

Repeated Games and Finite Automata

by Robert E. Marks

from *Recent Developments in Game Theory*,
ed. by J. Creedy, J. Eichberger, and J. Borland
(London: Edward Elgar, 1992), pp. 43–64.

Repeated Games and Finite Automata†

. . . why may we not say that all Automata . . . have an artificial life.

Hobbes, *Leviathan* (1651)

GAME THEORY—usually thought of as the framework par excellence for analysing strategic interactions—has also been characterised as a means of analysing the meaning of “rational” behaviour. One source of interest in rational behaviour flows from such games as the Prisoner’s Dilemma, in which the Nash equilibrium is not Pareto-optimal, in the one-shot game. Can the efficient, Pareto-optimal equilibrium be supported if the game is played repeatedly? The Folk Theorem (Aumann 1981) asserts that in the repeated game, the individually *rational* outcome may support the Pareto-optimal outcome of mutual coöperation instead of costly mutual defection.

A second source of interest in rational behaviour is as a datum against which “irrational” behaviour can be measured, and as a description of ways in which irrational behaviour is to be avoided.

But *irrational* behaviour is important in game theory, too. How robust is an equilibrium to apparently irrational behaviour? To what extent is apparently irrational behaviour the rational response to a coarse information partition, to incorrect information, to unobserved payoffs, to a mistaken action? As Aumann & Sorin (1989, p.37) put it:

The work on equilibrium refinements since Selten’s “trembling hand” (1975) indicates that rationality in games depends critically on *irrationality*. In one way or another, all refinements work by assuming that irrationality cannot be ruled out, that the players ascribe irrationality to each other with a small probability. True rationality needs a “noisy”, irrational environment; it cannot grow in sterile soil, cannot feed on itself only.¹

This issue is discussed at length in Binmore (1988).

Lest game theoreticians fall into ad hoc characterisations of these apparently necessary irrationalities, it is important to consider how best to model such phenomena. One method that has proved very productive is to model a player as a stimulus–response machine, in which the stimulus of the other players’ previous actions maps into a response. Of course, such machines can have no expectations, and can have no intentions.

Stimulus–response machines are in general finite: they have finite memory; they accept a finite number of input signals; and they have a finite set of

† The author wishes to thank Larry Samuelson for his assistance.

1. The irrationality required here is quite different from the apparent need for irrationality—or bounded rationality—identified by Simon (1984) in all schools of economic thought, which he characterises as ad hoc, casual appeals to limited rationality in order to explain such phenomena as business cycles and apparently involuntary unemployment, in the absence of exogenous shocks.

responses.² (The memory of previous moves or actions in general may constitute the stimulus to which the machine responds.) The machine may be modelled to respond to a set of stimuli with a particular response, that is, to use the previous moves as a basis for choosing its response, which obviates the need to postulate expectations. That is, the machines model forward induction; they cannot anticipate, and so cannot engage in backwards induction, as such.

Apart from modelling degrees of irrationality, or “bounded rationality”, to use Simon’s phrase (1972), stimulus–response machines have been used (*a*) in formal proofs of behaviour with players who exhibit irrational or bounded-rational behaviour, (*b*) in simulations of such behaviour, and (*c*) to formalise measures of strategic complexity (Marks 1990). A special example of such simulations has been what Binmore and Dasgupta (1986) call “descriptive” game theory, which can be modelled by an evolutionary process, in which a search algorithm from artificial intelligence machine learning, the Genetic Algorithm, mimics the evolution of “successful” machines, as measured by their payoffs in repeated games against other machines or against a “niche” of strategies, a weighted average of other machines (Marks 1989*a*, 1989*b*).

In general, stimulus–response machines have been modelled as responding with pure strategies. Mixed strategies can be modelled by positing a distribution over the (deterministic) machines; this may be a probability distribution in selection or a frequency distribution across a population. Moreover, it may be possible to construct “Markov machines”, which select a mixed strategy, a probability distribution over a set of pure strategies at each stage of the repeated game.

This paper is in several parts. Section 2.1 discusses “rationality” and “bounded rationality” in game theory. Section 2.2 introduces finite automata and Turing machines, and discusses how automata can model various forms of bounded rationality. Section 2.3 discusses the game-theoretical literature which uses the notion of players as finite automata to explore and prove existence, uniqueness, necessity, and sufficiency. Section 2.4 discusses the selection of finite automata, both theoretically and using Genetic Algorithm simulations.

2.1 BOUNDED RATIONALITY

In order to describe limits to rationality, it is first necessary to define rationality formally. Inevitably, such a discussion must rely on Herbert Simon’s writings, since he has been at the forefront of the Behaviouralist school in arguing for less ad-hockery and more empirical consistency in the modelling and use of “bounded rationality” in economic theorising. Although the postulate of rationality can take many forms,

for a wide range of assumptions, rationality implies that, in equilibrium,

2. Finite automata are just that, but infinite machines exist: Turing machines have infinite tapes, permitting more complicated behaviour than finite automata can exhibit (Megiddo and Wigderson 1986).

people will have no motivation to modify their behaviours, and resources will be fully employed. The equilibrium need not, of course, be static. (Simon, 1984, p.37)

In economic theory, out-of-equilibrium paths are usually assumed to be the result of exogenous shocks. In game theory, however, such shocks are not in general modelled, and out-of-equilibrium behaviour, if it exists, must be the result of irrationalities. (Equilibrium concepts have been developed to deal with behaviour flowing from imperfect or incomplete information, so by definition out-of-equilibrium behaviour cannot be due to this.) Binmore (1988) analyses the rôle of bounded rationality in economics in general and in game theory in particular; he concludes with a programme of research into the thinking processes of the players to better model rationality.

Although many might agree that to model *Homo œconomicus* as an all-powerful computing machine—*Homo calculans*—with unlimited abilities to determine actions necessary to maximise expected utility is a far-from-realistic assumption, I would argue that the profession has not adopted Simon’s concept of bounded rationality with great enthusiasm—especially in operationalising it—because of the absence of a consistent framework for modelling it, even if just how the human “machine” exhibits bounded rationality could be agreed on.

By any definition, there can be no limits to the complexity of response available to the unbounded-rational player. But bounded rationality implies limited complexity. We should like to be able to characterise the complexity of implementing strategies, using a cardinal measure of complexity. Borrowing from the mathematics of computing machines provides a means of definition and measurement of the complexity of the responses of players in repeated games, and by extension of all strategic actions.³

2.2 FINITE AUTOMATA AND MOORE MACHINES

The use of stimulus–response machines in repeated games derives from Aumann (1981),⁴ and has since been used by several authors (Neyman 1985; Radner 1986; Rubinstein 1986; and others). The most commonly used machines have been finite automata, although infinite machines (including Turing machines) have also been discussed (Megiddo and Wigderson 1986; Gilboa and Schmeidler 1989). Originally, economists’ interest in finite automata theory was to develop theoretical results about strategies in repeated games with limits on strategic complexity, but finite automata also provide a way of using techniques of machine learning to examine the processes of out-of-equilibrium behaviour and to search for robust strategies in repeated games, as discussed in Section 2.4, below.

We can formalise a *finite automaton*, and provide some examples. Let Q_i

3. Megiddo (1986) raises some objections to the characterisation of bounded rationality that focuses on time constraints for information processing, as captured with finite automata.

4. The notion of “machine models” had earlier been mentioned by Selten (1978), and according to Radner (1986) by T.A. Marschak and C.B. McGuire in unpublished lecture notes in 1971.

be a finite set, called the set of possible internal states of player i 's automaton, and let S_i and S_j denote the finite sets of actions or moves for players i and j , respectively. If in round t the state of player i 's machine is $q_i(t) \in Q_i$ and player j 's move is $s_j(t) \in S_j$, then at round $t+1$ the state of player i 's machine, $q_i(t+1)$, will be

$$q_i(t+1) = \delta_i[q_i(t), s_j(t)],$$

and player i 's move (or action) in round $t+1$, $s_i(t+1) \in S_i$, will be

$$s_i(t+1) = \lambda_i[q_i(t+1)].$$

The quadruple $\langle Q_i, \bar{q}_i, \lambda_i, \delta_i \rangle$ constitutes player i 's automaton,⁵ where $\bar{q}_i \in Q_i$ is the initial state of the machine, where λ_i is the action function, $\lambda_i : Q_i \rightarrow S_i$, and where δ_i is the next-state (or transition) function, $\delta_i : Q_i \times S_j \rightarrow Q_i$. The number of elements in Q_i is called the size of the automaton. In order to rank automata by size, care must be taken to compare minimal machines of behavioural equivalence (Harrison 1965), that is, to compare the sizes of the reduced forms (Moore 1956).

Rubinstein (1986) describes a world in which players select Moore machines (Moore 1956) instead of explicit strategies. A *Moore machine* is a finite automaton in which the player's next move (the machine's output) is contingent on the existing state of the machine, which in turn is a function of the previous state of the machine (at the previous round) and the other player's previous move (the machine's input), through a transition (or next-state) function. (The initial state and the set of all feasible internal states of each machine must be defined at the outset, along with the set of all feasible moves and the transition function and the "action"—or output—function.) If both players in a two-person game have chosen Moore machines, then the game can continue between the machines, which will generate moves (and states) as the repeated game progresses.

It is possible to depict Moore machines as *transition diagrams*, directed graphs whose vertices or nodes correspond to the states, q_i , of the machine represented and whose edges correspond to the possible transitions between those states. One of the nodes is the "Start," \bar{q}_i . Below, we present transition diagrams of strategies in the repeated Prisoner's Dilemma. The letters C or D immediately beneath each node show the machine's move (the output) associated with that node; the letters C and/or D immediately above each arc correspond to the other player's move (the input), after which the machine moves to the new node at the arrowed end of the arc.

For instance, a machine which plays C constantly (Always Coöperate) can be described as

$$Q = \{ q^* \}, \bar{q} = q^*, \lambda(q^*) = C, \text{ and } \delta(q^*, \cdot) \equiv q^*.$$

This is depicted in Figure 2.1.

5. Strictly (Hopcroft and Ullman 1979), the description should also include the sets of input and output symbols, but since we are modelling games in which both players face the same action sets, we omit these.

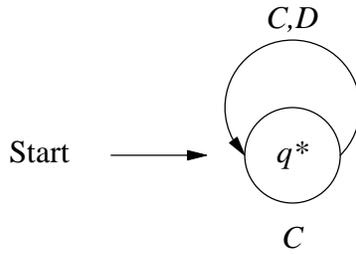


Figure 1. The “Always Coöperate Moore Machine

Rapoport’s strategy, Tit for Tat, can be described as

$$Q = \{q^C, q^D\}, \bar{q} = q^C, \lambda(q^s) = s \text{ and } \delta(q, s) = q^s \text{ for } s = C, D.$$

Its transition diagram is given by Figure 2.2.

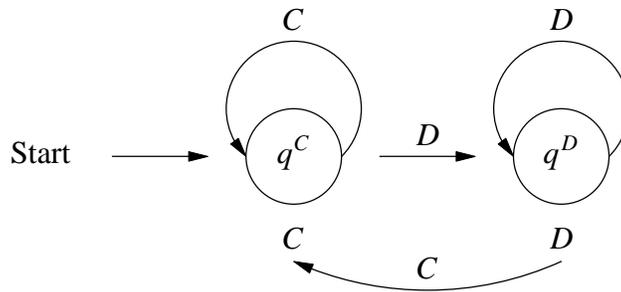


Figure 2. The “Tit for Tat” Moore Machine

The strategy of playing C until the other player plays D and then punishing him for three periods regardless of what moves he makes in the meantime before returning to coöperation requires at least a four-node machine, as depicted in Figure 2.3.

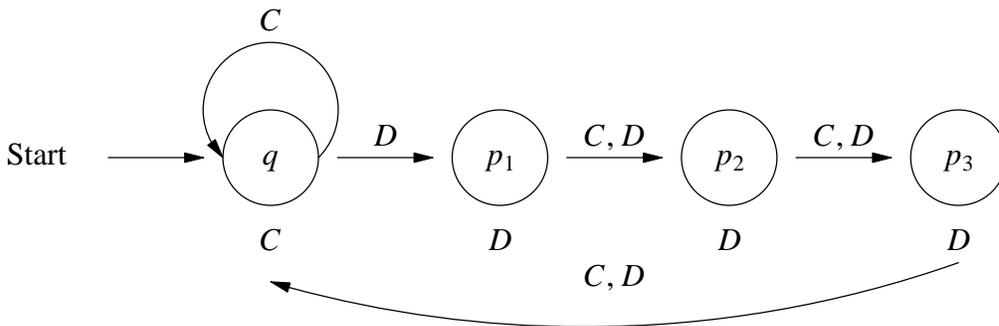


Figure 3. A Four-Node Moore Machine

Each of the states reached by an unconditional transition (that is, regardless of the opponent's move) is called a *counting state* (Miller 1988), and the number of counting states or strings of connected counting states in the minimal finite automaton provides additional information on the behaviour of the machine. The machine of Figure 2.3 can be described as

$$\begin{aligned} Q &= \{q, p_1, p_2, p_3\}, \bar{q} = q, \lambda(q) = C, \lambda(p_h) = D, (h = 1, 2, 3), \\ \delta(q, C) &= q, \delta(q, D) = p_1, \delta(p_h, \cdot) \equiv p_{h+1}, \text{ and } \delta(p_3, \cdot) \equiv q. \end{aligned}$$

It is possible to model a *trigger strategy* (Radner 1980), in which a pattern of play on the part of the opponent triggers the machine's moves into (usually) the punishment of continual defection. This is shown in Figure 2.4, in which q^D is the *trapping state*:

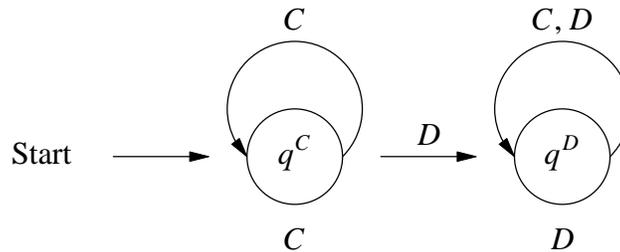


Figure 4. A Trigger-Strategy Moore Machine

the first play of D by the opponent triggers the move to q^D , and the machine remains in that state for the rest of the game, playing D . The number of trapping or terminal states in a minimal finite automaton is of interest, since at least one is required for each trigger strategy (Miller 1988).

It is possible to think of the succession of opponent's moves as constituting symbols on an input tape read by the automaton, in response to which the machine changes state and produces a succession of moves of its own. That is, the state of the automaton, and hence its own moves, is a function of the concatenation of the input symbols it has received since the start (Hopcroft and Ullman 1979).

Gilboa and Samet (1989) define a *connected finite automaton* (CFA) as follows: given an automaton $\langle Q, \bar{q}, \lambda, \delta \rangle$ (we drop the subscripts for clarity), and given two states $q, \hat{q} \in Q$, we say that \hat{q} is accessible from q (and write $q \rightarrow \hat{q}$) if there exists a history h_r such that $\delta(q, h_r) = \hat{q}$. (A history of player r is the concatenation of player r 's moves since the start of the repeated game, and $\delta(\cdot)$ is the transition function; player r is the opponent in the two-person game.) Two states, q and \hat{q} , are mutually accessible (written $q \leftrightarrow \hat{q}$) if both $q \rightarrow \hat{q}$ and $\hat{q} \rightarrow q$. The automaton is said to be connected if all states belonging to Q are mutually accessible.⁶ A connected automaton cannot describe trigger strategies; connectedness rules out what Gilboa and Samet call "vengeful" strategies:

6. Marks (1990) describes how a finite automaton can be modelled algebraically, specifically as a general non-negative matrix, and how these propositions are related to the matrix structure.

however “angry” the automaton may be, it can always be appeased.

It is convenient for using the Genetic Algorithm (Section 2.4) to represent these machines by strings, together with rules describing the transition and action functions. Each locus (of one or more characters) on the string corresponds uniquely to a state. The action function is simply a mapping from the locus on the string to the output character (or characters) (in the case of the Prisoner’s Dilemma the single characters *C* or *D*). The transition function will result in a new locus (or state), contingent on the previous locus and the input of the other player’s previous move.

For instance, in the Always Coöperate machine of Figure 2.1, there is only one node, which always results in *C*. Thus, the string representation of this machine might be the string *C*. Then, whatever the previous move of the other player, the machine’s response would be an unchanging *C*. For Tit for Tat there must be at least two elements in the string, one corresponding to the other player’s coöperating in the previous round, and the other corresponding to his defecting. The first results in the machine’s responding with *C*, the second with *D*. Thus, the string representation of Tit for Tat might be, say, *CD*, where *C* corresponds to node 1 and *D* corresponds to node 2, as in Figure 2.2. The algorithm would tell us to look at node 1 for our next move if the other player’s previous move was *C*, and to look at node 2 for our next move if the other player’s previous move was *D*. The four-node strategy of Figure 2.3 might be represented by the string *CDDD*; this strategy is not as simple as the previous two; the transition function, for instance, is not simple, although the transition diagram can be followed without too much difficulty. This machine recalls up to three moves ago—only after three *D*s does it revert to a *C*, a kind of Three Tits for a Tat.

It might be concluded that a strategy which has no memory (such as Figure 2.1) requires one node, that a 1-round memory (Figure 2.2) requires two nodes, and that a 3-round memory (Figure 2.3) requires four nodes. A moment’s thought, however, will reveal that (a) the number of states must be a function of the number of possible inputs and outputs, and (b) in a two-person game with *s* possible symmetric moves there are s^2 possible combinations of play per round, so that to recall all possible moves for the last *r* rounds a machine will require s^{2r} states. For a specific strategy, however, not all of these states will be connected, which is the reason for comparing the sizes of minimal machines, which are behaviourally equivalent to their unreduced originals.

When using finite automata to simulate play in a repeated game, or when selecting finite automata to play more successfully in a repeated game as discussed in Section 2.4, we face an engineering problem. As Harrison (1965, p.299) puts it:

The trouble with computing the behaviour of a machine directly from its definition is that the concept is not finitary in nature. In principle one cannot feed all possible tapes [successions of opponent’s moves] into the machine to decide which input words [ditto] cause the machine to go into a final state.

It is possible, however, to define the behaviour of a finite automaton and to use finite experiments to determine whether two machines are behaviourally equivalent. This hastens solution of the analysis problem, which consists of

describing the behaviour, or “emergent properties”, of a given finite automaton. A second problem is to design a finite machine which has a specific behaviour. With a solution to this problem, we can attempt to find a “best design”, where “best” might be the least complex machine.⁷

The size of an automaton can be defined as the number of states it has. The complexity of a strategy is defined by Ben-Porath (1987) as the minimal size of the automaton that can implement it. From the transition diagrams above, it appears that Always Coöperate is of lowest strategic complexity, followed by Tit for Tat, and that Figure 2.3 depicts a strategy of higher complexity. Kalai and Stanford (1988) note that for any machine this complexity measure is equivalent to the number of distinct strategies induced by the original strategy in all possible subgames, so that the trigger strategy automaton of Figure 2.4 has complexity two, since it induces only itself or the constant D strategy. As Radner (1986) notes, this measure does not take account of the complexity of the action function and the transition function—what Gottinger (1983, p.127) calls the tradeoff between structural complexity and computational complexity. Banks and Sundaram (1990) develop a complexity measure that takes into account both the size (number of states) and transitional structure of an automaton.

Nonetheless, Ben-Porath’s measure of strategic complexity raises the question: Given any level of strategic complexity, what is the most successful strategy in competing against a given environment of strategies? Tit for Tat has proved itself to be, at a low level of strategic complexity, extremely robust against a wide range of opponents. This raises another question: With no limit on strategic complexity, can Tit for Tat be soundly bettered? We shall return to these questions in Section 2.4.

Of the three measures of the characteristics of finite automata mentioned above—the numbers of states, counting states, and trapping states—the last is by far the most significant: with no trapping states, a finite automaton will eventually forget; with trapping states, a connected finite automaton may eventually “trigger,” never to forget. Let us call finite automata with no trapping states *bounded recall finite automata* or BRFA; let us call finite automata with trapping states *trigger finite automata* or TFA. Gilboa and Samet (1989) assert that the set of connected-automaton (CFA) strategies is (strictly) larger than that of bounded-recall strategies (those associated with BRFA). There is a special class of TFA, those automata which possess a single state, which must therefore be a trapping state. These are like the Moore machine of Figure 2.1: they exhibit

7. This problem—of designing or choosing a machine to play the game—is a complex pure-strategy choice (Ben-Porath 1988), more complex than the actual game-playing decisions, as we see in Section 2.4, below. Binmore (1988) posits metaphorical meta-players, who make the machine choice, analogous with Walras’ auctioneer in tâtonnement.

unchanging behaviour, and so memory and forgetting are irrelevant.

2.3 REPEATED GAMES

In a one-shot Prisoner's Dilemma (PD) game, the dominant (pure) strategy is to defect,⁸ despite a higher payoff for coöperation, because of the reward of cheating and the penalty of being cheated.

In a repeated PD game of unknown length, however, the higher payoff to coöperation may result in strategies different from the Always Defect of the single game, because of the opportunity to punish defection provided by later rounds. By breaking the logical imperative of mutual defection inherent in the static, one-shot PD, the repeated PD—in which the players repeatedly face each other in the same situation—can admit the possibility of *learning* on the part of the players, which may result in mutual coöperation or some mixed strategy on their part, as they learn more about the type of behaviour they can expect from each other and build up a set of beliefs of behaviour.

An early analysis of successful strategies in the repeated PD (Luce and Raiffa 1957, pp.97–102) suggested that continued, mutual coöperation might be a viable strategy, despite the rewards from defection, but for twenty years no stronger analytical results were obtained for the repeated PD.

As is now widely known, Axelrod's tournaments (1984) revealed that one very simple strategy is difficult to better in the repeated PD: Rapoport's Tit for Tat. When pitted against a "nasty" strategy, such as Always Defect, it does almost as well, itself defecting on every round but the first, but at the cost of the aggregate score. When played against itself, each player's aggregate score is a maximum, since every round will then be mutual coöperation, a result which resembles collusion, although each player's decisions are made independently of the other's.

In the one-shot PD game the Cournot–Nash non-coöperative equilibrium dominates the Pareto-superior coöperative solution. This result generalises to n -player games and provides a rationale for price wars when there are a small number of sellers of differentiated products, as the in MIT tournaments (Fader and Hauser 1988), and in other cases (Eaton and Slade 1989). With a simple game played between two opponents for more than a single round, the opportunity of responding to an opponent's defection in the previous round with a defection in this and later rounds raises the possibility that the threat of defection may induce mutual coöperation. But for games of finite duration with low discount rates (we can use the "limit of means" or the discounted payoffs for the game score) this hope is dashed by the end-game behaviour, or what Selten (1975) called the "chain-store paradox". There is a discontinuity for infinitely repeated games (or supergames): the Folk Theorem (Aumann 1989) tells us that any individually rational payoff vector can be supported in infinitely repeated games, for sufficiently low discount rates. (For high discount rates the threat of future punishment may not be sufficiently great to offset the gain from defecting now.)

8. Although Aumann and Sorin (1989) use the terms "friendly" and "greedy" play instead of the more usual "coöperate," "defect", or "fink", we shall stay with the familiar, if imprecise, words.

In order to explain the apparent evidence of coöperative behaviour among oligopolists in the real world, among experimental subjects in clinical trials, and among strategy simulation tournaments—all of them examples of finite repetitions—researchers have sought relaxation of the underlying assumptions in the finite game.

Kreps et al.—the so-called gang of four—(1982) assumed incomplete information: they relaxed the assumption that rationality is common knowledge (Aumann 1976) among the players. This allowed them to perturb a finitely repeated Prisoner's Dilemma by assuming that with a small probability one of the players is playing Tit for Tat rather than maximising as a perfectly rational player. They showed that with a sufficiently long repetition all sequential equilibrium outcomes are close to coöperative. But, as Aumann and Sorin (1989) point out, this result could be stronger: because Tit for Tat is the only perturbation allowed, in a sense it is the input as well as the output. The coöperative sequential equilibrium is not really endogenously coöperative, as might be concluded if the perturbation admitted of all possible alternative strategies. (See Aumann and Sorin's (1989) result below.)

The literature on finite automata in repeated games can be categorised into two distinct branches: the analysis of the theoretical equilibrium properties of machine games, and the effect of finite computational abilities on supporting coöperative outcomes (the Folk Theorem and its relatives). Rubinstein (1986), Abreu and Rubinstein (1988), and Banks and Sundaram (1990) fall into the first category, in which the level of strategic complexity is endogenous; Neyman (1985), Megiddo and Wigderson (1986), and others fall into the second, in which the level of strategic complexity is exogenous.

Using the number of states as their measure of the complexity of implementing a strategy, Abreu and Rubinstein (1988) consider the tradeoff between the cost of this complexity and the repeated-game payoffs in the players' choices of Moore machines. This generalises the earlier work of Rubinstein (1986), in which the level of complexity of the strategies—modelled as finite automata—was a lexicographic ordering of average payoff above machine complexity. (Rubinstein had introduced a dynamic concept of *automaton equilibrium*: at no time during the infinite-length game would the players want to alter their machines. The earlier work demonstrated that opposing machines will coördinate their actions, which sharply reduces the set of equilibrium outcomes from the game, and that coöperation cannot be the outcome of a solution of the infinitely repeated Prisoner's Dilemma.) Players simultaneously choose Moore machines to implement their strategies, the complexities of which are measured by the number of states in the minimal automaton necessary to play the strategy. Abreu and Rubinstein analyse Nash equilibrium in the machine game, and derive necessary conditions on the form of equilibrium strategies and plays, rather than the more frequent results concerning equilibrium payoffs. They show that in any Nash equilibrium of the machine game, “the two machines have an equal number of states, and maximise repeated game payoffs against one another”. That is, in equilibrium, players' choices are fully optimal, despite the complexity considerations explicitly introduced. Their results suggest that the introduction of

implementation costs—through the complexity of the strategies—results in a “striking” discontinuity in the Nash equilibrium set in terms of strategies, plays, and payoffs, as with the chain-store paradox.

Banks and Sundaram (1990) attempt to capture the transitional complexity of machine strategies in the repeated game by considering the number of edges in the transition diagram of the Moore machine representation of the automaton. They find that the one-shot Nash equilibrium is invariably supported in the repeated PD—*only* mutual cooperation.

Neyman (1985) investigated what happens when fully rational players are replaced by automata in finitely repeated games. Neyman showed that when the players are restricted to finite automata, no matter how much larger these machines are than the number of repetitions, there exist equilibria with payoffs that are on average close to the cooperative payoff. That is, automaton players enable—but do not ensure—cooperation that is impossible with full rationality.

This is also the conclusion reached by Radner (1986), who explored three departures from full rationality: uncertainty about the degree of cooperativeness of the other player in a two-person game; the epsilon-equilibrium concept, in which each player is satisfied to approach the payoffs of the other player’s strategy; and (following Neyman) machine strategies implemented by finite automata of limited size (complexity). Radner found in the first case that, under certain conditions, the larger the total number of stages in the repeated game, the longer the players remain cooperative; in the second case that as the number of stages increases the corresponding sets of equilibria include those with longer and longer cooperation; and in the case of finite-automata strategies that if the number of stages is sufficiently large compared to the size of the automaton, then there are equilibria in which the players cooperate throughout the repeated game. Harrington (1987) found that limited complexity of players’ beliefs—instead of players’ strategies—could result in the emergence of cooperation. Friedman (1971) and Sorin (1986) showed that a sufficiently high discount rate was sufficient. Fudenberg and Maskin (1986) extended the proofs in the infinitely repeated case to games of three or more players.

Megiddo and Wigderson (1986) model a finitely repeated Prisoner’s Dilemma game played by Turing machines, each with a symmetrically restricted number of internal states, using unlimited time and space. Their results strongly suggest that Folk Theorem holds: the cooperative outcome of the game can be approximated in equilibrium; that is, even if the machines memorise the entire history of the game and are capable of counting the number of stages, the cooperative play can be approximated. Their Turing machines differ from Neyman’s finite automata in several ways: (a) they consider machines with unlimited memory, whereas automata have no memory besides their states; (b) their machines are uniform, and can play any number of rounds, announced at the start of the game; and (c) they consider pure-strategy choices of machines, rather than Neyman’s mixed-strategy choices.

Lehrer (1988) addresses repeated games played by asymmetric players with bounded recall who do not know the stage of the infinite game at which they are currently playing. In a non-zero-sum game, he finds that the set of Nash-

equilibrium payoffs tends to the set of all the individually rational and feasible payoffs. (He also examines the asymptotic behaviour of the set of equilibrium payoffs as the capacity of the memories of both players grow to infinity.) Although not explicitly modelled as finite automata, his bounded-recall strategies can be so modelled (Marks 1989a).

Aumann and Sorin (1989) define *common interests* in a two-person game if there exists a single payoff pair that strongly Pareto-dominates all other payoff pairs, such as (C,C) in the Prisoner's Dilemma. They model a perturbation in which during repetitions of a game with common interests each player attaches a small but positive probability to the other's playing some bounded-recall fixed-strategy automaton. (This is their irrationality in the search for coöperative outcomes.) They find that this perturbation of the repeated game possesses pure-strategy equilibria, and that all such equilibria are close (in payoff) to the unique coöperative (efficient, Pareto-optimal) pair of payoffs of the game with common interest. That is, coöperation is ensured under their conditions, not merely possible, as in the Folk Theorem.

They report that they first conjectured that it might be sufficient to perturb the game with strategies that could be played by automata of bounded complexity, but found that bounded recall is essential. As they put it (Aumann and Sorin, 1989, p.8):

People must be willing to forget past grievances; remembering the distant past is not a good means for fostering coöperation. More accurately, in a culture in which irrational people have long memories, rational people are less likely to coöperate.

Moreover, the set of possible automata must be sufficiently rich: it must contain at least all the zero-recall strategies. Their result is a powerful theoretical justification for the coöperation that Axelrod (1984) was able to evolve in his computer tournaments, and which Miller (1988) and Marks (1989a) also obtain with their Genetic Algorithm simulations.

Kalai and Stanford (1988) follow Ben-Porath's (1988) work on the relationship between the structure of strategies and equilibria, as opposed to the characterisation of equilibrium payoffs of Neyman, and others considering exogenous, or uniform, strategic complexity. They assert that their finite automata are richer than the Moore machines described in Section 2.2 above, since they use Mealy machines (Mealy 1955), which include their own actions as inputs, as well as their opponent's. This, Kalai and Stanford assert, enables their automata to deal with every history of past plays and not merely self-consistent histories, which in turn allows *subgame perfection* to become a relevant solution concept. Since Moore and Mealy machines are behaviourally equivalent (Assmus and Florentin 1968), the basis for their assertion is unclear.

Combining finite complexity of automaton players with epsilon equilibrium, Kalai and Stanford find that every subgame-perfect equilibrium of the repeated game can be approximated (with regard to payoffs) by a subgame-perfect epsilon equilibrium of finite complexity. They also prove necessary relationships among the complexities and memories of players' strategies for certain classes of subgame-

perfect equilibria in two-person games.

Gilboa and Samet (1989) consider two-person repeated games in which a player of bounded rationality (modelled as a connected finite automaton CFA) chooses pure strategies against an unbounded rational player (leaving the issue of the existence of such an animal unresolved). They determine that the rational player has a dominant strategy; that in some cases the weaker, bounded CFA player may exploit this fact to “blackmail” the rational player: the “tyranny of the weak”. This analysis formalises the idea of “stubbornness”: the CFA player does not have to announce his choice, he simply has to play it and let the rational player learn it through experimentation. This is a dominant strategy for the rational player. Since the automaton is connected, it has no trapping states and cannot therefore implement trigger strategies, which would be costly, perhaps fatally so, to its opponent, if triggered by experimentation. The results hold even if the automaton player is allowed to randomise over CFAs.

Gilboa and Schmeidler (1989) introduce three assumptions to the theoretical literature: (a) infinite histories, which means that there is no period zero to begin forward induction from; (this models institutional interactions which continue without beginning or end—or may do); (b) Turing machines with memory: they show that with infinite histories a decision-maker’s Turing-machine strategy, implementable by a Turing machine which always halts, is no more than a finite-recall strategy; this enables them to strengthen the computational model by endowing the machines with external memory to allow them to carry over some memory from one stage to the next;⁹ (c) what they call non-strategic players, who do not speculate on others’ strategies but rather treat the history of play as a stimulus to generate the next action. This describes machine players, of course, but is also close in spirit to the evolutionary modelling to be described in the next section. With these assumptions, the authors define a solution concept for the one-shot game, called “steady orbit”. They determine that the closure of the set of steady-orbits payoffs strictly includes the convex hull of the Nash equilibria payoffs, and is strictly included in the correlated equilibria payoffs (Aumann 1974). This can be viewed as an attempt to formulate the “repeated game” interpretation of Nash equilibrium in the one-shot game.

As Binmore and Dasgupta (1986) suggest, an evolutionary competition among game-playing programs provides an avenue for linking prescriptive game theory with descriptive game theory: in the long run not quite all of us are dead, only those who were unsuccessful in the repeated game—some genes of those who scored well survive in their descendents. This provides a *learning* model in which it is the generations of populations of strategies that learn, not individuals, which are immutable. Samuelson (1988) provides a theoretical framework for examining the processes of the evolution of strategies, at least for finite, two-person normal-form games of complete information. He proves that, under certain properties of the evolutionary process, equilibrium strategies will be supported that are

9. Whereas finite automata use their states to remember information—previous plays—from one stage of the repeated game to the next, Turing machines in an infinite-history game require additional “external” memory to do this, since they use their states for computation alone.

“trembling-hand perfect” (Selten 1975, 1983; Binmore and Dasgupta 1986), a subset of Cournot–Nash equilibrium.

Early work by biologists on the emergence of coöperation in animal populations (Maynard Smith 1982) was also concerned with the *evolutionary stability* of strategies (or genetically determined behaviour traits): their ability to survive in the face of an “invasion” by other strategies. Simulation (Marks 1989*b*) allows precise and unambiguous examinations to be made of such occurrences by use of a non-random initial population of strategies that has been seeded with any desired ratio of specific invaders to incumbents. The invaders can be any of the strategies possible within the particular formulation used.

Binmore and Dasgupta (1986, pp.16–19) argue that the equilibrium concept that Selten (1975) calls *perfect* equilibrium but that they call *trembling-hand* equilibrium¹⁰ is relevant to the discussion of stability to invasion. Roughly speaking, a Nash equilibrium for any game is a trembling-hand equilibrium if each of its component strategies remains optimal even when the opponents’ hands “tremble” as they select their equilibrium strategies. This concept models out-of-equilibrium behaviour, perhaps due to a mistake, or perhaps due to incorrect information.¹¹

2.4 SELECTING FINITE AUTOMATA

In the previous section we focused on equilibrium concepts. We now turn to the questions of selection and design mentioned above. Until the end of the section we restrict discussion to the problem of selecting a best-response automaton in a two-person repeated game when there is uncertainty about the machine selected by the other player. In an analysis of the complexity of selection—as opposed to the strategic complexity of the machine—Ben-Porath (1988) shows that both versions of the selection problem—finding a best-response automaton, or deciding whether a given automaton is a best-response—are “difficult” (that is, not polynomial).¹² Gilboa (1988) had previously shown that when players select pure strategies (that is, select a single machine and not a distribution across machines), the problem of finding a best-response automaton is polynomial if the number of players is known in advance, but NP otherwise. Ben-Porath shows that when players use mixed strategies (that is, select from a distribution across automata), the selection problem is NP even in a two-person game.

10. They prefer *trembling hand* to *perfect* in order to distinguish the concept clearly from another of Selten’s: *subgame-perfect* (Binmore and Dasgupta 1986, fn.18). All trembling-hand equilibria are subgame perfect, but the converse is not true. See also Selten (1983).

11. Binmore and Samuelson (1990) regard the choice of automaton of Abreu and Rubinstein (1988) as the outcome of an evolutionary process. They define a modified evolutionarily stable strategy (Maynard Smith 1982) and examine the circumstances under which the only evolutionarily stable outcome in an infinitely repeated game is “utilitarian”, in which the sum of the players’ payoffs is maximised.

12. In the computer science literature, problems are categorised as either polynomial or non-polynomial (NP). Polynomial problems are considered “simple”, non-polynomial problems “difficult”.

As Ben-Porath puts it (1988, p.2):

[T]here is an interpretation of Nash equilibrium in which it is not necessary to assume that the players can compute a best-response strategy. This is known as the evolutionary interpretation. Each player in the game corresponds to a group of a certain type in a population, and a mixed strategy represents the fractions of individuals that play different actions. A Nash equilibrium corresponds to a steady state in the following sense: If a population is not in a Nash equilibrium, over time some individuals will find (by error or by experimenting but not necessarily by calculation) a profitable deviation and will stick to it. Others will mimic them, or if they are not capable of doing even that, will eventually join them by the same process.

This is a good description of the process, first used by Axelrod (1987), of simulating the evolution of strategies as stimulus–response machines in a repeated game by means of the process of machine learning known as the Genetic Algorithm (Holland 1975; Goldberg 1988).¹³

Given the rules of the game and the payoff matrix in normal form, and given an upper bound on the complexity of possible strategies as measured by the number of rounds of the game “recalled” by the machine, the process of simulated evolution searches the large space of available machines to derive those behaviourally equivalent machines which are “best”, as measured by average payoff or discounted payoff across the repeated game. In Axelrod’s (1987) study, in a game of perfect information the machines were playing against a “niche” of strategies derived from his earlier (1984) computer tournaments. He did not characterise his derived strategies as machines or automata; it was left to Marks (1989*a*) to attempt to replicate his work, and to present the generated strategies as Moore machines.

Miller (1988) uses the Genetic Algorithm to generate strategies as explicit finite automata, that is, in his formulation the strategies are not simply interpreted as finite automata after the selection process, which is what Marks (1989*a*) does, but are available from a family of Moore machines only. He argues that there are two advantages of finite automata over the n -round-recall machines of Axelrod (1987) and Marks (1989): finite automata can embody a greater range of strategies, such as trigger strategies, which require trapping states, which are unavailable to n -round-recall strategies, which eventually forget; and, he asserts, finite automata are analytically richer.

Miller’s automata are two-round recall machines, modelled as bit-strings of length 148 ($4 + 16 \times 9$ bits). Miller’s study includes games of imperfect information, as well as perfect information, by modelling symmetric noisy

13. Fujiki and Dickinson (1987) describe using the GA to generate programs written in Lisp to “solve” the repeated PD—this is much more complex than our modelling. Chess (1988) describes simulations to generate best-response strategies in the iterated Prisoner’s Dilemma, and generates simple algorithms, but the set of possible machines is small and he does not use the Genetic Algorithm.

Marimon et al. (1990) use a Genetic Algorithm classifier system to model “artificially intelligent” agents learning to trade in an economy with money as a medium of exchange.

reporting of the opponent's actual moves: for each round there is a finite probability, in the repeated Prisoner's Dilemma, that the opponent's move is wrongly reported. His results suggest that the level of noise in the system has a fundamental effect on the outcome: higher levels of imperfect information are associated with less coöperation and lower payoffs. The effect of noise is apparently not continuous—phase transitions are evident in his results.

In a second study, Marks (1989*b*) uses the Genetic Algorithm to examine the extent to which repetition supports coöperation in repeated games, both two- and three-person, of perfect information. He models one-, two-, and three-round-memory strategies. In what he dubbed *bootstrapping evolution*, he allows the evolution of both players to occur by pitting each individual strategy in a population of strategies against all other strategies (or combinations of strategies in three-person games) to obtain a fitness score for each strategy. This bootstrap breeding, together with the Genetic Algorithm's search properties, should result in "evolutionary" convergence to the optimum optimum of all possible strategies. (There is some doubt whether all loci will be optimally selected for: an individual emerging into a population of similar strategies will not experience much opportunity to respond to hugely different strategies, and over time there may be genetic drift, as the descendents lose some traits previously strongly selected for. The consequences of this kin-selection for the possibility of invasions examined in Marks (1989*b*)).

As a consequence of the GA's processes, we speak of convergence to *behaviour*, not to *structure*: when, amongst themselves, the population of strategies all play the same action for the duration of each repeated game and for all possible combinations, we say that the population has *converged*. That is, we are searching for behaviourally equivalent strategies. Marks (1989*b*) examines the resistance of these converged populations to the introduction or invasion of new strategies from outside, in a simulation of trembling-hand equilibrium, as discussed by Binmore and Samuelson (1990).

Marks' simulations relax three of the assumptions of simple models: (*a*) strategies with longer than one-round memories, (*b*) games with more than two possible actions per player, and (*c*) games with more than two players. For those games for which theoretical results had been derived, he was able to simulate them using bounded-recall automata and the Genetic Algorithm.

Eaton and Slade (1989) demonstrate analytically and using evolutionary simulations with the Genetic Algorithm that small deviations from Axelrod's (1984) setup break the link that enables coöperation to emerge in the repeated Prisoner's Dilemma. In particular, they show that allowing players to change strategies without announcing this change to opponents drastically changes the result, and they demonstrate that the unique evolutionary equilibrium of the infinitely repeated Prisoner's Dilemma without discounting is observationally equivalent to infinite repetition of the Nash equilibrium of the one-shot game, that

is, mutual defection.

2.5 CONCLUSION

This paper has attempted to do several things. First, it has attempted to review the growing literature on the use of stimulus–response machines as players in repeated games. It will be seen that finite automata and bounded-recall strategies are more frequently used, while two papers have also used the more powerful Turing machines of computer science. We have derived a beginner’s taxonomy of finite automata: connected finite automata, bounded-recall finite automata, trigger-strategy finite automata, and the trivial constant-behaviour automata (in the repeated Prisoner’s Dilemma: “always coöperate”, and “always defect”).

Furthermore, we have shown how stimulus–response machines of various kinds (bounded-recall, finite automata) have been used in the beginnings of a study of what Binmore (1988) calls the evolutive study of the adjustment process, in which the value of the machines is that various forms of bounded rationality can be explicitly modelled and examined by the evolutionary simulations possible with the Genetic Algorithm. Examples of this literature are Axelrod (1987), Miller (1988), Marks (1989*a*, 1989*b*), and Eaton and Slade (1989). Future extensions of the use of finite automata in game theory include the possibility of modelling the Markov processes which may occur in non-deterministic games, but this area is virtually untouched; simulation may prove equally valuable in this application.

The importance of machines in game theory is to allow us to introduce forms of irrationality in a gentle way, by means of various bounds on the computational power of the automata. This may accelerate Simon’s hope to introduce a Behaviouralist approach to economics in general and game theory—the study of strategy—in particular.

REFERENCES:

- Abreu, D. and Rubinstein, A. (1988) The structure of Nash equilibrium in repeated games with finite automata. *Econometrica*, 56, pp.1,259–1,282.
- Assmus, E.F., Jr., and Florentin, J.J. (1968) Algebraic machine theory and logical design. In *Algebraic Theory of Machines, Languages, and Semigroups* (edited by M.A. Arbib), pp.15–35. New York: Academic Press.
- Aumann, R. (1974) Subjectivity and correlation in randomized strategies. *J. Math. Econ.*, 1, pp.67–95.
- Aumann, R. (1976) Agreeing to disagree. *Annals Stat.*, 4, pp.1,236–1,239.
- Aumann, R. (1981) Survey of repeated games. In *Essays in Game Theory and Mathematical Economics in Honor of Oskar Morgenstern* (by R.J. Aumann et al.), pp.11–42. Zurich: Bibliographisches Institut.
- Aumann, R. (1989) Game theory. In *The New Palgrave: Game Theory* (edited by J. Eatwell, M. Milgate, P. Newman), pp.1–53. London: Macmillan.
- Aumann, R.J. and Sorin, S. (1989) Coöperation and bounded recall. *Games & Econ.*

- Behav.*, 1, pp.5–39.
- Axelrod, R. (1984) *The Evolution of Coöperation*, New York: Basic Books.
- Axelrod, R. (1987) The evolution of strategies in the iterated Prisoner's Dilemma. In *Genetic Algorithms and Simulated Annealing* (edited by L. Davis), London: Pittman.
- Banks, J.S. and Sundaram, R.K. (1990) Repeated games, finite automata, and complexity. *Games & Econ. Behav.*, 2, pp.97–117.
- Ben-Porath, E. (1987) Repeated games with finite automata. Stanford University Institute for Mathematical Studies in the Social Sciences, Tech. Report No. 515, August.
- Ben-Porath, E. (1988) The complexity of computing a best response automaton in repeated games with mixed strategies. Mimeo., Grad. School of Bus., Stanford Univ.
- Binmore, K. (1988) Modeling rational players, Part II. *Economics and Philosophy*, 4, pp.9–55.
- Binmore, K. and Dasgupta, P. (1986) Game theory: a survey. In *Economic Organizations as Games* (edited by K. Binmore and P. Dasgupta), pp.1–45. Oxford: B. Blackwell.
- Binmore, K. and Samuelson, L. (1990) Evolutionary stability in repeated games played by finite automata. Mimeo.
- Chess, D.M (1988) Simulating the evolution of behaviour: the Iterated Prisoners' Dilemma. *Complex Systems*, 2, pp.663–670.
- Eaton, B.C. and Slade, M.E. (1989) Evolutionary equilibrium in market supergames. Mimeo., November.
- Fader, P.S., and Hauser, J.R. (1988) Implicit coalitions in a generalized Prisoner's Dilemma. *J. Conflict Resol.*, 32, pp.553–582.
- Friedman, J.W. (1971) A non-coöperative equilibrium of supergames. *Rev. Econ. Stud.*, 38, pp.1–12.
- Fudenberg, D., and Maskin, E. (1986) The Folk Theorem in repeated games with discounting or incomplete information. *Econometrica*, 54, pp.533–554.
- Fujiki, C., and Dickinson, J. (1987) Using the genetic algorithm to generate Lisp source code to solve the Prisoner's Dilemma. In *Genetic Algorithms & Their Applications, Proc 2nd. Intl. Conf. Gen. Alg.* (edited by J.J. Grefenstette), pp.236–240. Hillsdale, N.J.: Lawrence Erlbaum Assoc.
- Futia, C. (1977) The complexity of economic decision rules. *J. Math. Econ.*, 4, pp.289–299.
- Gilboa, I. (1988) The complexity of computing best-response automata in repeated games. *J. Econ. Theory*, 45, pp.342–352.
- Gilboa, I. and Samet, D. (1989) Bounded versus unbounded rationality: the tyranny of the weak. *Games and Econ. Behav.*, 1, pp.213–221.
- Gilboa, I. and Schmeidler, D. (1989) Infinite histories and steady orbits in repeated games. Mimeo., August.
- Goldberg, D.E. (1988) *Genetic Algorithms in Search, Optimization, and Machine Learning*. Reading, Mass.: Addison-Wesley.
- Gottinger, H.W. (1983) *Coping with Complexity: Perspectives for Economics, Management and Social Sciences*. Dordrecht: D. Reidel.
- Harrington, J.E., Jr. (1987) Finite rationalizability and coöperation in the finitely repeated Prisoner's Dilemma. *Econ. Lett.*, 23, pp.233–237.
- Harrison, M.A. (1965) *Introduction to Switching and Automata Theory*. New York: McGraw-Hill.
- Holland, J.H. (1975) *Adaptation in Natural and Artificial Systems*. Ann Arbor: Univ. Michigan Press.
- Hopcroft, J.E. and Ullman, J.D. (1979) *Introduction to Automata Theory, Languages, and*

- Computation*. Reading: Addison-Wesley.
- Kalai, E. and Stanford, W. (1988) Finite rationality and interpersonal complexity in repeated games. *Econometrica*, 65, pp.397–410.
- Kreps, D., Milgrom, P., Roberts, J., and Wilson, R. (1982) Rational cooperation in the finitely repeated Prisoner's Dilemma. *J. Econ. Theory*, 27, pp.245–252.
- Lehrer, E. (1988) Repeated games with stationary bounded recall strategies. *J. Econ. Th.*, 46, pp.130–144.
- Luce, R.D. and Raiffa, H. (1957) *Games and Decisions: Introduction and Critical Survey*. New York: Wiley.
- Marimon, R., McGrattan, E., Sargent T. J. (1990) Money as a medium of exchange in an economy with artificially intelligent agents. *J. of Econ. Dynamics and Control*, 14, pp.329–373.
- Marks, R.E. (1989a) Niche strategies: the Prisoner's Dilemma computer tournaments revisited. *AGSM Working Paper 89-009*.
- Marks R.E. (1989b) Breeding hybrid strategies: optimal behaviour for oligopolists. In *Proceedings of the Third International Conference on Genetic Algorithms*, George Mason University, June 4–7, 1989 (edited by J.David Schaffer), pp.198–207. San Mateo: Morgan Kaufmann.
- Marks R.E. (1990) Measures of strategic complexity. Mimeo. Presented at the Sixth World Congress of the Econometric Society, Barcelona.
- Maynard Smith, J. (1982) *Evolution and the Theory of Games*. Camb.: Camb. Univ. Press.
- Mealy, G.H. (1955) A method of synthesizing sequential circuits. *Bell System Tech. J.*, 34, pp.1,045–1,079.
- Megiddo, N. (1986) Remarks on bounded rationality. IBM Research Report, RJ 5270 (54310). Yorktown Heights: IBM Research Division.
- Megiddo, N. and Wigderson, A. (1986) On play by means of computing machines. In *Reasoning About Knowledge* (edited by J.Y. Halpern) pp.259–274. Los Altos: Kaufmann.
- Miller, J.H. (1988) The evolution of automata in the repeated Prisoner's Dilemma. Mimeo., Dept. Econ., Univ. Mich., Aug.
- Moore, E.F. (1956) Gedanken-experiments on sequential machines. In *Automata Studies* (edited by C.E. Shannon and J. McCarthy), pp.129–153. Princeton: Princeton Univ. Press.
- Neyman, A. (1985) Bounded complexity justifies cooperation in the finitely repeated Prisoners' Dilemma. *Econ. Lett.*, 19, pp.227–229.
- Radner, R. (1980) Collusive behaviour in noncoöperative epsilon-equilibria of oligopolies with long but finite lives. *J. Econ. Theory*, 22, pp.136–154.
- Radner, R. (1986) Can bounded rationality resolve the Prisoners' Dilemma? In *Contributions to Mathematical Economics in Honor of Gérard Debreu* (edited by W. Hildenbrand and A. Mas-Colell), pp.387–399. Amsterdam: North-Holland.
- Rubinstein, A. (1986) Finite automata play the repeated Prisoners' Dilemma. *J. Econ. Theory*, 39, pp.83–96.
- Samuelson, L. (1988) Evolutionary foundations of solution concepts for finite, two-player, normal-form games. Mimeo., Dept. Econ., Penn. State Univ.
- Selten, R.C. (1975) Reëxamination of the perfectness concept for equilibrium points in extensive games. *Inter. J. Game Theory*, 4, pp.25–55.
- Selten, R. (1978) Chain-store paradox. *Theory and Decision*, 9, pp.127–159.
- Selten, R.C. (1983) Evolutionary stability in extensive two-person games. *Math. Soc. Sci.*,

5, pp.269–363.

Simon, H.A. (1972) Theories of bounded rationality. I: *Decision and Organization* (edited by C.B McGuire and R. Radner), pp.161–188. Amsterdam: North Holland.

Simon, H A. (1984) On the behavioral and rational foundations of economic dynamics. *J. of Econ. Behavior and Organization*, 5, pp.35–55.

Sorin, S. (1986) On repeated games with complete information. *Math. of O. R.*, 11, pp.147–160.

BIOGRAPHY

Robert Marks lectures at the Australian Graduate School of Management in the University of New South Wales, where he was a foundation lecturer. Previously, he had been an instructor in the Department of Engineering-Economic Systems, Stanford University, which he later visited as an Assistant Professor. He has also visited the Energy and Resources Group at UC Berkeley and the M.I.T. Energy Laboratory. His major research interests include game theory (in 1987 he was the winner of the Second M.I.T. Competitive Strategy Computer Tournament), learning models in economics, energy policy, and drug policy. His publications include the book, *Nonrenewable Resources and Disequilibrium Macrodynamics*, (New York: Garland Publishing, 1979).

Repeated Games and Finite Automata

Robert E. Marks

Australian Graduate School of Management,
University of New South Wales,
P.O. Box 1,
Kensington NSW 2033

Phone: (02) 662-0271

Internet: bobm@agsm.unsw.oz.au

Presented at the two-day seminar,
Recent Developments in Game Theory,
the University of Melbourne,
June 7-8, 1990.

CONTENTS

2.1 BOUNDED RATIONALITY	3
2.2 FINITE AUTOMATA AND MOORE MACHINES	4
2.3 REPEATED GAMES	10
2.4 SELECTING FINITE AUTOMATA	15
2.5 CONCLUSION	18
REFERENCES:	18
BIOGRAPHY	22

LIST OF FIGURES

Figure 1. The “Always Coöperate Moore Machine	6
Figure 2. The “Tit for Tat” Moore Machine	6
Figure 3. A Four-Node Moore Machine	6
Figure 4. A Trigger-Strategy Moore Machine	7