

4. The Uncertain Fate of Self-Defeating Systems

David Batten

Abstract

Complex patterns of human behaviour are difficult to capture in agent-based simulations of socio-ecological systems. Even knowing each individual agent's strategy at one point in time may not help when trying to predict the collective behaviour of certain systems—if it is in each agent's best interest to do the opposite of most other agents. In *self-defeating* situations, the collective population of agents may exhibit a panorama of simple or complex behaviour, depending on the extent to which information is shared. A well-known example is the El Farol bar problem, in which a simulated population of bar attendees oscillates in a seemingly random manner around a critical congestion level. In this paper, it is suggested that a family of resource management problems involving human interactions with ecosystems may possess a self-defeating character. This poses several new challenges for integrated resources management. A case in point is the potential over-fishing of fisheries, which is addressed in the paper and likened to a minority game.

Introduction

The use of agent-based simulation to assist in environmental management and decision-making has increased markedly in recent years. Simulation models representing various facets of human behaviour have emerged (Hare and Deadman 2004). Social scientists use this kind of simulation for several purposes, including the discovery of collective regularities. Some are built on strong behavioural foundations, where mental models of the silicon agents are based on empirical data or stakeholders' views. In other cases, agents' strategies are selected stochastically or even arbitrarily and the outcomes tested in computational experiments.

What is arbitrarily decided in such simulations is how to drive the artificial learning process that enables each agent to adapt to other agent strategies interactively and to decide whether and how to modify its own strategy over time. Imitation of other (e.g. nearby/visible) agents is a simple and common method of introducing adaptation over time. In other simulations, agents choose from a suite of heuristics or mental models made available to some or all of them. More sophisticated learning methods (such as neural networks and genetic algorithms) have also been adopted.

Largely overlooked in the simulation work to date is the fact that the intrinsic structure of the system being simulated can have a direct influence on ways in which the human agents are likely to behave, i.e. react and adapt over time. Some social systems that fall into this category are called *self-referential* problems—situations in which the forecasts made by agents serve to create the very world they are trying to forecast. In these complex systems, the *best* thing to do depends on what everyone else is doing.

The purpose of this paper is to look briefly at a special class of self-referential problems called *self-defeating* systems. A typical example is the bar problem, in which a simulated population of attendees oscillates in an apparently random manner around a critical congestion level. It is suggested that several resource management problems involving human interactions with ecosystems have a self-defeating character, posing new challenges for integrated resources management. An example is the potential over-fishing of fisheries, behavioural aspects of which are discussed in the paper and likened to a minority game. Such common problems will require new institutional arrangements to overcome the self-defeating character of these situations.

Self-referential systems

Definition

A *self-referential* situation is one in which the forecasts made by the human agents involved serves to create the world they are trying to forecast. Such co-evolutionary systems have also been called *reflexive* and are generally *self-reinforcing*. They have attracted attention in several socio-economic contexts (Arthur 1994; Casti 1999; Batten 2000). A key feature of these systems is that the *best* thing to do (such as *to go* or *not to go*) depends on what everyone else is doing. As no individual agent knows that, the best thing that any agent can do is to apply the strategy or mental model that has worked best so far.

In self-referential systems, the mix of possible strategies is co-evolving incessantly over time, so much so that the choices made by individual agents seem to have little impact on collective outcomes. Sometimes, collective populations of agents exhibit emergent regularities. At other times, they display different types of complex behaviour. In the next section, we describe the El Farol bar problem, an early agent-based simulation of a complex adaptive system reported by Brian Arthur (1994).

The bar problem defined

Consider a system of $N = 100$ agents deciding independently each week whether or not to go to their favourite bar (called the El Farol) next Thursday. Space is limited, so the evening is enjoyable only if the bar is not too crowded

(say $N_{\max} = 60$). There is no collusion or prior communication among agents. Knowing the bar attendance over the past few weeks, each bar-loving agent simply decides independently to go if he expects less than N_{\max} to attend or stay home if he expects more than N_{\max} to go.

This problem is a metaphor for a broad class of self-referential situations: e.g. urban traffic congestion, canteen crowding or queue lengths at big sporting events. It has some interesting properties. First, if a decision model existed that agents could rely upon to forecast attendance, then a deductive solution would be possible. No such model exists. Irrespective of past attendance figures, many plausible hypotheses could be adopted to predict future attendance. Because agents' rationality is bounded, they are forced to reason *inductively*. Second, any shared expectations will be broken up. If all agents believe *most* will go, then *nobody* will go. By staying home, that common belief will be destroyed. If all agents believe *few* will go, then *all* will go, thus undermining that belief. Because any mental model that is shared by most of the agents will be *self-defeating*, agents' expectations must always differ.

Perplexed by the intractability of this problem, Arthur created a computer simulation in which his agents were given attendance figures over the past few months. He also created an *alphabetic soup* of several dozen predictors replicated many times. After randomly ladling out k of these to each agent, each kept track of his k different predictors and decided whether to go or not according to a preferred predictor in his set. This preferred predictor could be chosen in a variety of ways, although Arthur adopted the most accurate current predictor for each agent in his simulations.

The bar problem simulated

Once decisions have been made in Arthur's simulated bar, agents learn the new attendance figure, updating the accuracy of their own set of predictors. Then decisions are made for the following week. In this kind of problem, the set of predictors acted upon by agents—called the set of *preferred* predictors—determines the attendance. But the attendance history also determines the set of preferred predictors. We can think of this set as forming a kind of *ecology* (John Holland's term). Of interest is how this *ecology* evolves over time.

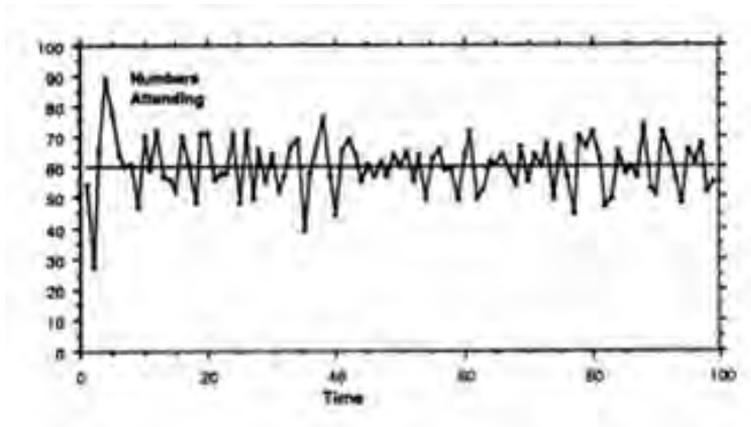
The simulations show that weekly attendance fluctuates unpredictably, but mean attendance always converges to 60 in the long run. The predictors self-organise into an equilibrium pattern or ecology in which (on average) 40 per cent of the preferred predictors forecast above 60, and 60 per cent below 60. This 40/60 split remains although the population of preferred predictors keeps changing in membership. The emergent ecology is rather like a forest whose

contours do not change, but whose individual trees do. Similar results appeared throughout Arthur's experiments, robust to changes in the types of hypotheses.

There is another intriguing result. Although the computer-generated attendance results look more like the outcome of a random process than a deterministic one (see Figure 4.1), there is no inherently random factor governing how many people attend. Weekly attendance is a deterministic function of the individual predictions, themselves being deterministic functions of the past attendance figures.

Thus the bar problem is a relatively simple example of an emergent, self-defeating system. It is a situation in which a system of interacting agents can develop collective properties that are not at all obvious from our knowledge of the agents themselves. Even if we knew all agents' individual idiosyncrasies, we are no closer to anticipating the emergent outcome. Under the influence of a sufficiently strong *attractor*, individually subjective, boundedly rational expectations self-organise to produce a kind of *collectively rational* behaviour (Arthur 1994).

Figure 4.1. A simulated 100-week record of attendance at El Farol



Source: Arthur, 1994

The bar problem extended

Given that agents do communicate with each other in the real world, we wonder what may happen if they are permitted to exchange information? Bruce Edmonds' work allows communication among agents before they make their final decisions whether or not to go to the bar (Edmonds 1999). Using a genetic programming algorithm to simulate adaptive learning, he allows each agent to competitively develop its models of what the other agents are going to do. Although the beliefs and goals of other agents are not known or represented by each agent, heterogeneity among the agents emerges in the form of non-uniform tactics and role-playing identities. These collective properties are features that emerge purely from the micro-dynamics.

Like the Prisoners' Dilemma, the bar problem is currently receiving more attention outside economics—as a metaphor for learning and bounded rationality. It has inspired a new literature in statistical physics on a closely related problem known as the minority game. A minority game is a repeated game in which N (odd) players have to choose one out of two alternatives (say A and B) at each time step. Those who happen to be in the minority win.

Seemingly simple at first glance, the game is subtle in the sense that if all players analyse the situation in the same way, all will choose the same alternative and therefore all will lose. Thus, players behave heterogeneously over time. Moreover, there is a frustration since not all players can win at the same time (this is an essential mechanism for modelling competition).

Since its introduction, this game has generated incredible interest. As there are only three parameters, the game is suitable for detailed numerical studies and analytical descriptions. Although each alternative is unspecified, for A or B one could read 'I am going to the bar', or, 'I am staying home'; 'I am choosing the motorway', or, 'I am going by the scenic route'; 'I am going to fish at the usual place', or, 'I am going to find a better spot'. Perhaps the most striking properties of the minority game are that:

- it is a model that addresses the interaction between agents and information;
- agents are able to cooperate and make use of the available information;
- there is a second order phase transition between a symmetric phase (in which no information is available to agents) and an asymmetric phase (in which information is available to agents); and
- if agents take their impact on the game into account, there is no phase transition and this case has an exact solution.

The minority game is an abstraction of the bar problem. Like simulations of the bar problem, numerical simulations of this game have displayed a remarkably rich set of emergent, collective behaviours (Challet and Zhang 1998).

Both the bar problem and the minority game contain the key elements of a *complex adaptive system*: firstly, a *medium* number of agents—a number too large for hand-calculation or intuition, but too small to use statistical methods applicable to very large populations; secondly, these agents are *intelligent* and *adaptive*, making decisions on the basis of rules of thumb or heuristics, like the bar predictors (needing to modify these rules or come up with new ones if necessary, they reason *intuitively*); and, finally, no single agent knows what all the others are (thinking of) doing, because each only has access to limited information.

Human ecosystems

In human ecosystems, human agents interact with one another and other life in a natural environment. Self-referential problems of the binary variety arise in these systems, but are seldom recognised as such. Since Hardin's influential article in *Science* (Hardin 1968), most are treated purely as commons dilemmas in which agents over-exploit a scarce resource in common. Examples include the over-fishing of fisheries, the degradation of national parks and the destruction of coral reefs. The question we address in the next section is to what extent the fisheries problem may be viewed as a complex adaptive system of the self-defeating variety.

Fisheries

In an insightful paper almost two decades ago, Allen and McGlade (1986) stressed that the key elements of fishing (like ancestral hunting) are *discovery* and *exploitation*, not simply the latter as is currently the case in western agriculture. As well as a desire to avoid congestion (like in the bar problem), fishing involves communication, adaptive learning and the emergence of non-uniform tactics and role-playing among fishing agents. Within such a socio-ecological system, key questions on our relationship with nature and the general problem of the management of a complex adaptive system can be explored.

Largely hidden from human view, the marine ecosystem is complex with respect to its species composition as well as to the processes occurring in it. As exploitation increases, Beddington and May (1977) showed that such an ecosystem moves towards instability and greater risk of collapse. Human responses tend to amplify random fluctuations, thereby further increasing the risk of collapse.

Ignoring the complex issue of ecosystem dynamics for the moment, let us concentrate on the strategies of the owners of the fishing vessels. From an owner's selfish viewpoint, if there are no institutional rules imposed on him, he aims to maximise his own vessel's catch. Obviously, the *best* thing for him to do will depend on what the other fishing vessels in the vicinity of his vessel are doing. Thus our fishing problem is a self-referential one.

Allen and McGlade (1986) developed dynamic, multi-species, multi-fleet models to explore the implications of different fishing strategies and information flows among fishing vessels in a Nova Scotia fishery. They identified two fishing strategies, calling them *Cartesians* and *Stochasts*. Stochasts search randomly for better sites or use their own intuition, without resorting to information that is shared between associates. They are risk-takers, seeking higher returns commensurate with the higher risks they take. At the other extreme we have the Cartesians, skippers who (for various reasons) are unwilling to take any risk and who only go to the zone promising the best *known* return. It does not matter

how long it takes for a Cartesian to reach their chosen site, as long as the final catch is guaranteed.

Allen and McGlade found that less information exchange among the fishing vessels ensures a more random response on the part of boats (i.e. more Stochasts among the fleet). Because the Stochasts continue to explore less visited parts of the system, the fishery as a whole has a greater chance of survival. If boats refuse to take risks and go only to where they know there are fish, the end result can be disastrous for everyone. Discovery involves risk, but to abandon it invites disaster.

There is a powerful message here that goes beyond the bounds of fisheries alone. In any mobile population, for example, we find some risk-takers and some who are risk-averse. Like the number of fishing boats turning up at the same site, the number of vehicles turning up on a specific road each day is unpredictable. Of interest are the adaptive strategies of drivers exposed to regular traffic jams. Anthony Downs (1962) identified two behavioural classes of driver: those with a low propensity to change their mode or route strategy, called *sheep*, and those with a propensity to change, called *explorers*. Explorers search for alternative options to save time. They are quick to learn and hold several heuristics in mind simultaneously. Sheep are more conservative and prone to following the same option. Empirical work in North America has confirmed the presence of sheep and explorer behaviour in real traffic (Conquest et al. 1993).

The parallels between Cartesian or Stochast fishing strategies and sheep or explorer driving strategies may seem striking. Yet they are less surprising if thought of as symptomatic of a more general phenomenon. In the world of technology, risk-aversers and risk-takers appear under different guises: *imitators* and *innovators*. If we allow for the coexistence of imitative and innovative mechanisms in a population, both must be treated as co-evolutionary variables dependent on the unfolding of events.

Evolutionarily stable strategies

Is the coexistence of Cartesian (imitative) and Stochast (innovative) strategies in a human population an *Evolutionarily Stable Strategy* (ESS)? An ESS is a strategy (or mix of strategies) which, if most members of a population adopt it, cannot be bettered by an alternative strategy. The reason is that the fitness of individuals adopting an ESS strategy is higher than the fitness of individuals adopting other strategies (Maynard-Smith 1982). Another way of putting it is that the best strategy for any individual depends on what the majority of the population is doing (Dawkins 1976). This last perspective helps us to grasp the idea that an ESS may be an important property of self-referential systems.

There are three classes of ESS: pure, mixed and conditional. A pure ESS is a strategy that is consistently exhibited by individuals throughout their lifetimes. A mixed ESS is a complex of two or more strategies varying either within or among individuals over time. In a conditional ESS, each individual's strategy varies under different conditions in the social or physical environment. If the evolutionary payoff for each strategy depends on what other individuals are doing and decreases as a greater proportion of the population adopts that strategy, other strategies may also be an ESS (frequency-dependent selection).

To apply the ESS idea to our fishery problem, consider the following variant of one of Maynard Smith's simplest hypothetical cases (hawks and doves). Here we draw extensively on the simulation results in Allen and McGlade (1986). Suppose that the only two types of fishing strategy in a fleet of vessels are Stochasts and Cartesians. We want to know whether Pure-Stochast or Pure-Cartesian is an evolutionarily stable strategy. As Stochasts and Cartesians compete for returns, we must estimate payoffs to the different fleet strategies in order to find an ESS.

Serving as the *eyes* of the fleet, Stochasts search for fishing zones of high return. A fleet with a high proportion of risk-taking Stochasts will discover successive zones of high return, eventually fishing out the high-return zones in a patchwork manner. But Stochasts ignore fish from nearby zones of intermediate returns. As it takes time for the high-return zones to recover, pure Stochast strategies tend to lead to boom-and-bust series of good and bad years. However, this patchwork approach guarantees that the fishery system survives rather than running the risk of collapsing.

Based on the results in Allen and McGlade (1986), we can assign a Pure-Stochast a cumulative payoff of 76 points (equivalent to 6 large circles plus 8 small ones) over a ten-year fishing period. This includes several bad years during which the stocks must recuperate.

Cartesians refuse to take risks and rely solely on the information passed on by the Stochasts. Thus they go only to zones where they know in advance that there are some fish. A fishing fleet dominated by risk-averse Cartesians will simply direct all the boats in the short term to a few of the *best* zones for fishing. Because the information about which are the best zones is generated by a small number of Stochasts, the tendency for Cartesians is to *lock* onto one particular location as the best. In the long run, this leads to the fleet exploiting a single location instead of the whole area, leading to a small catch and a small fleet. Based on figures found in Allen and McGlade, a Pure-Cartesian strategy yields a payoff of 45 points over the same ten-year period.

The above figures confirm our suspicions that Pure-Stochasts do better than Pure-Cartesians over time. They are rewarded for taking more risks. Pure-

Stochasts achieve better returns because they can maintain fishing activities over the whole area, obtaining a higher catch in the good years and maintaining a larger fishing industry overall. Pure-Cartesians become trapped in a few locations instead of spreading out across the whole area. They are limited each year to a smaller catch and a declining fishing fleet and industry over time.

Table 4.1. Strategies and payoffs for homogeneous fleets

<i>Fleet strategy 10-year payoff</i>	
Pure-Stochasts	76 points
Pure-Cartesians	45 points

But is a Pure-Stochast strategy evolutionarily stable on its own? To answer this question, let us suppose that we have a fleet consisting only of Pure-Stochasts. Despite the bad years, they seem to do very nicely, earning the payoff given in Table 4.1. Now suppose that a mutant Cartesian arises in the fleet. Being the only Cartesian, must he compete aggressively with the Pure-Stochast? No. Instead he could follow the Pure-Stochasts, letting them show him where it is best to fish. Then, by fishing at the edges of sites discovered by Pure-Stochasts, he does not irritate them too much. Remaining mostly unnoticed, he can *free-ride*, enjoying above-average returns for very little search effort.

As long as there are sites offering intermediate returns that are left untouched by Stochasts, there are niche opportunities for Cartesians to fill. Based on figures in Allen and McGlade (1986), 10-year payoffs to a mixed fleet of Stochasts and Cartesians under various levels of information exchange are given in Table 4.2.

Table 4.2. Fleet strategies and payoffs when information is shared

<i>Fleet strategy 10-year payoff</i>	
If catch information is shared equally:	
Stochasts	104 points
Cartesians	83 points
If Cartesians fail to inform Stochasts:	
Stochasts	95 points
Cartesians	81 points
If Stochasts fail to inform Cartesians:	
Stochasts	109 points
Cartesians	30 points

When information is shared equally between both groups, the payoff to each is higher than to a Pure-Stochast or Pure-Cartesian fleet (see Table 4.1). Thus a Pure-Stochast strategy is not evolutionarily stable, since it can be bettered by a strategy in which Stochasts and Cartesians communicate and cooperate with each other.

What happens if either group cheats by not transmitting reliable information about their catches? As Allen and McGlade noted, there is an interesting asym-

metry. If Cartesians fail to inform Stochasts, there is very little effect (see Table 4.2). But if the information possessed by Stochasts is not transmitted to Cartesians, the latter perish (or are compelled to become Stochasts). Thus Cartesians will try to obtain catch information and Stochasts will try to withhold it. Such effects have evolved in real fishing fleets. For example, Vignaux (1996) found that trawlers fishing for hoki in waters off New Zealand's coastline do not share catch information, instead basing their own decisions partly on watching where other vessels fish. This amounts to spying. Other tactics include *listening in* on radios, lying about catches and spreading other misleading information.

Using agent-based simulation, a richer suite of possibilities can be explored. Recent Agent-Based Modelling (ABM) work involving Bayesian belief networks has shown that various kinds of information flow among fishing vessels has important effects on the dynamics and resource exploitation of a simulated fishery (Little et al. 2004). As stated earlier, information flow tends to benefit the Cartesians at the expense of the Stochasts. This asymmetry confirms the importance of realistically representing the rich variety of possible fisher behaviour in any modelling framework that aims to assist with integrated management of fishery resources.

A modified minority game?

The superior payoff achieved when catch information is shared reliably is in line with observed properties of the minority game. If all agents in a fleet have access to public information about catches and zones of highest return over a period of time, then the agents interact only through this public information and the system has a *mean-field* character (in the sense that no short-range interactions exist). Self-organisation in such a system is achieved by allowing each agent to have several strategies from which he selects the one that seems best (to him).

In the literature on the minority game, there is a second order phase transition between a symmetric phase (in which no information is available to agents) and an asymmetric phase (in which information is available). It is easy to see why this phase change might occur in our fisheries context. First, if reliable catch information is passed between Stochasts and Cartesians, this may ensure that a boom-and-bust series of good and bad years is avoided. Stochasts can direct some Cartesians to high-return zones and others to medium-return zones, in such a way that no high-return zone will be in great danger of over-fishing. Second, the two strategies are complementary. Cartesians make good use of information, while Stochasts generate it. Together they can exploit the resource more efficiently because of their complementarity.

Ideally, each fleet needs some Stochasts (researchers) and Cartesians (producers) who cooperate within a fleet but not with competing fleets. To survive, each

needs the other. What, then, is the ideal ratio of Stochasts to Cartesians? This is a difficult question to answer without considering the dynamics of the fish population and other environmental factors. The higher payoff that a Stochast may enjoy is tempered by the higher risks involved, whereas the risks taken by a Cartesian are minimal.

In a highly dynamic world like fisheries, one may expect the ratio of Stochasts to Cartesians in a fleet to vary in response to the search success of the Stochasts and their willingness to inform their Cartesian partners. Like music lovers at the El Farol bar, the ratio of Stochasts to Cartesians can oscillate forever. But, in this case, there may be no emergent regularity as the congestion level is a dynamic variable. The situation is a self-defeating one, since the *best* thing for each vessel to do depends on what everyone else is (thinking of) doing.

Discussion

Previous simulation work referred to in this paper has shown the importance of information sharing and communication strategies (Allen and McGlade 1986; Maury and Gascuel 2001; Little et al. 2004). Irrespective of whether vessel owners within fishing fleets share reliable catch information, agent-based models can help to clarify the potential value of information-sharing. Earlier work has shown that lower returns will be experienced by Cartesians in the absence of information flow from Stochasts. In the long run, however, Stochasts suffer boom-and-bust years by fishing alone because they fish out the high-return zones in a patchwork manner. Thus the evolutionarily stable strategy for a fishing fleet is a mixed one in which Stochasts and Cartesians communicate and cooperate with each other.

When the economic viability of a fishery is under serious threat, it is worthwhile adopting a cooperative approach—not only to encourage information sharing, but also to overcome such common pool dilemmas. If conditions are suitably conducive, a sustainable solution can even be found by the participants themselves. One successful example is the inshore fishery at Alanya in Turkey (Ostrom 1990). Members of the local cooperative found an ingenious rotation system allocating fishing sites to local fishers that builds upon reliable and mutually beneficial information exchange.

As mentioned near the outset, the purpose of this paper was to look at a class of self-referential problems called *self-defeating* systems. A typical example is the bar problem, in which the simulated population of attendees oscillates in a seemingly random manner around a critical congestion level. Another self-defeating problem is urban traffic congestion. Without a degree of cooperation and information sharing between the sheep and explorers (Cartesians and Stochasts) on our roads, the collective behaviour of these complex adaptive systems will continue to defy prediction or control on a daily basis.

The fishery situation is of a similar ilk. As well as a desire to avoid congestion, fishing involves communication, adaptive learning and the emergence of role-playing and non-uniform tactics among the agents involved. Human behaviour in national parks and coral reefs raises similar collective problems. In each of these ecosystems, it is impossible to achieve an eco-efficient outcome unless a threshold level of cooperation and information sharing can be achieved. Stochasts and Cartesians need to work together in order to achieve sustainability and thus avoid over-exploiting the resources therein.

Within such socio-ecological systems, the use of agent-based simulation allows key questions about our relationship with nature and the more general problem of the management of a complex adaptive system to be explored qualitatively. Rule-change experiments—such as those that led to the successful fishing regime at Alanya—fall into the scientific domain of participatory, agent-based modelling. The great thing about this kind of simulation is that we can assess the likely outcomes of such rule changes ahead of their implementation.

As individual humans—in our roles within families, communities, firms, institutions and regions—we must decide how to divide our time and effort between doing what we know (with known values and payoffs) and searching for new opportunities and roles that may have superior payoffs in the future. We can choose to explore new pathways and connections or try to minimise such deviations from those pathways that we are accustomed to in our lives to date. Thus Stochasts and Cartesians are two important extremes in human society. The first group take more risks, venturing into the unknown, be it by hunches, tactics or entrepreneurial activities. Their discoveries nourish a society in the long run, expanding it into the adjacent possible and assuring its long-term survival in some form. The second group prefers to devote themselves to tasks already assigned to them, seeking to undertake them as quietly and efficiently as possible. They constitute the *backbone* of a society and come much closer to our definition of *normality*. The survival and sustainable prosperity of any individual, community, firm or nation, or even a human ecosystem, requires a mixture of both types of behaviour.

Acknowledgments

Helpful comments from Paul Atkins, Marco Janssen, participants at the joint CSIRO Agent-Based Modelling/Human Ecosystems Modelling with Agents (CABM/HEMA) Workshop, Canberra (17–18 May 2004) and at the IEMSS Congress, Osnabrück (14–17 June 2004) are gratefully acknowledged. The usual caveat regarding author's responsibility applies.

References

- Allen, P.M. and J.M. McGlade (1986) Dynamics of discovery and exploitation: the case of Scotian Shelf fisheries. *Canadian Journal of Fisheries and Aquatic Sciences* 43: 1187–200.
- Arthur, W.B. (1994) Inductive behaviour and bounded rationality. *The American Economic Review* 84: 406–11.
- Batten, D.F. (2000) *Discovering Artificial Economics: How Agents Learn and Economies Evolve*. New York: Westview Press.
- Beddington, J.R. and R.M. May (1977) Harvesting natural populations in a randomly fluctuating environment. *Science* 197: 463–5.
- Casti, J. (1999) Would-be worlds: the science and surprise of artificial worlds. *Computers, Environment and Urban Systems* 23: 193–203.
- Challet, D. and Y-C. Zhang (1998) On the minority game: analytical and numerical studies. *Physica A* 256: 514–518.
- Conquest, L., J. Spyridakis, M. Haselkorn and W. Barfield (1993) The effect of motorist information on commuter behaviour: classification of drivers into commuter groups. *Transportation Research C* 1: 183–201.
- Dawkins, R. (1976) *The Selfish Gene*. Oxford: Oxford University Press.
- Downs, A. (1962) The law of peak-hour expressway congestion. *Traffic Quarterly* 16: 393–409.
- Edmonds, B. (1999) Gossip, sexual recombination and the El Farol bar: modelling the emergence of heterogeneity. *Journal of Artificial Societies and Social Simulation* 2(3): 1–21.
- Hardin, G. (1968) The tragedy of the commons. *Science* 162: 1243–8.
- Hare, M. and P. Deadman (2004) Further towards a taxonomy of agent-based simulation models in environmental management. *Mathematics and Computers in Simulation* 64: 25–40.
- Little, L.R., S. Kuikka, A.E. Punt, F. Pantus, C.R. Davies and B.D. Mapstone (2004) Information flow among fishing vessels modelled using a Bayesian network. *Environmental Modelling and Software* 19: 27–34.
- Maury, O. and D. Gascuel (2001) Local over-fishing and fishing tactics: theoretical considerations and consequences in stock assessment studied with a numerical simulator of fisheries. *Aquatic Living Resources* 14: 203–10.
- Maynard-Smith, J. (1982) *Evolution and the Theory of Games*. Cambridge: Cambridge University Press.

Ostrom, E. (1990) *Governing the Commons: the Evolution of Institutions for Collective Action*. Cambridge: Cambridge University Press.

Vignaux, M. (1996) Analysis of vessel movements and strategies using commercial catch and effort data from the New Zealand hoki fishery. *Canadian Journal of Fisheries and Aquatic Sciences* 53: 2126–36.