Playing Games with Genetic Algorithms

Robert E. Marks[1]

Australian Graduate School of Management, Universities of Sydney and New South Wales, Sydney 2052, Australia. bobm@agsm.edu.au

Abstract. In 1987 the first published research appeared which used the Genetic Algorithm as a means of seeking better strategies in playing the repeated Prisoner's Dilemma. Since then the application of Genetic Algorithms to game-theoretical models has been used in many ways. To seek better strategies in historical oligopolistic interactions, to model economic learning, and to explore the support of cooperation in repeated interactions. This brief survey summarises related work and publications over the past thirteen years. It includes discussions of the use of game-playing automata, co-evolution of strategies, adaptive learning, a comparison of evolutionary game theory and the Genetic Algorithm, the incorporation of historical data into evolutionary simulations, and the problems of economic simulations using real-world data.

1 Introduction

Over the past twenty-five years, non-cooperative game theory has moved from the periphery of mainstream economics to the centre of micro-economics and macro-economics. Issues of information, signalling, reputation, and strategic interaction can best be analysed in a game-theoretic framework. But solution of the behaviour and equilibrium of a dynamic or repeated game is not as simple as that for a static or one-shot game. The multiplicity of Nash equilibria of repeated games has led to attempts to eliminate many of these through refinements of the equilibrium concept. At the same time, the use of the rational expectations assumption to cut the Gordian Knot of the intractability of dynamic problems has led to a reaction against the super-rational *Homo calculans* model towards boundedly rational models of human behaviour.

In the late 'eighties the Genetic Algorithm (GA) was first used to solve a dynamic game, the repeated Prisoner's Dilemma (RPD) or iterated Prisoner's Dilemma (IPD). Mimicking Darwinian natural selection, as a simulation it could only elucidate sufficient conditions for its Markov Perfect equilibria (MPE), rather than the necessary conditions of closed-formed solution, but over the past twenty-odd years, its use in economics has grown to facilitate much greater understanding of evolution, learning, and adaptation of economic agents. This brief survey attempts to highlight emergence of the marriage of GAs and game theory.

2 Deductive Versus Evolutionary Approaches to Game Theory

Ken Binmore and Partha Dasgupta [?] argued that the use of repeated games in economics can be viewed from two perspectives; first, deductively determining what the equilibrium outcome of a strategic interaction will be, using the Nash equilibrium concept as a solution method, and, second, evolutively, asking how the strategic players will learn as they play, which will also result in equilibrium. In general, the possible equilibria which follow from the second approach are a subset of those that follow from the first. Indeed, "learned" equilibria are one attempt to "refine" the possible Nash equilibria of the deductive approach. A good summary of the rationale and techniques of the "learning" approach to solving game equilibria is given in Fudenberg and Levine [?], who mention the GA as a means of exploring the space of strategies to play repeated games.

3 The Repeated Prisoner's Dilemma

Several years earlier, Robert Axelrod [?] had set up a tournament among computer routines submitted to play the RPD. Pairs of routines were matched in the repeated game, and their scores used to increase or reduce the proportion of each routine in the total population of potential players. Although Axelrod was a political scientist, he was influenced by William Hamilton, a biologist, and this simulated "ecology" mimicked in some way the natural selection undergone by species competing for survival in an ecological niche. It is well known that Anatol Rapoport's Tit for Tat emerged as a robust survivor among the many routines submitted. Foreshadowing the analytical tool of replicator dynamics (see below), however, the number of strategies in the game was fixed at the start: no new strategies could emerge.

4 Boundedly Rational Players

Axelrod's 1984 tournament pitted routines that were perforce boundedly rational. That one of the simplest, Tit for Tat, emerged as the best, as measured by its high score (although not always: see [?]), might only have been because of the bounded nature of its strategy (start off "nice", and then mimic your opponent's last move).

Deductive game theory has not assumed any limits to agents' abilities to solve complex optimisation problems, and yet Herbert Simon has been arguing for many years what applied psychologists and experimental economists are increasingly demonstrating in the laboratory: human cognitive abilities are bounded, and so descriptive models should reflect this. But one virtue of the assumption of perfect rationality is that global optima are reached, so that there is no need to ask in what manner our rationality is bounded. It may be no coincidence that two recent monographs on bounded rationality in economics (Thomas Sargent [?], and Ariel Rubinstein [?]) are written by authors who have previously been concerned with repeated games played by machines [?] and one of the first published uses of GAs in economics [?]. What is the link between these two areas? Some background will clarify.

5 Game-Playing Automata

The GA was developed by John Holland [?,?] and his students, including David Goldberg [?]. The original applications were in engineering and were predominantly optimisation problems: find $x = \arg\max f(x)$. The comparative advantage of GAs was that $f(.)$ was not required to be continuous or convex or differentiable or even explicitly closed-form – for a recent overview, see [?]. So the original focus was on optimisation of closed-form functions.

But the GA can search for more general optima. Axelrod [?] recounts how his colleague at Michigan, John Holland, mentioned that there was this new technique in Artificial Intelligence (or what has become known as Machine Learning) which, by simulating natural selection, was able to search for optima in extremely non-linear spaces. Axelrod [?], the first published research to use the GA with repeated games, was the result: against the same niche of submitted routines that he had used in his 1984 study, Axelrod used the GA to see whether it could find a more robust strategy in the RPD than Tit for Tat. I came across a reference to this then-unpublished work in Holland et al. [?] and obtained both the draft and the code in early 1988 from Axelrod. I had been interested in solutions to the RPD since a routine of mine had won the second M. I. T. competitive strategy tournament [?], an attempt to search for a generalisation for Tit for Tat in a three-person game with price competition among sellers of imperfect substitutes. What could this new technique tell me about my serendipitous routine?

Axelrod (with the programming assistance of Stephanie Forrest [?]) modelled players in his discrete RPD game as stimulus-response automata, where the stimulus was the state of the game, defined as both player's actions over the previous several moves, and the response was the next period's action (or actions). That is, he modelled the game as a state-space game [?,?], in which past play influences current and future actions, not because it has a direct effect on the game environment (the payoff function) but because all (or both) players believe that past play matters. Axelrod's model focused attention on a smaller class of "Markov" or "state-space" strategies, in which past actions influence current play only through their effect on a state variable that summarises the direct effect of the past on the current environment (the payoffs). With state-space games, the state summarises all history that is payoff-relevant, and players' strategies are restricted to depend only on the state and the time.

Specifically, Axelrod's stimulus-response players were modelled as strings, each point of which corresponds to a possible state (one possible history) and decodes to the player's action in the next period. The longer the memory, the longer the string, because the greater the possible number of states. Moving to a game with more than the two possible moves of the RPD will lengthen the string, holding the number of players constant. Increasing the number of players will also increase the number of states. Formally, the number of states is given by the formula $a^{m\,p}$, where there are a actions, p players, each remembering m periods [?].

Although the implicit function of Axelrod's 1987 paper is non-linear and open-form, in one way this pioneering work was limited: the niche against which his players interact is static, so the GA is searching a niche in which the other player's behaviour is determined at time zero: these routines do not learn. This type of problem is characterised as open-loop, and is essentially non-strategic [?].

6 Co-Evolution of Automata

The interest of game theory is in strategic interaction, in which all players respond to each other, and can learn how better to respond. I thought that GAs playing each other would be more interesting than Axelrod's static niche (which I replicated in [?]) and extended Axelrod's work to co-evolution of stimulus-response players in [?], later published as [?], although I termed the simultaneous adaptation of each player in response to all others' actions as "boot strapping", being unaware of the biologist's term.

Another approach to modelling players as co-evolving stimulus-response machines was taken by John Miller, in his 1988 Michigan PhD thesis (later published as [?]). He modelled the finite automaton explicitly, and let the GA derive new automata. This is explained further in [?]. Miller's approach has the advantage that offspring of the genetic recombination of the GA are functioning finite automata, whereas offspring in the Axelrod approach may include many irrelevant states which – via the curse of dimensionality [?] – dramatically increases the computation costs. Holland and Miller [?] argued for the merits of simulation in general (anticipating some of the arguments made by Judd [?]), and bottom-up artificial adaptive agents (such as the population of strategies playing the RPD to derive a fitness score for the GA) in particular.

As discussed at length in the two monographs, Sargent [?] and Rubinstein [?], modelling players as boundedly rational poses important questions for the modeller: which constraints are realistic? which ad hoc? Deterministic finite automata playing each other result in Markov perfect equilibria. But as soon as the researcher uses a (necessarily finite) automaton to model the player – or its strategy – the question of boundedness must be addressed.

Eight years ago, the literature on finite automata in repeated games could be categorised into two distinct branches: the analysis of the theoretical equilibrium properties of machine games, and the effect of finite computational abilities on the support of cooperative outcomes [?]. Since then, as discussed in the section, Empirical Games, below, there is a growing literature which uses such techniques as the GA and neural nets to explore historical market behaviour, especially in oligopolies.

Rubinstein [?] was using automata as players to explore some theoretical issues associated with the RPD. Others have followed, without necessarily using the GA to solve the strategy problem. Megiddo and Wigderson [?] used Turing machines (potentially of infinite size) to play the RPD; Chess [?] used simulations to generate best-response strategies in the RPD, and generated simple routines, but he did not use the GA; Cho [?,?,?] and Cho and Li [?] used "perceptrons," neural nets, to demonstrate support of the Folk Theorem in RPD games, with simple strategies, with imperfect monitoring, and with nets that are of low complexity. Fogel and Harrald [?] and Marks and Schnabl [?] compare neural nets and the GA for simulating the RPD, and Herbrich et al. [?] survey the use of neural networks in economics.

7 Learning

The earliest published paper to use the GA to solve the RPD is Fujiki and Dickinson [?], but instead of explicitly using automata, they modelled the players as Lisp routines, and allowed the GA to search the space of LisP routines for higher scoring code. (Lisp is a programming language often used in Artificial Intelligence applications.) The earliest publication in a peer-reviewed economics journal of a GA used in economics modelling was Marimon et al. [?], but it was not explicitly about using the GA to solve games. Instead, in a macro-economic model it used the GA to model learning, the first of many papers to do so.

I must admit that in an early 1989 conversation with Tom Sargent, I discounted the mechanisms of the GA – selection, crossover, and mutation – as models of the learning process, but Jasmina Arifovic, a student of Sargent's, in a continuing of series of papers over the past decade which model GA learning, has shown that this conclusion was wrong. Arifovic [?] uses the GA to simulate the learning of rational expectations in an overlapping-generations (OLG) model; her first published paper [?] was on learning in that canonical model of equilibrium determination: the Cobweb Model, in which she introduces the election operator, the first contribution by an economist to the practice of GAs, previously dominated by, first, engineers, and, second, mathematicians; Arifovic [?] shows that using the GA in an OLG model results in convergence to rational expectations and reveals behaviour on the part of the artificial adaptive agents which mimics that observed by subjects in experimental studies; [?] finds a two-period stable equilibrium in an

OLG model; Arifovic and Eaton [?] use the GA in a coordination game. The Cobweb Model has proved a popular subject of investigation using the GA: Dawid and Kopel [?] explore two models – one with only quantity choices and one with the prior decision of whether to exit or stay in the market; Franke [?] explores the stability of the GA's approximation of the moving Walrasian equilibrium in a Cobweb economy, and provides an interpretation and extension of Arifovic's election operator [?].

One of the traits of Tit for Tat as characterised by Axelrod [?] is that it is easily recognised. That is, other players can soon learn the type of player they are facing. A more complex strategy would not necessarily be as readily recognised. Dawid [?,?,?] argues that GAs may be an economically meaningful model of adaptive learning if we consider a learning process incorporating imitation (crossover), communication (selection), and innovation (mutation).

There are several other relevant papers on this topic. Bullard and Duffy [?] explore how a population of myopic artificial adaptive agents might learn forecasting rules using a GA in a general equilibrium environment, and find that coordination on steady state and low-order cycles occurs, but not on higher-order periodic equilibria. Bullard and Duffy [?] use a GA to update beliefs in a heterogeneous population of artificial agents, and find that coordination of beliefs can occur and so convergence to efficient, rational expectations equilibria. Riechmann [?] uses insights from the mathematical analysis of the behaviour of GAs (especially their convergence and stability properties) to argue that the GA is a compound of three different learning schemes, with a stability between asymptotic convergence and explosion. Riechmann [?] demonstrates that economic GA learning models can be equipped with the whole box of evolutionary game theory [?], that GA learning results in a series of near-Nash equilibria which then approach a neighbourhood of an evolutionary stable state. Dawid and Mehlmann [?], and Chen et al. [?] also consider genetic learning in repeated games. The recent volume edited by Brenner [?], on computational techniques of modelling learning in economics, includes several papers of interest, in particular Beckenbach [?] on how the GA can be re-shaped as a tool for economic modelling, specifically to model learning.

8 Replicator Dynamics

But the GA is not the only technique to come from biology into economics. Convergence of a kind was occurring as economics in general, and game theory in particular, was borrowing insights from a second source in biology. For the past fifty years there has been some interest in the Darwinian idea of natural selection as applied to firm survival in the (more or less) competitive marketplace, as summarised in Nelson and Winter [?]. Borrowing game theoretical ideas from economics, John Maynard Smith [?] introduced the concept of the evolutionarily stable strategy (ESS) to biology, and this in turn was used by economists eager to reduce the number of possible equilibria in repeated games. ESS is concerned with the invasion of a new species (or strategy) into an ecology (market) in which there is an existing profile of species (strategies). Axelrod [?] had been mimicking these changes as he allowed the proportions of his RPD-playing routines to alter in response to their relative success. The most widely discussed paper, and one of the earliest to use GA's in examining the changing profiles of strategies in the RPD, is Lindgren [?]. Two recent investigations of the RPD are by Ho [?], with information processing costs, and by Wu and Axelrod [?], with noise.

Axelrod [?] had argued that simple strategies such as Tit for Tat were collectively evolutionarily stable, and so could sustain cooperation in the finitely RPD. In an unpublished paper, Slade and Eaton [?] argue that Axelrod's conclusion is not robust to small deviations from the Axelrod setup, in particular, allowing agents to alter their strategies without announcement. Similarly, Nachbar [?,?] questions the robustness of Axelrod's findings, arguing that prior restrictions on the strategy set result in a convergence to cooperation with Tit for Tat, and that without these restrictions cooperation would eventually be exhausted.

Replicator dynamics [?,?] exhibit similar behaviour, as the profile of initial strategies changes, but unlike the GA as applied to stimulus-response machines (strategies) replicator dynamics cannot derive completely new strategies. Replicator dynamics have provided the possibility of closed-form investigation of evolutionary game theory [?,?,?], but evolutionary game theory in general, and replicator dynamics in particular, despite the word "evolutionary", are not directly related to GAs (see [?,?] for comparisons).

Chen and Ni [?] use the GA to examine Selten's [?] chain store game, in a variant due to Jung et al. [?], where there are "weak" and "strong" monopolists, in the context of answering questions of Weibull's on evolutionary game theory [?]. They characterise the phenomenon of coevolutionary stability, although it may be that the simulated behaviour observed are actually Markov Perfect equilibria [?] with a very long periodicity.

To what extent can interactions in oligopolies be modelled as the play of a repeated Prisoner's Dilemma? Yao and Darwen [?] examined the emergence of cooperation in an n-person iterated Prisoner's Dilemma, and showed that the larger the group size, the less likely its emergence. Chen and Ni [?] examine an oligopoly with three players, which they characterized as a time-variant, state-dependent Markov transition matrix, and conclude that, owing to the path dependence of the payoff matrix of the oligopoly game, this is not the case. Further, they argue that the rich ecology of oligopolistic behavior can emerge without the influence of changing outside factors (such as fluctuating income or structural changes), using GAs to model the adaptive behavior of the oligopolists.

9 Other Refinements

The literature on the RPD cited above restricts the player's decisions to its next action, but does not model an earlier decision: whether or not to play. This constraint is relaxed – players get to decide with whom to interact – in a series of papers by Leigh Tesfatsion and co-authors [?,?], culminating in her simulation of trading [?].

The GA has been used to search for the emergence of risk-averse players, in a simple model [?], and more recently in a robust model of the farmer's decision faced with high- and low-risk options [?]. Huck et al. [?] use replication and mutation (but not crossover or the GA) to search for the emergence of risk aversion in artificial players of repeated games.

Tomas Başar has been prominent in the mathematics of game theory: the book [?] solved many issues in dynamic, continuous, differential games before many economists had become aware of the problems. Recently, he has used GAs to solve for Stackelberg solutions, in which one player moves first [?].

Following in Başar's tradition, Özyildirim [?] applied the GA to search for approximations of the non-quadratic, non-linear, open-loop, Nash Cournot equilibria in a macro-economic policy game. She goes beyond many earlier studies in using historical data to explore the interactions between the Roosevelt administration and organised labour during the New Deal of the 'thirties. In a second paper [?], the GA is used to search for optimal policies in a trading game with knowledge spillovers and pollution externalities.

10 Empirical Games

The two papers by Özyildirim, discussed above [?,?], show the future, I believe: moving from stylised games, such as the RPD, to actual market interactions. Slade [?] provides a clear structure for this project, as well as reporting preliminary results. She finds that the Nash equilibrium of the one-shot game is emphatically not a reasonably approximation to real-world repeated interactive outcomes, and that most games in price or quantity appear to yield outcomes that are more collusive than the one-shot outcome, a finding which is consistent with the Folk Theorem.

Marks and his coauthors have been involved in a project to use GAs to solve for MPE using historical data on the interactions among ground-coffee brands at the retail level, where these players are modelled as Axelrod/Forrest finite automata in an oligopoly. This work is a generalisation of Axelrod [?] and Marks [?], and uses the ability of the GA to search the highly disjoint space of strategies, as Fudenberg and Levine [?] suggest.

The first two papers [?,?] build on a "market model" that was estimated from historical supermarket-scanner data on the sales, prices, and other marketing instruments of up to nine distinct brands of canned ground coffee, as well as using brand-specific cost data. The market model in effect gives the one-shot payoffs (weekly profits) for each brand, given the simultaneous actions of all nine (asymmetric) brands. There was natural synchrony in brands' actions: the supermarket chains required their prices and other actions to change together, once a week, at midnight on Saturdays.

Modelling the three most rivalrous historical brands as artificial adaptive agents (or Axelrod strings), the authors use the GA to simulate their coevolution, with the actions of the other six brands taken as unchanging, in a repeated interaction to model the oligopoly. After convergence to an apparent MPE in their actions, the best of each brand is played in open-loop competition with the historical data of the other six in order to benchmark its performance. The authors conclude that the historical brand managers could have improved their performance given insights from the response of the stimulus-response artificial managers, although this must be qualified with the observation that closed-loop experiments [?], allowing the other managers to respond, would be more conclusive.

Marks et al. [?] refine the earlier work by increasing the number of strategic players to four, by using four distinct populations in the simulation (against the one-string-fits-all single population of the earlier work), and by increasing the number of possible actions per player from four to eight, in order to allow the artificial agents to learn which actions (prices) are almost always inferior and to be avoided. Marks [?] explores some issues raised by the discrete simulations and the curse of dimensionality, and uses entropy as a measure of the loss of information in partitioning the continuous historical data (in prices etc.). Slade [?,?] considers two markets with similar rivalrous behaviour – gasoline and saltine crackers – but does not use a GA.

Curzon Price [?] uses a GA to model several standard industrial-organization models in order to demonstrate that the GA performs well as a tool in applied economic work requiring market simulation. He considers simple models of Bertrand and Cournot competition, a vertical chain of monopolies, and an electricity pool.

11 Conclusion

When John Holland [?] invented the GA, his original term for it was an "adaptive plan" [?], which looked for "improvement" in complex systems, or "structures which perform well." Despite that, most research effort, particularly outside economics, has been on its use as a function optimiser. But, starting with Axelrod [?], the GA has increasingly been used as an adaptive search procedure, and latterly as a model of human learning in repeated situations. In the 1992 second edition of his 1975 monograph, Holland expressed the wish that the GA be seen more as a means of improvement and less on its use as an optimiser.

This survey, although brief, has attempted to show that use of the GA in economics in general, and in game theory in particular, has been more and more focusing on its use as an adaptive search procedure, searching in the space of strategies of repeated games, and providing insights into historical oligopolistic behaviour, as well as into human learning processes.

12 Acknowledgements

The work in this study was partly supported by an Australian Research Council grant and the Australian Graduate School of Management. I wish to thank the editor, Shu-Heng Chen, for his encouragement in the face of my recent bereavement, and my colleague, Robin Stonecash, for her support.

References

1. Alemdar N. M., Özyildirim S. (1998) A genetic game of trade, growth and externalities, Journal of Economic Dynamics and Control **22**(6): 811–32.
2. Arifovic J. (1990) Learning by Genetic Algorithms in economic environments, Santa Fe Institute Econ. Res. Prog. W.P. 90-001.
3. Arifovic J. (1994) Genetic Algorithm learning and the Cobweb Model, Journal of Economic Dynamics and Control **18**(1): 3–28.
4. Arifovic J. (1995) Genetic Algorithms and inflationary economies, Journal of Monetary Economics **36**(1): 219–243.
5. Arifovic J. (1998) Stability of equilibria under Genetic Algorithm adaptation: an analysis, Macroeconomic Dynamics **2**(1): 1–21.
6. Arifovic J., Eaton B. C. (1995) Coordination via genetic learning, Computational Economics **8**(3): 181–203.
7. Arifovic J., Eaton B. C. (1998) The evolution of type communication in a sender/receiver game of common interest with cheap talk, Journal of Economic Dynamics and Control **22**(8-9): 1187–1207.
8. Ashlock D., Smucker M. D., Stanley E. A., Tesfatsion L. (1996) Preferential partner selection in an evolutionary study of Prisoner's Dilemma, BioSystems **37** (1–2): 99–125.
9. Axelrod R. (1984) The Evolution of Cooperation, N. Y.: Basic Books.
10. Axelrod R. (1987) The evolution of strategies in the iterated Prisoner's Dilemma, in Genetic Algorithms and Simulated Annealing, L. Davis (ed.), London: Pittman. pp.32–41.
11. Axelrod R. (1997) The Complexity of Cooperation, Princeton: P. U. P.
12. Başar T., Oldser G. J. (1982) Dynamic Non-Cooperative Game Theory, New York: Academic Press.
13. Beckenbach F. (1999) Learning by Genetic Algorithms in economics? in Brenner (1999) op. cit., pp. 73–100.
14. Binmore K., 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.
15. Birchenhall C. R. (1996) Evolutionary games and Genetic Algorithms In Computational Economic Systems: Models, Methods Econometrics, (edited by M. Gilli), in the series Advances in Computational Economics 5, pp. 3–23. Dordrecht: Kluwer Academic Publishers.
16. Brenner T. (ed.) (1999) Computational Techniques for Modelling Learning in Economics, in the series Advances in Computational Economics 11, Dordrecht: Kluwer Academic Publishers.
17. Bullard J., Duffy J. (1998) Learning and the stability of cycles, Macroeconomic Dynamics **2**(1): 22–48.
18. Bullard J., Duffy J. (1999) Using Genetic Algorithms to model the evolution of heterogeneous beliefs, Computational Economics **13**(1): 41–60.
19. Cacho O., Simmons P. (1999) A Genetic Algorithm approach to farm investment, Australian Journal of Agricultural and Resource Economics, **43**(3): 305–322.
20. Chen S.-H., Duffy J., Yeh C. H. (1996) Equilibrium selection via adaptation: using genetic programming to model learning in a co-ordination game mimeo.

21. Chen S.-H., Ni C.-C. (1997) Coevolutionary instability in games: an analysis based on Genetic Algorithms. In Proceedings of 1997 IEEE International Conference on Evolutionary Computation, Piscataway, N.J.: Institute of Electrical and Electronics Engineers, pp. 703–708.
22. Chen S.-H., Ni C.-C. (2000) Simulating the ecology of oligopolistic competition with Genetic Algorithms, Knowledge and Information Systems, forthcoming.
23. Chess D. M. (1988) Simulating the evolution of behavior: the Iterated Prisoners' Dilemma, Complex Systems, **2**: 663–670.
24. Cho I.-K. (1995) Perceptrons play the repeated Prisoner's Dilemma, Journal of Economic Theory, **67**: 266–284.
25. Cho I.-K. (1996a) On the complexity of repeated principal agent games, Economic Theory, **7**(1): 1–17.
26. Cho I.-K. (1996b) Perceptrons play repeated games with imperfect monitoring, Games and Economic Behavior **16**(1): 22–53.
27. Cho I.-K., Li H. (1999) How complex are networks playing repeated games? Economic Theory **13**(1): 93–123.
28. Curzon Price T. (1997) Using co-evolutionary programming to simulate strategic behaviour in markets, Journal of Evolutionary Economics **7**(3): 219–254.
29. Dawid H. (1996a) Genetic Algorithms as a model of adaptive learning in economic systems, Central European Journal for Operations Research and Economics **4**(1): 7–23.
30. Dawid H. (1996b) Learning of cycles and sunspot equilibria by Genetic Algorithms, Journal of Evolutionary Economics **6**(4): 361–373.
31. Dawid H. (1999) Adaptive Learning by Genetic Algorithms: Analytical Results and Applications to Economic Models, Lecture Notes in Economics and Mathematical Systems, vol 441. Heidelberg: Springer-Verlag, 2nd ed.
32. Dawid H., Kopel M. (1998) On economic applications of the Genetic Algorithm: a model of the cobweb type, Journal of Evolutionary Economics **8**(3): 297–315.
33. Dawid H., Mehlmann A. (1996) Genetic learning in strategic form games, Complexity **1**(5): 51-59.
34. Fader P. S., Hauser J. R. (1988) Implicit coalitions in a generalised Prisoner's Dilemma, Journal of Conflict Resolution, **32**: 553–582.
35. Fogel D. B., Harrald P. G. (1994) Evolving continuous behaviors in the iterated Prisoner's Dilemma in: Sebald, A., Fogel, L. (eds.) The Third Annual Conference on Evolutionary Programming, Singapore: World Scientific, pp. 119–130.
36. Forrest S. (ed.) (1991), Emergent Computation: Self-Organizing, Collective, and Cooperative Phenomena in Natural and Artificial Computing Networks, Cambridge: M.I.T. Press.
37. Franke R. (1998) Coevolution and stable adjustments in the cobweb model, Journal of Evolutionary Economics **8**(4): 383–406.
38. Friedman D. (1991) Evolutionary games in economics, Econometrica, **59**(3): 637–666.
39. Fudenberg D., Levine D. K. (1998) The Theory of Learning in Games, Cambridge: M.I.T. Press.
40. Fudenberg D., Tirole J. (1991) Game Theory, Cambridge: M.I.T. Press.
41. Fujiki C., Dickinson J. (1987) Using the genetic algorithm to generate Lisp source code to solve the Prisoner's Dilemma. In: Grefenstette J. J. (ed.) Genetic Algorithms and their Applications, Proceedings of the 2nd. International Conference on Genetic Algorithms. Hillsdale, N. J.: Lawrence Erlbaum.

42. Goldberg D. E. (1988) Genetic Algorithms in Search, Optimization, and Machine Learning, Reading, Mass.: Addison-Wesley.
43. Herbrich R., Keilbach M., Graepel T., Bollmann-Sdorra P., Obermayer K. (1999) Neural networks in economics, in Brenner (1999), op. cit., pp. 169–196.
44. Ho T. H. (1996) Finite automata play repeated Prisoner's Dilemma with information processing costs, Journal of Economic Dynamics and Control **20**(1–3): 173–207.
45. Hofbauer J., Sigmund C. (1988) The Theory of Evolution and Dynamic Systems, Cambridge: C. U. P.
46. Hofbauer J., Sigmund C. (1998) Evolutionary Games and Population Dynamics, Cambridge: C. U. P.
47. Holland J. H. (1975) Adaptation in Natural and Artificial Systems. Ann Arbor: Univ. of Michigan Press. (A second edition was published in 1992: Cambridge: M.I.T. Press.)
48. Holland J. H. (1992) Genetic algorithms, Scientific American, pp. 66–72, July.
49. Holland J. H., Holyoak K. J., Nisbett R. E., Thagard P. R. (1986) Induction: Processes of Inference, Learning, and Discovery, Cambridge: M.I.T. Press.
50. Holland J. H., Miller J. H. (1991) Artificial adaptive agents in economic theory, American Economic Review: Papers and Proceedings, **81**(2): 365–370.
51. Huck S., Müller W., Strobel M. (1999) On the emergence of attitudes towards risk, in Brenner (1999), op. cit., pp. 123–144.
52. Jung Y. J., Kagel J. H., Levin D. (1994) On the existence of predatory pricing: an experimental study of reputation and entry deterrence in the chain-store game, Rand Journal of Economics, 25(1): 72–93.
53. Judd K. L. (1998) Numerical Methods in Economics, Cambridge: M.I.T. Press.
54. Lindgren K. (1991) Evolutionary phenomena in simple dynamics, pp. 295–324 in C. Langton, C. Taylor, J. Farmer, and S. Rasmussen, (ed.), Artificial Life II, Vol. 10, Reading: Addison-Wesley.
55. Mailath G. J. (1998) Do people play Nash equilibrium? Lessons from evolutionary game theory, Journal of Economic Literature, **36**(3): 1347–74.
56. Marimon R., McGrattan E., Sargent T. J. (1990) Money as a medium of exchange in an economy with artificially intelligent agents. Journal of Economic Dynamics and Control **14**: 329–373.
57. Marks, R. E. (1989a) Niche strategies: the Prisoner's Dilemma computer tournaments revisited. AGSM Working Paper 89–009. `<http://www.agsm.edu.au/~bobm/papers/niche.pdf>`
58. Marks R. E. (1989b). Breeding optimal strategies: optimal behavior for oligopolists. In J. David Schaffer (ed.), Proceedings of the Third International Conference on Genetic Algorithms, 198–207. San Mateo, CA.: Morgan Kaufmann.
59. Marks R. E. (1992a) Repeated games and finite automata, in Recent Developments in Game Theory, ed. by J. Creedy, J. Borland, and J. Eichberger, Aldershot: Edward Elgar, 43–64.
60. Marks R. E. (1992b) Breeding hybrid strategies: optimal behaviour for oligopolists, Journal of Evolutionary Economics, **2**: 17–38.
61. Marks R. E. (1998) Evolved perception and behaviour in oligopolies, Journal of Economic Dynamics and Control, **22**(8–9): 1209–1233.
62. Marks R. E., Midgley D. F., Cooper L. G. (1995) Adaptive behavior in an oligopoly, in Evolutionary Algorithms in Management Applications, ed. by J. Biethahn and V. Nissen, Berlin: Springer-Verlag, pp.225–239.

63. Marks R. E., Midgley D. F., Cooper L. G. (1998). Refining the Breeding of Hybrid Strategies, Australian Graduate School of Management Working Paper 98–017, Sydney. <http://www.agsm.edu.au/~bobm/papers/wp98017.html>
64. Marks R. E., Midgley D. F., Cooper L. G. (2000). Breeding better strategies in oligopolistic price wars, submitted to the IEEE Transactions on Evolutionary Computation, special issue on Agent-Based Modelling of Evolutionary Economic Systems, INSEAD Working Paper 2000/65/MKT, <http://www.agsm.edu.au/~bobm/papers/ieee-marks.html>
65. Marks R. E., Schnabl H. (1999) Genetic Algorithms and Neural Networks: a comparison based on the Repeated Prisoner's Dilemma, in Brenner (1999), op. cit., pp. 197–219.
66. Maynard Smith J. (1982) Evolution and the Theory of Games, Cambridge: C. U. P.
67. Megiddo N., 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.
68. Midgley D. F., Marks R. E., Cooper L. G. (1997) Breeding competitive strategies, Management Science, **43**(3): 257–275.
69. Miller J. H. (1996) The coevolution of automata in the repeated Prisoner's Dilemma. Journal of Economic Behavior and Organization, **29**: 87–112.
70. Mitchell M. (1996) An Introduction to Genetic Algorithms, Cambridge: M.I.T. Press.
71. Nachbar J. H. (1988) The evolution of cooperation revisited, mimeo., Santa Monica: RAND Corp., June.
72. Nachbar J. H. (1992) Evolution in the finitely repeated Prisoner's Dilemma, Journal of Economic Behavior and Organization **19**(3): 307–326.
73. Nelson R. R., Winter S. G. (1982) An Evolutionary Theory of Economic Change, Cambridge: Belknap Press of Harvard University Press.
74. Özyildirim S. (1997) Computing open-loop noncooperative solution in discrete dynamic games, Journal of Evolutionary Economics **7**(1): 23–40.
75. Riechmann T. (1999) Learning and behavioural stability: an economic interpretation of Genetic Algorithms, Journal of Evolutionary Economics **9**(2): 225–242.
76. Riechmann T. (2001) Genetic Algorithm learning and evolutionary games, in this volume.
77. Rubinstein A. (1986) Finite automata play the repeated Prisoner's Dilemma, Journal of Economic Theory, **39**: 83–96.
78. Rubinstein A. (1998) Modeling Bounded Rationality, Cambridge: M.I.T. Press.
79. Sargent T. J. (1993) Bounded Rationality in Macroeconomics, Oxford: O. U. P.
80. Selten R. (1978) The chain store paradox, Theory and Decision **9**: 127–59.
81. Selten R. (1991) Evolution, learning, and economic behavior, Games and Economic Behavior **3**: 3–24.
82. Slade M. E. (1992) Vancouver's gasoline-price wars: an empirical exercise in uncovering supergame strategies, Review of Economic Studies, **59**: 257–274.
83. Slade M. E. (1995) Empirical games: the oligopoly case, Canadian Journal of Economics, **28**(2): 368–402.
84. Slade M. E., Eaton B. C. (1990) Evolutionary equilibrium in market supergames, University of British Columbia Department of Economics Discussion Paper: 90-30.

85. Stanley E. A., Ashlock D., Tesfatsion L. (1994) Iterated Prisoner's Dilemma with choice and refusal of partners, in C. Langton, Artificial Life III, Vol. 17, Santa Fe Institute Studies in the Sciences of Complexity, Redwood City: Addison-Wesley, pp. 131–175.
86. Szpiro G. (1997). The emergence of risk aversion, Complexity **2**: 31–39.
87. Tesfatsion L. (1997) A trade network game with endogenous partner selection, in H. M. Amman, B. Rustem, and A. B. Whinston (eds.), Computational Approaches to Economic Problems, Dordrecht: Kluwer Academic Publishers, pp. 249–269.
88. Vallee T., Başar T. (1999) Off-line computation of Stackelberg solutions with the Genetic Algorithm, Computational Economics, **13**(3): 201–209.
89. Weibull J. W. (1995) Evolutionary Game Theory, Cambridge: M.I.T. Press.
90. Wu J., Axelrod R. (1995) How to cope with noise in the Iterated Prisoner's Dilemma, Journal of Conflict Resolution **39**(1): 183–189.
91. Yao X., Darwen P. J. (1994) An experimental study of n-person iterated prisoner's dilemma games, Informatica **18**: 435–450.