DISCOVERING ARTIFICIAL ECONOMICS

How Agents Learn and Economies Evolve

DAVID F. BATTEN

Revised Manuscript

December 1999

To be published by Westview Press

2

On the Road to Know-Ware

"The heart has its reasons
that reason does not know."
PASCAL

{A}What is Knowledge?{/A}

In a novel study of the American economy, Fritz Machlup claimed that the "knowledge industry" accounted for about 29 percent of the U.S. Gross National Product *by 1958*.[1] He also found that the growth rate of "knowledge-producing" occupations had exceeded all other job classes since the turn of the century. Studies elsewhere have confirmed that knowledge-producing occupations are growing rapidly. After subdividing the workforce force into four occupations – knowledge-handling jobs, administration and information jobs, personal services, and goods-handling jobs – a Swedish study found that the share of knowledge-handling jobs grew from about 10 percent in 1960 to 18 percent in 1980, and was expected to reach about 30 percent by the turn of the millenium.[2]

Who are these knowledge producers? Machlup combined people who create new knowledge (e.g. research scientists) with those who communicate existing knowledge to others (e.g. teachers, managers, air traffic controllers). On reflection, this seems too generous. Why? Because the *creation* of new knowledge is a more complicated and demanding undertaking than routine tasks like passing on existing knowledge or information to others. Although all of us possess some degree of creative talent, the most creative knowledge-producers are well-known. They include researchers, inventors, entrepreneurs, authors, artists and composers.

[1] For a fuller account of the distribution of knowledge-handlers in the American economy, see Machlup (1962).

[2] The Swedish study and its results are discussed in Andersson (1986).

Some intriguing questions then arise. What kind of new knowledge do they create? How is this knowledge acquired and improved? Does it differ from information? One kind of new knowledge comes from scientific research. Scientific knowledge is produced by universities, research institutes and other science-oriented organizations, where it often mixes with associated tasks (e.g. teaching, consulting). Much of this collective knowledge appears in books, scientific journals and other technical publications. Some of it gets applied by industry in the design and production of new technology and products. Occasionally, it finds its way into the popular press.

Traditionally, most scientists used to work alone or in small groups led by a senior scientist. Today, however, there's a growing trend towards collaboration between researchers in many different locations.[3] Globalization, aided by the advent of the Internet and rapid advances in communication technology, seem to have spawned greater collaboration over longer distances. Nevertheless, there's also evidence that the frequency of research collaboration between researchers who speak the same first language decreases exponentially with the distance separating them.[4] Distance still matters, despite the information technology revolution.

The explanation for this frictional effect is that informal, face-to-face contact is thought to be an essential ingredient for generating new ideas by collaborating partners, but that factors such as geographical proximity become significant because of the additional travel cost and time needed to bring the partners together. Thus researchers who are closer geographically can meet and exchange ideas more easily and more often than distant partners. The debatable point is that face-to-face meetings cannot be replaced by the use of telecommunications, the Internet, or related technologies when it

[3] One way of monitoring research collaboration is to measure the number of internatonally co-authored scientific articles. Work of this kind has been reported recently by Swedish researchers. For example, Andersson and Persson claim that, in recent decades, such collaboration has been growing at an average of 14% per year; see Andersson and Persson (1993).

[4] For an interesting analysis of the effect of distance on collaboration between the UK, Canada and Australia, see Katz (1993). He normalized the distance variable as a fraction of the largest distance beteen pairs of universities in each country, and found that the frictional impedance effect was exponential. A similar result can be found in Beckmann (1994), whose theoretical analysis weighed the advantages of collaboration against its cost. His results suggest that we should expect an exponential distance effect.

comes to the generation of new knowledge.[5] One reason for this is that novel and surprising ideas tend to emerge more spontaneously from face-to-face exchanges than from other kinds of contacts. We'll return to the issue of face-to-face contacts later in this chapter, when we come to distinguish knowledge from information.

An increasing number of scientists work on research and development (R&D) in private sector corporations or public agencies. This R&D leads to new products like "high-tech" pharmaceuticals, chemicals, aeroplanes, space products, instruments, electronics, computer-based hardware and software. Because chance plays such a major part in the process of scientific discovery through R&D, once again we find ourselves plunged into the world of morphogenesis. Although knowledge-producers can be found in all sectors of the economy, successful R&D is a crucial ingredient in the "get ahead, stay ahead" world of high technology. The potential payoffs are enormous for the winners. As we learnt in the first chapter, high-tech scientists live in that very chancy world of increasing returns.

How do scientists and others acquire new knowledge? There's no magic elixir involved here. Like all of us, scientists must acquire their knowledge from learning. Yet for most of the twentieth century, economists have treated knowledge and learning as exogenous variables when it comes to explaining economic growth and development. Ironically, the idea that increasing returns could arise from the accumulation of knowledge is almost as old as economics itself. In his *Principles of Economics*, Alfred Marshall noted that an increase in "trade-knowledge" that cannot be kept secret is a form of external economy. Yet very few models of economic change adopted this suggestion.

Perhaps the foremost advocate of knowledge as the endogenous engine of growth and technological change is Paul Romer, an economist at the University of Chicago. For several decades now, there's been a frantic rush to absorb learning into the main corpus of economic theory; ever since the appearance of Kenneth Arrow's pioneering paper on "learning-by-doing." Much of the discussion in this and later chapters will focus on learning as an adaptive process, drawing partly on the domain of psychology.

[5] This highly contentious issue can be relied upon to raise a storm of debate among researchers claiming expertise in the field. But independent research has shown that the strongest proponents supporting each

When we probe learning problems, we find that psychologists are no more in agreement than economists. But one empirical generalization seems to stand firm, having been accepted by all schools of thought: Learning is a product of *experience*. In other words, learning can only occur when we're attempting to *do* something - like reading a book, talking to someone, playing a game or solving a problem. We already know that learning associated with repetition of much the same problem is subject to diminishing returns. We're trapped in the world of stasis again. According to Arrow, the stimulus situations must themselves be steadily evolving rather than merely repeating, if we're to build on our experience and enjoy steadily increasing returns. This is a vital distinction because it means that learning belongs to the world of morphogenesis. And that makes it evolutionary.

Before we launch into a discussion of the evolutionary aspects of learning, let me round off this discussion with a few words about knowledge and information. Superficially, we tend to think of knowledge as data or facts which we can organize into predefined categories - compartments in the brain, if you like. For example, a financial analyst might notice that the rate of exchange between the yen and the dollar is running at 125. He stores this fact in the "yen-dollar" compartment in his memory bank; he now *knows* this. The next day he notices that the same exchange rate has dropped to 123. Therefore he revises this entry in his memory bank. Having updated his yen-dollar data string, he has *learnt* something new.

Undoubtedly this is a kind of learning, but it's a very simple learning model. In fact, it's too simple. Knowing means nothing more than being aware of data and assigning them to the correct memory compartments. Learning means nothing more than revising these data in the light of fresh information from outside. While this may be fine for some purposes, to understand what learning is really about we need to go deeper than this. Certainly the world of economic decision-making is not that shallow. Knowledge is not always spoonfed to us in pre-packaged forms. Rarely do we enjoy the luxury of infallible conceptual models that we can use to forecast, analyze and act upon with

argument tend to have vested interests in their particular point of view. For a hierarchical treatment of the knowledge exchange issue, see Batten, Kobayashi and Andersson (1989).

absolute confidence. Our personal stock of knowledge is a very individual thing, being a unique product of our own experiences, constructs and memories.

Like the empiricist philosophers before him, Immanuel Kant believed that our knowledge of the world comes from our sensations. But he also believed that *how* we see the world depends on the particular "glasses" we're wearing. We can never have certain knowledge of things "in themselves," just how things appear to us. Kant believed that there are clear limits to what we can know. Each mind's glasses set these limits. When we ask questions about totality, such as whether the Universe is finite or infinite, we're asking about a totality of which we're but a small part. There's simply too much for a single person to handle, so we can never know this totality completely. We'll return to the issue of what may be knowable and unknowable in Chapter 8.

Knowledge becomes a much fuzzier concept in Kant's world. If there are clear limits to what we can know, then our ability to reach identical conclusions under similar conditions should not be taken for granted. Erwin Schrödinger sums it up nicely in his book *Mind and Matter*:

{EXT}The world is a *construct* of our sensations, perceptions, memories. It is convenient to regard it as existing objectively on its own. But it certainly does not become manifest by its mere existence. Its becoming manifest is conditional on very special goings-on in very special parts of this very world, namely on certain events that happen in a brain.{/EXT}[6]

Each of us is a unique product of our own brain and our uniquely individual experiences. Our personal knowledge is honed by the concepts, notions and models which we choose to use to represent it. All of this has to be *created*, put together over time by us as well as by others in society as a whole. Learning is a cumulative process which can be frustratingly slow, partly because some of us are stubbornly resistant to change. Thus some of our knowledge stocks turn out to be surprisingly durable.

An obvious distinction can be made between the simpler, more objective kinds of knowledge, such as the yen-dollar exchange rate, and the more complex, subjective kinds which result from our own mental "gymnastics" and contacts with others. We can

[6] See Schrödinger (1956).

usefully think of the simpler kinds as *information* and our use of them as *information-processing*. Basic information -- like an exchange rate or the maximum daily temperature – is readily quantifiable. It's dispersed easily over geographical space and changes rather quickly over time. Most information has limited value, in the sense that its usefulness erodes rather quickly. Last Thursday's maximum temperature is of much less interest than today's high. Furthermore, this kind of information is *not* a product of our own thinking and experience. It's external to us, bestowed upon us like "manna from heaven." For this reason, information can be gathered and transferred in many different ways. Moreover, it's exchange doesn't require face-to-face interactions.

[Table 2.1 near here]

Thus simpler tasks like information-processing should be distinguished from the more complex processes that make up learning. Whereas information-processing tends to be routine and repetitive, learning is more intuitive and adaptive. A key point is that learning cannot take place in a vacuum, isolated from others. It relies heavily on interactions between individuals, or between individuals and their environment. For example, individuals learn to make a living by creating and selling goods and services that only make economic sense in the niches afforded by other goods and services. As Stuart Kauffman suggests, an economy resembles an ecosystem because it consists of a web of coevolving agents. The learning effects of the interactions between agents can be beneficial to some, detrimental to others; or mutually beneficial. As we'll see shortly, the true story of technological evolution is actually one of adaptive coevolution.

{A}Finding the Road to Know-Ware{/A}

Biological evolution also teaches us that information is a relatively trivial concept in comparison with knowledge. For example, the divided cell or the fertilized egg possess a basic kind of *know-how*. Two fertilized eggs can have the same amount of information in terms of bits, but one knows how to make a hippopotamus and the other knows how to make a giraffe.[7] An ant or rat knows how to find its way to food in a maze. A trained racing pigeon knows how to find its way home from any distant origin. This special "sixth-sense" or know-how is part of any living creatures behavioural repertoire. Furthermore, it gives the impression of being responsive, perhaps even *purposive*.

By giving the impression that it's purposive, know-how stretches the dimensions of knowledge beyond a simple information concept. The problem is that it's devilishly difficult to quantify. Because it's a multi-dimensional part of a person's knowledge stock, for example, it's more qualitative than quantitative. We can only measure know-how vaguely, using coarse measures like "little," "basic" or "extensive."

We run into a similar kind of problem in economics. Two different countries can have the same level of GNP, for example, but one knows how to make the world's best watches while the other knows only how to grow olives. One may have a fairly equal income distribution, the other a very unequal one. Yet they're given equal scores in comparative studies of economic performance! The problem doesn't go away by decomposing GNP into its constituent parts, such as firms and their employees. Although two watchmaking firms may have the same output or earnings, one can make highly sophisticated watches using leading-edge technology, while the other knows only how to make old-fashioned faces with leather bands. Even at the microscopic level – say of individual watchmakers -- the problem persists. No two individuals possess identical know-how.

Know-how is not the only component of human knowledge. There is also *know-what* and *know-that*. Returning to the biological example, a fertilized egg definitely has the know-how to make a hippotamus, but it's unlikely to know *what* it's doing, i.e. *that* it's making a hippopotamus rather than a giraffe. Knowing how to reach food inside a maze does not mean that a rat knows *what* food is there, or even *that* there's food there. Knowing how to fly home does not imply that a pigeon knows *that* a particular path will get it home.

When it comes to economic agents, however, the nice thing is that we do have know-what and know-that. In fact, we use them to create know-how. By concentrating

[7] This suggestion came from Kenneth Boulding.

on know-what and know-that, for example, scientists have greatly increased our stocks of know-how. Until science discovered the chemical elements and the periodic table, there could be never be a chemical industry. Until they discovered the silicon chip, developing the know-how to create integrated circuits was out of the question. Perhaps all we need to spawn some new know-how is a mixture of know-what and know-that, together with the added "spice" of a little contact with some other clever thinkers.[8]

Regrettably, it's not even that "simple". We must continue our little journey along this road to know-ware, going beyond know-what and know-how to *know-whether*. Human decision making involves a series of choices. Know-whether involves evaluating the implications of alternative decisions in order to find out if the chosen course of action was the best decision. Before choosing from among several alternative courses of action, we harbour expectations about the likely consequences of each different course. Once we've made our choice, however, we're "locked-in" to it. Only at some later stage do we get additional feedback telling us about the wisdom of our choice. In other words, to know-whether requires feedback from the decision environment in which the choices are made.

This subjective ability is a more sophisticated part of the behavioural repertoire of humans. It's one of those instincts that seems to set us apart from other living species. Although rats and pigeons may act expectantly, they cannot explicitly state an expectancy. Nor do we believe that they can really think about it. Different environmental conditions lead us not just to act differently, but also to think differently. For example, being caught up repeatedly in traffic jams not only encourages us to travel off-peak or take the train instead. It also makes us think about the future of the whole traffic system, and whether our politicians are doing enough to improve it.

Our economy needs science to develop know-what and know-that, primarily as ingredients for improving our know-how. Production always begins with know-how.

---

[8] Both animal and human behaviour can be thoughtful, and learning depends on some kind of thinking. It's easy to regard thinking as something we *do* and knowledge as something we *have*. But this may be an oversimplification. We don't just store our knowledge, we also use and improve it. Ryle (1949) points out that we can also use it without necessarily being able to explain what it is that we know. He introduced a distinction between "knowing that" and "knowing how", partly to reinforce his own idea that intelligence is a characteristic of performance. For a summary of Ryle's thesis, and an interesting analysis of thinking as an activity that "unfolds" in time, see Eiser (1994).

There were no plastics one hundred years ago because we didn't know how to make them. Most managers and economic agents, however, rely on know-whether to make sound decisions. We'll reconfirm the need for careful distinctions between these elements of our personal know-ware when we look at an example of adaptive learning on networks; which we call *learning-by-circulating*. For the moment, though, let's return to some conventional wisdom: how knowledge and learning have been treated in traditional economic theory.

{A}The Age of Deception{/A}

We've emphasized that the conventional world of economics is mostly confined to the world of stasis. What does this imply in terms of human reasoning and behaviour? Nothing very sophisticated. Among the most popular set of simplifying assumptions, a particular one tends to proliferate. Most theoretical reasoning in economics assumes that economic agents behave in a perfectly rational manner, that is, they possess perfect, logical, deductive rationality.

Deduction is reasoning from the general to the particular.[9] A perfectly logical deduction yields a conclusion that must be true provided that its premises are true. Thus deduction involves specifying a set of axioms and proving consequences that can be derived from those premises. Sounds straightforward enough, doesn't it? The catch is that the premises must be complete, consistent and well-defined. As such, it's pretty easy to run into problems. While deduction is handy for solving a host of theoretical problems, it's much less helpful for tackling practical problems. Why? Because for premises to be complete, consistent and well-defined, the problem must be relatively simple. In an economic setting, for example, the problem must be simple enough for agents to know what's in their self-interest, to act in their self-interest, and to perform the

[9] Reasoning from the general to the particular is a top-down approach. It leads to the concept of a representative agent like the "average citizen." General rules of behaviour are determined at the macrolevel and then assigned unilaterally to all individual agents. Thus the population of agents is regarded as being homogeneous.

calculations needed to know the implications of alternative decisions. In other words, they need to have the brainpower to figure out the optimal decision.

A case in point is education. Far-sighted parents plan ahead and deduce the amount of education that's economically feasible, both for their children and for themselves. Another example is housing. Potential buyers plan ahead and deduce the location, type and amount of housing that they can afford. This is the kind of reasoning, deduction and analysis that is assumed in most areas of economics. It certainly can help in family planning and house purchases. Also it's well researched. Psychologists have accumulated almost a century's worth of experiments based on deductive reasoning.[10] But is a deductive approach always sufficient to solve the full range of problems confronted in economics? To answer this question, let's see what's required in terms of know-ware.

To reason deductively, economic agents need to have a complete set of know-ware at their disposal. First, they need to know-what serves their best interest. Second, they need the know-how to act in their best interest. Third, they need know-whether in order to evaluate the implications of alternative decisions and be sure that they have chosen wisely. In brief, they must have perfect *know-what*, perfect *know-how* and perfect *know-whether*. Their know-ware must be honed to perfection. This, of course, is a rather tough condition.

From the above, it's hardly surprising to find that deduction only works well on relatively simple problems. As Brian Arthur puts it: "If we were to imagine the vast collection of decision problems economic agents might conceivably deal with as a bottomless sea or ocean, with the easier ones on top and more complicated ones at increasing depth, then deductive rationality would describe human behaviour only within a foot or two of the surface."[11] Before we tackle economic problems, let's visualize where various games might be found as we dive to various depths. Simple games – like Tic-Tac-Toe – are readily solved by a deducible minimax solution. In everyday terms, this means that the human brain is quite capable of figuring out the

[10] For a summary of some of these experiments in deductive reasoning, see Johnson-Laird and Byrne (1991).

[11] See Arthur (1994a), page 406.

"best" moves on a game board consisting of only nine squares. The best moves are the ones which leave your opponent in the worst possible situation. Thus a way of testing for goodness is to pretend you've made the move, then evaluate the board from your opponent's viewpoint. Meanwhile, your opponent is doing the same. He or she mentally runs through all possible moves and evaluates them from what he thinks is *your* viewpoint.

Note that we've defined our best move *recursively*, using the maxim that what's best for one side is worst for the other. It's recursive because it operates by trying a move, and then calling on itself in the role of opponent. Since recursion can go on several moves ahead, it's possible to figure out the best strategy to adopt and the likely result in a game of Tic-Tac-Toe. Each move generates its own "look-ahead tree", with the move itself as the trunk, your opponent's responses as main branches, your counter-responses as subsidiary branches, etc. In Figure 2.1, I've shown the look-ahead tree corresponding to the first few moves of a game.

[Figure 2.1 near here]

Let's see how the minimax solution can be deduced. If you move first, your best move is to choose the central square, thus limiting your opponent's opportunities of scoring three noughts in a row to just four possibilities - two horizontally and two vertically. At the same time, you've secured four ways of winning in two additional moves. Choosing any other opening move offers your opponent more opportunities to win and secures fewer ways for you to win. Your opponent's best response to your opening move is to choose any one of the four corner squares. Then the game will finish in a draw. However. if your opponent doesn't choose a corner square, then you'll win the game.[12]

If each player always chooses their best move, the game's outcome can be deduced in advance. This is a rather pleasing result. In exchange for a well-defined decision problem, we get back a well-defined solution. The recursive solution invokes a

[12] These two outcomes assume that all remaining moves are the best possible ones.

logic that is relentless and consistent. It acts step by step on premises that are well-defined. It's also self-consistent and self-enforcing, in the sense that if your opponent behaves according to the deductive solution (i.e. chooses his/her best move), then it would not be in your best interest to do otherwise. If the best moves are implemented properly, the drawn outcome confirms the deductions that went into it. It's a rational expectations equilibrium.

Because we can deduce the likely outcome to a game of Tic-Tac-toe after the first few moves, we can think of it's logic as lurking just below the surface. We'd need to dig a little deeper, however, to catch a glimpse of games like Checkers or Quads.[13] Owing to larger board sizes (Quads = 36 squares, Checkers = 64 squares), it takes longer to figure out the likely outcome. In other words, we can't deduce the perfectly rational solution. Not only are the choices of moves more numerous, but our best moves depend more and more on the moves of our opponent. Neither our best strategy nor the likely outcome are deducible in advance. There are simply too many possible branches in the corresponding look-ahead trees.

[Fig 2.2 near here]

Deduction has no chance whatsoever when we finally reach the game of Chess. Something else is needed in even greater doses at these deeper levels. It's really an art to figure out how to avoid exploring every branch of a look-ahead tree out to its very tip. Good chess players seem to excel at this art. Or do they? The funny thing is that top-level players look ahead relatively little, especially if compared to chess programs. Until Deep Blue's success against Gary Kasparov, people were superior as chess strategists. We'll come back to the issue of the best Chess strategy shortly.

Why, then, does deductive rationality fail us when we're faced with more complicated decision problems? Three reasons spring to mind.. The obvious one is that

[13] Quads is a game between two players, each of whom attempt to place 18 different pieces on a game board of 36 squares. Players take turns, choosing and placing one of their pieces on the board in such a way that it borders at least one other piece. The sides facing each other must be identical. The game ends as soon as one of the players cannot place another piece on the board. The other player is the winner.

beyond a certain degree of complicatedness, our logical apparatus ceases to cope. In other words, our rationality is *bounded*. Social scientists have been aware of this problem ever since Herbert Simon suggested “satisficing” as a way of describing less-than-logical behaviour in some decision situations. But Simon's notion of satisficing is too vague to be adopted as a practical method for solving complicated problems. Of late, economists have joined the search for something to put in place of deductive rationality. When we look into the growing literature on bounded rationality, however, there seems to be little agreement on a suitable successor.

The second reason for the deductive mode to break down is more ominous. In *interactive* decision situations, where the rationality of one agent's decision is dependent on the strategy of other agents, there are no guarantees that each agent will toe the line, i.e. behave with perfect rationality. Instead agents may be forced to guess the behaviour of other agents. Suddenly they’re plunged into a world of subjective beliefs, and subjective beliefs about subjective beliefs. Complete, consistent, well-defined premises are impossible under these conditions. Deductive reasoning breaks down for one very significant reason: the problem has become *ill-defined*.

A third reason is just as devastating. Even if one agent guesses the behaviour of others correctly at one point in time, there are no guarantees that this success will ever be repeated. In complex economic situations, agents learn and adapt differently. Thus future guesswork becomes more difficult. The evolutionary paths traced out by each agent are not familiar to other agents. Nevertheless, agents still make decisions in situations that are fuzzy or ill-defined. What’s even more surprising is that we seem to make them quite comfortably under such conditions. Perhaps we don't realize that the problem is ill-defined. But it's clear that we no longer reason deductively. A different kind of decision making process comes to the rescue. To find out what it is, we must look into psychology.

Let's begin with some basics from modern psychology which touch upon our area of interest. Aside from suggesting that we’re creatures of habit, psychologists claim that we make use of three varieties of reasoning: calculation, deduction, and induction. As we’ve just learnt, deductive logic is only useful in simple circumstances; like the family planning or residential choice examples we mentioned. What we're very good at, however, is recognizing or matching patterns.

When things get too complicated for our deductive powers, we seem to undergo a cognitive shift to the other side of our brain. Psychologists tell us that the right hand side handles pattern recognition, as well as intuition, synthesis, and creative insights. By putting a combination of these processes to work, we use the perceived patterns to fashion temporary constructs in our mind. We can call these constructs *mental models* or *hypotheses*. Once we have a set of such hypotheses firmly in our mind, they assist us to carry out "localized" deductions and act upon them.

But the whole process doesn't finish there. Our observations and experiences provide us with feedback from the decision environment; feedback that may alter the patterns we perceive, strengthening or weakening our confidence in our current set of hypotheses. What we're doing, of course, is trying to improve our ability to make prudent decisions; upgrading our know-whether, so to speak. We discard hypotheses that have proved to be unreliable, replacing them with new ones. We retain others. Wherever we lack full definition of the problem, we devise simple hypotheses to paper over the gaps in our understanding and we act on the best of these. This kind of behaviour is not deductive. It's *inductive*.[14]

If all of this sounds a little complicated, that's hardly surprising. Inductive reasoning *is* complicated. Even the psychologists don't fully understand it. Luckily, we can picture the inductive mind at work in a setting we've mentioned already: a Chess game. In the 1940s, the Dutch psychologist, Adrian de Groot, studied how Chess novices and Grand Masters perceive a Chess situation. He found that Grand Masters don't simply look further ahead than novices. Instead, they sharpen their intuition by studying the board's configuration and trying to discern chunks or patterns. In other words, they develop their own mental model of the board.[15] Then they use the perceived patterns and

---

[14] For an extensive discussion of inductive reasoning, and its prominent role in the learning process, see Holland, Holyoak, Nisbett and Thagard (1986).

[15] A revealing finding was that Grand Masters' mistakes involved placing whole *groups* of pieces in the wrong place, which left the game almost the same to the Master but bewildering to the novice. Even more revealing was the fact that, when pieces were assigned randomly to the squares on the board, instead of

their mental model to form hypotheses about their opponent's likely motives and strategies.

Chess openings are definitive patterns of play formed in the first dozen or so moves.[16] The leading players can even recall their opponent's favourite openings from previous games. "He's offering me the Queen's Gambit again." "Isn't this the Catalan opening?" "That looks like the modern Dragon Variation of the Sicilian defence." Good players carry out local deductions based on these mental models, analyzing the implications of alternative moves and their subsequent responses. As play proceeds, they hold onto the most plausible hypotheses and toss away the others, replacing them with new ones as the state of the game dictates.

{Figure 2.3 near here}

Clearly then, Chess players engage in a sequence of reasoning that's inductive. This includes pattern formation, pattern recognition, hypothesis formation, deduction using the currently-held hypotheses, and hypothesis replacement as dictated by the patterns that unfold (see Figure 2.3). Seasoned players build on their experiences from earlier games. Chess is a strongly-interactive game because players only learn *during the game* which of their hypotheses work best. It involves *adaptive learning* rather than information-processing. Each player's strategy evolves partly in response to the evolutionary path chosen by his opponent. Neither player can afford to adopt a fixed strategy. They must be flexible and "roll with the punches," so to speak. In a word, strategies need to be *coevolutionary*..

It's interesting to note that Hermann Haken, founder of the field of *synergetics*, believes that the kinds of pattern recognition used in Chess are closely associated with pattern formation. Synergetics is a general theory of self-organization. Although its generality comes from timescale separation and a slaving principle, the formalism of

---

being copied from actual games, Grand Masters were found to be no better than novices in reconstructing such random boards.

[16] For a wide range of chess openings, see Horowitz (1964).

synergetics allows us to calculate evolving patterns, provided the microscopic laws for the formation of patterns are known.[17] Take another look at Figure 1.1 (near the start of the book). Interestingly, neither birds nor antelopes can be perceived all the time. After a while, one image fades away, allowing the brain to perceive the other interpretation. Some kind of oscillatory process sets in.

Haken believes that decision making can be regarded as pattern recognition. Like Grand Masters staring at a Chessboard, all of us search for a resemblance between a situation that we now confront and one that we have met before. To do this, we establish a "similarity measure" in our minds, which allows us to choose a course of action that is the best under the given information. Haken and his colleagues have mimicked this pattern formation process inside their "synergetic computer." Thus they're also adopting an inductive approach to decision making.

Because it's closely connected with learning and adaptation, inductive behaviour sits firmly in the world of morphogenesis. In economics, the popular interpretation of "rational" behaviour connotes behaviour that's sensible or sound-minded. Deductive reasoning is only sensible or sound-minded in fairly simple, well-defined problems. Once a situation gets too complicated or ill-defined, like in a Chess game, an intelligent person begins to reason inductively. Induction enters the scene whenever someone has to derive a whole solution from partial information. In the social sciences, for example, induction is widely used in the analysis of opinion surveys and macroeconomic data.

Surprisingly, there's now a third way of doing social science. It corresponds to the third mode of reasoning cited by psychologists: calculation. But the calculations are done by machines instead of humans. This growing focus on calculation goes by the name of *agent-based* computer modeling or *simulation*. Like deduction, agent-based simulation starts off with a set of assumptions. Unlike deduction, however, it doesn't prove theorems. Instead, an agent-based model generates simulated data that can be analyzed inductively. But the simulated data come from a rigorously specified set of rules rather than from direct measurements of the real world.

---

[17] For an introduction to synergetics, see Haken (1977).

Whereas the purpose of deduction is to find consequences of assumptions, and that of induction is to find patterns in data or real-world experiences, agent-based modeling is a way of doing thought experiments that help to sharpen our intuition.[18] The examples to be discussed in this and later chapters have the same collective properties as those we found in Schelling's segregation model. Locally interacting agents can produce large-scale effects, most of which turn out to be far from obvious.

Economists have pushed the assumption of deductive reasoning beyond its limits. In doing so, they've locked most of us into the world of stasis. But this frozen world is only a tiny part of the whole universe of economic behaviour. From an intellectual viewpoint, we've been trapped in an era of tomfoolery. We might even call it the "Age of Deception". Thankfully a new light is glowing in the distance. The source of this light is a group of social scientists who just happen to believe that human agents reason inductively and adaptively. Furthermore, they also believe that agent-based simulation represents a promising new way of doing social science, one that can help us to unravel some of the complexities of human behaviour. The search for a new age of human enlightenment is now underway.

{A}Seeing the Light at the El Farol{/A}

El Farol is a bar on Canyon Road in Sante Fe which offers Irish music every Thursday evening. Being born in Belfast, Brian Arthur was fond of going to hear the music and to enjoy a few beers in a relaxed atmosphere once a week. But soon he encountered a thorny problem. If the bar was too crowded, the chances of brushing up against a few too many pushing-and-shoving bar louts were high. This would spoil the night and cause him to think twice about going the following Thursday. Arthur realized that he needed a more reliable method of deciding whether the bar was likely to be overcrowded each coming Thursday night.

[18] The growing importance of computer simulations can be gauged from the improved performance of silicon "agents" in high-level Chess games. Since 1997, when IBM's Deep Blue defeated Garry Kasparov in their second round of games, this third way of reasoning in open-ended situations has been recognized as full of promise.

The reader might like to ponder this problem for a moment. As we'll see shortly, it turns out to be an instructive example of a complex adaptive system. To make it more concrete, let's suppose that there are 100 people in Sante Fe who, like Arthur, are keen to go to the El Farol on Thursdays. Space is limited and everyone enjoys themselves if the bar is not too crowded. A crowd beyond sixty is thought to be excessive. The tricky thing is that there's no way of telling how many will come beforehand. A person simply goes if he expects fewer than sixty to turn up or stays home if he expects more than sixty to show.

Arthur has highlighted two interesting aspects of this problem. First, if there was an obvious model that all agents could use to forecast bar attendance, then a deductive solution would be possible. But there's no such model. Irrespective of past attendance figures, a wide range of plausible hypotheses could be adopted to predict future attendance. This dastardly multiplicity of possibilities means that nobody can choose in a well-defined manner. The problem becomes ill-defined and all the potential bar attendees are catapulted into a world of induction.

Second, any shared expectations will tend to be broken up. If all music lovers believe *most* will go, then *nobody* will go. But by all staying home, that common belief will be destroyed immediately. On the other hand, if all of them believe *few* will go, then *all* will go, thereby undermining that belief. The net result of this diabolical state of affairs is that expectations must differ.

Perplexed yet fascinated by this intractable problem, Arthur decided to turn his computer loose on it.[19] By creating a surrogate El Farol bar inside his machine to study how electronic music lovers would act in this situation, he stepped into the exciting new realm of agent-based computer simulation. All of his music-loving "agents" were given Thursdays' bar attendance over the past few months. For example, typical attendance figures might be:

... 44, 78, 56, 15, 23, 67, 84, 34, 45, 76, 40, 56, 22, 35

[19] For a full description of his computer experiments, see Arthur (1994a).

With this information at hand, each electronic agent has to keep track of a different subset of predictors (or hypotheses). He opts to go or stay home each Thursday according to the currently most accurate predictor in his set. Typical predictors might include the following:

- the same number as last week's (35)
- a mirror image around 50 of last week's (65)
- a rounded average of the last four weeks (38)
- the same as two weeks ago (22)

Once decisions have been taken, surrogates converge on the silicon bar and a new attendance figure is recorded. Each person reexamines the accuracy of his set of predictors, replacing the poorer ones with more reliable predictors. Then the whole decision process is repeated.

Fondly enough, the set of predictors deemed most credible and acted upon by each person - which Arthur calls the set of *active* predictors - determines the attendance. But the attendance history also determines the set of active predictors. This process of cumulative causation creates what John Holland has called an *ecology* of predictors. One of the aims of Arthur's computer experiments was to find out how this ecology evolves over time. So he created an "alphabetic soup" of several dozen predictors and randomly "ladled" out various mixes of these to each of the 100 persons.

As long as the predictors are not too simplistic, the simulations show that the weekly attendance will fluctuate, but mean attendance always converges to sixty. The predictors self-organize themselves into an equilibrium "ecology" in which 40% of the active predictors forecast above sixty and 60% of them below sixty. This happens despite the fact that the population of active predictors keeps changing in membership forever. Such an emergent ecology is more like a forest whose contours do not change, but whose individual trees do.

[Fig. 2.4 near here]

But there’s also another intriguing result. The computer-generated attendance results look more like the outcome of a random process rather than a deterministic one (see Figure 2.4). Yet there’s no inherently random factor governing how many people show up. Weekly attendance is a purely deterministic function of the individual predictions, which themselves are deterministic functions of the past attendance figures. Curiously, the existence of a statistical regularity might be attributed to the deterministic nature of chaos. In other words, the time-series of attendance figures might well be a deterministically random (i.e. chaotic) process. The irritating thing is that we have no mathematical formalism with which to either prove or disprove this conjecture.

What does this kind of emergent simplicity tell us? It confirms that a system of interacting people (in this case, bar attendees) can "spontaneously" develop collective properties that aren’t obvious from our knowledge of each of the individuals themselves. These statistical regularities are large-scale features that emerge purely from the microdynamics. As Jack Cohen and Ian Stewart have stressed: "Emergent simplicities collapse chaos; they bring order to a system that would otherwise appear to be wallowing hopelessly in a sea of random fluctuations."[20]

The El Farol problem contains all the essential elements of a *complex adaptive system*. First, it involves a "largish" number of agents, where "largish" denotes a number too large for hand-calculation or intuition but too small to call upon statistical methods applicable to very large populations. Second, it involves agents who are *adaptive* and *intelligent*. Such agents can take decisions on the basis of mental models (like the El Farol predictors) and are willing to modify these mental models or come up with new ones where necessary. In other words, they can reason *inductively*. Third, no single agent knows what all the others are thinking of doing, having access to a limited amount of information only. The El Farol case is extremely tight as each agent only knows what he or she is thinking of doing.

Conventional economic wisdom would tell us that these agents have only one reasoning skill: the ability to process the information available to them in a purely logical, deductive manner to arrive at the best decision in a given situation. But this is useless when the best thing to do - to go or not to go - depends on what everyone else is doing. There's no optimal predictor. The best each agent can do is to apply the predictor that has worked best so far, to be willing to revealuate the effectiveness of his set of predictors, and to adopt better ones as new information about bar attendance becomes available. The latter is the inductive part of the decision process. This is the way that Arthur's surrogate music lovers behave in his silicon world, which is our world of morphogenesis.

Three results of the El Farol problem are significant. First, the computer experiments show that inductive reasoning can be modelled. Second, they show that agents' belief systems should be thought of as evolving and coevolving. Third, they suggest that under the influence of a sufficiently strong "attractor", individual expectations that are boundedly-rational can self-organize to produce collectively "rational" behaviour. But there's also an even stronger message for economists. Learning and adaptation should not be addenda to the central theory of economics. They should be at its core, especially in problems of high complexity.

What's more daunting is the idea that those silicon agents in Arthur's computer experiments may have developed superior intuition (i.e. know-whether) to the real bar-lovers on which they're based. Ever since Deep Blue defeated Kasparov in their second round of chess games, the frightening idea that "calculative" reasoning might be able to outperform human reasoning (based on deduction or induction) in open-ended situations has been recognized.

One perplexing issue remains. How and why do the predictors self-organize so that sixty emerges as the mean attendance in the long run? Was it simply because Arthur picked sixty as the crowding threshold? If the threshold had been seventy instead of sixty, the simulations may have showed that the mean attendance would have converged to seventy. What would have happened if the assumption of a uniform crowding threshold had been dropped completely? After all, the heterogeneities of human thinking and decision making suggest that music lovers could never agree on the same figure to define a bar's crowding threshold. Would sixty still emerge under these more realistic conditions?

---

[20] See Cohen and Stewart (1994), page 232.

Arthur's explanation for his result is that sixty may be a natural "attractor" for the microdynamics in this bar problem. If we view the problem as a pure prediction game, then a mixed strategy of forecasting above 60 with probability 0.4 and below it with probability 0.6 corresponds to what game theorists call a *Nash equilibrium*. Could this be the reason the forecasts split into a 60/40 ratio? Although this explanation illuminates the end of the journey, it fails to explain the means of achieving that end. Given each agent's subjective reasoning, it's still impossible to explain the collective outcome. In the next section, therefore, we'll look more closely at Nash equilibria in both a static *and* a dynamic context.

{A}The Emergence of Cooperation{/A}

Life is literally teeming with perplexing problems, dilemmas and paradoxes - not all of which are abstract and philosophical. Rather than being a source of frustration, some paradoxes are superbly enlightening. We savour the moment when the truth dawns upon us. One paradox which hinges on the quirks of human nature is the game-theorist's favourite game, the Prisoner's Dilemma. This tantalizing puzzle was discovered in the 1950s by Merrill Flood and Melvin Drescher of the RAND Corporation, two early game theorists who were testing some bargaining theories with experimental games.[21]

Social scientists have become quite fond of the game, seeing Prisoner's Dilemma situations arising everywhere in socio-economic interactions. It's a surprisingly ubiquitous metaphor. As Russell Hardin notes, if the dilemma had been called "exchange" originally, then everyone would've expected it to be ubiquitous.[22] In economics, we view exchange as a two-party affair. But exchange situations can involve

---

[21] The Prisoner's Dilemma game first appeared in 1952, in an unpublished memorandum from the Rand Corporation discussing some experimental games. It remained unnamed in those early days, until A.W. Tucker called it the Prisoner's Dilemma. Tucker saw in it a perverse analog of the American criminal justice system, where prosecutors extract confessions on the promise of reduced sentences (plea bargaining).

[22] For a comprehensive discussion of the Prisoner's Dilemma and the insights it can provide for various problems of collective action in social and economic contexts, see Axelrod (1984) and Hardin (1982). These two authors have recently updated their original work; see Axelrod (1997) and Hardin (1995).

more than two parties. Think of the bidding by competing parties at an auction, verbal exhanges between participants at a public meeting, or the debate between political parties leading up to an election. In its many-person or collective guise, exchange is a very interesting problem. Although less tractable than the traditional two-party problem, it captures the perversity of the logic of collective action.[23] Under this logic, a group of people with a common interest that requires a common action may share an interest collectively but not individually.

Before we get immersed in a debate over collective versus individual behaviour, it's worthwhile getting more closely acquainted with the Prisoner's Dilemma. Let's sidestep the original formulation about prisoners and gaol terms because it can baffle the uninitiated. Instead we'll look at this puzzling paradox in the form of an economic metaphor: the "Trader's Dilemma." The scene for this metaphor is Medieval Europe. Imagine yourself as an enterprising merchant from Venice in the early days of Mediterranean trade. Venetian salt is what you have in abundance. Instead of trading it for grain and linen, on your next voyage you opt for something different. The latest rage is precious silk from Thebes (near Athens), so you stop there on your way from Venice to Constantinople.

You find a silk dealer whose terms are acceptable. But disappointment follows when you learn that he has a binding agreement with another Italian merchant. He's willing to trade, but only if the trade takes place secretly. He can't risk being seen dealing with you. So you agree to leave your salt consignment in bags at a well-concealed spot in a nearby forest and to pick up the bags of silk at the silk dealer's designated place. Of course you'll have to leave more bags than him because salt has a much lower value than silk.

Given that the silk dealer is nervous about jeopardizing his existing agreement, it's pretty clear to both of you that this will be a "once-off" exchange. You're unlikely to meet again or have any further dealings with each other. Suddenly you realize that there's something for each of you to fear: namely that *the bags which you get could be empty*. Of

---

[23] See Olson (1965).

course there's no risk if you both leave full bags. But getting something for nothing would be even more rewarding. So you're tempted to leave empty bags.

Here's how you might think this through:

"If the silk dealer brings full bags, then I'll be better off leaving empty bags - because I'll get all the silk I want and keep all my salt. Even if the silk dealer brings empty bags, I'll still be better off leaving nothing - because that way I can never be cheated. *No matter what the silk dealer does*, my smartest move is to leave empty bags. So I'll leave empty bags."

By similar reasoning, however, the silk dealer reaches the same conclusion. So you both leave empty bags and come away empty-handed.

The result is obviously disappointing. In the jargon of the Prisoner's Dilemma, both of you chose *defection* over *cooperation.* If you'd both cooperated, you could have sailed away with precious silk and the silk dealer would own a cellar full of salt. There would have been smiles all round. Instead of that you have no silk. What went wrong? Why did logical reasoning rule out cooperation? This is the perplexing aspect of the Prisoner's Dilemma.

A revealing theorem by John Nash, an early pioneer in game theory, threw some light on this surprising outcome. He showed that there's always at least one "Nash" strategy for each player, with the property that if each player chooses that strategy, he or she will be better off than with any other strategy. But this situation holds only if all the other players opt for their Nash strategies. A choice of Nash strategies among all the players is called a Nash equilibrium. No doubt you've guessed the connection by now. The decision of both traders to leave empty bags corresponds to a Nash equilibrium.

To pinpoint a Nash equilibrium precisely, we need to assign some numbers to our problem. But how do we quantify it? According to game theory, we define a set of payoffs to each of the players. You and the silk dealer have a pair of strategies to choose from: you can leave either full or empty bags. The payoff for each of you depends on the strategy chosen by the other. To display these various alternatives, we define a *payoff matrix* containing point values for the different pairs of strategies. Some typical point values are shown in Figure 2.5.

[Fig. 2.5 near here]

How do we interpret the figures? Payoff doesn't mean money since the exchange involves goods only. The numbers indicate the *degree of satisfaction* associated with each strategic outcome. For example, mutual cooperation scores two points to both of you. In this problem, two points means "quite happy". Both of you would be quite happy if all bags were full and you both got what you wanted. Mutual defection scores zero, and zero means you're "indifferent" to the idea of gaining and losing nothing.[24] If your worst fears were realized, and you got empty bags after leaving full ones, then you'd score -1 and be "feeling upset," while the silk dealer would score 4 and be feeling "very happy". These scores would reverse if you found full bags after leaving empty ones.

The Nash equilibrium solution to this game lies at P = (0,0). But this means zero payoff to both of you. Why should you choose zero reward? Because defection risks at worst indifference (0), whereas cooperation could leave you upset (-1). No matter what the silk dealer does, your safest *individual* strategy is to defect. Note that the payoff for mutual defection is lower than for mutual cooperation. Thus your best *joint* strategy is to cooperate. Now you can see the dilemma. Because you're unlikely to meet again, the best solution is to defect - despite the seemingly paradoxical outcome that it would be *collectively* superior for you both to cooperate.

Stuart Kauffman sums up the Nash concept as a penetrating one: "The concept of Nash equilibria was a remarkable insight, for it offers an account of how independent selfish agents might coordinate their behaviour without a master choreographer."[25] Despite its allure, however, a Nash equilibrium "solution" to the Trader's Dilemma has some major weaknesses. First, it sits stubbornly in the world of stasis. In the case where the Trader's Dilemma is played only once, a Nash solution is stable, predictable and resistant to change. No other strategy can invade the strategy of pure defection. Second, it relies on all traders thinking and acting

[24] If your perilous journey to reach Thebes was taken into account, however, you may feel that zero is far too generous.

[25] See Kauffman (1995), page 218.

rationally and identically. Like the deductively rational economic agent we discussed earlier, traders are assumed to have the brainpower and enough information to figure out their optimal strategy. Third, it may not be the solution in which the payoff to traders is particularly good.

Even if these weaknesses could be overlooked (which they can't), there's still another problem. What happens in a large game with many players? How would you approach a market boasting a dozen or more silk traders and just as many salt merchants desperate to out-perform you? There could be many more possible strategies and many more Nash equilibria. How would you find all these Nash strategies and then pick the best Nash equilibrium? Would this Nash payoff be worth pursuing in any case? Like Thurday nights at the El Farol bar, your problem quickly becomes complicated and ill-defined.

Setting these problems aside for the moment, let's return to the scene of our Trader's Dilemma. By now your ship has been reloaded with the untraded salt. What have you learnt from this abortive experience? What might you do next? You've discovered that for a *single* exchange conducted in secret, the temptation for both parties to cheat is irresistible. Now you recall your normal recipe for success: long-term, bilateral agreements. You go in search of another silk dealer who shares your desire for lifelong exchanges. Eventually you find a young dealer who accepts your long-term proposal. You agree to exchange fixed amounts every quarter but are unlikely to meet again face-to-face.

What do you do on the occasion of your first exchange? Leaving empty bags would hardly be a friendly way of fostering goodwill with a new trading partner. So you leave full bags. So does your silk dealer. What a relief! Three months pass and then you must go again. Empty or full? Every quarter you must take the decision whether to cooperate or defect. Two years later, the silk dealer defects unexpectedly. What will you do now? Can he ever be trusted again or will you call a halt to all future exchanges with him?

In the literature, the game you're now playing is known as the *Iterated* Prisoner's Dilemma. Just as adding more players creates complications, so does allowing *repeated* exchanges. But it also adds more realism. The trading world has always featured long-term agreements, cartels, price-fixing and other multilateral trading arrangements. Furthermore, traders engaging in repeated exchanges have always shown some degree of cooperation. And they also tend to review their strategies regularly over time. So it makes alot of sense to study the iterated version of the Prisoner's Dilemma in order to understand the exact conditions under which cooperation might emerge.

The million-dollar-question is: *Can cooperation ever evolve out of noncooperation*? Well, the answer turns out to be a resounding **yes**. Emergent cooperation has been demonstrated by a novel method: a computer tournament organized by Robert Axelrod, a political scientist at the University of Michigan. Cooperation won out among a diverse population of computer programs playing repeated games of the Prisoner's Dilemma with one another. After the tournament was over, Axelrod spotted the salient principles and proved theorems that could explain cooperation's rise from nowhere. His findings have been published in many papers and two thought-provoking books.[26]

Axelrod sent out invitations to game theorists in economics, sociology, political science, and mathematics. The rules implied a similar payoff matrix to the one shown earlier for the Trader's Dilemma. Submitted programs were designed to respond to the "cooperate" or "defect" decisions of other programs, taking into account the remembered history of previous interactions with that particular program. Fourteen entries duly arrived. Axelrod added another called RANDOM, which simply flipped a coin each time it met another program. Each program was made to engage each other program 200 times.

Although some of the entries were quite sophisticated, the winning program turned out to be the *simplest* of all the strategies submitted. Known as TIT FOR TAT, it simply cooperates on the first move and then does whatever the other player did on the preceding move. That's all there is to it. TIT FOR TAT was written by the psychologist and philosopher, Anatol Rapoport, who turns out to be an old hand at the Prisoner's

[26] See Axelrod (1984, 1997).

Dilemma game.[27] Rapoport's program was also the shortest of all the programs submitted. Small is beautiful after all!

Given this surprising outcome, the full results were circulated and entries for a second tournament were solicited. Axelrod provided a few hints this time. He pointed out that many of the losing strategies suffered from self-punishment because such a possibility was not perceived by their decision rules. He also stressed that TIT FOR TAT was a strategy of cooperation based on reciprocity and that many of the other strategies were not forgiving enough.

This time Axelrod received 62 entries from 6 countries. Most came from computer hobbyists (including one ten year old). Rapoport's TIT FOR TAT was there again. So was a variation on the same theme called TIT FOR TWO TATS, which tolerates two defections before getting mad (but still only strikes back once). Axelrod already knew that TIT FOR TWO TATS would have won the first tournament if it had been in the line-up. So it was sure to be hard to beat, and was entered by one of the world's experts on game theory and evolution, John Maynard Smith, professor of biology at the University of Sussex.

Can you guess what happened? Amazingly, TIT FOR TAT won again. What a remarkable result! How on earth could such a simple decision strategy defeat so many other strategems devised by all those whizzkids? Axelrod attributes TIT FOR TAT's success to its being:

nice (not the first to defect);
provokable (responding to the other player's defection);
forgiving (punishing and then cooperating after a defection);
clear (easy for other players to understand).

So, nice guys, or more precisely, nice, provokable, forgiving and clear guys, can indeed finish first. But success in two computer tournaments is hardly enough to prove

[27] For an overview of the Prisoner’s Dilemma Game, along with some of the background to TIT FOR TAT, see Rapoport and Chammah (1965).

that TIT FOR TAT would do well as an evolutionary strategy. To test this possibility, Axelrod conducted a series of *ecological* simulations, with various tournament entries and other strategies as his starting population. He found that TIT FOR TAT quickly became the most common strategy. Thus the conclusions seemed clear. A small cluster of players who opt for cooperation based upon reciprocity could establish themselves in a population of noncooperative players. Once established, they could become immune from re-invasion by any other strategies and could thus take over such a population.

The Prisoner's Dilemma game captures the tension between the advantages of selfishness in the short term versus the need to elicit cooperation to be successful in the long run. But the jury's still out on whether such voluntary cooperation is sustainable over the long run. Many analysts (like Axelrod) believe that TIT FOR TAT could be a robust and stable mutant; even an *evolutionarily stable strategy* under certain conditions. But others have challenged this idea. One *Nature* article showed that no pure strategy can be evolutionarily stable in the iterated Prisoner's Dilemma.[28] Another *Nature* article reported that the most successful strategy was one that repeats its previous choice when it gets one of the two highest payoffs.[29] This ongoing debate suggests that there's much more to be learnt about the evolution of cooperation, especially under different socio-economic conditions.

Nevertheless, it's nice to know that benign cooperation among selfish agents can emerge despite the constant temptation to defect. In Axelrod's words:[30] "Mutual cooperation can emerge among a world of egoists without central control, by starting with a cluster of individuals who rely on reciprocity." and also those of Hardin:[31] "But coordination can come about without intent,

[28] Boyd and Lorberbaum's (1987) claim that no pure strategy is evolutionarily stable was disputed recently by Bendor and Swistak (1998), who found that strategies that are nice and retaliatory (like TIT FOR TAT) are the most stable against possible invasions.

[29] TIT FOR TAT is sensitive to the occurrence of mistakes or misunderstandings, commonly called noise. Noise played a part in the simulation study reported in *Nature* by Nowak and Sigmund (1993). Unlike TIT FOR TAT, their WIN-STAY, LOSE-SHIFT strategy defects after the other player is exploited and cooperates after a mutual defection. Axelrod's reaction to their study was to feel a little protective of TIT FOR TAT. With the help of a postdoctoral fellow, he found that adding either generosity or contrition to TIT FOR TAT was an effective way of coping with noise. Furthermore, the WIN-STAY, LOSE-SHIFT strategy did not perform as well in this variegated environment. For further details, see Axelrod (1997).

[30] See Axelrod (1984).

without overcoming contrary incentives. It can just happen." In other words, cooperation can *self-organize* within a population, despite its members' biologically determined egoism. This has ramifications for both the economic and political arenas.

Many of the challenging problems facing humanity relate to globalization and international relations, where independent nations often refuse to cooperate, instead exhibiting stark hostility. Many of these problems closely resemble the iterated Prisoner's Dilemma. In addition to trade negotiations, arms races, crisis bargaining nuclear proliferation and environmental pollution fall into this category. By understanding the process of mutual cooperation a little better, perhaps we could use our foresight to speed up its evolution.

There's another interesting message for economics from the TIT FOR TAT story. A strategy's success depends entirely on the environment in which it's swimming. At the beginning of Axelrod's ecological tournaments, poor programs *and* good programs were well represented. But as they "swam" through generation after generation of interactions, this environment changed. The poorer strategies began to drop out and the better ones flourished. So it is with technologies. Their rank order changes because their "goodness" is remeasured alongside a different field of competitors. Success breeds further success, but only when that success comes from interaction with other reasonably successful technologies.

Doesn't this ring a bell or two? Remember those high-tech firms we discussed in Chapter 1. In the economic world of morphogenesis, firms that get ahead get even further ahead. Like good strategies in the Prisoner's Dilemma, they're self-reinforcing. TIT FOR TAT didn't win all those tournaments by *beating* the other players, but by eliciting behaviour from the others which allowed *both* to do well. This kind of mutual learning process isn't just evolutionary, it's *coevolutionary*.

Coevolution has been found in evolutionary versions of the Prisoner's Dilemma. Instead of starting like Axelrod did, with random selection from a rich set of strategies, Kristian Lindgren began with the simplest possible strategies. Using an extension of the genetic algorithm to evolve more and more complex possibilities, and allowing for the

[31] See Hardin (1995, page 45)

possibility of mistakes, he found that selection favors the evolution of cooperation and unexploitable strategies.[32] The result was long periods of stasis alternating with periods of instability, as one dominant pattern of strategies was invaded by another. Various kinds of evolutionary phenomena, like coexistence, punctuated equilibria, exploitation, the coevolution of mutualism, and evolutionarily stable strategies, were encountered in the model simulations.

{A}Coevolutionary Learning{/A}

How stunning it was when biology revealed that organisms don't just evolve, they *coevolve*. Their adaptation over time isn't shaped merely by their encounters with other organisms, it's also honed by the environment in which they live; and this environment isn't fixed but is also adapting to the behaviour of its changing inhabitants. What might this imply for agents in an economy? In behavioural terms, it suggests that what agents believe affects what happens to the economy <u>and</u>, in turn, what happens to the economy affects what agents believe. This, in fact, is the hypothesis explored in this book. The agents, goods and services in the economy coevolve, because those that are present must always make sense in the context of all the others that already exist. Diversity breeds more diversity, thereby fueling the growth of complexity.

There's another way of recognizing the fundamental difference between evolution and coevolution. Stuart Kauffman couches it in terms of fitness landscapes.[33] Evolution occurs on a fixed landscape where the attractors are local optima in the form of single points. This kind of landscape is a familiar one to economists, who are taught that optimization is a simple hill-climbing procedure. All that we thought we needed to know was the topography of the hill.

In a coevolutionary process, however, the landscape isn't fixed. Instead, it's adapting incessantly. In an economic environment, for example, we can picture the landscape of one agent heaving and deforming incessantly as other agents make their own

---

[32] For further details, see Lindgren (1992).

[33] For a comprehensive discussion of fitness landscapes, see Kauffman (1993;1995, especially Chapter 8).

adaptive moves. The resulting environment has a very unstable surface. In fact, it's more like a seascape than a landscape. Have you ever tried to surf on an incessantly choppy and undulating ocean surface? That's the kind of decision environment that exists when agents must adapt incessantly. Nobody quite knows where they're going next. It's even difficult to define the problem. For example, technological evolution is a process attempting to optimize a system riddled with conflicting constraints. In such an uncertain environment, the various behavioural regimes that may coevolve are mindboggling. The likelihood of reaching local optima in the form of point attractors is very remote. *Coevolving systems may not be optimizing anything*.

What *are* they doing then? One insight comes from Axelrod's results. Sometimes coevolution allows TIT FOR TAT-style cooperation to emerge and thrive in a world full of treacherous defectors. The iterated Prisoner's Dilemma game simultaneously provides an abstract model for the evolution of cooperation and the setting for a very complex evolutionary landscape.[34] Another insight comes from Arthur's inductive economic models in which agents learn and adapt. Inductive agents, who persistently alter their mental models of other agents' behaviour, will decide and behave differently. They're forever changing their mental images of others. These mental images are often nothing more than subjective expectations or half-hoped anticipations. They can be mutually cooperative or mutually competitive. They can arise, get a solid footing, gain prominence, fall back, and disappear. Arthur regards them as the DNA of an economy.

Whenever beliefs form a complex ocean of interacting, competing and cooperating, arising and decaying entities, it's possible that they may simplify into an ordered equilibrium now and then. But most of the time they'll be found in complex, unsettled, ever-changing states. Beliefs about beliefs are mostly volatile. There's no evidence to suggest that such adaptive behaviour ever settles down into a steady, predictable pattern. This is a another signature of *coevolutionary learning*.

The light is beginning to shine at last! When economic agents interact, when they must think about what other agents might be thinking, their coevolving behaviour can take a wide variety of forms. Sometimes it might look chaotic, sometimes it might appear to be ordered, but more often than not it will lie somewhere in between. At one distant end of the spectrum, chaotic behaviour would correspond to *rapidly* changing models of other agents' beliefs. If beliefs change too quickly, however, there may be no clear pattern at all. Such a volatile state could simply appear to be random. At the other end of the spectrum, ordered behaviour could emerge, but only if the ocean of beliefs happens to converge onto a mutually consistent set of models of one another. One familiar example is that classical pillar of the world of stasis, a state of equilibrium among a set of deductively rational agents.

For most of the time, however, we'd expect that mental models of each other's beliefs would be *poised* somewhere in between these two extremes, tending to change, poised ready to unleash avalanches of small and large changes throughout the whole system of interacting agents. Why should we expect this? Given more data, we would expect each agent to improve his ability to generalize about the other agents' behaviour by constructing *more complex* models of their behaviour. These more complex models would also be more sensitive to small alterations in the other agents' behaviour. Thus as agents develop more complex models to predict better, the coevolving system of agents tends to be driven away from the ordered regime toward the chaotic regime. Near the chaotic regime, however, such complexity and changeability would leave each agent with very little reliable data about the other agents' behaviour. Thus they would be forced to simplify, to build *less complex* models of the other agents' behaviour. These less complex models are less sensitive to the behaviour of others and live in calmer oceans.

Thus we can picture a constant struggle between the need to simplify and the need to "complexify" our thinking. This is tantamount to a struggle between the two halves of our brain. Being more objective, rational and analytical, the left hand side houses the simpler confines of convergent thinking. It's also responsible for the deductive metaphor among economic agents. By way of contrast, the right hand side is more subjective, intuitive and holistic. This kind of divergent thinking produces multiple outcomes and creative ideas. Thus the right hand side is responsible for inductive reasoning among economic agents. We can picture these two modes of thought in constant interplay,

[34] See Lindgren (1992).

driving the whole ocean of beliefs back and forward, from chaos to order and back again. In general, therefore, we might expect to find most economic agents hovering somewhere in between. Poised, if you like, near "the edge of chaos".

Believe it or not, such poised states proliferate in our everyday world. Here's one example from my own backyard. For several years, my ten-year-old daughter, Sofie, and I have enjoyed the habit of cycling together on a bayside circuit near our Melbourne home. While pedalling around this predetermined circuit, we often play that simple guessing game which kids everywhere seem to adore: "I spy with my little eye, something beginning with......." For the first few months, we played the standard game: we tried to guess each other's word given its first letter. Then Sofie began to ask for a clue if the word turned out to be elusive. A few months later, we allowed two clues under certain conditions. Then came a major "mutation". Our population of potentially guessable "things" expanded dramatically when we allowed two-word descriptions to join single words. First, it was just nouns needing two words to define them (e.g. ice cream). But very soon the set of possibles was expanded again. We included adjective-noun pairs (e.g. red car). Can you guess what was happening? Our little game of "I spy" was *coevolving* all the time. It seemed to be forever poised on the verge of an avalanche of smaller and larger changes. A bit like those sandpiles we discussed in Chapter One.

Progress in science is a fertile example of coevolutionary learning. Scientists develop their own models or hypotheses about a particular phenomenon of interest. Then they test these ideas in a variety of ways: performing experiments in laboratories, discussing their results with colleagues, lecturing about them at seminars and conferences, and publishing them in journals and books. From each of these critical audiences, they get feedback. Usually they're obliged to revise their ideas in the light of this feedback, to discard some of their old models and replace them with new ones. Like Chess players, scientists discern patterns, build temporary mental models, test them in a competitive environment, revise them in the light of new information, and come up with improved hypotheses. They're immersed in an ocean of ideas, where each one is evolving and coevolving. Scientists are experts at coevolutionary learning.

As we noted earlier, researchers understand the importance of oral communication when their own ideas are exposed for critical comment. One of the great advantages of face-to-face contact is that it fosters the unexpected! Unsolicited comments from thoughtful peers can be pearls of wisdom. Interestingly, face-to-face contact is dependent on the transportation system. And the world of transport is another arena where coevolutionary learning is hard at work. A traffic system supports a large number of interacting vehicles and drivers, and the behaviour of both can be complex and unpredictable. As we shall learn in a later chapter, busy traffic may be poised near the edge of chaos.

TABLE 1.1:
Two Economic Worlds - The Simple and the Complex

NECESSITY	CHANCE
Stasis	Morphogenesis
Resource-Based	Knowledge-Based
Unique Outcome	Multiple Outcomes
Equilibrium	Path-Dependent
Mechanistic	Organic
Predictable	Unpredictable
Diminishing Returns	Increasing Returns
Convex	Nonconvex
Easy to Model	Difficult to Model
A SIMPLE WORLD	A COMPLEX WORLD

TABLE 2.1:
Information and Knowledge

Characteristic	Information	Knowledge
Source	External	Internal
Nature	Weakly-interactive	Strongly-interactive
Primary exchange mode	Interface	Face-to-face
Learning rate	Fast	Slow
Usefulness	Temporary	Longlasting
Exchange process	Simple	Complex
Unit of measurement	Quantitative (e.g. bits)	Qualitative (e.g. deep)

TABLE 4.1
Population Growth in Europe

Date	European Population	Margin of Error (%)
200	48	35
500	36	30
800	32	30
1000	39	20
1300	75	20
1500	76	10
1700	102	8

TABLE 4.2:
The Ten Largest Cities in Europe by Population, 1000-1400

1000	1100	1200	1300	1400
Cordova	Constantinople	Constantinople	Paris	Paris
Constantinople	Fez	Palermo	Granada	Bruges
Seville	Seville	Seville	Constantinople	Milan
Palermo	Palermo	Paris	Venice	Venice
Kiev	Cordova	Venice	Genoa	Genoa
Venice	Granada	Cordova	Milan	Granada
Thessalonika	Venice	Granada	Sarai	Prague
Ratisbon	Kiev	Milan	Seville	Constantinople
Amalfi	Salerno	Cologne	Florence	Rouen
Rome	Milan	London	Cologne	Seville

Table 5.1: Changes in Rank of Selected American Cities, 1810-1910

City	Rank in 1810	Rank in 1860	Rank in 1910
New York	1	1	1
Philadelphia	2	2	3
Baltimore	3	3	7
Boston	4	4	5
New Orleans	6	5	14
Cincinatti	42	6	13
St. Louis	-	7	4
Chicago	-	8	2
Buffalo	-	9	10
Louisville	-	10	22
Albany	17	11	44
Washington	12	12	16
San Francisco	-	13	11
Providence	8	14	21
Pittsburgh	28	15	8
Rochester	-	16	23
Detroit	-	17	9
Milwaukee	-	18	12
Cleveland	-	19	6
Charleston	4	20	77

Table 5.2: Similarities Between CAs and Socio-Economic Dynamics

	Cellular Automata	Socio-Economic Dynamics
Basic elements	Cells are the basic units or "atoms" of a CA	Individual agents are the basic units of an economy
Possible states	Cells assume one of a set of alternative states	Agents form mental models which enable them to make choices from alternatives
Interdependence	The state of a cell affects the state of its closest neighbors	The choices made by agents affect the choices made by other agents
Applications and tasks	Modeling the emergence of order, macro outcomes explained by micro rules, and the path dependence of dynamic processes	Important tasks include: understanding the emergence of order, macro to micro relationships, and economic dynamics

Bibliography

Albin, P. (1975) The Analysis of Complex Socio-Economic Systems. Lexington, MA: Lexington Books.

Allen, P. and M. Sanglier (1979) A Dynamic Model of Growth in a Central Place System, Geographical Analysis, vol. 11, pp. 256-272.

______ and ______ (1981) Urban Evolution, Self-Organization, and Decisionmaking, Environment and Planning A, vol. 13, pp. 167-183.

Andersson, Å.E. (1986) "Presidential Address: The Four Logistical Revolutions," Papers of the Regional Science Association, vol. 59, pp. 1-12.

______ (1995) "Economic Network Synergetics", in Batten, D.F., Casti, J.L. and R. Thord (eds.), Networks in Action, Berlin: Springer Verlag, pp. 309-318.

Andersson, Å.E. and O. Persson (1993) "Networking Scientists," The Annals of Regional Science, vol. 27, pp. 11-21.

Arnott, R., A. de Palma and R. Lindsay (1993) "A Structural Model of Peak-Period Congestion: a Traffic Bottleneck with Elastic Demand," American Economic Review, vol. 83, pp. 161-179.

Arrow, K. (1962) "The Economic Implications of Learning by Doing," Review of Economic Studies, vol. 29, pp. 155-173.

Arthur, W.B. (1994a) "Inductive Behaviour and Bounded Rationality", American Economic Review, vol. 84, pp. 406-411.

______ (1994b) Increasing Returns and Path Dependence in the Economy. Ann Arbor: University of Michigan Press.

______ (1995) "Complexity in Economic and Financial Markets," Complexity, vol. 1, pp. 20-25.

Arthur, W.B., Holland, J.H., LeBaron, B., Palmer, R. and P. Tayler (1997) "Asset Pricing under Endogenous Expectations in an Artificial Stock Market," in Arthur, W.B., Durlauf, S.N. and D.A. Lane, eds. The Economy as an Evolving Complex System II, Reading, Ma: Addison-Wesley, pp. 15-44.

Auerbach. F. (1913) "Das Gesetz der Bevölkerungskonzentration," Petermanns Geographische Mitteilungen, no. 59, pp. 74-76.

Axelrod, R. (1984) The Evolution of Cooperation. New York: Basic Books.

______ (1997) The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration. Princeton: Princeton University Press.

Bachelier, L (1900) Théorie de la Spéculation. Doctoral Dissertation in Mathematical Sciences, Faculté des Sciences de Paris, defended 29 March.

Bairoch, P. (1988) Cities and Economic Development: From the Dawn of History to the Present. (Translated by Christopher Braider) Chicago: The University of Chicago Press.

Bak, P. (1996) How Nature Works: The Science of Self-Organized Criticality. New York: Springer-Verlag.

Bak, P., Chen, K., Scheinkman, J.A. and M. Woodford (1993) "Aggregate Fluctuations from Independent Shocks: Self-Organized Criticality in a Model of Production and Inventory Dynamics," Ricerche Economiche, vol. 47, pp. 3-24.

Barrett, C.L., Thord, R. and C. Reidys (1998) "Simulations in Decision Making for Socio-Technical Systems," in M.J. Beckmann, B. Johansson, F. Snickars and R. Thord, eds. Knowledge and Networks in a Dynamic Economy. Berlin: Springer-Verlag, pp. 59-82.

Batten, D.F. (1982) "On the Dynamics of Industrial Evolution," Regional Science and Urban Economics, vol. 12, pp. 449-462.

______ (1995) "Network Cities: Creative Urban Agglomerations for the Twenty-First Century," Urban Studies, vol. 32, pp. 313-327.

______ (1998) "Coevolutionary Learning on Networks," in M. Beckmann, B. Johansson, F. Snickars and R. Thord, eds., Knowledge and Networks in a Dynamical Economy. Berlin: Springer-Verlag, pp. 311-332.

Batten, D.F. and B. Johansson (1989) "The Dynamics of Metropolitan Change," Geographical Analysis, vol. 19, pp. 189-199.

Batten, D.F., Kobayashi, K. and Å.E. Andersson (1989) "Knowledge, Nodes and Networks: An Analytical Perspective", in Å.E. Andersson, D.F. Batten and C. Karlsson, eds. Knowledge and Industrial Organization, Berlin: Springer-Verlag, pp. 31-46.

Batty, M., Couclelis, H. and M. Eichen (1997) "Urban Systems as Cellular Automata," Environment and Planning B, vol. 24, pp. 159-164.

Beckmann, M.J. (1958) “City Hierarchies and the Distribution of City Size,” Economic Development and Cultural Change, vol. 6, pp. 243-248.

______ (1994) “On Knowledge Networks in Science: Collaboration among Equals,” The Annals of Regional Science, vol. 28, pp. 233-242.

Bell, D. (1981) “Models and Reality in Economic Discourse,” in Bell, D. and I. Kristol (eds.) The Crisis in Economic Theory, New York: Basic Books.

Ben-Akiva, M., A. de Palma and I. Kaysi (1991) "Dynamic Network Models and Driver Information Systems," Transportation Research A, vol. 25, pp. 251-266.

Bendor, J. and P. Swistak (1998) “The Evolutionary Advantage of Conditional Cooperation,” Complexity, vol. 4, pp. 15-18.

Benhabib, J., ed. (1992) Cycles and Chaos in Economic Equilibrium. Princeton: Princeton University Press.

Berry, B.J.L. (1961) “City-Size Distributions and Economic Development, Economic Development and Cultural Change, vol. 9, pp. 573-588.

Biham, O., Middleton, A. and D. Levine (1992) "Self-Organization and a Dynamical Transition in Traffic-Flow Models", Physical Review A, vol. 46, pp. R6124-7.

Bloch, M. (1962). Feudal Society. (Translated by L.A. Manyon), London: Routledge & Kegan Paul.

Bossomaier, T. and D. Green (1998) Patterns in the Sand: Computers, Complexity and Life, Sydney: Allen and Unwin.

Bower, G.H. and E.R. Hilgard (1981) Theories of Learning, Englewood Cliffs: Prentice Hall.

Boyd, R. and J.P. Lorberbaum (1987) “No Pure Strategy is Evolutionarily Stable in the Iterated Prisoner’s Dilemma Game,” Nature, vol. 327, pp. 58-59.

Braess, D. (1968) “Über ein Paradoxon aus der Verkehrsplanung,” Unternehmensforschung, vol. 12, pp. 258-268.

Braudel, F. (1982) The Wheels of Commerce. London: William Collins.

Brock, W., Hsieh, D. and B. LeBaron (1991) Nonlinear Dynamics, Chaos, and Instability: Statistical Theory and Economic Evidence. Cambridge, MA: MIT Press.

Bunge, W. (1966) Theoretical Geography. Lund Studies in Geography, The Royal University of Lund, Series C, No.1.

Casti, J.L. (1989) Paradigms Lost, New York: Avon.

______ (1994) Complexification, London: Abacus.

______ (1997) Would-Be Worlds, New York: Wiley.

Chandler, A.D. Jr. (1965) The Railroads: The Nation's First Big Business. New York: Harcourt Brace Jovanovich.

Christaller, W. (1933) Central Places in Southern Germany. (Translated by Carlisle W. Baskin, Englewood Cliffs, N.J.: Prentice Hall, 1966).

Cohen, J. and I. Stewart (1994) The Collapse of Chaos, New York: Penguin.

Conquest, L., Spyridakis, J., Haselkorn, M. and W. Barfield (1993) "The Effect of Motorist Information on Commuter Behaviour: Classification of Drivers into Commuter Groups," Transportation Research C," vol.1, pp.183-201.

Cronon, W. (1991) Nature's Metropolis: Chicago and the Great West. New York: W.W. Norton.

Dafermos, S. and A. Nagurney (1984) "On Some Traffic Equilibrium Theory Paradoxes," Transportation Research B, vol. 18, pp. 101-110.

Daganzo, C. and Y. Sheffi (1977) "On Stochastic Models of Traffic Assignment," Transportation Science, vol. 11, pp. 253-274.

Darley, V. (1995) "Emergent Phenomena and Complexity," in Brooks, R.A. and P. Maes (eds.), Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems. Cambridge, MA.: The MIT Press, pp. 411-416.

Day, R.H. and J.L. Walter (1995) "Economic Growth in the Very Long Run: on the Multiple-Phase Interaction of Population, Technology, and Social Infrastructure," in Barnett, W.A., Geweke, J. and K. Shell (eds.), Economic Complexity: Chaos, Sunspots, Bubbles, and Nonlinearity. Cambridge: Cambridge University Press, pp. 253-288.

De Bondt, W. and R. Thaler (1985) "Does the Stock Market Overreact?" Journal of Finance, vol. 60, pp. 793-805.

Delong, J.B., Schleifer, A., Summers, L.H. and J. Waldmann (1990) “Positive Feedback and Destabilizing Rational Speculation,” Journal of Finance, vol. 45, pp. 379-395.

Dillard, D. (1967) Economic Development of the North Atlantic Community. New Jersey: Prentice-Hall.

Downs, A. (1962) "The Law of Peak-Hour Expressway Congestion," Traffic Quarterly, vol. 16, pp. 393-409.

Dunlap, R.A. (1997) The Golden Ratio and Fibonacci Numbers. Singapore: World Scientific.

Dyckman, T.R. and D. Morse (1986) Efficient Capital Markets and Accounting: A Critical Analysis. Englewood Cliffs, N.J.: Prentice-Hall.

Eiser, J.R. (1994) Attitudes, Chaos and the Connectionist Mind. Oxford: Blackwell.

Eldredge, N. and S. Gould (1972) “Punctuated Equilibria: an Alternative to Phyletic Gradualism,” in Schopf, T.J.M. (ed.) Models in Paleobiology, San Francisco: Freeman Cooper, pp. 82-115.

Elliott, R.N. (1946) Nature’s Law: the Secret of the Universe (Reprinted in R.N. Elliott, R.N. Elliott’s Masterworks: The Definitive Collection, Gainesville, GA: New Classics Library, 1994).

Epstein, J.M. and R. Axtell (1996) Growing Artificial Societies. Cambridge, MA.: The MIT Press.

Erdos, P. and A. Renyi (1960) “On the Evolution of Random Graphs,” Institute of Mathematics, Hungarian Academy of Sciences, Publication No.5.

Fama, E. (1970) “Efficient Capital Markets: a Review of Theory and Empirical Work,” Journal of Finance, vol. 25, pp. 383-417.

Farmer, R. A. (1993) The Macroeconomics of Self-Fulfilling Prophecies. Cambridge, MA.: MIT Press.

Fogel, R.W. (1964) Railroads and American Economic Growth. Baltimore, Maryland: Johns Hopkins Press.

Forrester, J.W. (1987) "Nonlinearity in High-order Models of Social Systems," European Journal of Operational Research, vol. 30, pp. 104-109.

French, R.M. and A. Messinger (1995) “Genes, Phenes and the Baldwin Effect: Learning and Evolution in a Simulated Population,” in Brooks, R.A. and P.

Maes (eds.), Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems. Cambridge, MA.: The MIT Press. pp. 277-282.

Frost, A.J. and R.P. Prechter (1985) Elliott Wave Principle: Key to Stock Market Profits. Gainesville, GA: New Classics Library.

Fujita, M. (1996) "On the Self-Organization and Evolution of Economic Geography," The Japanese Economic Review, vol. 47, pp. 34-61.

Gell-Mann, M. (1995) "What is Complexity?" Complexity, vol. 1, pp. 16-19.

Goodin, S.H. (1851) "Cincinatti – It's Destiny," in C. Cist, ed. Sketches and Statistics of Cincinatti in 1851. Cincinatti: William Moore.

Gorman, M.E. (1992) Simulating Science: Heuristics, Mental Models, and Techno-Scientific Thinking. Bloomington: Indiana University Press.

Greising, D. and L. Morse (1991) Brokers, Bagmen, and Moles: Fraud and Corruption in the Chicago Futures Markets. New York: Wiley.

Griffin, K. (1998) "Friction in Economics," in Åkerman, N. The Necessity of Friction. Boulder: Westview Press, pp. 119-131.

Grimmett, G. (1998) Percolation. Berlin: Springer.

Haag, G. (1994) "The Rank-Size Distribution of Settlements as a Dynamic Multifractal Phenomena," Chaos, Solitons, and Fractals, vol. 4, pp. 519-534.

Haag, G and H. Max (1993) "Rank-Size Distribution of Settlement Systems: A Stable Attractor in Urban Growth," Papers in Regional Science, vol. 74, pp. 243-258.

Haken, H. (1977) Synergetics: An Introduction, Berlin: Springer-Verlag.

______ (1998) "Decision Making and Optimization in Regional Planning," in M.J. Beckmann, B. Johansson, F. Snickars and R. Thord, eds. Knowledge and Networks in a Dynamic Economy. Berlin: Springer-Verlag, pp. 25-40.

Hardin, R. (1982) Collective Action, Baltimore: Resources for the Future.

______ (1995) One for All: The Logic of Group Conflict, Princeton, New Jersey: Princeton University Press.

Harker, P. (1988) "Multiple Equilibrium Behaviours on Networks", Transportation Science, vol. 22, pp. 39-46.

Hegselmann, R. (1996) “Understanding Social Dynamics: The Cellular Automata Approach,” in K.G. Troitzsch, U. Mueller, G.N. Gilbert and J.E. Doran, eds. Social Science Microsimulation. Berlin: Springer, pp. 282-306.

Hinton, G.E. and S.J. Nowlan (1987) “How Learning Can Guide Evolution,” Complex Systems, vol. 1, pp. 495-502

Holland, J.H (1988) “The Global Economy as an Adaptive Process,” in Anderson, P.W., Arrow, K.J. and D. Pines (eds.) The Economy as an Evolving Complex System. Reading, MA: Addison-Wesley, pp. 117-124.

Holland, J.H., Holyoak, K.J., Nisbett, R.E. and P.R. Thagard (1986) Induction: Processes of Inference, Learning, and Discovery, Cambridge, MA: The MIT Press.

Horowitz, I.A. (1964) Chess Openings: Theory and Practice. London: Faber and Faber.

Hume, D. (1957) An Inquiry Concerning Human Understanding. New York: Library of Liberal Arts.

Jacobs, J. (1969) The Economy of Cities. New York: Random House.

Johansson, B. and L-G. Mattsson (1995) "Principles of road pricing", in Johansson, B. and L-G. Mattsson (eds.) Road Pricing: Theory, Empirical Assessment and Policy, Boston: Kluwer.

Johnson, J. (1995) "The multidimensional networks of complex systems", in Batten, D.F., Casti, J.L. and R. Thord (eds.) Networks in Action, Berlin: Springer Verlag, pp. 49-79.

Johnson-Laird, P.N. and R.M.J. Byrne (1991) Deduction, Hove: Lawrence Erlbaum Associates.

Katz, J.S. (1993) “Geographical Proximity and Scientific Collaboration.” Scientometrics, vol. 31 (1): pp. 31-43.

Kauffman, S. (1993) The Origins of Order: Self-Organization and Selection in Evolution, New York: Oxford University Press.

______ (1995) At Home in the Universe: The Search for Laws of Complexity, London: Penguin.

Kobayashi, K. (1993) "Incomplete Information and Logistical Network Equilibria, in Andersson, Å.E., Batten, D.F., Kobayashi, K. and K. Yoshikawa (eds.) The Cosmo-Creative Society, Berlin: Springer, pp. 95-119.

Kobayashi, K., Kunihisa, S. and K. Fukuyama (2000) "The Knowledge Intensive Nature of Japan's Urban Development, in Batten, D.F., Bertuglia, C.S., Martellato, D. and S. Occelli (eds.) Learning, Innovation and Urban Evolution, Boston: Kluwer, in press.

Krugman, P. (1993) "On the Number and Location of Cities," European Economic Review, vol. 37, pp. 293-298.

______ (1994a) "Complex Landscapes in Economic Geography," American Economic Association, Papers and Proceedings, vol. 84, pp. 412-416.

______ (1994b) The Age of Diminished Expectations. Cambridge, Ma.: The MIT Press.

______ (1994c) Peddling Prosperity: Economic Sense and Nonsense in the Age of Diminished Expectations. New York: W.W. Norton.

______ (1996) The Self-Organizing Economy. New York: Blackwell.

Kuhn, T. (1962) The Structure of Scientific Revolutions. Chicago: University of Chicago Press.

Lane, D. (1993) "Artificial Worlds and Economics, Part I," Journal of Evolutionary Economics, vol. 3, pp. 89-107.

______ (1997) "Is What is Good for Each Best for All? Learning from Others in the Information Contagion Model," in Arthur, W.B., Durlauf, S.N. and D.A. Lane (eds.) The Economy As An Evolving Complex System II, Reading, Ma: Addison-Wesley, pp. 105-127.

Langton, C.G. (1996) "Artificial Life," in Boden, M.A. (ed.) The Philosophy of Artificial Life, Oxford: Oxford University Press, pp. 39-94.

LeRoy, S.F. and R.D. Porter (1981) "Stock Price Volatility: Tests based on Implied Variance Bounds," Econometrica, vol. 49, pp. 97-113.

Lindgren, K. (1992) "Evolutionary Phenomena in Simple Dynamics," in Langton, C.G., Farmer, J.D., Rasmussen, S. and C. Taylor (eds.) Artificial Life II, Redwood City, CA: Addison-Wesley, pp. 295-312.

Lindsay, C.L. (1991) Trident: A Trading Strategy. Brightwaters, N.Y.: Windsor Books.

Lösch, A. (1944) The Economics of Location (Translated from the 2nd Revised Edition by William H. Woglom and published by Yale University Press, New Haven, 1954).

Machlup, F. (1962) The Production and Distribution of Knowledge in the United States, Princeton, N.J.: Princeton University Press.

Mandelbrot, B. (1963) "The Variation of Certain Speculative Prices," Journal of Business, vol. 36, pp. 394-419.

______ (1997) Fractals and Scaling in Finance. New York: Springer-Verlag.

______ (1999) "A Multifractal Walk Down Wall Street," Scientific American, February, pp. 50-53.

Marshall, A. (1920) Principles of Economics. London: Macmillan.

Maruyama, M. (1963) "The Second Cybernetics: Deviation-Amplifying Mutual Causal Processes," American Scientist, vol. 51, pp. 164-179.

Marx, K. (1973) Grundrisse: Foundations of the Critique of Political Economy. (Translated by Martin Nicolaus). Harmondsworth, UK: Penguin.

de la Maza, M. and D. Yuret (1995) "A Futures Market Simulation with Non-Rational Participants," in Brooks, R.A. and P. Maes, Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems. Cambridge, MA.: The MIT Press. pp. 325-330.

Mees, A. (1975) "The Revival of Cities in Medieval Europe", Regional Science and Urban Economics, vol. 5, pp. 403-425.

Monod, J. (1971) Chance and Necessity. London: Penguin.

Nagel, K. and S. Rasmussen (1995) "Traffic at the Edge of Chaos," in Brooks, R.A. and P. Maes, Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems. Cambridge, MA.: The MIT Press. pp. 222-235.

Nagel, K. and M. Schreckenberg (1992) "A Cellular Automaton Model for Freeway Traffic," Journal de Physique I, vol. 2, p. 2221.

Nelson, R.R. and S.G. Winter (1982) An Evolutionary Theory of Economic Change, Cambridge, MA.: Harvard University Press.

Neumann, J. von (1966) Theory of Self-Reproducing Automata (edited and completed by Arthur Burks), Urbana: University of Illinois Press.

Nicolis, G. and I. Prigogine (1977) Self-Organization in Nonequilibrium Systems, New York: Wiley.

______ and ______ (1989) Exploring Complexity: An Introduction, New York: W.H. Freeman.

North, D.C. and R.P. Thomas (1973) The Rise of the Western World: A New Economic History, Cambridge: Cambridge University Press.

Nowak, M. and K. Sigmund (1993) "A Strategy of Win-Shift, Lose-Stay that Outperforms Tit-for-Tat in the Prisoner's Dilemma Game," Nature, vol. 364, pp. 56-58.

Ohmae, K. (1991) The Borderless World: Power and Strategy in the Interlinked Economy. New York: HarperPerennial.

Olson, M. (1965) The Logic of Collective Action, Cambridge, Ma.: Harvard University Press.

Pareto, V. (1896) Oeuvres Completes. Geneva: Droz.

Pigou, A.C. (1927) Industrial Fluctuations, London: Macmillan.

Pirenne, H. (1925) Medieval Cities: their Origins and the Revival of Trade. (Translation by F.D. Halsey, Princeton: Princeton University Press, 1952).

______ (1936) Economic and Social History of Medieval Europe. (Translation by I.E. Clegg, London: Routledge & Kegan Paul, 1965).

Portugali, J., Benenson, I. and I. Omer (1994) "Sociospatial Residential Dynamics: Stability and Instability within a Self-Organizing City," Geographical Analysis, vol. 26, pp. 321-340.

______, ______ and ______ (1997) "Spatial Cognitive Dissonance and Sociospatial Emergence in a Self-Organizing City," Environment and Planning B, vol. 24, pp. 263-285.

Pred, A. (1966) The Spatial Dynamics of U.S. Urban-Industrial Growth, 1800-1914. Cambridge, Ma.: Harvard University Press.

Prigogine, I. and R. Herman (1971) Kinetic Theory of Vehicular Traffic, New York: Elsevier.

Puu. T. (1997) Mathematical Location and Land Use Theory, Berlin: Springer.

Rapoport, A. and A.M. Chammah (1965) Prisoner's Dilemma: A Study in Conflict and Cooperation, Ann Arbor: University of Michigan Press.

Rasmussen, S. and C.L. Barrett (1995) "Elements of a Theory of Simulation," in ECAL 95, Lecture Notes in Computer Science. New York: Springer.

Ray, T.S. (1992) "An Approach to the Synthesis of Life," in Langton, C., Taylor, C., Farmer, J.D. and S. Rasmussen, eds. Artficial Life II. Redwood City, Ca: Addison-Wesley, pp. 371-408.

Rhea, R. (1932) The Dow Theory. New York: Barron's.

Riising, A. (1952) "The Fate of Henri Pirenne's Thesis on the Consequences of Islamic Expansion," Classica et Medievalia, vol. 13.

Rostow, W.W. (1960) The Stages of Economic Growth. New York: Cambridge University Press.

Rothstein, M. (1982) "Frank Norris and Popular Perceptions of the Market," Agricultural History, vol. 56, pp. 50-66.

Rouse, W.B. and Morris, N.M. (1986) "On Looking into the Black Box: Prospects and Limits in the Search for Mental Models," Psychological Bulletin, vol. 100, pp. 349-363.

Rubinstein, M. (1975) "Securities Market Efficiency in an Arrow-Debreu Economy," American Economic Review, vol. 65, pp. 812-814.

Ryle, G. (1949) The Concept of Mind. London: Hutchinson.

Saari, D. (1995) "Mathematical Complexity of Simple Economics," Notices of the American Mathematical Society, vol. 42, pp. 222-230.

Sakoda, J.M. (1971) "The Checkerboard Model of Social Interaction," Journal of Mathematical Sociology, vol. 1, pp. 119-132.

Samuelson, P. (1976) Economics. New York: McGraw Hill.

Sanders, L., Pumain, D., Mathian, H., Guérin-Pace, F. and S. Bura (1997) "SIMPOP: A Multiagent System for the Study of Urbanism," Environment and Planning A, vol. 24, pp. 287-305.

Sargent, T.J. (1993) Bounded Rationality in Macroeconomics, New York: Oxford University Press.

Scheinkman, J.A. and M. Woodford (1994) "Self-Organized Criticality and Economic Fluctuations, American Journal of Economics, vol. 84, pp. 417.

Schelling, T.S. (1969) "Models of Segregation," American Economic Review, Papers and Proceedings, vol. 59 (2), pp. 488-493.

______ (1978) Micromotives and Macrobehaviour. New York: W.W. Norton and Company.

Schroeder, M. (1991) Fractals, Chaos, Power Laws: Minutes from an Infinite Paradise. New York, W.H. Freeman.

Schrödinger E. (1956) Mind and Matter (Reprinted together with What is Life, Cambridge: Cambridge University Press, 1967).

Schumpeter, J. (1934) The Theory of Economic Development. Cambridge, MA.: Harvard University Press.

Scott, J.W.(1876) A Presentation of Causes Tending to Fix the Position of the Future Great City of the World in the Central Plain of North America: Showing that the Centre of the World's Commerce, Now Represented by the City of London, Is Moving Westward to the City of New York, and Thence, within One Hundred Years, to the Best position on the Great Lakes. Toledo.

Sendut, H. (1966) "City-Size Distributions of South-East Asia," Asian Studies, vol. 4, pp. 165-172.

Sheffi, Y. (1985) Urban Transportation Networks, New Jersey: Prentice Hall.

Shiller, R.J. (1981) "The Use of Volatility Measures in Assessing Market Efficiency," Journal of Finance, vol. 36, pp. 291-304.

______ (1989) Market Volatility, Cambridge, MA: The MIT Press.

Simon, H. (1987) "Giving the Soft Sciences a Hard Sell," Boston Globe, 3 May.

Singer, H.W. (1936) "The 'Courbes des Populations': A Parallel to Pareto's Law," Economic Journal, vol. 46, no. 182, pp. 254-263.

Soros, G. (1994) The Alchemy of Finance: Reading the Mind of the Market, New York: Wiley.

Spence, A.M. (1981) "The Learning Curve and Competition," The Bell Journal of Economics, vol. 12, pp. 49-70.

Stauffer, D. (1985) Introduction to Percolation Theory. London: Taylor and Francis.

Taylor, G.R. (1951) The Transportation Revolution: 1815-1860. New York: Rinehart.

Turner, F.J. (1920) The Frontier in American History. New York: Holt, Rinehart and Winston.

Vickrey, W.S. (1969) "Congestion Theory and Transport Investment", American Economic Review, vol. 59, pp. 251-260.

Thünen, J.H. von (1826) Der Isolierte Staat in Beziehung auf Landtschaft und Nationalökonomie. Hamburg. (English translation by C.M. Wartenburg: von Thünen's Isolated State, Oxford: Pergamon Press, 1966).

Vries, J. de (1984) European Urbanization: 1500-1800. London: Methuen.

Ward, M. (1999) Virtual Organisms. London: Macmillan.

Yang, H. and M.G.H. Bell (1998) "A Capacity Paradox in Network Design and How to Avoid It," Transportation Research, vol. 32 (7), pp. 539-545.

Zhang, W-B. (1993) Synergetic Economics. Berlin: Springer-Verlag.

Zipf, G.K. (1941) National Unity and Disunity. Bloomington, Indiana: Principia Press.

______ (1949) Human Behaviour and the Principle of Least Effort. New York: Hafner.

Åkerman, N. (1998) The Necessity of Friction. Boulder: Westview Press.