handshakes. As Banks and Sundaram [7] have shown, however, other measures of complexity may lead to different conclusions. Moreover, we look at the extreme case in which costs of complexity are ranked lexicographically behind profits. We should stress that this allows us much freedom in constructing possible mutants when testing the stability of populations. Successful mutants may be vastly more complex than the population they invade provided only that they succeed in achieving a slightly higher profit.

We follow the standard practice of evading a study of the dynamics of the evolutionary process by appealing to the idea of an evolutionarily stable strategy. Our modification of this definition to suit our special circumstances, though essential to our results by allowing us to avoid the difficulties arising out of the nonexistence of an evolutionarily stable strategy, will probably not be controversial. However, it is important to stress that evolutionarily stable strategies are designed for use in a biological context. Mutations are then *rare*, so that after each mutant invasion, the system has time to attain a new equilibrium before the next mutant invasions. Is such a model appropriate in our context? We think that socioeconomic evolution would be better modeled by supposing that mutations are sufficiently frequent that the systems do *not* have time to adjust to the last mutation before the next appears. In particular, we plan a future paper in which there is a steady stream of mutations consisting of simplified versions of those machines that are currently present in the population.¹²

In pointing out this last possible inadequacy in our model, we are agreeing with Fudenberg's [12] criticism that our model is overly restrictive in the type of trembles that equilibrium strategies must confront. Note, however, that our remedy does not lie in appealing to Selten's trembling hand but in expanding the set of mutations to be considered.

5. FORMALITIES

Only two-player games are considered. The *underlying game* G is specified by a quadruple (S_1, S_2, π_1, π_2) in which S_1 and S_2 are strategy sets and the payoff functions are $\pi_i: S_1 \times S_2 \rightarrow \mathbb{R}$ ($i = 1, 2$). The Prisoners' Dilemma on the left of Fig. 1 will be the principal example.

A repeated game $G^\infty = \langle R_1, R_2, P_1, P_2 \rangle$ with the underlying game G as its stage-game, is constructed in the usual way. The payoff functions P_1 and P_2 are defined as limits of the means.¹³ (Discounting is discussed in Appendix C.) Thus

$$P_i(r_1, r_2) = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \pi_i(r_1(h_t), r_2(h_t)),$$

¹² We believe that such a model may lead to markedly different conclusions. Linster [19] provides some support for such a conjecture while pursuing Nachbar's [21] criticisms of Axelrod's "Olympiad" simulations. He finds that GRIM does exceedingly well, much better than TIT-FOR-TAT, in simulations of evolutionary competition between one and two state machines provided that the inflow of mutants is sufficiently rich and varied.

¹³ Since only strategies implementable by finite automata are considered, $P_i(r_1, r_2)$ will always be defined.

where h_t denotes the history of the game up to and including time t , and $r_i(h_t)$ is the action in S_i taken by player i at time $t + 1$. The number $P_i(r_1, r_2)$ will be called a *profit*.

An automaton selection game $G^\#$ is defined as in Abreu and Rubinstein [1]. The strategy space A_i consists of a set of finite automata (Moore machines) described formally in Appendix B. An automaton a in the set A_i is capable of occupying the role of player i in the repeated game G^∞ . No confusion will result from using a to denote the strategy implemented by the automaton a in G^∞ as well as the automaton a itself.

The automaton selection game $G^\# = \langle A_1, A_2, U_1, U_2 \rangle$ is played by two metaplayers. The payoff functions U_i reflect the fact that the metaplayers are assumed to care, not only about profits in G^∞ , but also about the complexity $|a|$ of the machine to which they delegate the duty of playing G^∞ . The precise assumption to be made attributes lexicographic preferences to the metaplayers. More precisely

$$U_1(a, c) > U_1(b, c) \Leftrightarrow \{P_1(a, c) > P_1(b, c)\} \quad \text{or} \\ \{P_1(a, c) = P_1(b, c) \quad \text{and} \quad |a| < |b|\},$$

with a similar requirement for U_2 . (Recall that $|a|$ is the number of states in the machine a .)¹⁴

6. EVOLUTIONARY STABILITY

In Abreu and Rubinstein [1] a metaplayer's choice of an automaton is interpreted as an act of delegating responsibility for the play of G^∞ to a simple-minded underling. In this section, the choices attributed to the metaplayers are assumed to emerge from an evolutionary process.

The case when G is symmetric will be studied first. Then $G^\#$ is also symmetric and we may write $A = A_1 = A_2$ and $U_1(a, b) = U(a, b) = U_2(b, a)$. For the moment, forget that U represents lexicographic preferences. If, instead, U measures fitness, then the standard criteria (Maynard Smith [20]) for a to be an *evolutionarily stable strategy* (ESS) are

- (I) $U(a, a) > U(b, a)$, or
- (II) $U(a, a) = U(b, a)$ and $U(a, b) > U(b, b)$,

for all possible mutants b . These conditions imply that a is a best response to itself, and hence (a, a) is a symmetric Nash equilibrium for $G^\#$. In addi-

¹⁴ We choose to represent the lexicographic preferences here with a utility function (which exists because the set of finite automata is countable) in order to ease comparison of our modified definition of evolutionary stability with the standard definition.

tion, the ESS conditions incorporate a stability requirement. For any alternative best response b to a , a must be a better response to b than b is itself. The ESS requirements are intended to capture the idea that a population of replicas of a should be invulnerable to invasion by sufficiently small groups of mutant b 's.

With lexicographic preferences, (I) and (II) need to be modified. The modified requirements for a to be an ESS are

- (i) $P(a, a) > P(b, a)$, or
- (ii) $P(a, a) = P(b, a)$ and $P(a, b) > P(b, b)$, or
- (iii) $P(a, a) = P(b, a)$ and $P(a, b) = P(b, b)$ and $|a| < |b|$,

for all possible mutants b . The principle is that profits are always counted before complexity.

Various authors have commented on the existence problem for an evolutionarily stable strategy in the repeated Prisoners' Dilemma, including Boyd and Lorberbaum [9], Kim [16], and Robson [23]. Existence is also a problem for the modified definition of an ESS given above. Since it is easy to see that COOPERATE and DEFECT can both be invaded, the following result shows that no automaton satisfies this definition of an ESS when G is the Prisoners' Dilemma.

LEMMA 6.1. *For any G , an automaton a that satisfies (i), (ii) or (iii) for all automata b has only one state.*

Proof. This is obvious if G has only one action. If G has two actions, consider an automaton a with more than one state that satisfies (i), (ii) and (iii). Suppose that a uses action x in its initial state. Now construct a mutant b that is identical to a except that, if its opponent plays something other than x when b is in its initial state, then b switches to a different state from that to which a would switch. Then $|a| = |b|$ and $P(a, a) = P(b, a) = P(a, b) = P(b, b)$. The reason for the latter set of equations is that a and b always use x in their initial state, and a and b are indistinguishable in play provided that their opponent uses x in its initial state. Thus the hypothesis that a has more than one state leads to the conclusion that none of (i), (ii), or (iii) hold. ■

Existence problems of this type seem to us to be an artificial construct arising from a definition of an ESS that is not entirely appropriate to the situation.¹⁵ The standard definition of an ESS demands that any sufficiently small invading group of mutants be eventually eradicated. One certainly

¹⁵ The existence problem does not arise because automata are used. Any pure strategy will always yield a host of out-of-equilibrium paths in a repeated game, and invading mutants can be constructed by altering behavior on these paths.

would wish to use a definition that precludes mutant invasions in which the original small mutant bridgehead expands at the expense of the original normal population. But what of mutant invasions after which the original bridgehead neither expands nor contracts? The original normal population and the mutant invaders will then survive together in a state of peaceful co-existence. In particular, the observed behavior of both natives and invaders may be precisely the same, their differing strategies leading to different behavior only under unrealized contingencies. As regards putative applications to human societies, who cares if different people would behave slightly differently under certain unrealized contingencies if, in fact, they all cooperate in sustaining a utilitarian outcome?

To accommodate this issue we introduce a further modification into the evolutionary viability requirement to be studied. A machine a will be called a *modified evolutionary stable strategy* (MESS) if, for all possible mutants b ,

- (1) $P(a, a) > P(b, a)$, or
- (2) $P(a, a) = P(b, a)$ and $P(a, b) > P(b, b)$, or
- (3) $P(a, a) = P(b, a)$ and $P(a, b) = P(b, b)$ and $|a| \leq |b|$.

As will become clear, one cannot dispense with the definition of a MESS simply by considering a mixed ESS or a polymorphous extension of the ESS idea.¹⁶ On the other hand, polymorphous MESSES will be important for the reasons given in Section 7. However, in this and the following section, the polymorphy issue is put to one side in an attempt to ease the exposition.

Complexity is important in the requirements for a MESS because its inclusion means that a MESS a can have no states that are not used when a plays G^* against itself.¹⁷ If an unused state existed, an automaton b could be constructed that is similar to a but which dispenses with the unused state so that $|b| < |a|$. The automaton b plays G^∞ exactly as a plays G^∞ provided that its opponent has always played G^∞ exactly as a plays G^∞ in the past. It follows that $P(a, a) = P(b, a) = P(a, b) = P(b, b)$, and hence Condition (3) for a MESS is violated. This is one of the easier arguments of Abreu and Rubinstein [1] adapted to our purpose. Appendix B (Lemma B.1) provides the formal details.

Note that it follows from the fact that all states must be used in equi-

¹⁶ A mixed ESS occurs when no sufficiently small group of mutants can invade a population whose members all play the same mixed strategy. A polymorphy occurs when no member of the population randomizes, but the effect is the same for an invading mutant because different members of the original population use different pure strategies.

¹⁷ Rubinstein's [24] original work on automata used a criterion that called for all states to be used *infinitely often*. We see no grounds for imposing such a constraint here.

librium that neither TIT-FOR-TAT nor GRIM can be a MESS (although this does not preclude their being constituents of a polymorphous MESS).

So far, only the case of a *symmetric* underlying game G has been considered. This section concludes by describing how asymmetric games may be addressed. In the symmetric case, automata do not “know” whether they are player 1 or player 2. Even if the payoff matrix for G is symmetric, this lack of information may be a handicap. For example, if the 3s in the version of the Prisoners’ Dilemma of Figure 1 are replaced by 6s, it ceases to be true that the sum of the payoffs is maximized at (C, C) . The sum is maximized instead at (C, D) or (D, C) (or some mixture of these).

If this problem arises or the payoff matrix of G is asymmetric, we will follow the standard practice in evolutionary game theory of introducing a symmetrized version of G^* to be denoted by G^{**} . The first event in G^{**} consists of the automata receiving a signal indicating whether they are to be player 1 or player 2, where each signal is received with probability $1/2$. An automaton a can then be seen as a pair (a_1, a_2) of simpler automata in which a_1 is used when a acts as player 1, and a_2 is used when a acts as player 2.

The profit for player 1 in the symmetrized game G^{**} is given by

$$\mathcal{P}(a, b) = \frac{1}{2}(P_1(a_1, b_2) + P_2(b_1, a_2)).$$

The complexity of the machine $a = (a_1, a_2)$ in the symmetrized game is taken to be $|a| = |a_1| + |a_2|$. These definitions allow a utility function V to be defined lexicographically by writing

$$\begin{aligned} V(a, c) > V(b, c) &\Leftrightarrow \{\mathcal{P}(a, c) > \mathcal{P}(b, c)\} \quad \text{or} \\ &\{\mathcal{P}(a, c) = \mathcal{P}(b, c) \quad \text{and} \quad |a| < |b|\}. \end{aligned}$$

One may then take G^{**} to be the quadruple

$$\langle A_1 \times A_2, A_1 \times A_2, V_1, V_2 \rangle,$$

in which V_1 and V_2 are defined by $V_1(a, b) = V(a, b) = V_2(b, a)$.

7. UTILITARIAN AUTOMATA

When a plays itself in G^{**} , the largest value its profit can be is

$$P(a^*, a^*) = \max_{(s_1, s_2) \in S_1 \times S_2} \frac{1}{2}(\pi_1(s_1, s_2) + \pi_2(s_1, s_2)). \quad (1)$$

An automaton a^* that achieves this maximum will be called a *utilitarian automaton* for G^{**} because it acts to maximize the sum of payoffs in the underlying game G .¹⁸

A utilitarian automaton a^* in G^* is defined to be one that satisfies (1), but with the maximization subject to the additional constraint that $s_1 = s_2$. This constraint reflects the fact that no mechanism is provided in G^* for breaking the symmetry of G . Of course, when G is the Prisoners' Dilemma of Fig. 1, the utilitarian profits for G^* and G^{**} are the same. Unless this is the case, the definition of a utilitarian automaton for G^* is of no great interest.

THEOREM 7.1. *An automaton a can be a MESS in G^{**} only if a is utilitarian.*

Proof. Let a^* be a utilitarian automaton and, for the purposes of contradiction, let a be an automaton which is a MESS but is not utilitarian. It will be shown that a population consisting entirely of replicas of a can be displaced by a mutant b . The mutant b will be more complex than a , but complexity counts after profit in the lexicographic preferences described by U .

The mutant b is constructed to have the following properties:

- Initially, b_i outputs a strategy for G that differs from a_i 's initial output.
- If the initial output of the opposing machine differs from a_{-i} 's initial output, then b_i "knows" that it is playing a replica of itself. It therefore continues by mimicking the utilitarian automaton a_i^* .
- If the initial output of the opposing machine is the same as a_{-i} 's initial output, then b_i "knows" that it is playing a . It therefore emits an output in the next period that "convinces" a that its opponent is a replica of itself, and then continues by mimicking a_i , and so achieves a profit of $P(a, a)$.

If such a machine b can be constructed, then it follows from the third property that $P(b, a) = P(a, a) = P(a, b)$. But $P(b, b) = P(a^*, a^*) > P(a, a)$. Thus none of the criteria (1), (2), or (3) for a MESS are satisfied and the theorem follows.

It remains to be confirmed that the third property for the mutant b can be satisfied, and so $P(b, a) = P(a, a) = P(a, b)$. This is not entirely obvious. For example, if a were the GRIM strategy of Fig. 2 and G were the Prisoners' Dilemma of Fig. 1, then b 's playing differently from a at the first

¹⁸ But one must recall that P is a "limit of means." Thus a^* need not maximize the sum of payoffs in every stage game.

stage of G^∞ would push a into its punishment phase for ever. Thus $P(b, a) = 0 < 2 = P(a, a)$. Note, however, that the GRIM strategy is not a MESS because it has a state that is unused in equilibrium. This observation is the key in establishing that $P(b, a) = P(a, a) = P(a, b)$.

Suppose that b_i 's initial action switches a_{-i} from its initial state q_0 to another state q . Since all of a_{-i} 's states are used in equilibrium, a_{-i} would eventually reach state q when playing a_i . At that time, let a_i be in state q^* generating output r . The automaton b should therefore be constructed so that, if the opponent's initial output is the same as a_{-i} 's, then b switches to a state which we will label q^* and emits the output r . After that, b can mimic a_i without difficulty. Appendix B (Lemma B.2) provides the details. A machine designed in this way satisfies the third of the three properties attributed to b . In particular, $P(b, a) = P(a, a) = P(a, b)$. ■

Theorem 7.1 provides a necessary condition for an automaton to be a MESS in G^{**} . The same necessary condition, of course, holds for G^* .

Sufficient conditions for the existence of a MESS are easily obtained using Folk Theorem arguments. We again quote only the result for G^{**} .

THEOREM 7.2. *Let \bar{m} denote the pure strategy minimax point of the underlying game G . Let u be its utilitarian outcome. Then a sufficient condition for the existence of a MESS in G^{**} is that $u \geq \bar{m}$.*

Note that the sufficient condition of Theorem 7.2 need not be satisfied. Its failure will sometimes be a consequence of the exclusion of mixed strategies from the analysis. Our results can be extended to machines that output mixed strategies for G rather than just pure strategies.¹⁹ One can then replace \bar{m} in Theorem 7.2 by its mixed strategy equivalent $\bar{v} \leq \bar{m}$. However, there remains no guarantee that $u \geq \bar{v}$. If this inequality fails, a MESS fails to exist.

8. POLYMORPHOUS POPULATIONS

The definition of an evolutionarily stable strategy is designed for the case when mutations come at sufficiently infrequent intervals that the system can fully recover from the effects of repelling one mutation before the next appears. The definition of a MESS is inconsistent with such a story because it allows for the possibility that certain mutations may not be repelled at all. After an invasion by such a mutant, the system will therefore not revert to its original state. This section illustrates the significance of this feature of the definition.

¹⁹ An alternative would be to allow the machines access to random *inputs*, but to continue to insist on pure outputs. This approach creates difficulties for our analysis.

Appendix A summarizes the relevant properties of all one and two state machines that are capable of playing the infinitely repeated Prisoners' Dilemma. Observe that a population consisting entirely of TIT-FOR-TAT machines will be displaced by an invasion of mutant COOPERATE machines. After the invasion, all machines will be getting the utilitarian profit, but COOPERATE has an advantage in that it is less complex. With the exception of a population consisting only of copies of the COOPERATE machine, the same goes for any population of nice machines.

This argument shows that no nice machine can be a MESS.²⁰ A machine can be an optimal reply to itself in our framework only if it exhibits some nasty behavior in equilibrium. The simplest such machine is DEFECT, but DEFECT is not a MESS because an invasion by TIT-FOR-TAT will expand from its initial bridgehead.

The simplest machine that is a MESS is TAT-FOR-TIT (Fig. 2). But the fact that no invading mutant can *expand* at the expense of TAT-FOR-TIT does not imply that TAT-FOR-TIT cannot be *invaded*. In the notation of Appendix A, TAT-FOR-TIT can be invaded by "cc", "cd", "AA", "AC", or "CA". (Small letters indicate machines that begin by cooperating; large letters indicate machines that begin by defecting.) A tiny invading group of one of these machines will not expand, but neither will it be driven to extinction.

The inability of a MESS to repel all invaders implies that the purity of a MESS in which the whole population consists of copies of a single machine will eventually be sullied. This makes it necessary to study polymorphous populations. The next section offers a definition of a polymorphous MESS, and proves an appropriate generalization of Theorem 7.1. This section continues with a simple example that is intended to indicate some of the complexities that this analysis evades.

In a population consisting of positive fractions of TAT-FOR-TIT, TWEEDLEDUM (cc), and TWEEDLEDEE (cd) all machines will always earn the utilitarian profit. The population cannot be invaded by any one or two state machine except TWEETYPIC (CA). Nor can any machine of complexity greater than two invade. However, after an attempted invasion has been repelled, the proportions of the three types of machine may be altered. For example, an attempted invasion by "AC" will increase the relative numbers of TWEEDLEDUM and an attempted invasion by "AA" will reduce the relative numbers of TWEEDLEDEE. In the long run, the proportions of the different types in the population will therefore be subject to drift in the face of mutations.

²⁰ This argument also shows that, with the definition of a polymorphous MESS given in the next section, no population consisting entirely of "nice" machines can be a MESS.

Not only this, none of the three machines are equipped to repel TWEETYPiE. TWEETYPiE can therefore invade the population. After the invasion, all machines will still always earn the utilitarian profit. However, once TWEETYPiE gets established, its toothlessness allows suitable machines (of complexity greater than two) to gain a temporary foothold. Such machines will get more than the utilitarian payoff when playing TWEETYPiE who will consequently get less than the utilitarian payoff. Eventually, the invading machines will lose their edge when their exploitation of TWEETYPiE results in its being eliminated. They will then, in turn, be eliminated themselves because their greater complexity no longer generates any extra advantage. This restores a population of TAT-FOR-TIT, TWEEDLEDUM, and TWEEDLEDEE in which all machines earn the utilitarian payoff.

A population consisting of suitable fractions of TAT-FOR-TIT, TWEEDLEDUM, TWEEDLEDEE, and transient numbers of TWEETYPiE can therefore co-exist in a relationship whose stability is somewhat precarious. The proportions in which they are present will drift depending on the shocks the system receives as different mutants appear. Although the symbiotic relationship may persist for long periods of time, it will be stressed to the point of collapse if a sufficiently adverse sequence of mutations is encountered.

Perhaps more satisfactory stability situations exist at higher complexity levels. At the two state level, little more can be said without information about how mutations arise.

9. POLYMORPHOUS MESSAGES

It remains to confirm that the conclusion of Theorem 7.1 is valid in the presence of polymorphy. To this end, let F denote a population in which the automaton a^n occurs with frequency $f_n > 0$ ($n = 1, 2, \dots, N$). Define

$$P(a, F) = \sum_{n=1}^N f_n P(a, a^n).$$

A population F will be said to be a *polymorphous* MESS if, for all possible mutants b and all a^n in F ,

- (1) $P(a^n, F) > P(b, F)$, or
- (2) $P(a^n, F) = P(b, F)$ and $P(a^n, b) > P(b, b)$, or
- (3) $P(a^n, F) = P(b, F)$ and $P(a^n, b) = P(b, b)$ and $|a^n| \leq |b|$.

The conditions for a polymorphous MESS ensure that each machine in the population does at least as well as any potential invader, provided the

latter appears in sufficiently small numbers. Some examples for the case when the underlying game G is the Prisoners' Dilemma of Fig. 1 may help to clarify the definition. We have seen that a population consisting entirely of TIT-FOR-TAT is not a MESS. However, a population consisting entirely of TAT-FOR-TIT is a MESS. So is any population consisting of specimens of TAT-FOR-TIT, TWEEDLEDUM, and TWEEDLEDEE. But the addition of small numbers of TWEETYPIC generates a population that is not a MESS.

Since any member of the population can itself appear as a mutant b , a polymorphous MESS must have the property that $P(a^m, F) = P(a^n, F)$ for all a^m and a^n in F . It follows from part (2) of the definition of a polymorphous MESS that, for all a^m and a^n in F ,

$$P(a^m, a^n) \geq P(a^n, a^n). \tag{2}$$

We define

$$z = \min_n P(a^n, a^n) \equiv P(a^k, a^k). \tag{3}$$

It follows from parts (2) and (3) of the definition of a polymorphous MESS that, for all a^m and a^n in F ,

$$P(a^m, a^n) \geq z. \tag{4}$$

A population F will be said to be *utilitarian* if, for all a^m and a^n in F , $P(a^m, a^n) \geq P(a^*, a^*) = u$, where a^* is a utilitarian automaton.²¹

THEOREM 9.1. *If the underlying game G is symmetric, then a population F can be a polymorphous MESS for G^* only if F is utilitarian.*

Proof. Let F be a population that is a polymorphous MESS. The theorem will follow from (4) if it can be shown that $z \geq u$.

The first step is to prove that, for any a^m in F ,

$$P(a^m, a^k) \geq u \quad \text{or} \quad (P(a^k, a^m), P(a^m, a^k)) = (z, z), \tag{5}$$

where a^k satisfies (3).

Suppose that (5) fails. Then there exists a^j in F for which

$$P(a^j, a^k) < u \quad \text{and} \quad (P(a^k, a^j), P(a^j, a^k)) \neq (z, z). \tag{6}$$

²¹ If the automaton selection game is G^{**} , then it is also true that $P(a^m, a^n) \leq P(a^*, a^*)$ for all a^m and a^n in F . This also holds when G is symmetric and the automaton selection game is G^* provided that, as in the Prisoners' Dilemma of Figure 1, the utilitarian profit in G^* coincides with that of G^{**} . If this proviso does not hold, the definition of a utilitarian population is not very interesting for G^* .

It will be shown that (6) allows the construction of a mutant b that does better than a^j . The mutant plays as follows:

- Initially, b mimics a^k , and continues to do so unless a cycle is established that generates a profit pair of (z, z) . (Because each automaton in F is finite, b can be programmed to detect the occurrence of one of the finitely many cycles generated by the automata in F which yields the profit pair (z, z)).

- If such a cycle is established, b sends an identifying signal. If this is reciprocated, b "knows" it is playing itself and thereafter mimics a^* .

- If the signal is not reciprocated, b experiments until it learns the identity of the automaton a^n in F that it is playing.

- This will leave a^n in some state q . For familiar reasons, a^n must use q in playing against at least one automaton a^m in F . Thus b may now mimic a^m thereafter, and hence, by (4), secure a profit of at least z .

To see that the mutant b so constructed does better than a^j , first observe that, for all a^n in F ,²²

$$P(b, a^n) \geq P(a^k, a^n). \quad (7)$$

It follows that $P(b, F) \geq P(a^k, F) = P(a^j, F)$. The first of the conditions for a polymorphous MESS is therefore not satisfied. The second and third conditions are not satisfied either because $P(b, b) = P(a^*, a^*) = u$, and $P(a^j, b) = P(a^j, a^k) < u$, by (6). Thus, $P(a^j, b) < P(b, b)$. It follows that (5) must hold if F is a MESS. It remains to complete the proof by showing that $z \geq u$ follows from (5) and the fact that F is a MESS. A mutant b is constructed exactly as in the first part of the proof, but we now examine how this mutant performs in comparison to a^k rather than a^j . As before, the mutant satisfies (7) for all a^n in F .

When playing a^k , the mutant b eventually behaves like some machine a^m in F . Suppose that $P(b, a^k) = P(a^m, a^k) \geq u$. If $u > z = P(a^k, a^k)$, then (7) holds with *strict* inequality when $n = k$. It follows that $P(b, F) > P(a^k, F)$. This contradicts the assumption that F is a MESS. Thus, either $z \geq u$ as we are trying to prove, or else $P(a^m, a^k) < u$.

We therefore consider the case when $P(a^m, a^k) < u$. Then (5) implies that $P(a^k, b) = P(a^k, a^m) = z$. Since $P(b, F) \geq P(a^k, F)$, the second and third conditions for a polymorphous MESS yield that $P(a^k, b) \geq P(b, b) = u$. On combining these results, we therefore obtain again that $z \geq u$. ■

²² To establish (7), notice that if a^m is the automaton that b mimics against a^n in the last part of the description of b , then $P(b, a^n) = P(a^m, a^n) \geq P(a^k, a^n)$, since otherwise $P(a^k, a^n) > P(a^m, a^n) \geq z$ (from (2)–(3)), precluding the ability of b and a^n to yield a cycle of profits (z, z) and yielding a contradiction.

THEOREM 9.2. *A population F can be a polymorphous MESS for $G^{*#}$ only if F is utilitarian.*

Proof. We note only that, in adapting the proof of Theorem 9.1, one need not construct a mutant that signals its identity under certain contingencies both when in the role of player 1 and in the role of player 2. It is enough to do better than a machine in the current population in one of the two roles. ■

10. CONCLUSION

We have examined circumstances under which the only evolutionarily viable outcome in an infinitely repeated game is utilitarian. The argument depends on the players being boundedly rational, but no uniform bound is imposed on the complexity of their thinking processes.

Our notion of evolutionary viability, captured by the *modified evolutionarily stable strategy* (MESS) permits players operating different strategies to co-exist in a symbiotic relationship that sustains the utilitarian outcome against any *single* potential disrupting mutant invasion. However, questions remain about the viability of such populations in the presence of *repeated* invasion attempts. We feel that the examination of more stringent stability requirements than those studied in this paper is an important area for further research. In particular, it is not clear that one should follow the biologists in supposing that the interval between successive invasion attempts is sufficiently large that a nonviable mutant will be entirely eliminated before the appearance of a new mutant. Overlapping invasions would seem more suited to a social context.

APPENDIX A

This appendix contains a listing of all 26 one and two state Moore machines capable of playing the infinitely repeated Prisoners' Dilemma. The entry in row r and column c in the Figs. 3 and 4 give the profit (the limit-of-the-means payoff) that machine r gets when it plays machine c . Machines that cooperate in their initial state are labeled with lower case letters. Machines that defect in their initial state are labeled with upper case letters.

APPENDIX B

This appendix discusses some formal issues in the theory of finite automata that are neglected in the text. For more details, see Hopcraft and

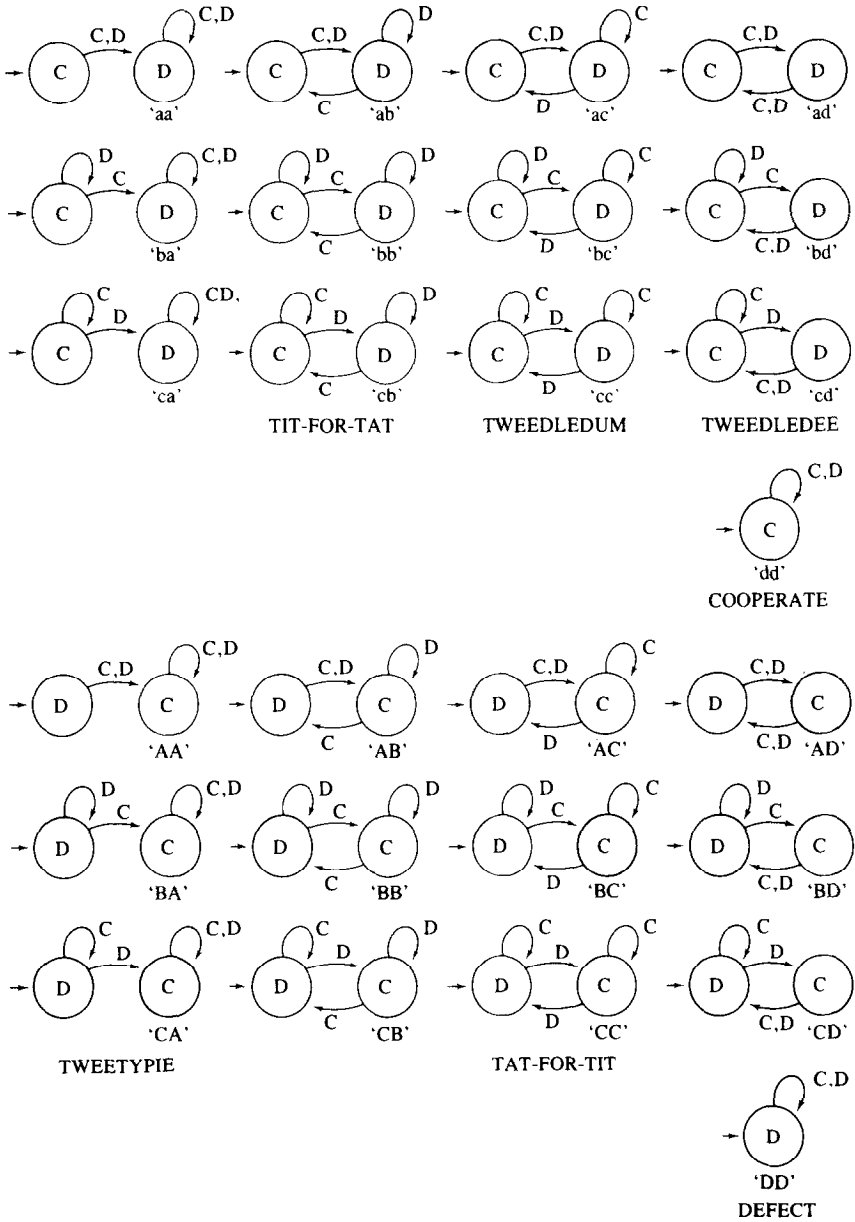


FIG. 3. All one and two state machines

	aa	ab	ac	ad	ba	bb	bc	bd	ca	cb	cc	cd	dd
aa	0	0	$\frac{2}{3}$	$\frac{2}{3}$	0	0	3	3	0	0	$\frac{2}{3}$	$\frac{2}{3}$	3
ab	0	0	$\frac{2}{3}$	1	0	0	$\frac{5}{3}$	$\frac{5}{3}$	0	1	$\frac{2}{3}$	1	$\frac{5}{3}$
ac	$-\frac{1}{2}$	$\frac{2}{3}$	1	1	$-\frac{1}{2}$	3	1	1	$-\frac{1}{2}$	$\frac{2}{3}$	$\frac{5}{3}$	$\frac{5}{3}$	3
ad	$-\frac{1}{2}$	1	1	1	$-\frac{1}{2}$	1	1	1	$-\frac{1}{2}$	1	1	1	$\frac{5}{3}$
ba	0	0	$\frac{2}{3}$	$\frac{2}{3}$	0	0	3	3	0	0	$\frac{2}{3}$	$\frac{2}{3}$	3
bb	0	0	-1	1	0	0	$\frac{5}{3}$	$\frac{5}{3}$	-1	$\frac{4}{3}$	-1	$\frac{4}{3}$	$\frac{5}{3}$
bc	-1	$\frac{1}{3}$	1	1	-1	$\frac{1}{3}$	1	1	-1	1	$\frac{2}{3}$	$\frac{2}{3}$	3
bd	-1	$\frac{1}{3}$	1	1	-1	$\frac{1}{3}$	1	1	-1	$\frac{4}{3}$	-1	$\frac{4}{3}$	$\frac{5}{3}$
ca	0	0	$\frac{2}{3}$	$\frac{2}{3}$	0	0	3	3	2	2	2	2	2
cb	0	0	$\frac{2}{3}$	1	0	$\frac{4}{3}$	1	$\frac{4}{3}$	2	2	2	2	2
cc	$-\frac{1}{2}$	$\frac{2}{3}$	1	1	$-\frac{1}{2}$	3	1	3	2	2	2	2	2
cd	$-\frac{1}{2}$	1	1	1	$-\frac{1}{2}$	$\frac{4}{3}$	1	$\frac{4}{3}$	2	2	2	2	2
dd	-1	$\frac{1}{2}$	-1	$\frac{1}{2}$	-1	$\frac{1}{2}$	-1	$\frac{1}{2}$	2	2	2	2	2

	AA	AB	AC	AD	BA	BB	BC	BD	CA	CB	CC	CD	DD
aa	3	3	$\frac{2}{3}$	$\frac{2}{3}$	3	3	0	0	3	3	$\frac{2}{3}$	$\frac{2}{3}$	0
ab	$\frac{5}{3}$	$\frac{5}{3}$	1	1	$\frac{5}{3}$	0	1	1	$\frac{5}{3}$	$\frac{5}{3}$	$\frac{2}{3}$	$\frac{2}{3}$	0
ac	3	3	$\frac{5}{3}$	1	3	3	$\frac{2}{3}$	$\frac{2}{3}$	3	1	$\frac{5}{3}$	1	$-\frac{1}{2}$
ad	$\frac{5}{3}$	1	1	1	$\frac{5}{3}$	1	1	1	$\frac{5}{3}$	1	1	1	$-\frac{1}{2}$
ba	3	3	$\frac{2}{3}$	$\frac{2}{3}$	3	0	0	0	-1	-1	-1	-1	-1
bb	$\frac{5}{3}$	$\frac{5}{3}$	$\frac{4}{3}$	1	$\frac{5}{3}$	0	$\frac{4}{3}$	0	-1	-1	-1	-1	-1
bc	3	1	$\frac{4}{3}$	1	3	$\frac{1}{3}$	1	$\frac{1}{3}$	-1	-1	-1	-1	-1
bd	$\frac{5}{3}$	1	$\frac{4}{3}$	1	$\frac{5}{3}$	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{3}$	-1	-1	-1	-1	-1
ca	3	3	$\frac{2}{3}$	$\frac{2}{3}$	3	3	0	0	3	3	$\frac{2}{3}$	$\frac{2}{3}$	0
cb	2	$\frac{4}{3}$	1	1	2	$\frac{4}{3}$	1	1	2	1	$\frac{4}{3}$	$\frac{4}{3}$	0
cc	3	3	2	1	3	3	$\frac{2}{3}$	$\frac{2}{3}$	2	1	2	2	$-\frac{1}{2}$
cd	2	$\frac{4}{3}$	1	1	2	$\frac{4}{3}$	1	1	2	$\frac{4}{3}$	2	1	$-\frac{1}{2}$
dd	2	$\frac{1}{2}$	2	$\frac{1}{2}$	2	$\frac{1}{2}$	2	$\frac{1}{2}$	-1	-1	-1	-1	-1

	AA	AB	AC	AD	BA	BB	BC	BD	CA	CB	CC	CD	DD
AA	-1	$\frac{1}{2}$	-1	$\frac{1}{2}$	-1	$\frac{1}{2}$	-1	$\frac{1}{2}$	-1	2	-1	2	2
AB	-1	$\frac{1}{2}$	-1	$\frac{1}{2}$	-1	$\frac{1}{2}$	-1	$\frac{1}{2}$	-1	$\frac{4}{3}$	-1	$\frac{4}{3}$	$\frac{5}{3}$
AC	$-\frac{1}{2}$	1	$\frac{1}{3}$	1	$-\frac{1}{2}$	$\frac{4}{3}$	$\frac{4}{3}$	1	2	1	2	1	2
AD	$-\frac{1}{2}$	1	1	1	$-\frac{1}{2}$	1	1	1	$-\frac{1}{2}$	1	1	1	$\frac{5}{3}$
BA	-1	$\frac{1}{2}$	-1	$\frac{1}{2}$	-1	$\frac{1}{2}$	-1	$\frac{1}{2}$	-1	2	-1	2	2
BB	-1	0	-1	1	0	$\frac{5}{3}$	$\frac{5}{3}$	-1	$\frac{4}{3}$	-1	$\frac{4}{3}$	$\frac{5}{3}$	$\frac{5}{3}$
BC	0	1	$\frac{2}{3}$	1	0	$\frac{4}{3}$	1	$\frac{4}{3}$	0	1	$\frac{2}{3}$	1	2
BD	0	1	$\frac{2}{3}$	1	0	$\frac{2}{3}$	$\frac{2}{3}$	0	1	$\frac{2}{3}$	1	$\frac{2}{3}$	$\frac{2}{3}$
CA	-1	$-\frac{1}{3}$	-1	$-\frac{1}{3}$	3	3	3	-1	2	2	2	3	3
CB	-1	$-\frac{1}{3}$	1	1	3	3	3	-1	1	$\frac{5}{3}$	$\frac{5}{3}$	3	3
CC	$-\frac{1}{2}$	$\frac{2}{3}$	$\frac{1}{3}$	1	3	3	3	$-\frac{1}{2}$	$\frac{4}{3}$	2	2	3	3
CD	$-\frac{1}{2}$	$\frac{2}{3}$	1	1	3	3	3	$-\frac{1}{2}$	$\frac{4}{3}$	$\frac{5}{3}$	$\frac{5}{3}$	3	3
DD	0	0	$\frac{2}{3}$	$\frac{2}{3}$	3	3	3	0	0	$\frac{2}{3}$	$\frac{2}{3}$	$\frac{2}{3}$	3

	AA	AB	AC	AD	BA	BB	BC	BD	CA	CB	CC	CD	DD
AA	2	$\frac{1}{2}$	2	$\frac{1}{2}$	2	$\frac{1}{2}$	2	$\frac{1}{2}$	2	-1	2	-1	-1
AB	$\frac{5}{3}$	1	$\frac{4}{3}$	1	$\frac{5}{3}$	$\frac{1}{3}$	$\frac{4}{3}$	$\frac{1}{3}$	$\frac{5}{3}$	1	-1	1	-1
AC	2	$\frac{4}{3}$	2	1	2	$\frac{4}{3}$	1	1	2	$\frac{1}{3}$	2	$\frac{1}{3}$	$-\frac{1}{2}$
AD	$\frac{5}{3}$	1	1	1	$\frac{5}{3}$	1	1	1	$\frac{5}{3}$	1	1	1	$-\frac{1}{2}$
BA	2	0	2	$\frac{1}{2}$	0	0	0	0	2	-1	-1	-1	0
BB	$\frac{5}{3}$	$\frac{5}{3}$	$\frac{4}{3}$	1	0	0	0	0	$\frac{5}{3}$	$\frac{5}{3}$	-1	-1	0
BC	2	$\frac{4}{3}$	1	1	0	0	0	0	2	1	$\frac{4}{3}$	$\frac{4}{3}$	0
BD	$\frac{5}{3}$	$\frac{5}{3}$	1	1	0	0	0	0	$\frac{5}{3}$	$\frac{5}{3}$	$\frac{4}{3}$	$\frac{4}{3}$	0
CA	2	$-\frac{1}{3}$	2	$-\frac{1}{3}$	2	$-\frac{1}{3}$	2	$-\frac{1}{3}$	2	-1	2	-1	-1
CB	3	1	$\frac{5}{3}$	1	3	$-\frac{1}{3}$	1	$-\frac{1}{3}$	3	1	1	1	-1
CC	2	3	2	1	3	3	$\frac{2}{3}$	$\frac{2}{3}$	2	$\frac{1}{3}$	2	$\frac{1}{3}$	$-\frac{1}{2}$
CD	3	1	$\frac{5}{3}$	1	3	3	$\frac{2}{3}$	$\frac{2}{3}$	3	1	$\frac{1}{3}$	1	$-\frac{1}{2}$
DD	3	3	$\frac{2}{3}$	$\frac{2}{3}$	0	0	0	0	3	3	$\frac{2}{3}$	$\frac{2}{3}$	0

FIG. 4. Payoff for all one and two state machines

Ullman [14]. We work with a particular type of finite automaton called a Moore machine. A Moore machine capable of serving as player i in the repeated game G^∞ has the set S_i of player i 's actions in the underlying game G as its output (or response) set. The set S_{-i} of actions jointly available to the other players in G is its input (or stimulus) set.

Formally, a Moore machine with these input and output sets is a quadruple $\langle Q, q_0, t, u \rangle$. The set Q is finite. An element $q \in Q$ is interpreted as a *state* of the machine. The *output function* $t: Q \rightarrow S_i$ describes the action $t(q)$ used by the automaton when it is in state q . The first time that the machine plays G , it is in its *initial state* q_0 . The *transition function* $u: Q \times S_{-i} \rightarrow Q$ describes how other states are reached. If the machine is in state q , it will play $t(q)$ in the current repetition of G . The other players will use some joint action s_{-i} . This shifts the automaton to state $u(q, s_{-i})$, where it will remain until after the next repetition of G .

For simplicity, the following results are confined to the case when G is symmetric.

LEMMA B.1. *If an automaton a is a MESS for G^* , then every state of a is used when a plays G^* against itself.*

Proof. Suppose that $a = \langle Q, q_0, t, u \rangle$ has a state q^* that is not reached when it plays itself. A machine $b = \langle Q', q'_0, t', u' \rangle$ is then constructed as follows. Take $Q' = Q \setminus \{q^*\}$ and $q'_0 = q_0$ ($\neq q^*$). Define $u': Q' \times S_{-i} \rightarrow Q'$ so that $u'(q, s_{-i}) = u(q, s_{-i})$, unless $u(q, s_{-i}) = q^*$. In the latter case, take $u'(q, s_{-i}) = q$. The game G^∞ will then be played the same whether the players are two a machines, two b machines, or an a machine and a b machine. Thus, $P(a, a) = P(b, a) = P(a, b) = P(b, b)$. All the conditions for a to be a MESS then fail because $|b| < |a|$. ■

LEMMA B.2. *If an automaton a is a MESS for G^* and u is the utilitarian payoff, then there exists a machine b such that $P(b, b) = u$ and $P(b, a) = P(a, a)$.*

Proof. Suppose that $a = \langle Q, q_0, t, u \rangle$. Let $t(q_0) = x$. The lemma is immediate unless G admits a second action y . Let $u(q_0, y) = q^*$. Take $\tilde{Q} = Q \cup \{r, s\}$, and define $b = \langle \tilde{Q}, r, \tilde{t}, \tilde{u} \rangle$ so that $\tilde{t}(r) = y$ and $\tilde{t}(s) = z$, where z is the utilitarian action in G . Define $\tilde{t}(q) = t(q)$, for all $q \in Q$. Define $\tilde{u}(r, y) = s$ and $\tilde{u}(r, w) = q^*$ for all $w \neq y$. Define $\tilde{u}(s, w) = z$ for all w . Define $\tilde{u}(q, w) = u(q, w)$ for all w and all $q \in Q$.

Then $P(b, b) = u$ because, after observing the opponent play y , the automaton b switches from state r to state s , where it remains thereafter using action z . Why is $P(b, a) = P(a, a)$? After b plays y and a plays x in their respective initial states, both automata then switch to state q^* and

play like a thereafter. Since all of a 's states are used when a plays itself (by Lemma B.1), it follows that $P(b, a) = P(a, a)$. ■

The formalities for the corresponding results for polymorphous MESSes and for asymmetric underlying games G are similar but have longer proofs.

APPENDIX C

The main body of the paper looks only at limit-of-the-means profits in the repeated game G^∞ . This appendix indicates some extensions to the discounting case. Only pure populations and a symmetric underlying game G will be considered.

If an automaton a plays G^∞ with a second automaton b , denote the stage-game payoff to the former at time t by $\pi^t(a, b)$. Denote the "continuation payoff" at the T th stage of G^∞ by

$$P_\delta^T(a, b) = (1 - \delta) \sum_{t=T}^{\infty} \delta^{t-T} \pi^t(a, b),$$

where δ is a discount factor satisfying $0 < \delta < 1$. The function $P_\delta(a, b) = P_\delta^0(a, b)$ then replaces the profit function of the text. Complexity considerations remain unchanged.

More liberties will be taken with the definition of evolutionary stability. If $0 < \delta < 1$ and $\varepsilon > 0$, a population will be said to be (δ, ε) -viable if a mutant bridgehead consisting of a fraction ε of the population cannot expand when the discount factor is δ . An automaton a will be said to be a GUESS if, whenever $\varepsilon = O(1 - \delta)$ as $\delta \rightarrow 1$, a population consisting entirely of specimens of a is (δ, ε) -viable for all δ sufficiently close to 1.

The following theorem is a limiting result. Away from the limit, one can not guarantee that machines are utilitarian. One can only ask that they be approximately utilitarian.

THEOREM C.1. *If a is a GUESS, then*

$$P_\delta(a, a) \rightarrow u \text{ as } \delta \rightarrow 1,$$

where u is the utilitarian payoff.

Proof. If a is a GUESS, then it has no unused states for familiar reasons. Let $t = T$ be the first stage at which $P_\delta^t(a, a)$ achieves its minimum. Assume for the moment that $T > 0$. Construct a mutant b with the following properties:

- Until time $T - 1$, the mutant b mimics a .

- At time $T-1$, the mutant b outputs something different from a .
- If the signal is reciprocated, b continues by mimicking the utilitarian machine from period T onwards.
- If the signal is not reciprocated, b returns to mimicking a from T onward, but starting from the state to which a was switched by b 's deviation at time $T-1$.

The expected continuation profit from time $T-1$ to the a automaton after an invasion by a fraction ε of b mutants is

$$(1-\delta)\{(1-\varepsilon)\pi^{T-1}(a, a) + \varepsilon\pi^{T-1}(a, b)\} \\ + \delta\{(1-\varepsilon)P_\delta^T(a, a) + \varepsilon P_\delta^T(a, b)\}. \quad (8)$$

The expected continuation profit from time $T-1$ to the b automaton after the invasion is

$$(1-\delta)\{(1-\varepsilon)\pi^{T-1}(b, a) + \varepsilon\pi^{T-1}(b, b)\} \\ + \delta\{(1-\varepsilon)P_\delta^T(b, a) + \varepsilon P_\delta^T(b, b)\}. \quad (9)$$

Both a and b get the same profit flow up to time $T-1$, and so a will get displaced if (9) exceeds (8). If m and M are the minimum and maximum values in the payoff matrix for G , we obtain the following sufficient condition for (9) to exceed (8) and hence for a to be displaced by b :

$$(1-\varepsilon)P_\delta^T(a, a) + \varepsilon P_\delta^T(a, b) - (1-\varepsilon)P_\delta^T(b, a) \\ < \varepsilon P_\delta^T(b, b) - \left(\frac{1-\delta}{\delta}\right)(M-m).$$

A simpler sufficient condition is obtained by noting that $P_\delta^T(b, b) = u$, and that

$$P_\delta^T(a, b) = P_\delta^T(b, a) \geq P_\delta^T(a, a).$$

The latter observation follows from the fact that b mimics a from time T onward, and hence $P_\delta^T(a, b) = P_\delta^T(b, a) \geq P_\delta^t(a, a) \geq P_\delta^T(a, a)$, for some t . On substituting, we obtain our final sufficient condition that b be able to displace a . The sufficient condition is

$$\varepsilon P_\delta^T(a, a) < \varepsilon u - \left(\frac{1-\delta}{\delta}\right)(M-m). \quad (10)$$

Suppose that it is false that $P_\delta(a, a) \rightarrow u$ as $\delta \rightarrow 1$. Then there exists an $\eta > 0$ and a sequence $\delta_i \rightarrow 1$ such that $P_{\delta_i}(a, a) \leq u - \eta$. It follows that,

$$P_{\delta_i}^{T(\delta_i)}(a, a) \leq P_{\delta_i}(a, a) \leq u - \eta < u - \left(\frac{1-\delta_i}{\delta_i \varepsilon_i}\right)(M-m), \quad (11)$$

provided that

$$\left(\frac{1-\delta_i}{\delta_i \varepsilon_i}\right)(M-m) < \eta.$$

To obtain a contradiction, it remains to select

$$\varepsilon_i = \frac{2(1-\delta_i)}{\eta \delta_i}.$$

Then $\varepsilon_i = O(1-\delta_i)$, but (11) and hence (10) hold. Therefore b can displace a for each value of i . Thus a is not a GUESS.

This proves the theorem provided that it is never true that $T=0$. If $T=0$, b is constructed to screen at time 0. Then $P_\delta^1(b, a) = P_\delta^1(a, b) \geq P_\delta^0(a, a)$, and (10) must be replaced by

$$(1-\varepsilon)P_\delta^1(a, a) - (1-2\varepsilon)P_\delta^0(a, a) < \varepsilon u - \left(\frac{1-\delta}{\delta}\right)(M-m).$$

Since $P_\delta^1(a, a) = P_\delta^0(a, a) + O(1-\delta)$, the result still follows after a little extra algebra.

REFERENCES

1. D. ABREU AND A. RUBINSTEIN, The structure of Nash equilibrium in repeated games with finite automata, *Econometrica* **56** (1988), 1259–1282.
2. L. ANDERLINI, "Communication, Computability and Common Interest Games," Working paper, St. John's College, Cambridge, 1990.
3. R. AUMANN, Survey of repeated games, in "Essays in Game Theory and Mathematical Economics in Honor of Oskar Morgenstern," Wissenschaftsverlag, Bibliographisches Institut, Mannheim/Wein/Zurich, 1981.
4. R. AUMANN AND S. SORIN, Cooperation and bounded recall, *Games Econ. Behav.* **1** (1989), 5–39.
5. R. AXELROD, "The Evolution of Cooperation," Basic Books, New York, 1984.
6. R. AXELROD AND W. HAMILTON, The evolution of cooperation, *Science* **211** (1981), 1390–1396.
7. J. BANKS AND R. SUNDARAM, Repeated games, finite automata and complexity, *Games Econ. Behav.* **2** (1990), 97–117.
8. K. G. BINMORE, Social contract III: Evolution and utilitarianism, *Constitutional Polit. Econ.* **1** (1990), 1–26.
9. R. BOYD AND J. LOBERBAUM, No pure strategy is evolutionarily stable in the repeated Prisoners' Dilemma game, *Nature* **327** (1987), 58–59.
10. R. DAWKINS, "The Selfish Gene," Oxford Univ. Press, Oxford, 1976.
11. J. FARRELL AND R. WARE, Evolutionary stability in the repeated Prisoner's Dilemma game, *Theoret. Population Biol.* **36** (1988), 161–166.
12. D. FUDENBERG, "Repeated Game Explanations of Commitment and Cooperation," in "Advances in Economic Theory: Sixth World Congress of the Econometric Society" (J.-J. Laffont, Ed.), Cambridge Univ. Press, Cambridge, 1991.

13. D. FUDENBERG AND E. MASKIN, Evolution and cooperation in noisy repeated games, *Amer. Econ. Rev.* **80** (1990), 274–279.
14. J. HOPCRAFT AND J. ULLMAN, "Introduction to Automata Theory, Languages and Computation," Addison–Wesley, Reading, MA, 1979.
15. E. KALAI AND A. NEME, "The Strength of a Little Perfection," Discussion paper 858, Northwestern University, 1989.
16. Y. KIM, "Evolutionary Stable Strategies in the Repeated Prisoner's Dilemma," Working paper, University of Iowa, 1989.
17. D. KREPS, P. MILGROM, J. ROBERTS, AND R. WILSON, Rational cooperation in the finitely repeated prisoners' dilemma, *J. Econ. Theory* **27** (1982), 245–252.
18. D. LEWIS, "Counterfactuals," Basil Blackwell, Oxford, 1976.
19. B. LINSTER, "Essays on Cooperation and Competition," Ph.D. Thesis, University of Michigan, 1990.
20. J. MAYNARD SMITH, "Evolution and the Theory of Games," Cambridge Univ. Press, Cambridge, 1982.
21. J. NACHBAR, "The Evolution of Cooperation Revisited," Working paper, Washington University, St. Louis, 1989.
22. A. NEYMAN, Bounded complexity justifies cooperation in the finitely repeated prisoners' dilemma, *Econ. Lett.* **19** (1986), 227–229.
23. A. ROBSON, Efficiency in evolutionary games: Darwin, Nash, and the secret handshake, *J. Theoret. Biol.* **144** (1990), 379–396.
24. A. RUBINSTEIN, Finite automata play the repeated prisoners' dilemma, *J. Econ. Theory* **39** (1986), 83–96.
25. R. SELTEN, Reexamination of the perfectness concept for equilibrium points in extensive-games, *Int. J. Game Theory* **4** (1975), 25–55.