Ragupathy Venkatachalam   *Editor*

Artificial Intelligence, Learning and Computation in Economics and Finance

Springer

Springer Complexity

Springer Complexity is an interdisciplinary program publishing the best research and academic-level teaching on both fundamental and applied aspects of complex systems—cutting across all traditional disciplines of the natural and life sciences, engineering, economics, medicine, neuroscience, social and computer science.

Complex Systems are systems that comprise many interacting parts with the ability to generate a new quality of macroscopic collective behavior the manifestations of which are the spontaneous formation of distinctive temporal, spatial or functional structures. Models of such systems can be successfully mapped onto quite diverse "real-life" situations like the climate, the coherent emission of light from lasers, chemical reaction-diffusion systems, biological cellular networks, the dynamics of stock markets and of the internet, earthquake statistics and prediction, freeway traffic, the human brain, or the formation of opinions in social systems, to name just some of the popular applications.

Although their scope and methodologies overlap somewhat, one can distinguish the following main concepts and tools: self-organization, nonlinear dynamics, synergetics, turbulence, dynamical systems, catastrophes, instabilities, stochastic processes, chaos, graphs and networks, cellular automata, adaptive systems, genetic algorithms and computational intelligence.

The three major book publication platforms of the Springer Complexity program are the monograph series "Understanding Complex Systems" focusing on the various applications of complexity, the "Springer Series in Synergetics", which is devoted to the quantitative theoretical and methodological foundations, and the "Springer Briefs in Complexity" which are concise and topical working reports, case studies, surveys, essays and lecture notes of relevance to the field. In addition to the books in these two core series, the program also incorporates individual titles ranging from textbooks to major reference works.

Indexed by SCOPUS, INSPEC, zbMATH, SCImago.

Series Editors

Understanding Complex Systems

Ragupathy Venkatachalam
Editor

Artificial Intelligence, Learning and Computation in Economics and Finance

Springer

Editor
Ragupathy Venkatachalam
Institute of Management Studies
Goldsmiths, University of London
London, UK

Dedicated to
Prof. Shu-Heng Chen

Acknowledgements

Contents

Contributors

Aliyev Nihad University of Technology Sydney, Finance Department, UTS, Ultimo, Australia

Aruka Yuji Chuo University, Hachioji, Tokyo, Japan

Chie Bin-Tzong Department of Industrial Economics, Tamkang University, New Taipei City, Taiwan

He Xue-Zhong International Business School Suzhou (IBSS) Business Building (BS), South Campus, Xi'an Jiaotong-Liverpool University, Suzhou, People's Republic of China

Kao Ying-Fang Experimentation Team, Machine Learning and AI Division, Just Eat, London, UK

Kokensparger Brian J. Computer Science, Design and Journalism Department, Creighton University, Omaha, NE, USA

Kumar Sunil Mitra India Institute and Department of International Development, King's College London, London, UK

Lawless W. F. Professor of Math and Psychology, Paine College, Augusta, GA, USA

Luna Francesco Office of Innovation and Change, International Monetary Fund, Washington DC, NW, USA

Lux Thomas Department of Economics, University of Kiel, Kiel, Germany

Marks Robert Economics, University of New South Wales, Sydney, Australia

Pyka Andreas Institute of Economics, University of Hohenheim, Stuttgart, Germany

Rosser J. Barkley Jr. JMU College of Business, James Madison University, Harrisonburg, VA 22807, USA

Shen Dehua School of Finance, Nankai University, Tianjin, P.R. China

Tai Chung-Ching Southampton Business School, University of Southampton, Southampton, UK

Tesfatsion Leigh Research Professor and Professor Emerita of Economics, Courtesy Research Professor of Electrical and Computer Engineering, Iowa State University, Ames, IA, USA

Velupillai K. Vela Solna, Sweden

Venkatachalam Ragupathy Institute of Management Studies, Goldsmiths, University of London, London, UK

Vermeulen Ben IQIB, Bad Neuenahr-Ahrweiler, Germany

Wang Chen College of Management and Economics, Tianjin University, Nankai District, Tianjin, P.R. China

Wang Shu G. Department of Economics (retired), National Chengchi University, Taipei, Taiwan

Zambelli Stefano Department of Economics and Management, University of Trento, Trento, Italy

Chapter 15
Thirty-Five Years of Computational Economics

Robert Marks

Abstract This chapter describes the evolution of my work in computational economics, from 1987 to 2020, as I submitted algorithms to play a generalisation of the iterated Prisoner's Dilemma (IPD), applied a new method of machine learning (the genetic algorithm, or GA) to this, extended these techniques to trying to understand and improve on asymmetrical seller behaviour in historical oligopoly pricing, generalised this to fully fledged agent-based models (ABMs), applied the GA to exploring the best method of decision making in uncertain situations, and finally as I used simulation to search exhaustively among decision-making methods with uncertainty. This research program also led me to derive new techniques for validating the output of simulation models when the metric was *ordinal* (rather than the *interval* or *ratio* metrics usually encountered, especially in physics and engineering). At least since 1995, while undertaking this work, I have been involved with Distinguished Professor Shu-Heng Chen at many conferences, and with several publications and lectures. This chapter recounts this history.

Keywords Decision-making · Strategic oligopoly · Simulation · Validation of patterns

R. Marks (✉)
Economics, University of New South Wales, Sydney, Australia
e-mail: robert.marks@gmail.com

15.1 Tournaments to Explore the Iterated Prisoner's Dilemma

In 1987 I won the Second MIT Competitive Strategy Tournament.[1] This was a three-person pricing competition among imperfect substitutes, in which pricing low would increase sales but cut into profits. Just what my win led to, including my connection with Shu-Heng Chen, is recounted below.

My 1978 Ph.D. thesis had been a theoretical study of the interactions of markets for output, labour, and non-renewable resources in the macro economy.[2] This led to publications on energy and the environment. It did not lead directly to the work described below. From my teaching[3] and interest in environmental issues, I was, however, taken with the Prisoner's Dilemma (Marks 1998), thanks to Hardin's (1968) "The tragedy of the commons" and Schelling's (1978) *Micromotives and Macrobehavior.* As computers and computing had become more than a means of number crunching, and as the Internet had spread and enabled world-wide communication (even before the World Wide Web), the notion that it was possible to pit submitted algorithms to compete in a tournament, and that researchers from around the world could submit algorithms to compete was new.

Political scientist Robert Axelrod had been the first to use this possibility seriously, the first on-line instance of what we now call crowdsourcing. In 1979 he invited submission of algorithms via email to play an IPD game in silico, in order to see whether the temptation to defect (in the one-shot game) could be avoided in the iterated game.[4]

The Prisoner's Dilemma is an example of a strategic situation where, in the one-off interaction, mutual defection D, D is the likely outcome, whereas mutual cooperation C, C is Pareto superior (Marks 1998). That is, the payoff to each player from C, C is higher than that from D, D. Could repetition result in C, C? This was a contention of Axelrod's, who spelt this out in his 1984 book which asked how cooperation might have evolved in political systems.

His research had begun in the late 1970s, with an investigation of the emergence of cooperative behaviour and social norms in Hobbesian societies. In Axelrod's first tournament, Anatol Rapoport submitted a simple algorithm: start by cooperating and then mimic the other player's action in the previous round. This strategy, now known as Tit for Tat, was very successful, outperforming all other strategies but for Always Defect (Axelrod 1980a).

---

[1] I thank my marketing colleague at the Australian Graduate School of Management (AGSM), John Roberts, a recent MIT graduate, for alerting me to the Tournaments.

[2] It has recently been republished, with a new preface by Robert Solow (Marks 2018 and https://www.agsm.edu.au/bobm/papers/2018prefaces.pdf).

[3] From 1977 to 1982 I taught AGSM 713 Management in Society on the MBA Program. See the 1982 outline at https://www.agsm.edu.au/bobm/teaching/MinS82.pdf.

[4] Fourteen entries were submitted (by email) from three countries and five disciplines (Axelrod 1980a).

Axelrod announced the outcome of Tit for Tat in his first tournament and announced a second. Despite entrants' knowledge that Tit for Tat was very likely a competitor, no other submitted strategies outperformed it in the second tournament.[5] It was robust. But it is not always a winner: how well it does depends on the collection of competing algorithms in the round robin.

As is now widely known, Axelrod's tournaments revealed that one very simple strategy is robust in the IPD: Rapoport's Tit for Tat, which cooperates on the first round of the iterated game, and then mimics its opponent. Tit for Tat can be characterised as nice (start off cooperating), but easily provoked (defect after a single defect by its opponent), forgiving (a single cooperate by its opponent leads to its cooperating), and easily identified.

I found Axelrod's work fascinating, and, as an early user of the Internet myself,[6] I was intrigued by his use of it to research social interactions.

15.2   Generalising Axelrod's Iterated Prisoner's Dilemma Tournaments

In November 1984, MIT marketing professors Pete Fader and John Hauser decided to see whether Tit for Tat could be generalised in a three-person game with continuous actions. They ran one and then another competitive strategy tournaments (Fader and Hauser 1988). The Tournaments modelled competitors' decisions to price their outputs, which were imperfect substitutes: like a generalised Prisoner's Dilemma, without collusion among sellers, in a once-off interaction, pricing low (that is, defecting) is likely to be the outcome, at some cost to their profits, whereas pricing high (that is, cooperating) would result in higher profits for all, if no-one priced low.

I had been following Axelrod's tournaments and, given the chance, I submitted a FORTRAN algorithm for the First Tournament, but it was trigger-happy: if any of the players priced low (equivalent to defecting), then it would price low and keep pricing low. It was not quite Always Defect, but close to it.

The First Tournament was promising, but the organisers decided that it was flawed: their set-up had been too amenable to gaming, and so ran the Second Tournament with a different profit function (additive, not multiplicative[7]). Before the September 1986 deadline for this, John Roberts and I ran a local tournament, the AGSM Double Auction Computer Tournament, with a prize of $500, using the software from the Second MIT Tournament. The winner was Tony Haig (University

---

[5] There were 62 entrants from six countries and eight disciplines, recruited through announcements in journals for users of small computers (Axelrod 1980b).

[6] My web site has not changed in appearance at all since June 1996. The earliest record of my use of the Internet is ten years earlier, on 5 June 1986, at https://groups.google.com/forum/#!topic/net. text/rvi_FPmQBTE (Thirty-five years later I still use the Free Software Foundation's groff package for my text-processing needs: comes with all Macintoshes).

[7] Marks (1992b) presents these two profit functions.

of Western Australia). I then submitted an algorithm to the MIT Tournament; my algorithm won.

How did my algorithm win, people would ask. Some insights into the IPD, I would answer, but also luck: the ecology of the Tournament was comprised of all the other algorithms submitted; this meant that there was no easy way of coming up with a winning algorithm, especially since the other algorithms were unknown at the start. But, given that all algorithms were attempting to maximise their profits, a Nash outcome was likely.

The MIT Competitive Strategy Tournaments presented players with market conditions which generalised from the RPD. The environment resembled a more complicated version of the Iterated Prisoner's Dilemma, but there were more than two players, and there were more than two moves, since the algorithms chose a price in a range, while knowing the previous prices chosen by all players, but without collusion over the prices in this period.

As with the most productive research, there was an unanswered question: how had my algorithm won the 1987 MIT Tournament. It nagged at me.

15.3   Enter Machine Learning: The Genetic Algorithm

Mathematically, the problem of generating winning strategies in these interactions is equivalent to solving a multi-dimensional, non-linear optimisation with many local optima. In biological-evolution terms, it is equivalent to selecting for "fitness."

Indeed, in a footnote, Cohen and Axelrod (1984, p .40) suggest that:

> One possible solution may lie in employing an analogue of the adaptive process used in a pool of genes to become increasingly more fit in a complex environment. A promising effort to convert the main characteristics of this process to an heuristic algorithm is given by Holland (1975 [1992]). This algorithm has had some striking preliminary success in the heuristic exploration of arbitrary high dimensionality nonlinear functions.

In 1987 I read a colleague's copy of *Induction*, by Holland et al. (1986), which made a passing reference to some more recent work of Axelrod's which found strategies for playing the IPD Tournament which resembled Tit for Tat.[8]

Perhaps this work would shed light on my winning algorithm in the MIT Tournament. I emailed Axelrod and he replied with a program (written in Pascal-VS by Stephanie Forrest, his R.A.) using a version of John Holland's Genetic Algorithm, written in C. In an early example of a trans-Pacific code-sharing, I received the C

---

[8] I thank another AGSM colleague, applied psychologist Bob Wood, who brought the book back to Sydney from California and fortuitously lent it to me. John's and Bob's suggestions were very much helped by the cross-disciplinary atmosphere of the AGSM, sadly now lost after it was destroyed by incoming Vice Chancellor, Fred Hilmer.

code for the GA from the the U.S. Naval Research Laboratories in Washington DC and compiled it on our Unix machine.[9]

Mitchell and Forrest (1994) report what Axelrod did:

> Axelrod (1987) performed a series of experiments to see if a GA could evolve strategies to play [the IPD] game successfully. Strategies were encoded as look- up tables, with each entry (C or D) being the action to be taken given the outcomes of three previous turns.
>
> In Axelrod's first experiment, the evolving strategies were played against eight human-designed strategies, and the fitness of an evolving strategy was a weighted average of the scores against each of the eight fixed strategies. Most of the strategies that evolved were similar to TIT FOR TAT, having many of the properties that make TIT FOR TAT successful. Strikingly, the GA occasionally found strategies that scored substantially higher than TIT FOR TAT.
>
> To study the effects of a dynamic environment, Axelrod carried out another experiment in which the fitness was determined by allowing the strategies in the population to play with each other rather than with the fixed set of eight strategies. The environment changes from generation to generation because the strategies themselves are evolving. At each generation, each strategy played an IPD with the other members of the population, and its fitness was the average score over all these games. In this second set of experiments, Axelrod observed the following phenomenon. The GA initially evolves uncooperative strategies, because strategies that tend to cooperate early on do not find reciprocation among their fellow population members and thus tend to die out. But after about 10–20 generations, the trend starts to reverse: the GA discovers strategies that reciprocate cooperation and that punish defection (i.e., variants of TIT FOR TAT). These strategies do well with each other and are not completely defeated by other strategies, as were the initial cooperative strategies. The reciprocators score better than average, so they spread in the population, resulting in more and more cooperation and increasing fitness.

What I learnt was that, following up his own suggestion, Axelrod used Holland's Genetic Algorithm (GA) to "breed" strategies in the two-person IPD game (Axelrod 1987). He reported that the GA evolved strategy populations whose median member was just as successful as Tit for Tat, whom they closely resembled. I proceeded to teach myself Pascal-VS (I had learnt FORTRAN IV as an undergraduate and C later) and wrote a paper (Marks 1989a)[10] that attempted to replicate what Axelrod had done. In particular, I "bred" strategies playing in six distinct niches: (a) a niche of Always Defect, (b) a niche of Always Cooperate, (c) a niche of Tit for Tat, (d) a 5-rule niche described by Axelrod (1984, p. 199) that approximates his second IPD tournament, (e) an 8-rule niche from Axelrod (1987) that is a better approximation, and (f) the 5-rule Axelrod niche but with "noise" added, to simulate Nalebuff's (1987) IPD game with imperfect information. I experimented with strategies of different

---

[9] John Grefenstette's GENESIS 4.5. In doing so, I was probably the first economist in Australia to compile a C program obtained via the internet: I thank the first AGSM guru, Iain Johnstone, for choosing Unix as our OS.

[10] An earlier version of Marks (1989a) was presented at the 1988 Australasian meetings of the Econometric Society, Canberra, 30 August 1988, under the title, "Breeding hybrid strategies: the Prisoner's Dilemma computer tournaments revisited" (See *Econometrica* 57: 240, 1989, for the program).

complexities (depth of memory of past plays).[11] Marks (1989a) was the one of first papers presented by an economist.

using the genetic algorithm.[12]

The six niches in Marks (1989a) are static. It was suggested to me by David Schaffer (I think) that the this work would be more interesting if the strategies (agents) competed not against static niches but against themselves, as they evolved.[13] In Marks (1989b) I modelled this process.[14] What to call it? Unaware.

that the biologists had preceded me, I dubbed it "bootstrapping evolution;" but biologists (Ehrlich and Raven 1964) had called it "coevolution." My experiments showed that coevolution would result in convergence to cooperative behaviour in several strategic interations (the simple IPD, an extended IPD, and a simple version of the second MIT Competitive Strategy Tournament).

The genetic algorithm turned out to be a powerful computational method for searching for optima in spaces that were not amenable to calculus-based solutions, such as the oligopolistic pricing of the MIT Tournaments.[15]

With David Midgley and Lee Cooper, I wrote several more papers using the GA to examine strategies in various strategic interactions, usually market-related, and latterly with heterogeneous agents (with multi-population GAs) (Marks 1989b, 1992b, 2002a). We wrote several papers using the GA to explore optimal oligopolistic pricing (including Midgley et al. 1997, which, Shu-Heng Chen tells me, inspired Chen and Ni 2000).

An issue arose when I attempted to compare the weekly prices obtained from simulating the asymmetric agents derived using the GA with the historical data.[16] How to measure a distance between the dynamic historical data and different sets of simulated agents interacting? This can be thought of as a means of verifying the simulation models derived from the GA. Trying to answer this question has led to a series of papers and presentations, including Marks (2013) and culminating in Marks (2019). This has turned out to be a fruitful line of research.

So, from my win in the MIT Tournament, to my use of Holland's pioneering GA, to exploring historical oligopolistic reactions with David Midgley and other marketing colleagues, to the puzzle of verifying simulation output which is patterns (ordinal metrics) rather than simple (interval or ratio) metrics, I believe I have contributed to the theory and practice of simulation.

---

[11] Note that Tit for Tat is a one-round-memory strategy: no need for a deeper memory.

[12] Engineers had used the GA as a technique for optimising static functions. This was different.

[13] Another paper, Marks (1989b), was presented at the Allied Social Science Associations meetings under the auspices of the Econometric Society, on 30 December 1988, in New York. John Miller, a student at Michigan, wrote his Ph.D. on coevolution of automata playing the IPD; he published Miller (1986, 1989), and eventually Miller (1996). Other early economics papers to use GAs are Marimon et al. (1989), Miller and Andreoni (1990a, b), Holland and Miller (1991), Arifovic (1994).

[14] This paper was first presented at the North American Winter Meeting of the Econometric Society, New York, on 30 December 1988 (See *Econometrica* 57: 757, 1989, for the program).

[15] Oligopolies are markets with small numbers of sellers, each of whose actions in general affect the outcomes for other sellers, as well as for itself.

[16] See Marks (1992a) for discussion of the automata derived by the GA, and their simulations.

15.4   From Genetic Algorithms to Agent-Based Models

Another line of research has been agent-based (AB) models: When the individual agents modelled by the GA are competing against each other, the GA is modelling the process of coevolution. GAs were originally used as means of seeking optimal solutions to static problems; Marks (1989b) and others adapted them to seek solutions of coevolutionary strategic problems, such as the IPD and oligopolies with asymmetric players, where the fitness of an agent depends on the individual agents' actions, that is, the state of the whole population of agents.

When the interacting players face identical payoff sets and choose from identical action sets, a single population is satisfactory, since the GA processes that model learning among the individuals and between generations of the population are focused to the same end: faced with the same state of the interaction, any of the players would behave identically, and fitness is average (or discounted) profit.

A single population was acceptable when the players were not differentiated and when the flow of information from parents to offspring at the genotype level was not an issue, but when the players are modelling heterogeneous actors—in realistic coevolution, for instance—each player requires a separate population, not least to prevent the modelling of illegally collusive, extra-market transfers of information. This is discussed in Marks (2012).

For instance, Marks et al. (1999) develops an oligopolistic model with three (or four) interacting asymmetric sellers. The general model is still one of using a GA to search for automata (or mappings from past marketing actions to future actions for each seller, but now we use separate populations of agents, one population for each of the asymmetric sellers. This paper is one of the first to use the GA in an agent-based model (separate populations in the GA).

From separate populations of asymmetric agents, it is a simple step to develop agent-based models (ABMs). With David Midgley (see Midgley et al. 2007), we discuss AB models in marketing, specifically the complex interactions among three types of agents—consumers, retailers, and manufacturers—that lead to market and economic outcomes such as consumer satisfaction, and retailer and manufacturer profits. We argue that AB modelling is more appropriate than previous methods, but that it requires "assurance:" verification and validation of the model. The paper discusses this at length.

The work on AB models in economics led to an invitation from Peter McBurney and the editors of the journal to edit a special issue of *The Knowledge Engineering Review*, on agent-based computational economics. I asked Nick Vriend to be a co-editor (promising him little extra work), and we approached economists and financial economists to contribute. The special issue included papers from many of the pioneers in ACE.[17] One of the contributors was Shu- Heng Chen.

---

[17] June 2012; Editors: Robert Marks and Nick Vriend. Papers by: Anufriev and Hommes (2012), Arifivic and Ledyard (2012), Cheng Chan and Du (2012), Fagioli and Roventini (2012), Ladley (2012), Marks (2012), Page (2012), Richiardi (2012), Wilhite and Fong (2012).

From 1997 to 2010 I was the General Editor of the *Australian Journal of Management*, but, apart from three editorials (Marks 2002b, 2003, 2006b) and the invited papers Hailu and Schilizzi (2004) and Byde (2006), I did not focus on the topics of this chapter.

15.5   The Issue of the Best Risk Profile for Risky Decision Making

From my graduate studies, I had developed an interest in Decision Analysis, the use of mathematical tools to examine and prescribe decision making under uncertainty (Howard 1968). Given my publications using the GA, and my visits to the Santa Fe Institute in 1993 and earlier, in 1995 I was asked by John Casti to review a submission by George Szpiro to *Complexity* which used the GA to explore the best risk profile for a decision maker faced with uncertainty. To my MBA students, I had always taught that a slightly risk-averse profile was best: too risk-averse and the decision maker would pass up risky prospects ("nothing ventured, nothing gained"), but too risk preferring and the agent would sooner or later bankrupt itself.

After I gave him some guidance about how the GA found its solutions (not so frequently on the boundaries of the feasible space), Szpiro found that slightly risk averse was optimal, at least according to the simulation of the GA, and the paper was published (Szpiro 1997).[18]

But Szpiro had used an indirect method, despite my suggestion as referee that he explore using the GA to search more directly in the risk-profile space of utility functions.[19] Instead, Szpiro's agents are characterised by a variable (beta) which determines how many shares (in a simple market) they buy (or sell) in any period. For any model's parameters (see Szpiro's paper), there is a threshold dividend rate which determines whether agents should buy or sell. His GA model maximises each agent's wealth as a function of the dividend rate, and examines the evolved beta as uncertainty rises. He argues that betas that jump from zero to maximum at the threshold dividend rate reflect risk-neutral agents (and this occurs with no uncertainty), but that betas that respond to an increasing dividend rate more gradually (with uncertainty) from zero to maximum exhibit risk aversion. His agents cannot exhibit risk preferring, they are only risk neutral or risk averse. He argues that his model produces optimised agents whose risk aversion rises as uncertainty rises.

---

[18] Szpiro was later invited by Shu-Heng Chen to publish a chapter (Szpiro 2002) in the same volume that also included a chapter of mine (Marks 2002a). In later correspondence, after I had sent him a copy of Marks (2016), Szpiro told me that he was no longer involved in this line of research.

[19] Indeed, in his Footnote 6, Szpiro acknowledged my suggestion as referee to model agent automata as utility functions with explicit risk profiles, as an alternative to his approach (They could then, he allowed, be characterised as risk preferring, as well as risk neutral or risk averse).

I decided that one day I would follow my own advice and use the GA to try to confirm Szpiro's conclusion. I did try, but ultimately found a different result: risk neutrality.

Although I had previously (Marks 1989a and other papers) used a GA written in C, I decided to use the higher-level language NetLogo, with its support and graphic capabilities for the exploration of the best risk profile. Nigel Gilbert had implemented a genetic algorithm in NetLogo. The papers that followed presented the convergence results graphically.[20]

The title of the paper (Marks 2015a) I presented in London in 2014 was "Learning to be risk averse?" which reflected my persisting belief that Szpiro's results would hold with my more direct method of exploring the risk profile space for the best profile. Later that year I presented a revised paper (Marks 2015b) in Singapore, entitled "Searching for agents' best risk profiles," in which I concluded that for some utility functions (the wealth-dependent Constant Relative Risk Aversion function and the Dual-Risk-Profile function[21] from Prospect Theory) the best functions could be shown to be slightly risk averse, but for the Constant Absolute Risk Aversion function risk neutral was best.

In 2016, in a paper entitled "Risk neutral is best for risky decision making," (Marks 2016) I concluded that risk-neutral decision makers—whichever utility function was modelled—outperformed others when agents successively chose among three lotteries with randomly allocated probabilities and outcomes (two per lottery).[22]

Nonetheless, I realised that the simulation experiments in these papers were not very clear for readers, who might find the concept of the genetic algorithm difficult to be convinced by and the statistical arguments unfamiliar. That the GA is searching for an optimum risk coefficient at a flattish apex also clouds the findings. Furthermore, NetLogo has its uses, but it lacks a reputation for exact scientific work.

Although I had come to this line of enquiry by a quite independent route, as described above, I found that my path had crossed with Shu-Heng Chen's: his paper (Chen and Huang 2008) also looks at this question of the risk profile of decision makers under uncertainty; he had come to it from his work in developing agent-based artificial stock markets.

---

[20] I presented my evolving findings at the Complex Systems Research Summer School 2007 at Charles Sturt University in NSW and the 26th Australasian Economic Theory Workshop 2008 at Bond University in Queensland.

[21] For the DRP utility function, the risk profile is coded with two parameters, unlike the other two utility functions for which the search is in a single parameter space. See Sect. 3.3 in Marks (2020).

[22] The 2016 research used the NetLogo GA with the experimental setup: a population of 100 agents, each of which has a average winnings or a cumulative level of wealth, based on its risk profile and the successive outcomes of its 1000 choices among the lotteries; Each lottery is randomly constructed: the two payoffs ("prizes") are uniformly chosen in the interval $[-\$100, +\$100]$, and the probability is chosen uniformly from [0,1]; Each agent faces 1000 lottery choices, and its cumulative winnings is that agent's "fitness" for the GA. The processes are stochastic. For each model we perform a number of Monte Carlo simulation runs to obtain sufficient data to analyse the results.

15.6 Simulation, Not Optimisation

Is risk neutrality best for making decisions under uncertainty? That is, for choosing amongst lotteries, for each of which the possible prizes and the probabilities of those outcomes were known? The analysis using the GA suggested risk neutral was best, but the results were not convincing. I decided I would search the risk-profile parameter space exhaustively, using simulation rather than searching the space using the GA. Perhaps not as elegant, but I hoped for a clear answer.

I decided not to continue using NetLogo (and certainly not to start coding in C); instead, I taught myself R.[23] Using R I wrote code to simulate exploration in the risk-profile space for eight methods of choosing among risky outcomes.[24] The paper that resulted (Marks 2020) demonstrates very clearly that risk-neutral decision makers outperform non-risk averse decision makers, whether CARA or CRRA or DRP utility functions (Marks 2020).

Two of the eight methods for choosing the best of the eight lotteries, max- min and max–max, are straightforward: max–min chooses the lottery with the highest minimum possible payoff, while max–max chooses the lottery with the highest maximum. Think of them as the pessimist's and optimist's methods, respectively. But the simulations revealed a surprising result which demands further research: the (pessimist's method) max–min, is almost twice as profitable as the (optimist's method) max–max. Just why is not yet clear. But even max–min is only a fraction as profitable as the risk-neutral Expected-Value method.

The question of which decision-making method gives the highest payoff in cases of uncertainty (where the possible pay-offs and their probabilities are known) is not, in general, amenable to closed-form solution. The answer is clearly that risk-neutral methods are best, as exemplified by the Expected Value method. I believe that exploration of other experiments in decision making under uncertainty (with complete information) will confirm the generality of this conclusion. Will relaxing our assumptions of complete information about possible outcomes and their probabilities result in different conclusions? This awaits further work. Marks (2020) was chosen for the Award for Best Paper at DECON 2019: the International Conference on Decision Economics, held in Ávila, Spain, on June 26−28, 2019.

---

[23] R is a programming language and free software environment for statistical computing and graphics supported by the R Foundation for Statistical Computing. It is widely used among statisticians and data miners for developing statistical software and for data analysis. As of January 2021, R ranks 9th in the TIOBE index, a measure of popularity of programming languages.

[24] See the R code at https://www.agsm.edu.au/bobm/papers/riskmethods.r. The R code sets up eight six-prize lotteries, with the prizes chosen randomly between +$10 and −$10. For each of eight methods of choosing one of the eight lotteries, there were 10,000 repetitions. See Marks (2020) for more details.

15.7   Connection with Shu-Heng Chen

I believe I first met Shu-Heng Chen at the 1st International Conference on Computational Economics, held at the University of Texas, Austin, May 21−24, 1995. At my invitation, in 2000 he came to Sydney and gave a presentation, "Genetic programming in the agent-based modeling of artificial stock markets," at the AGSM, UNSW, on February 15, 2000.

In 2001 he invited me to submit two papers for volumes that he was editing. The papers are Drake and Marks (2002) and Marks (2002a), published in *Genetic Algorithms and Genetic Programming in Computational Finance* and *Evolutionary Computation in Economics and Finance*, respectively. In return, some years later, as guest editor with Nick Vriend of the *The Knowledge Engineering Review*, a special issue on agent-based computational economics (ACE) (Marks and Vriend 2012),[25] I invited Professor Chen to contribute a paper, Chen et al. (2012).[26]

In July 2005 we were both in Bielefeld (at the the International Workshop on Agent-Based Models for Economic Policy Design, ACEPOL05) when Akira Nakatame invited some of us (including Shu-Heng and me, and other contributors to this volume) to become members of the Editorial Board of the new journal, the *Journal of Economic Interaction and Coordination*, the official journal of the new association, the Society of Economic Science with Heterogeneous Agents.[27] We have remained on the Editorial Board since then, and Shu-Heng is now co-editor.

At Professor Chen's invitation, in October 2005 I presented the Fourth Herbert Simon Seminar Series, on Agent-Based Computational Economics and Market Design, at the Artificial Intelligence Economics Research Center, Department of Economics, National Chengchi University, Taipei, and the National Kaohsiung University of Applied Sciences, Taiwan, on October 23−28.[28] These lectures were published as Marks (2006a), which, Shu-Heng Chen tells me, inspired his later research on agent-based modeling of lottery markets (Chen and chie 2008). At his invitation in 2013, I delivered three tutorials on "Validating Simulation Models, and Multi-agent Systems in the Social Sciences" at the Computational Finance and Economics Technical Committee (CFETC) of the Computational Intelligence Society (CIS) of the Institute of Electrical and Electronic Engineers, Inc. (IEEE) in Singapore in April 2013.[29]

---

[25] I thank Peter McBurney and the editors of the *KER* for this opportunity.

[26] With over 270 cites in Google Scholar, this is Professor Chen's second most cited paper.

[27] In November 2005 *JEIC*'s Editors in Chief: Akira Namatame, Tomas Lux, and Robert Axtell; Editorial Advisory Board: Mauro Gallegati, Masanao Aoki, and Alan Kirman; Editorial Board: Hideaki Aoyama, Yuji Aruka, Damien Challet, Shu-Heng Chen, Silvano Cinotti, Robin Cowan, Giorgio Fagiolo, David Green, Shouta Hattori, Dirk Helbing, Cars Hommes, Neil Johnson, Taisei Kaizoji, Sheri Markose, Matteo Marsili, Robert Marks, Denis Phan, Massimo Ricottilli, Erico Scalas, Frank Schweitzer, Didier Sornette, Hideki Takayasu, Leigh Tesfatsion, Zoltan Toroczkai, Bernard Walliser, David Wolpert, and Hirosi Yoshikawa.

[28] See my web page for these lectures, at https://www.agsm.edu.au/bobm/teaching/Taiwan.html.

[29] These tutorials are at https://www.agsm.edu.au/bobm/papers/IEEESingapore/tut01pr-3.pdf and https://www.agsm.edu.au/bobm/papers/IEEESingapore/tut02pr-3.pdf.

At Professor Chen's urging, I submitted a paper to DECON 2019: the International Conference on Decision Economics, held in Ávila, Spain, on June 26−28, 2019. That paper, Marks (2020), was judged the Best Paper at the Conference. Professor Chen was one of the conference organisers, and co-edited the conference proceedings.

Professor Chen's simulation study (Chen and Huang 2008) examines the survival dynamics of investors with different risk preferences in an agent-based, multi-asset, artificial stock market. They find that investors' survival is closely related to their risk preferences. Examining eight possible risk profiles, they find that only CRRA investors with relative risk aversion coefficients close to unity (that is, log-utility agents) survive in the long run (up to 500 simulations). This is not consistent with my findings in Marks (2020), whence I would expect risk-neutral agents to survive longer; these results remain to be reconciled.

My records reveal that Professor Chen and I have met at conferences in Austin, Texas (1995), Geneva, Switzerland (1996), Orlando, Florida (1999), Lake Arrowhead, California (2003), Kyoto (2004 and 2006), Bielefeld, Germany (2005 and 2010), Sydney (2009), Singapore (2013 and 2014), London (2014), and Ávila, Spain (2019).[30] Given our peripatetic movements (at least before COVID), there might well have been other occasions.

15.8   Conclusion

I have outlined how my research has evolved since 1987, the techniques I have used, and the results I have obtained. Moving from simple algorithms submitted to an in silico tournament of oligopolistic price competitors, to early adoption of the genetic algorithm in the Iterated Prisoner's Dilemma, and then to examining historical, asymmetric oligopolistic pricing decisions with multi-population agent- based models, and then to using the GA to search for the best risk profile for decision

---

[30] Respectively, the 1st International Conference on Computing in Economics and Finance (CEF 1995), Austin, May 21−24; the 2nd International Conference of Computing in Economics and Finance (CEF 1996), Geneva June 26−28; GECCO-99: the Genetic and Evolutionary Computation Conference, Orlando, July 13−17; the 2nd Lake Arrowhead Conference on Human Complex Systems, March 19−22; the 3rd International Workshop on Agent-based Approaches in Economic and Social Complex Systems (AESCS'04), Kyoto, May 27−29; the First World Congress on Social Simulation, Kyoto, August 21−25; the International Workshop on Agent-Based Models for Economic Policy Design (ACEPOL05), Bielefeld, June 30−July 2; and Advances in Agent-Based Computational Economics (ADACE 2010), Bielefeld, July 5−7; the 15th International Conference on Computing in Economics and Finance, UTS, July 15−17; Towards Large Multiscale Simulations of Complex Socio-Economic Systems of Heterogeneous Interacting Agents, the Society for Economic Sciences of Heterogeneous Interacting Agents (ESHIA), Nanyang, November 18−19; the 18th Asia Pacific Symposium on Intelligent and Evolutionary Systems (IES'2014); the IEEE conference on Computational Intelligence for Finance Engineering & Economics (CIFEr'2014), London, March 27−28; and DECON 2019: the International Conference on Decision Economics, at the 17th International Conference on Practical Applications of Agents and Multi-agent Systems (PAAMS), Universidad de Salamanca, Ávila, Spain, June 26−28. As well as his presentation at UNSW Sydney in February 2000.

makers facing uncertainty, to exhaustive simulations of this issue, I have encountered my friend Shu-Heng Chen many times, and greatly benefitted from his advice, invitations, and suggestions.

If there is a moral for young researchers in all this, it is to follow one's nose or one's hunch. And listen to your students: several times students told me about research I was unaware of.[31] It also helps to talk with researchers in other disciples—marketing, computer science, political science, applied psychology—and I was fortunate to have such colleagues, including Shu-Heng Chen. The resulting papers might not appear in *Econometrica*, but it is now possible to publish them on-line, and there are also a growing number of cross-disciplinary outlets, as the Internet affects the core disciplines. Above all, use your imagination—technical skills are all very well (necessary) but hardly sufficient: a desire to answer previously unanswered questions (or perhaps to answer questions that have never previously been asked) requires a bold imagination.

References

Anufriev M, Hommes C (2012) Evolution of market heuristics. Knowl Eng Rev 27(2):255–271

Arifovic J (1994) Genetic algorithm learning and the cobweb model. J Econ Dyn Control 18:3–28

Arifovic J, Ledyard J (2012) Individual evolutionary learning with many agents. Knowl Eng Rev 27(2):239–254

Axelrod R (1980a) Effective choice in the Prisoner's Dilemma. J Conflict Resolut 24(1):3–25

Axelrod R (1980b) More effective choice in the Prisoner's dilemma. J Conflict Resolut 24(3):379–403

Axelrod R (1984) The evolution of cooperation. Basic Books, New York

Axelrod R (1987) The evolution of strategies in the iterated Prisoner's dilemma. In Davis LD (ed) Genetic algorithms and simulated annealing. Research notes in artificial intelligence. Morgan Kaufmann, Los Altos, CA, pp 32−41

Byde A (2006) Applying evolutionary search to a parametric family of auction mechanisms. Austr J Manag 31(1):1−16

Chen S-H (ed) (2002a) Evolutionary computation in economics and finance, studies in fuzziness and soft computing, vol 100, ed. by J Kacprzyk. Springer, New York

Chen S-H (ed) (2002b) Genetic algorithms and genetic programming in computational finance. Kluwer, Boston

Chen S-H, Chie B-C (2008) Lottery markets design, micro-structure, and macro-behavior: an ACE approach. J Econ Behav Organ 67(2):463–480

Chen S-H, Huang Y-C (2008) Risk preference, forecasting accuracy and survival dynamics: simulation based on a multi-asset agent-based artificial stock market. J Econ Behav Organ 67(3–4):702–717

Chen S-H, Ni C-C (2000) Simulating the ecology of oligopoly games with genetic algorithms. Knowl Inform Syst: Int J 2:310–339

Chen S-H, Chang C-L, Du Y-R (2012) Agent-based economic models and econometrics. Knowl Eng Rev 27(2):187–219

Cohen MD, Axelrod R (1984) Coping with complexity: the adaptive value of changing utility. Am Econ Rev 74:30–42

---

[31] Such as Hofstadter (1983), which is relevant here, and alerted many to the Axelrod tournaments.

Drake AE, Marks RE (2002) Genetic algorithms in economics and finance: forecasting stock market prices and foreign exchange—a review, in Chen (2002b), Chapter 2, pp 29−54. https://doi.org/10.1007/978-1-4615-0835-9_2.https://www.agsm.edu.au/bobm/papers/drake1.pdf

Ehrlich PR, Raven PH (1964) Butterflies and plants: a study in coevolution. Evolution 18:586−608

Fader PS, Hauser JR (1988) Implicit coalition in a generalised Prisoner's dilemma. J Conflict Resolut 32:553–582

Fagiolo G, Roventini A (2012) On the scientific status of economic policy: a tale of alternative paradigms. Knowl Eng Rev 27(2):163–185

Hailu A, Schilizzi S (2004) Are auctions more efficient than fixed price schemes when bidders learn? Aust J Manag 29:147–168

Hardin G (1968) The tragedy of the commons. Science 162(3859):1243–1248

Hofstadter D (1983) Computer tournaments of the Prisoner's Dilemma suggest how cooperation evolves. Sci Am, May 1983:6−26; reprinted in his Metamagical Themas: questing for the essence of mind and pattern. Basic Books, New York, pp 715−738

Holland JH, Miller JH (1991) Artificial adaptive agents in economic theory, Am Econ Revi 81:365–370

Holland JH (1992) Adaptation in natural and artificial systems. MIT Press, Cambridge, 2nd edn. (1975)

Holland JH, Holyoak KJ, Nisbett RE, Thagard PR (1986) Induction: processes of inference, learning, and discovery. MIT Press, Cambridge

Howard RA (1968) The foundations of decision analysis. IEEE Trans Syst Sci Cybern ssc-4:211−219

Ladley D (2012) Zero intelligence in economics and finance. Knowl Eng Rev 27(2):273–286

Marimon R, McGrattan E, Sargent TJ (1989) Money as a medium of exchange in an economy with artificially intelligent agents. J Econ Dyn Control 14:329–373

Marks RE (1989a) Niche strategies: the Prisoner's dilemma computer tournaments revisited, AGSM Working Paper 89–009 (April). An earlier version of this paper had been presented at the 1988 Australasian meetings of the Econometric Society, ANU, Canberra, 30 August 1988, under the title, "Breeding hybrid strategies: the Prisoner's Dilemma computer tournaments revisited" (See *Econometrica* 57:240, 1989, for the program). https://www.agsm.edu.au/bobm/papers/niche.pdf

Marks RE (1989b) Breeding hybrid strategies: optimal behavior for oligopolists. In: Schaffer JD (ed) Proceedings of the third international conference on genetic algorithms, George Mason University, June 4−7, 1989. Morgan Kaufmann Publisher, San Mateo, CA, pp 198−207. This paper had been presented at the Allied Social Science Associations meetings under the auspices of the Econometric Society, 30 December 1988, in New York (See *Econometrica* 57:757, 1989, for the program). https://www.agsm.edu.au/bobm/papers/1989P3ICGA.pdf

Marks RE (1992a) Repeated games and finite automata. In: Creedy J, Eichberger J, Borland J (eds) Recent developments in game theory. Edward Elgar, London, pp 43−64. https://www.agsm.edu.au/bobm/papers/creedychapter.pdf

Marks RE (1992b) Breeding hybrid strategies: optimal behaviour for oligopolists. J Evol Econ 2:17−38. https://doi.org/10.1007/BF01196459. An earlier version of this paper was presented at the Winter meetings of the American Economics Association in New York on 30 December 1988 under the auspices of the Econometric Society. https://www.agsm.edu.au/bobm/papers/JEE1992.pdf

Marks RE (1998) Competition and common propert. AGSM Working Paper 98−003 AGSM-UNSW Sydney, February 11. https://www.agsm.edu.au/bobm/papers/ccp.pdf

Marks RE (2002a) Playing games with genetic algorithms, Chapter 2, pp 31−44, in Chen (2002a). https://doi.org/10.1007/978-3-7908-1784-3_2. https://www.agsm.edu.au/bobm/papers/playing_games-shu.pdf

Marks RE (2002b) Editorial: simulating economics. Austr J Manag 27:i−vi, June. https://www.agsm.edu.au/bobm/editorials/0206edit.html

Marks RE (2003) Editorial: models rule. Austr J Manag 28:i−v, June. https://www.agsm.edu.au/bobm/editorials/0306edit.html

Marks RE (2006a) Market design using agent-based models, Chapter 27. In: Tesfatsion L, Judd KL (eds) Handbook of computational economics, Volume 2: Agent-based modeling. Elsevier Science, Amsterdam, pp 1339−1380. https://doi.org/10.1016/S1574-0021(05)02027-7. https://www.agsm.edu.au/bobm/teaching/SimSS/Marks.finalrev.pdf

Marks RE (2006b) Editorial: the threesome: TLA, BCP, and AFP. Austr J Manag 31:i−iv, June. https://www.agsm.edu.au/bobm/editorials/0606edit.html

Marks RE (2007) Validating simulation models: a general framework and four applied examples. Comput Econ 30(3):265−290, October. https://doi.org/10.1007/s10614-007-9101-7. https://www.agsm.edu.au/bobm/papers/s1.pdf

Marks RE (2012) Analysis and synthesis: multi-agent systems in the social sciences. Knowl Eng Rev 27(2):123−136. https://doi.org/10.1017/S0269888912000094. https://www.agsm.edu.au/bobm/papers/KER2012-pap.pdf

Marks RE (2013) Validation and model selection: three measures compared. Comp Econ 2:41−61, May. https://doi.org/10.7564/13-COEC10. https://www.agsm.edu.au/bobm/papers/parispaper.pdf

Marks RE (2015a) Learning to be risk averse? In: Serguieva A, Maringer D, Palade V, Almeida RJ (eds) Proceedings of the 2014 IEEE computational intelligence for finance engineering & economics (CIFEr), London, March 28−29. IEEE Computational Intelligence Society, USA, pp 1075−1079. https://doi.org/10.1109/CIFEr.2014.6924096. https://www.agsm.edu.au/bobm/papers/IEEEmarks2014.pdf

Marks RE (2015b) Searching for agents' best risk profiles. In: Handa H, Ishibuchi M, Ong Y-S, Tan K-C (eds) The proceedings of the 18th Asia Pacific symposium on intelligent and evolutionary systems (IES'2014), Chapter 24, vol 1. Springer, pp. 297−309. https://doi.org/10.1007/978-3-319-13359-1_24. https://www.agsm.edu.au/bobm/papers/marksIES2014.pdf

Marks RE (2016) Risk neutral is best for risky decision making, mimeo, 16 March 2016. https://www.agsm.edu.au/bobm/papers/rnisbest.pdf

Marks RE (2018) Non-renewable resources and disequilibrium macrodynamics. Routledge, London (Reprinted version of my 1978 Stanford dissertation, with a new Preface by Robert Solow.) (Routledge Library Editions: Environmental and Natural Resource Economics, vol 9)

Marks RE (2019) Validation metrics: a case for pattern-based methods. In: Beisbart C, Saam NJ (eds) Computer simulation validation.: Springer International Publisher, Cham, Chapter 13,pp 319−338. https://doi.org/10.1007/978-3-319-70766-2_13. https://www.agsm.edu.au/bobm/papers/ch12_marks_rev4.pdf

Marks RE (2020) Calibrating methods for decision making under uncertainty. In: Bucciarelli E, Chen S-H, Corchado JM (eds) Decision economics: complexity of decisions and decisions for complexity. Springer, pp 1−9. https://doi.org/10.1007/978-3-030-38227-8_1. https://www.agsm.edu.au/bobm/papers/marksspain.pdf

Marks RE, Vriend NJ (2012) The special issue: agent-based computational economics—overview. Knowl Eng Rev 27(2):115−122. https://doi.org/10.1017/S0269888912000082. https://www.agsm.edu.au/bobm/papers/00marksvriend.pdf

Marks RE, Midgley DF, Cooper LG, Shiraz GM (1999) Coevolution with the genetic algorithm: application to repeated differentiated oligopolies. In: Banzhaf W, Daida J, Eiben AE, Garzon MH, Honavar V, Jakiela M, Smith RE (eds) GECCO-99: Proceedings of the genetic and evolutionary computation conference, July 13–17, 1999, Orlando, Florida. Morgan Kaufmann, San Francisco, pp 1609−1615. https://www.agsm.edu.au/bobm/papers/Gecco1999-RW-766.pdf

Midgley DF, Marks RE, Cooper LG (1997) Breeding competitive strategies. Manag Sci 43(3):257−275, March. https://doi.org/10.1287/mnsc.43.3.257. https://www.agsm.edu.au/bobm/papers/MS-MMC97.pdf

Midgley DF, Marks RE, Kunchamwar D (2007) The building and assurance of agent-based models: an example and challenge to the field. J Bus Res 60(8):884−893, August. https://doi.org/10.1016/j.jbusres.2007.02.004. https://www.agsm.edu.au/bobm/papers/jbr-mmk.pdf

Miller JH (1986) A genetic model of adaptive economic behavior, University of Michigan Working paper

Miller JH (1989) The coevolution of automata in the repeated prisoner's dilemma. Working paper 89-003 Santa Fe Institute

Miller JH (1996) The coevolution of automata in the repeated Prisoner's dilemma. J Econ Behav Organ 29:87–112

Miller JH, Andreoni J (1990a) A coevolutionary model of free riding behavior: replicator dynamics as an explanation of the experimental results. Working paper at Santa Fe Institute and University of Wisconsin

Miller JH, Andreoni J (1990b) Auctions with adaptive artificially intelligent agents. Working paper 90-01-004, Santa Fe Institute

Mitchell M, Forrest S (1994) Genetic algorithms and artificial life. Artif Life 1(3):267–289

Nalebuff B (1987) Economic puzzles: noisy prisoners, Manhattan locations, and more. J Econ Perspect 1:185–191

Page SE (2012) Aggregation inagent-based models of economies. Knowl Eng Rev 27(2):151–162

Richiardi MG (2012) Agent-based computational economics: a short introduction. Knowl Eng Rev 27(2):137–149

Schelling TC (1978) Micromotives and macrobehavior. Norton, New York

Szpiro GG (1997) The emergence of risk aversion. Complexity 2:31–39

Szpiro GG (2002) Tinkering with genetic algorithms: forecasting and data mining in finance and economics, Chapter 15, pp 273−286, in Chen (2002)

Tesfatsion L (2021) Agent-based computational economics: overview and brief history, Chapter in this volume

Wilhite A, Fong EA (2012) Agent-based models and hypothesis testing: an example of innovation and organizational networks. Knowl Eng Rev 27(2):221–238