

UNIVERSITY OF OXFORD  
FACULTY OF SOCIAL STUDIES  
SUB FACULTY OF ECONOMICS

THE EVOLUTION OF TECHNOLOGY AND  
ADAPTIVE ECONOMIC BEHAVIOUR

BENEDICT C. COOPER

THESIS SUBMITTED FOR THE DEGREE OF DOCTOR OF PHILOSOPHY  
AT THE UNIVERSITY OF OXFORD



TRINITY TERM 1997

**ABSTRACT**

THE EVOLUTION OF TECHNOLOGY AND ADAPTIVE ECONOMIC  
BEHAVIOUR

Benedict C. Cooper, Wolfson College

Submitted for the degree of Doctor of Philosophy  
Trinity Term 1997

This thesis studies the role of learning as a mechanism of economic change. Two areas are considered where this would seem to be important. First, how firms learn about new technology; and secondly, how agents learn to behave in interactive situations.

A model of research and development is presented which models the process by which firms solve specific design problems. This may be by individual experimental search or by partial imitation. In the latter case, a close parallel is drawn between biological evolution, based on genetic reproduction, and technological evolution, based on firms blending existing technologies. Some economic implications of these processes are explored, including their application to stochastic learning curves, patent design and the transfer of technology to developing countries.

The thesis continues by critically assessing the analogy between biological and cultural evolution often used to model how agents learn to behave in interactive situations. It is argued that the methods used by economists exploiting this analogy are often ill-suited to an economic context. Models are presented which deal with specific issues in the transition from a biological context to an economic context, including models of partnership formation, models of imperfect imitation, and models without payoff-monotonic dynamics. The issue of imperfect imitation is expanded upon in an evolutionary model of the infinitely repeated prisoners' dilemma, where it is shown that the problem of inter-generational copying fidelity may allow one to restrict attention to strategies with a very simple stochastic structure.

## ACKNOWLEDGEMENTS

I am very grateful to Paul Klemperer, who has been a model supervisor. His support and encouragement has been matched by an incisive awareness of the issues and careful, constructive criticism of my work.

Chris Wallace has also been a great encouragement. Some of our work on partnerships is all too briefly described in chapter 4. It has been great fun working with him.

The material included in chapters 2, 3 and 5 has been presented at seminars as part of the ‘Gorman workshop’ series at Nuffield College, Oxford. I would like to thank the participants of those workshops for their questions and advice, and Meg Meyer and Philippe Aghion for chairing them. Thanks also to seminar participants at University College London and the International Conference of the Seventh Stony Brook Summer Festival on Game Theory, who listened to and commented on versions of chapter 5. The material in chapter 5 has also benefitted from the advice of Ken Binmore, Larry Samuelson, Peter Sørensen and Martin Nowak.

This research was funded by ESRC research award R00429424072. My thanks to the ESRC also for helping to fund my attendance at the Seventh Stony Brook Summer Festival on Game Theory, along with the managers of the George Webb Medley Fund, the Graduate Studies Committee of Oxford University and Wolfson College, Oxford.

I must also mention the great patience, encouragement and—in her capacity at the Science Museum, London—professional help of my wife, Catherine. Anyone who has been through the trials and tribulations of living with a D.Phil. student will know what she has had to put up with.

## *Contents*

<b>1</b>	<b>Introduction</b>	<b>6</b>
1.1	The Evolution of Technology . . . . .	7
1.2	Evolution and Adaptive Economic Behaviour . . . . .	9
<b>I</b>	<b>Technology</b>	<b>10</b>
<b>2</b>	<b>Modelling Research and Development</b>	<b>11</b>
2.1	Introduction . . . . .	12
2.2	An Example Design Problem . . . . .	14
2.3	Two Models of Research and Development . . . . .	18
2.3.1	Individual Learning by Experimental Search . . . . .	18
2.3.2	Social Learning by Partial Imitation . . . . .	22
2.4	Economic Applications . . . . .	27
2.4.1	Stochastic Learning Curves . . . . .	27
2.4.2	Patent Design . . . . .	29
2.4.3	Technodiversity . . . . .	33
2.5	Other Methods of Modelling Research and Development . . . . .	35
2.6	Conclusion . . . . .	36
	Appendix: The Genetic Algorithm . . . . .	38

<b>3</b>	<b>Technodiversity</b>	<b>40</b>
3.1	Introduction . . . . .	41
3.2	Partial Imitation . . . . .	43
3.3	Two Seed Technologies . . . . .	45
3.3.1	A Linear Model . . . . .	46
3.3.2	A Non-Linear Model . . . . .	50
3.4	Multiple Seed Technologies . . . . .	53
3.5	Placing More Weight on Indigenous Innovation . . . . .	55
3.6	Encouraging High Rates of Innovation . . . . .	57
3.7	A Calibrated Diffusion Model . . . . .	58
3.8	Conclusion . . . . .	59
	Appendix . . . . .	59
<b>II</b>	<b>Adaptive Economic Behaviour</b>	<b>62</b>
<b>4</b>	<b>Evolutionary Game Theory</b>	<b>63</b>
4.1	Introduction . . . . .	64
4.2	The Methods of Evolutionary Game Theory . . . . .	68
4.3	Notes on the Transition from a Biological Context to an Eco- nomic Context . . . . .	75
4.3.1	Beyond Random Matching: Who to Play With? . . . . .	75
4.3.2	Imperfect Imitation in a 'Mini' Ultimatum Game . . . . .	82
4.3.3	Payoff Monotonicity . . . . .	86
4.4	Conclusion . . . . .	91
<b>5</b>	<b>Copying Fidelity in the Repeated Prisoners' Dilemma</b>	<b>93</b>
5.1	Introduction . . . . .	94
5.2	The Model . . . . .	97
5.2.1	Basic Elements . . . . .	97
5.2.2	Population Dynamics . . . . .	98
5.2.3	Evolutionary Equilibrium . . . . .	100
5.2.4	Imitation and The Copycat Criterion . . . . .	101

---

5.3	Utilitarian Populations and the Prisoners' Dilemma . . . . .	102
5.4	Non-Utilitarian Populations . . . . .	106
5.5	Restoring Copying Fidelity with Clumsy Automata . . . . .	107
5.6	Discussion . . . . .	112
	Appendix . . . . .	116
<b>6</b>	<b>Conclusion</b>	<b>120</b>

## *Chapter 1*

### *Introduction*

‘Evolution’ is a word that has been used—or mis-used—in so many ways and contexts that it has acquired a variety of meanings. In its most general sense—which stems from its origin as a Latin word used to describe the unrolling of a scroll—it simply refers to how some dynamical process ‘unfolds’ over time. As such, virtually any dynamical process can be described as ‘evolutionary’. More specifically, in biological contexts it can mean the process by which an organism develops from a rudimentary to a complete and complex state. More specifically still, the word ‘evolution’ was adopted in the nineteenth century by those who, like Darwin (1859)<sup>1</sup>, claimed that the process by which *species* develop from rudimentary to complex is by the dual mechanisms of inter-generational variation and natural selection. It is in this final, somewhat special, sense that the word is most commonly used today.

Indeed, when social scientists use the word ‘evolution’ to describe long-run cultural change they are usually making an analogy with this special

---

<sup>1</sup>Although Darwin was certainly not the only one proposing these ideas. Darwin’s major contribution was to synthesise the great diversity of views on evolution into a coherent whole, and to present his theory in a very readable book that was accessible to the general public.

sense. The inter-generational variation in biological evolution is caused by mutation or genetic crossover during reproduction; the variation in a model of social evolution may be caused by experimentation or mistakes during updating. The selection in biological evolution is caused by competition for scarce resources resulting in differing reproduction rates; the selection in a model of social evolution could have a number of alternative causes, including reinforcement (in the context of individual learning), imitation of success (in the context of social learning) or even bankruptcy (in a market context). The analogy is helpful in so far that it allows those interested in social evolution to borrow vocabulary, methods and models developed for similar processes.

However, we must be careful not to mis-apply the analogy, to abuse it, or to stretch it too far. There was a tendency at the end of the nineteenth century to find ‘evolutionary’ explanations for virtually anything. This tendency reached its worst excesses in so-called ‘Social Darwinism’, which claimed that since selection was an apparently natural phenomenon, it *ought* to be encouraged or even imposed in the social sphere—although to what purpose was never very clear. Borrowing ideas from evolutionary biology stops us having to ‘re-invent the wheel’ in many models of adaptive processes, but does no more than that. Indeed, the fact that the ideas are borrowed rather than tailor-made can often cause problems, as we shall see.

The aim of this thesis is thus to show that application of the evolutionary analogy *can* add to our understanding of long-run economic processes if it done carefully and selectively. No claim is made that using the analogy is the only way to model long-run processes, nor that it is always the best way to model them. However, two areas do seem to benefit from an evolutionary approach, and these are *technological change* and *adaptive behaviour*.

## 1.1 THE EVOLUTION OF TECHNOLOGY

If we return to Darwin’s original argument in favour of his theory of evolution, we find it goes something like this: Consider the huge diversity of new breeds

generated by selective breeding over a relatively few number of generations. For example, the number of distinctive pigeon breeds descended from the common rock-pigeon simply by the careful mating or ‘blending’ of pigeons with distinctive characteristics by expert breeders. Just imagine the changes that nature could accomplish through the same process of ‘blending’, now driven by *natural* selection, over the very long time-scale suggested by the geological record.

Of course, Darwin’s theory has undergone many revisions of its own over the last 140 years. The major innovation has been to combine it with a greater understanding of the laws of genetics. This innovation has been accompanied by a shift in emphasis in *neo*Darwinism away from some form of ‘blending’ during reproduction (or, as we might say, genetic crossover) to genetic mutations as the primary source of inter-generational variation<sup>2</sup>.

Now the analogy between technological development and biological evolution has been recognised at least as far back as ‘The Book of the Machines’ in Samuel Butler’s *Erewhon* (Butler 1872). Modern treatments, however, have borrowed from the *neo*Darwinian conception of evolution and have therefore perhaps missed out on some of Darwin’s original insights. In particular, there seems to be a close analogy between the way firms updating technology draw on parts of designs from multiple sources and Darwin’s admittedly vague notion of ‘blending’ during sexual reproduction. This is a key theme in Part I of this thesis. In chapter 2, the implications of assuming that firms construct new designs and solve design problems by this process of *partial* imitation are studied. In chapter 3, this form of technological innovation is linked to technological diffusion, with implications for the problem of transferring technology and encouraging indigenous development.

---

<sup>2</sup>This has broadened the scope of the theory in that, by construction, any biological development can be claimed to be consistent with it—although this gain in generality has arguably been at the expense of explanatory power.

## 1.2 EVOLUTION AND ADAPTIVE ECONOMIC BEHAVIOUR

While the analogy between biological and cultural evolution has been recognised at least since the Social Darwinists, it is only relatively recently that the idea has been endowed with much substance from a modelling point of view. Seminal work by authors such as Dawkins (1976) and Boyd and Richerson (1985), suggesting that *imitation* in cultural evolution may perform a similar function to genetic reproduction in biological evolution, has generated a great deal of interest. Over the past decade there has been an explosion in the study of ‘evolutionary game theory’, developed in the seventies to study biological evolution, as a potential tool for modelling long-run economic behaviour. Although the rate of growth is slowing, the influence of evolutionary game theory on economic theory is still increasing, and its impact is likely to be long-lasting.

The aim of chapter 4 is to step back and consider just how useful evolutionary game theory is to the economic theorist. When is it—and when it is not—an appropriate modelling tool? Moreover, if it is ever an appropriate modelling tool, are not current methods too derivative of biological mechanisms that assume a biological context? To address this objection, the chapter includes some examples of how one can improve the transition from a biological to an economic context. These cover the issues of matching, imperfect imitation and payoff monotonicity in cultural learning.

Chapter 5 expands on the issue of imperfect imitation in the context of an evolutionary model of the repeated prisoners’ dilemma. In simple games, evolutionary game theory has thrown some light on the problem of equilibrium selection; the question is, can it throw any light on equilibrium selection in repeated games, where the problem is especially acute? The evolutionary analysis of repeated games raises a number of difficulties, and these are also addressed in the chapter.

The thesis is concluded in chapter 6, which also includes some suggestions for further research.

*Part I*  
*The Evolution of Technology*

## *Chapter 2*

# *Modelling Research and Development: How Do Firms Solve Design Problems?*

### SUMMARY

One way of thinking about research and development is to recognise that firms are trying to solve particular *design problems*. We often build these design problems into our models, but are forced to oversimplify them in order to make the models solvable. The approach taken in this chapter is to acknowledge that design problems are often insoluble using standard techniques and to model instead the *process* by which firms solve them. Two such processes are simulated in detail. The first, individual experimental search, is based on a problem-solving technique known as *simulated annealing*. The second, partial imitation, involves learning at a social level and is based on a problem-solving technique known as the *genetic algorithm*. Some economic implications of these processes are explored, including their application to stochastic learning curves, patent design and the importance of ‘technodiversity’ in the introduction of new technology to developing countries.

## 2.1 INTRODUCTION

The aim of this chapter is to model how firms conduct research and development using processes analogous to some well-known problem solving techniques from computer science. The motivation is an obvious one. Some decision about technology choice or research and development lies at the heart of many a model of industrial organisation and, of course, at the heart of the new models of ‘endogenous’ growth. The more our portrayal of research and development corresponds to reality, the better these models will be. A pleasing side-effect is that, by modelling research and development in a different way, we obtain new insights into the origins of learning curves, the design of patents, and the problem of introducing ‘appropriate’ technology to developing countries.

When thinking about how to model research and development, it is already common to recognise that firms are trying to solve *design problems*. In models of industrial organisation, for example, firms, at some early stage in the extensive-form description of the model, choose products to sell at some later stage from some ‘product’ or ‘design’ space. (Chapters 2 and 7 of Tirole (1988) abound in examples of this sort.) With careful construction, it is relatively straightforward to solve for a subgame perfect equilibrium—which will include a solution of the design problem facing each firm. Such design problems are what Simon (1973) calls ‘well-structured’. That is, there is a clearly defined problem space in which *all* potential solutions can be expressed, and a definite criterion for testing any proposed solution. The danger is that *by construction* such problems will typically be easy to solve, while design problems in the real world are obviously very much *ill-structured* and difficult or awkward to solve. After all, technology is a complex object. The space of all possible designs is poorly defined and evolves as scientific understanding advances. Real design spaces are much more like the ‘evolving, rugged fitness landscapes’ of Kauffman (1988) and Kauffman and Macready (1995), where a design at a given set of coordinates has a ‘fitness’, ‘performance’ or

'quality' given by the height at that point. However, much of the landscape is uncharted, and even those regions thought to be well understood seem to shift in the occasional technological earthquake. Moreover, once the managers and engineers of a firm have, to the best of their abilities, sufficiently tamed an ill-defined design problem to call it well-defined, the structure of the problem is likely to remain highly complex and it is unlikely to be easily solvable.

The new approach pursued in this chapter tackles the issue of research and development from a different direction and may be dubbed *pseudo-empirical*. It recognises that there already exist specialists in turning ill-structured design problems into well-structured ones and then solving them, and they are not economic theorists but the engineers employed by firms for that express purpose<sup>1</sup>. By observing and modelling how they do so (rather than attempting to model the problem itself) we may obtain models of research and development that more closely correspond to the complexities of reality. Here, we observe that firms may take an independent approach and tackle design problems by simple methods of trial-and-error, or they may take a more 'social' approach and look around at the designs of other firms, combining features that seem to work well. In modelling these research *processes* we can exploit the fact that the former is similar to a problem-solving technique known as *simulated annealing* while the latter has some affinity to a problem-solving technique known as the *genetic algorithm*.

This is not the first time that these problem-solving techniques developed by computer scientists have found their way into economic models. For example Lindgren (1990) and Axelrod (1987) have applied the genetic algorithm to learning in the repeated prisoners' dilemma, while Arifovic (1994) has

---

<sup>1</sup>Of course, economic theorists *are* engineers when it comes to the task of selecting economic models from the space of all possible models. There is an interesting analogy one can draw between designing an economic model and designing a new product, although the criteria for testing the final results may be quite different! If they pause to consider their own research 'strategies', readers may recognise some of the techniques of experimental search and partial imitation described in this chapter.

applied it to learning in a simple cobweb model. In the genetic algorithm, the potential 'answers' to a problem are expressed as binary strings. The interpretation of these strings given in these models is that they somehow represent behavioural rules. The difference in the present chapter is that a string represents a physical design. Given that the genetic algorithm involves combining or 'mixing' strings in a fairly mechanical way, it may be argued that an interpretation based on strings representing physical designs, which are clearly and unambiguously observable, is more plausible than one based on strings representing behavioural rules, which are not necessarily directly observable and may be expressed in many different ways. It is certainly an interpretation that yields interesting results in a number of applications concerned with the evolution of technology.

The chapter proceeds as follows. The following section explains more fully the concepts of a design problem and a design space. Next, we consider two models of research and development: individual learning by experimental search, based on simulated annealing; and partial imitation, based on the genetic algorithm. Finally, we consider three applications of these models: to the origin of learning curves, the design of patents to encourage social learning, and to the introduction of new technology to developing countries.

## 2.2 AN EXAMPLE DESIGN PROBLEM

Where do well-structured design problems come from? We must remember that all design problems are *prima facie* ill-structured and only become well-structured by construction. Faced with an ill-structured problem, where the criterion for testing solutions may not be obvious and where the design space is poorly defined, a firm will attempt to break up the problem into well-defined sub-problems as much as is possible. The whole firm will be involved in this: commercial experts will make guesses about what is marketable; engineers will make guesses about what is possible. They will probably rely heavily on intuition, rules-of-thumb derived from experience, and convention.




Figure 2.1. A well-structured but non-trivial design problem

In a general sense, this process divides the product market into fairly well defined sectors, each with its own idiosyncrasies. More specifically, it provides engineers with well-structured design problems to solve. They might be asked to improve the comfort of car seating subject to space constraints, to increase the signal to noise ratio of audio equipment or reduce the mass of an aircraft component. Or they might be asked to redesign a product to allow fewer and more simple processes in its construction. Although the evolution of product sector divisions is an interesting topic in its own right, it is the solution of these resulting *restricted* design problems which concerns us here.

For example, suppose there exist consumers who value a product that is able to support a 50kN load 4m from a solid wall. Thus far, the problem of designing such a product is ill-structured. However, suppose that experience shows that a pin-jointed frame like that of Figure 2.1 works adequately well. The cost of the frame to the firm is directly proportional to its mass. The mass can be reduced by decreasing the cross-sectional areas of frame's members; but if the cross-sectional area of any member gets too low, the stress in the member—that is, the force per unit area—may make the frame deform.

The problem is thus to find the seven cross-sectional areas that minimize the mass of the frame subject to these stress constraints. We now have a well-structured problem in that there is a definite test criteria (mass) and a clearly defined problem space (cross-sectional areas).

Many design problems, restricted in this manner to make them well-structured, may be solvable directly. In this example, there is a unique combination of cross-sectional areas that will solve the constrained optimization program and a competent engineer with plenty of time, a pile of scrap paper and a calculator might be able find the answer using standard optimization techniques.

However, note that the problem is not as straightforward as it looks. The stress constraints define an irregular polyhedron in the restricted design space that will be difficult to work with. The engineer also needs to know the force in each member. Calculating these forces is not as trivial as one might think, as the frame is *statically indeterminate*. That is, it has a redundant member (member AD or EB could be removed and it would still hold together rigidly) and the forces cannot be found by simply performing vector addition at each node, and must be calculated using a technique known as ‘virtual work’. One can easily imagine similar problems—perhaps in three dimensions, involving composite materials and subject to variable loads (for example, the landing gear of an aircraft)—where solving the design problem with conventional linear or non-linear programming techniques may not be feasible. The complexity of the technology involved in many products means that we must not suppose that such problems have nice analytical properties that make them solvable directly. In particular, many design problems are *combinatorial* problems. That is, the design space is not Euclidian but a large, discrete configuration space. For example, the famous ‘travelling salesman problem’ is a combinatorial problem. The salesman visits  $N$  cities with given locations in a two-dimensional space, finally returning to his or her city of origin. Each city is to be visited only once and the route to be made as short as possible. More practical examples include the design of complex in-

tegrated circuits, the plumbing in a large building or ergonomically efficient work layouts in factories, kitchens or laboratories. The number of elements in such configuration spaces are factorially large, so it is impractical to search them exhaustively.

We may conclude that engineers do use classical optimization techniques—but only up to a point. Unlike economic theorists designing models, they are no pressure to return to the drawing board when it transpires that the well-structured problems they have constructed cannot be solved simply, elegantly and conclusively. Rather, they are not ashamed to use whatever technique will get the job done. Often they have no choice. So by assuming that engineers cannot use conventional techniques when finding design specifications for the design problem of Figure 2.1, it will hopefully be possible to obtain some insights into how engineers use adaptive techniques to develop products where the design problem is more complex.

So how *do* engineers solve such design problems? If exhaustive or random search is impractically slow, such problems seem insuperable. However, note what seems to be a common feature of all design spaces: a correlated structure. That is, designs close to each other in a design space tend to have similar fitnesses; those further apart may have fitnesses that are widely different. We shall look at two techniques that exploit this feature of design spaces. First, there is *experimental search*, where one looks at a slight adaptation of an existing design. Secondly, there is *partial imitation*, where firms look around at their opponents' designs and, if they seem to be successful, incorporate features of these designs into their own products. Both these techniques have implications for the way we model research and development.

## 2.3 TWO MODELS OF RESEARCH AND DEVELOPMENT

### 2.3.1 *Individual Learning by Experimental Search*

Economic theorists know that the easiest way to come up with a workable new model is to take an existing model and adapt it<sup>2</sup>. Similarly with product design: a great many product innovations are simple variations of existing products. Informally, we call this experimental search of a complex space *trial and error*. Try a simple adaptation: if it works better, adopt it; if it fails to work as well, reject it and put it down to experience.

This approach to finding improved elements of a complex space is incorporated into a formal method of numerical or combinatorial optimization known as *simulated annealing* (Kirkpatrick, Gelatt, and Vecchi 1983). Annealing is the process by which metals are cooled *slowly* to leave them in a low energy state where all the atoms are arranged in a well-ordered, regular crystal structure. Put loosely, the atoms, moving around randomly with steadily decreasing thermal mobility, have time to ‘search’ for this low-energy arrangement. The result is a stronger, softer and less brittle material. Simulated annealing attempts to make use of this general idea for other than thermodynamic systems. It has proved effective not just for combinatorial problems that are difficult to solve any other way, but also for problems which have a global extremum hidden among many local extrema.

One needs to specify:

1. A description of the search space—*e.g.* a description of possible system configurations.
2. An objective function (analog of energy) whose optimization is the goal of the procedure.

---

<sup>2</sup>For example, one could take a one-stage model of a patent race and extend it into a multi-stage model.

3. A mutation model. That is, a generator of random changes presenting the system with new options. This generation of new points in the search space is governed by:
  - (a) A control parameter (analog of temperature)—*e.g.* this might specify the ‘size’ of a mutation: how large an adaptation to try.
  - (b) An *annealing profile*, which specifies how the control parameter varies as the search progresses.

The method of simulated annealing presents us with a workable model of design by experimental search that we may apply to design problems like that of Figure 2.1. While too unwieldy for direct economic analysis, investigating the properties of design problem solving at this very basic level should help us to construct better models of research and development. In particular, we can investigate these properties by simulating the search process.

The search space for the design problem of Figure 2.1 is simply a vector of seven cross-sectional areas, and the objective is to minimize the mass of the frame. However, in order to specify a simple mutation model, we can specify a design as a single binary string of finite length. That is, one can code the design as a 28 bit concatenated mapped fixed-point binary string, with each of the seven 4 bit sub-strings representing a member’s cross-sectional area mapped linearly between  $\underline{A}$  and  $\bar{A}$ . (There is little reason to suppose why any design specification could not, in principle, be coded in a similar way.) A simple mutation model is then to specify a single probability  $T$  for an individual bit in the string to change state.  $T = \frac{1}{2}$  thus samples from the search space randomly, while  $T$  closer to zero will sample designs closer to the current design. In addition to being very simple, this mutation model has the advantage that for all  $T > 0$  there is positive probability of sampling *any* design in the space<sup>3</sup>.

---

<sup>3</sup>One could use this feature of the mutation model to demonstrate that simulated annealing will in this case converge asymptotically to the optimal design. However, this is not a striking result: both exhaustive and random search have similar asymptotic prop-

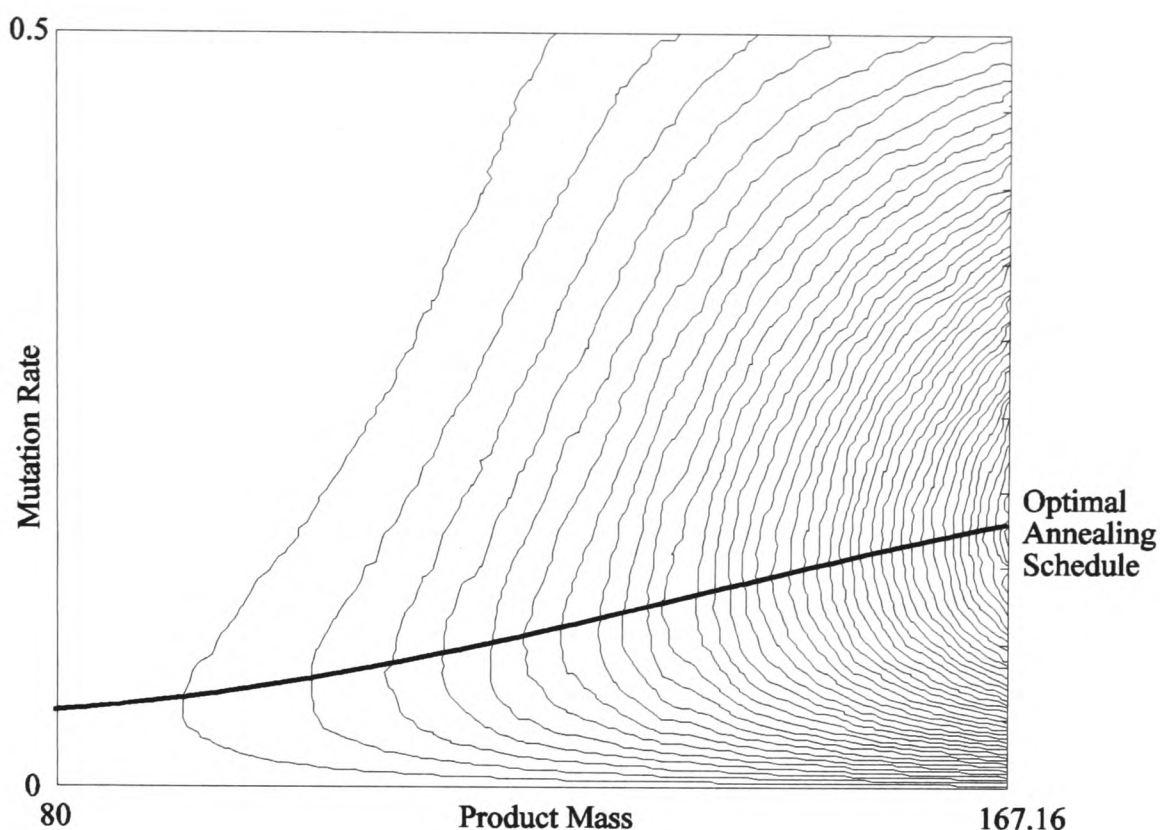


Figure 2.2. Expected mass reduction at a given starting mass and mutation rate. Each contour describes the locus of points giving the same expected mass reduction, with the highest expected mass reduction to the right of the figure. The 'optimal annealing schedule' is simply the locus of points at which the expected mass reduction is highest for a given starting mass. The figure was constructed by running the simulated annealing process many thousands of times at different mutation rates.

To complete the specification of a learning process based on simulated annealing, we need an annealing schedule. For any given problem, there is an optimal annealing schedule that maximizes the expected gain at each trial. The optimal schedule for the design problem of Figure 2.1 is shown in Figure 2.2. This plots the contours of a function giving the expected mass reduction at a given starting mass and mutation rate  $T$ . The locus of points on the 'ridge' of this function is the optimal schedule which, as can be seen, declines roughly exponentially as mass declines. Of course, there is

---

erties. The advantage of experimental search like simulated annealing is that it converges faster.

no guarantee that the firm involved knows what this schedule is. However, if it has experience solving similar design problems, then it will know the advantages of gradually reducing the control parameter  $T$  in this way. Note that we do observe a great deal of casual evidence that firms gradually reduce the size of adaptations they try over time. Bicycles, intended for the same purpose, were once wildly different in design; these days they are virtually identical.

It is now possible to investigate the sort of learning curves that this process generates. Simulated annealing is rather too complex to incorporate directly into an economic model; it is also very problem-specific. However, we may approximate experimental search by taking current mass as the state variable—rather than the current design vector. The process is thereby much simplified and more applicable to other problems.

In trial number  $\tau$ , the simulated annealing algorithm selects a new design. If this design deforms under the load because one or more of the stress constraints are violated, then the design is rejected. Panel (a) of Figure 2.3 shows that the stress constraints were violated approximately half the time in simulations of the process, regardless of the current mass. If the design passes the stress test, then the algorithm presents a new design with a new mass. In the simulations, the change in mass between this new mass and the current mass was distributed according to panels (b), (c) and (d) of Figure 2.3. As one would expect, panel (c) shows that the mean change in mass falls as the design mass declines. Indeed, as the process approaches the optimal mass, the new designs presented by the algorithm were typically of higher mass than the current design. Notice also, from panel (d), that the variance of mass change falls as current mass declines. This is largely attributable to the fact that the optimal annealing schedule shown in Figure 2.2 induces smaller search ‘jumps’ at lower masses. The new design is only adopted if it is of lower mass than the current design, so any ‘learning curve’ generated by this search process will only have downward changes in mass.

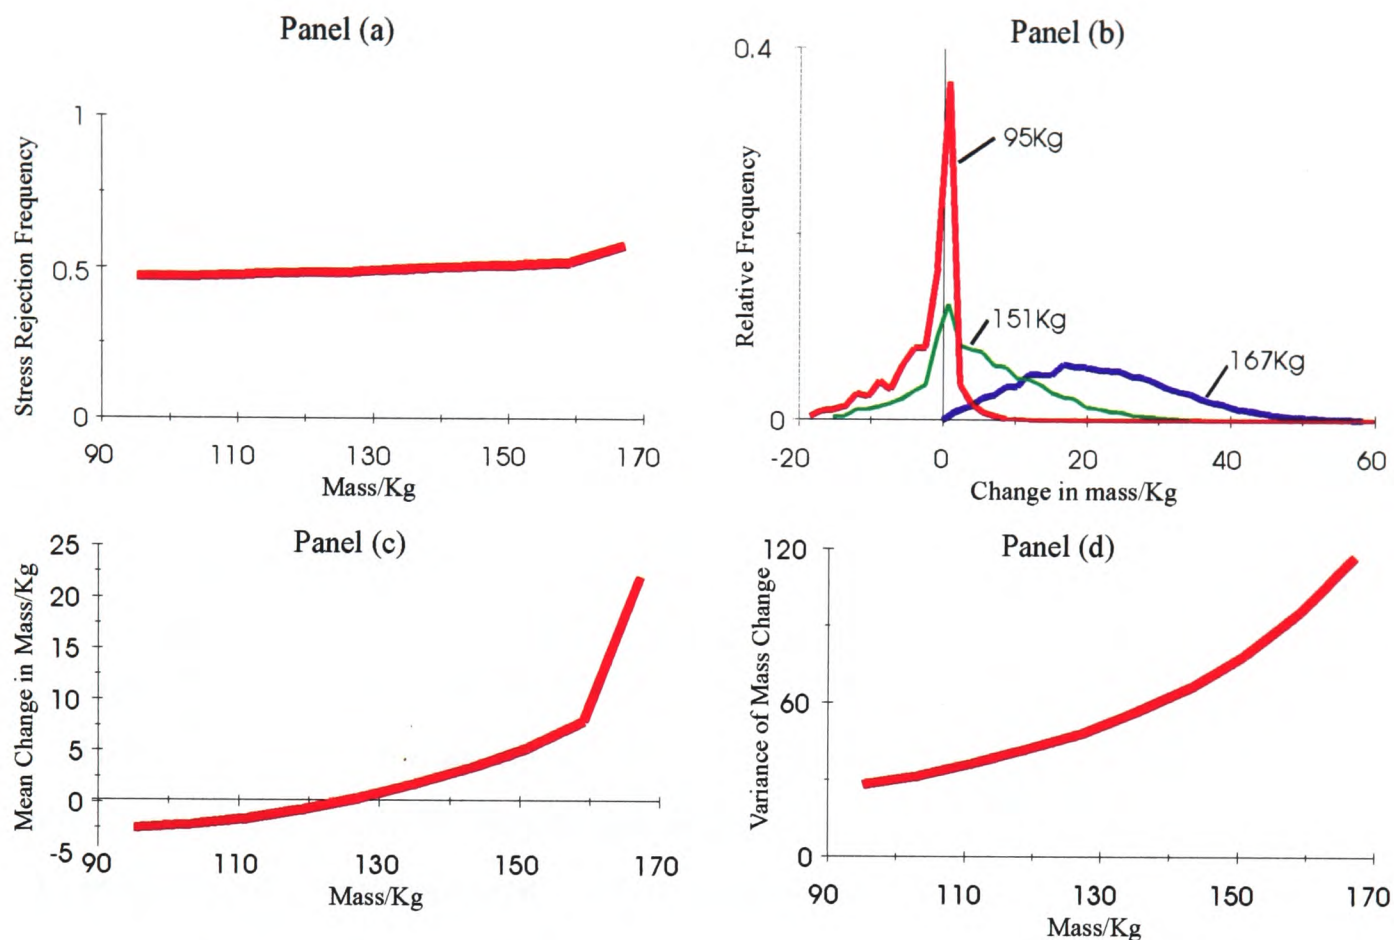


Figure 2.3. Features of experimental search based on simulated annealing. (Sample size 100000.)

### 2.3.2 Social Learning by Partial Imitation

Economic theorists also know, at least implicitly, that an effective method for creating innovative new models is to combine the successful features of two or more existing models<sup>4</sup>. We are perhaps a little sheepish about admitting that we do this—no-one likes to be accused of plagiarism, after all—but there is no doubt that we all do it and that the technique works.

This approach to finding improved elements of a complex space is incor-

<sup>4</sup>For example, the literature on endogenous growth originated by combining neoclassical savings models with microeconomic models of research and development. (Models of research and development which, if the conclusions of this chapter are correct, leave something to be desired.)

porated into a formal method of numerical or combinatorial optimization known as the *genetic algorithm*. The genetic algorithm is a general problem solving technique inspired by natural genetic processes. Its basic building block is a string of characters encoding some ‘answer’ to a certain problem. The string is generally written in binary form and the algorithm acts on a population of  $n$  strings. Each period, it evaluates the strengths of the strings—*i.e.* how ‘good an answer’ they are—and assigns each a probability weight in proportion to its strength. It selects  $p$  pairs of strings (where  $2p \leq n$ ) using this probability distribution and from them creates  $2p$  ‘offspring’, which replace the  $2p$  lowest strength strings in the population. The genetic algorithm ‘evolves’ better answers to a problem in a way similar to the natural processes that evolve fitter genes for an environment<sup>5</sup>.

However, while it may be an effective problem-solving technique, the genetic algorithm requires significant adaptation if it is to be used as a *model* of research and development. As it stands, the updating of strings in the genetic algorithm is a global process that does not readily correspond to individual firms making separate updating decisions<sup>6</sup>. Also, the genetic algorithm asserts that there are always two parents for every offspring (as in genetic reproduction); however, when it comes to the social learning envisaged here, there is no reason to suppose that firms are not influenced by multiple sources.

The model outlined here—which we shall call *partial imitation*—imagines a population of firms engaged in research and development. Each firm already has a working design but wishes to improve upon it. The firm does so by examining a number of high performance designs chosen from the ‘pool’ of existing designs (this may or may not include its own existing design). It blends these designs and randomly modifies the result—in a similar way to

---

<sup>5</sup>More details on the genetic algorithm, and how it differs from the ‘partial imitation’ described below, are given in the appendix to this chapter.

<sup>6</sup>As such, it is perhaps more appropriate for modelling the learning of a single firm rather than a population of firms.

the simulated annealing method described above—to create a prototype. If the prototype is superior to its existing design, then it adopts it.

Do firms really use such a mechanism to solve design problems? We certainly see many examples in the history of technology of *large scale* innovations taking place as the result of such a blending process. There are several good examples in James Burke's popular study of the history of technology, *Connections* (Burke 1978). In 1774 a technique was developed for constructing cannons less likely to explode in use, by boring a cylinder from a solid casting. By blending the resulting very strong and precise cylinders with the rudimentary steam technology then available, James Watt was able to build a steam engine enormously more powerful than anything that had come before. Burke describes the innovation of the 'kinetoscope'—the forerunner of modern cinema—as 'a fine example of how innovation sometimes occurs because at a certain point all the elements of a new development become available to one man, who fits them together like the pieces of a jigsaw'. The pieces were the electric light-bulb, a new material called celluloid (developed to make billiard balls) and the zoopraxiscope (which exploited persistence of vision to give the illusion of movement from a sequence of static images); the inventor who blended them together was Thomas Edison.

The role that blending plays in routine technological development is perhaps less obvious than it is in these dramatic innovations. However, designers are often highly influenced by the current designs around them. When Toyota was developing its luxury car, the 'Lexus', it blended the drawn back cabin of Mercedes (which apparently makes the format of the car like that of a coach and horses), the distinctive rear pillars of a BMW, and the shallow curves of up-market Audis. The result has been a steady commercial success, despite the difficulty of competing against such established names.

To be more explicit about this model of technology blending as it applies to the problem of Figure 2.1, let the frame masses in a population of  $n$  firms be given by  $(m_1, \dots, m_n)$ . Under partial imitation, firm  $i$  selects  $r_i$  designs as 'rôle models', where the probability of selecting the design of firm  $j$  in

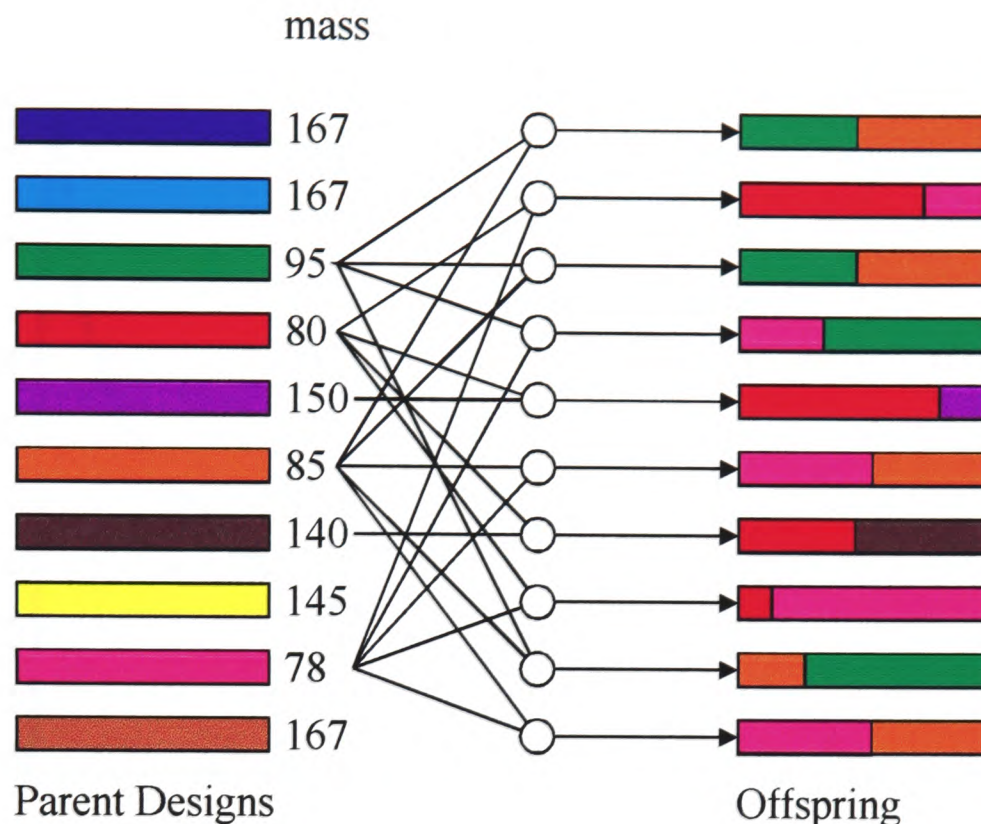


Figure 2.4. Reproduction and crossover with  $n = 10$  and  $r = 2$ .

any draw from the population is  $m_j^{-1} / \sum_{k=1}^n m_k^{-1}$ . (In the genetic algorithm literature, this is known as ‘roulette wheel reproduction’, as the operation is like a biased roulette wheel, where each design is assigned a slot sized in proportion to its performance.) From these selected rôle models it constructs a blended design by an operation known as ‘crossover’. The design strings of the rôle models and the prototype are divided into  $r_i$  randomly sized sub-strings. The first sub-string of the first rôle model is copied into the first sub-string of the prototype; the second sub-string of the second rôle model is copied into the second sub-string of the prototype, and so on. This process is illustrated in Figure 2.4 for  $n = 10$  and  $r_i = r = 2$  for all  $i = 1, \dots, n$ . The prototype is then randomly mutated according to the annealing schedule Figure 2.2 and adopted in the next generation if its mass is less than the mass of the firm’s current design.

Partial imitation offers a significant improvement in performance over the individual learning of the last section. For  $r_i = r = 1$  for all  $i = 1, \dots, n$ , where there is no crossover, this is, of course, a trivial observation. The cross-




Figure 2.5. The performance gains from social learning ( $r = 4$ ) relative to individual learning ( $r = 1$ ). (Ten firms; sample size 1000.)

imitation of successful designs keeps all firms close to the technology frontier. Each period there are thus  $n$  adaptations of designs close to the technology frontier rather than the one adaptation under individual learning. Learning is thus approximately  $n$  times faster.

However, it is apparent that there is an additional performance improvement from  $r > 1$ , where crossover or blending of designs takes place, as Figure 2.5 confirms. Why is this? First, blending two quite separate designs may be a very rapid way for a firm trapped in a local optima to make progress. Secondly, blending designs may be an analog of the 'golden section search' (Press, Vetterling, Teukolsky, and Flannery 1994, Chapter 10.1) method of numerical optimization, where one evaluates the objective function at a (weighted) intermediate point between two known points. Unlike such conventional numerical techniques, however, both the genetic algorithm and the partial imitation described here will work with combinatorial opti-

mization problems as well as smooth functions.

## 2.4 ECONOMIC APPLICATIONS

### 2.4.1 Stochastic Learning Curves

Arrow, in his model of ‘learning-by-doing’ (Arrow 1962), recognised that many observations of real-life research and development suggested a *learning curve* whereby cost decreased exponentially with ‘experience’, and built this observation into a model of growth. This semi-empirical approach depends, of course, critically on the quality and completeness of learning curve data. We avoid this problem here by taking the *processes* by which firms solve design problems as given, and then using computer simulations to generate high quality learning curve data artificially.

So what sort of learning curves do these processes generate? Consider first the stochastic learning curve generated by single-firm experimental search. Taking current mass  $m_\tau$  as a proxy for the state of the system, suppose that the learning process presents a new design each period which fails a stress test with probability  $h(m_\tau)$ . If the new design passes the stress test, assume that the mass reduction is distributed according to the cumulative distribution  $F$ , with density  $f$ , where the parameters of this underlying distribution depend only on current mass. The approximate model of the learning curve is thus given by:

$$Pr(m_\tau - m_{\tau+1} \leq x) = \begin{cases} 0 & : x < 0 \\ (1 - h(m_\tau))F(x; m_\tau) + h(m_\tau) & : x \geq 0 \end{cases} \quad (2.1)$$

In principle, we could use the data provided by simulation results to suggest functional forms—or even calibrate—this approximate model. For the example design problem discussed above, panel (a) of Figure 2.3 shows that  $h(m_\tau)$  is roughly constant; panel (b) shows that the underlying distribution

of mass reductions is fairly irregular even when the sample size is very large. However, it clearly would not be outrageous to model this as a normal distribution whose mean and variance vary with current mass as in panels (c) and (d).

However, before simply plugging a stochastic learning curve like that of equation (2.1) into a model, it might be prudent to consider the following:

1. What determines the rate at which trials are made? That is, how does a firm choose a profile  $\dot{\tau}(t)$ , where  $t \in [0, \infty)$ ? In the ‘learning-by-doing’ literature, of course, trials are tied directly to output. However, this need not necessarily be so in firms with a separate research and development capability that is able to create and test prototypes independently, at a rate related to research and development expenditure. Perhaps this depends on the industry in question. For relatively ‘heavy’ industry (*e.g.* aircraft manufacture, shipbuilding or power plants), where building and testing prototypes is an expensive business, the ‘learning-by-doing’ assumption might be quite appropriate.
2. Do the firms know how the parameters of  $F(x; m_\tau)$  vary with with state of technology, or do they have to learn these relationships? If the technology is relatively familiar, it might be safe to assume that they do know. However, a drastic innovation might lead to quite different research and development behaviour. For example, firms might forgo short-term profits by choosing a high trial rate to learn as much as they can about these underlying relationships<sup>7</sup>.
3. The model of *partial imitation* outlined above emphasized the potential impact that a ‘technological community’ may have on the rate of research and development. That is, the probability of stress rejection and

---

<sup>7</sup>Readers interested in doing a little partial imitation of their own might like to apply to this issue some of the tools developed in the now extensive literature on learning about a stochastic demand curve.

the mass distribution of new designs will also depend upon the number of firms making a product and the nature of the cross-imitation between them.

In practice, stochastic learning curves like that of equation (2.1) are likely to be difficult to incorporate into analytical models. For example, if the underlying density  $f$  were normal, the expected mass reduction at each trial would have to be calculated numerically. However, it may be possible to construct simple models of experimental search that give stochastic learning curves with features broadly consistent with those generated by the simulations above. To give an explicit example, suppose we specify an infinite product space in which all designs have a cost drawn independently from a uniform distribution on the unit interval. All research strategies in this case are effectively a random search of the space. It is relatively straightforward to show with such a product space that the expected cost after  $\tau$  trials (*i.e.* a firm's expected 'learning curve') is simply  $1/(1 + \tau)$ . That is, with respect to the number of trials, the expected learning curve is hyperbolic rather than logistic. However, we would expect the trial rate with respect to time to be smaller at the base of the learning curve than at the top, since the expected mass reduction here is small. The expected learning curve with respect to time may therefore have the logistic features that many authors claim to observe in real-life product development<sup>8</sup>.

#### 2.4.2 Patent Design

The literature on patent design looks at the trade-off between innovation-inducing incentives before the award of a patent and the costs of monopoly power or increased market concentration afterwards. The obvious way to

---

<sup>8</sup>A key feature of experimental search that this simple model fails to capture is the effect of an optimal annealing schedule causing the size of the 'jumps' in the search space to decline, resulting in a declining variance of underlying mass change, as in panel (d) of Figure 2.3.

perform this trade-off is to design a patent with a certain *length*: a long-lived patent gives strong incentives but high market power for its owner; short-lived patents give weaker incentives but a shorter period of high prices for consumers. Within a certain very simple innovation structure, Nordhaus (1969) explains how the optimal patent length may be calculated<sup>9</sup>. More recent work has recognised that an alternative to varying a patent's length is to vary its *breadth*. Defining patent breadth to be the ability of the patentee to raise price, Gilbert and Shapiro (1990) find conditions under which the optimal patent is of infinite length, with breadth used to perform the trade-off between market power and incentives. In a more sophisticated model, Klemperer (1990) shows conditions under which short, wide patents are to be preferred to long, narrow ones.

The model of partial imitation considered above suggests that there may be an additional cost after the award of a patent arising from a greatly restricted 'pool' of parent designs for further innovation. The model of partial imitation does not, however, suggest that patents are always a negative influence on innovation. On the contrary, it reveals a new way in which patents, quite apart from profit incentives, may encourage innovation.

For a given design problem and a given population size  $n$ , it seems there is an optimal symmetric number of 'parents'  $r^*$  that maximizes the performance of the social learning. Panel (a) of Figure 2.6 shows how  $r^*$  varied with  $n$  in the simulations. (This would seem to vindicate not taking the 'bisexual' genetic algorithm too literally!) Panel (b) shows how the improvement in performance translates into discounted social welfare in a Cournot model<sup>10</sup> with  $n = 10$ . We can see that the biggest improvement comes from moving from one parent per firm to two parents per firm.

An interesting question is: are the firms likely to choose the number of parents that maximizes social welfare? To investigate this issue, consider

---

<sup>9</sup>A more compact exposition is Scherer (1972)

<sup>10</sup>Inverse demand was given by  $P(Q) = 1600 - Q$ ; the marginal cost of firm was taken to be its product mass; a discount factor of 0.995 was used.

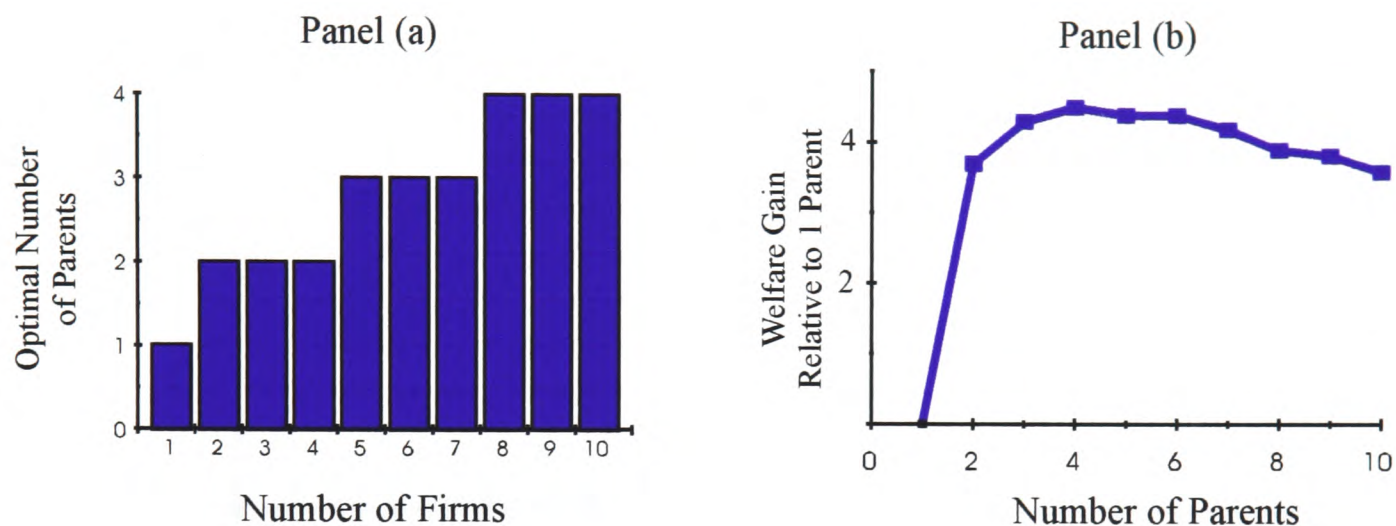


Figure 2.6. The optimal level of crossover.

a ‘meta-game’ where firms simultaneously choose  $r_1, \dots, r_n$ , and where the payoffs in the game are the expected, discounted sum of profits to each firm resulting from a strategy profile  $(r_1, \dots, r_n)$  within a certain market structure.  $(r^*, \dots, r^*)$  might not be a Nash equilibrium of this game, as each firm might have an incentive to deviate to  $r_i < r^*$  to maximize its expected payoff. This conjecture was borne out by the simulation results. For example, with ten firms and all firms choosing a socially optimal four parents, the mean gain for a single firm deviating to  $r_i = 1$  was significantly greater than zero<sup>11</sup>. The intuition for this is straightforward. The high performance of social learning based on partial imitation, relative to individual learning, depends upon there being a diverse ‘pool’ of alternative designs at any one time. This

<sup>11</sup>The mean gain in discounted profit was 2418 over 2000 samples. Given that the variance of the sampling distribution of differences was 903 032, this means that we can safely reject the hypothesis of no free-rider problem using a one-tailed test at a 99% confidence level.

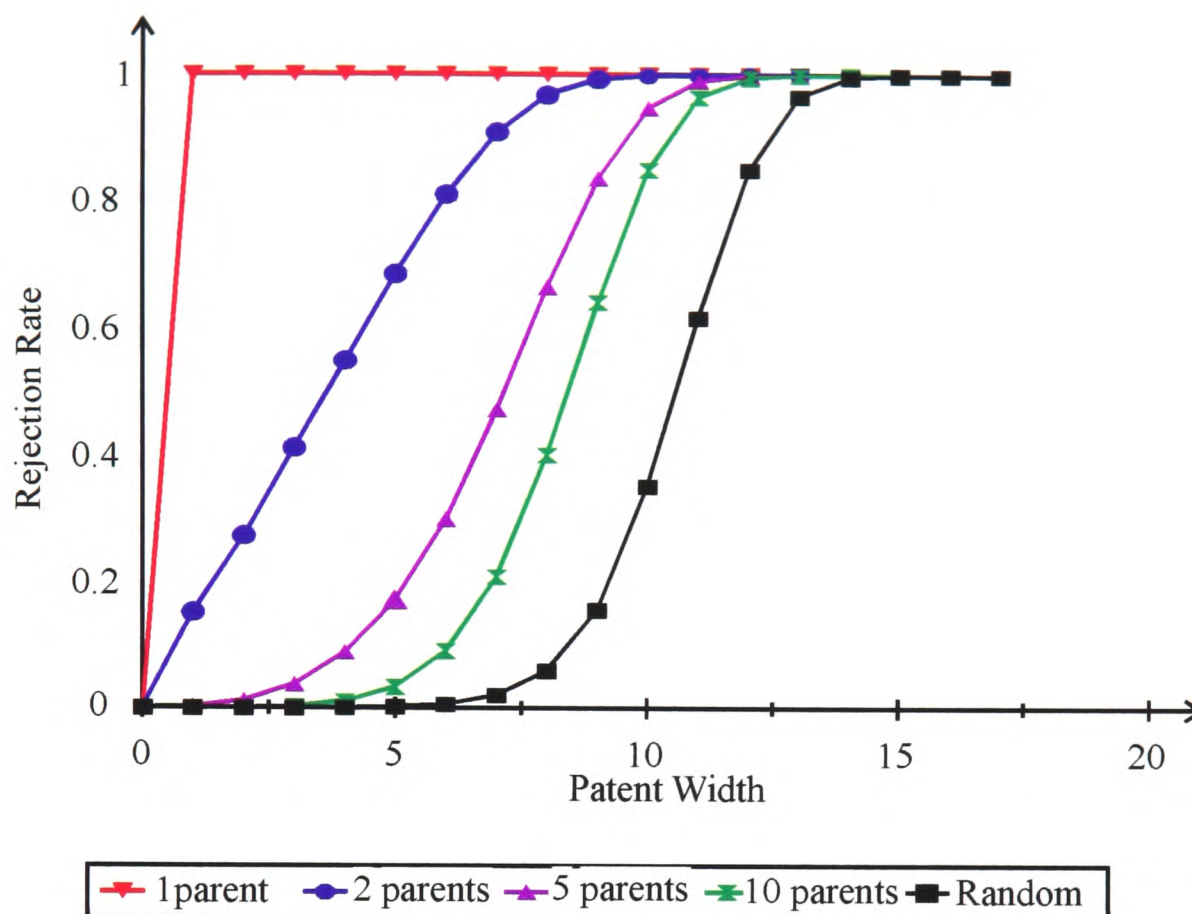


Figure 2.7. The impact of patent width on firms' R&D decisions. The results are for ten firms, with current designs taken from random points in the design space. The effect of mutation is ignored. (Sample size 100 000.).

diversity means that there will be a wide distribution of qualities within the population. The temptation facing each firm is to improve its chances of being consistently at the upper end of this distribution by directly copying the best current design, rather than blending designs.

How, then, can we design patents to encourage firms to choose at least two parents in cases where such a free-rider problem exists? We may start by noting that a natural definition of patent width in this context is to borrow a useful concept from the computer science literature and specify a minimum *hamming distance* between a protected design and any subsequent design. To find the hamming distance between two binary strings one simply compares the two strings and counts the number of bits that are different. Secondly, note that here the firms only look to current designs as a source for partial imitation, so that—if all we are interested in is encouraging efficient

learning—we may restrict attention to patents lasting just one period<sup>12</sup>.

Figure 2.7 shows the impact of awarding free one-period patent protection of width  $w$  to all existing designs on the success of different R&D strategies. A width of ten or more severely affected firms even choosing random points in the design space. However, a width between one and seven would seem to encourage partial imitation based on two to ten parents—resulting in the social benefits shown in Figure 2.6, panel (b). Two caveats should be noted. First, as we have not considered the full space of possible R&D methods, a patent width of one is unlikely to be sufficient to encourage a *uniform* mix of two designs: the temptation will exist to strongly bias the mix in one direction or the other. Secondly, the results of Figure 2.7 were conducted with current designs taken from random points in the design space. This may roughly correspond to an early stage in the social learning process (where, incidentally, we may note from Figure 2.5 the gains from crossover are highest), but as designs converge towards the optimal design the impact of patent protection will be much stronger. That is, at a later stage of social learning, the curves in Figure 2.6 will be shifted to the left. Taken together, these considerations suggest that patent width should begin fairly high to encourage a diverse ‘pool’ of potential parent designs, then gradually loosen as designs converge towards some optimum.

### 2.4.3 Technodiversity

‘Biodiversity’, the diversity of elements of life within an ecosystem, is considered worth preserving for a number of complex reasons. For example, the diversity might stabilise the ecosystem against external shocks. Or it may just be that a biodiverse ecosystem has a larger ‘genetic library’ and consequently greater evolutionary potential.

---

<sup>12</sup>In a more general model in which the ‘pool’ of designs supplying parents was cumulative in that it included *all* past designs, we would have to worry about patent width *and* length.

The discussion of partial imitation above suggests that a similar case could be made for what we might call *technodiversity*. We saw how a diverse ‘pool’ of designs at an early stage of social learning accelerated the learning process. The diversity stops the learning getting stuck on local optima, it provides multiple ‘sources of inspiration’ for further development. We saw in the last section how careful product protection might be used to encourage this diversity in the face of free-riding temptations that might reduce it.

We shall end by noting a further application of this idea, and that is to the introduction of new technologies to developing countries. Ever since Schumacher (1972), charities such as *Intermediate Technology* have been engaged in introducing ‘appropriate’ technology into deprived areas of the world. The principle is to find technology that suits the immediate needs of those it is given to, does not require external skills and does not unduly upset the social fabric of the communities who adopt it. A disadvantage of the approach is that it does little to encourage—indeed, may discourage—*indigenous* technological development. While the recipients may not depend on external skills to operate the technology, they *may* come to rely on external sources for the next innovation.

The discussion of technodiversity above suggests that rather than attempt to introduce the single-most ‘appropriate’ technology to a deprived community, a better approach might be to introduce a range of technologies even if some of these are obviously inferior. The diversity of designs should encourage the community itself to develop new and better designs using the simple techniques of partial imitation. For example, in the refugee camps of Zaire it was recently noted that modern stoves combined with charcoal pellets produced by technology originally introduced to make animal feed were able to substantially reduce smoke pollution and consumption of the extremely scarce local supply of wood. Not only will such social learning probably result in designs and methods more ‘appropriate’ than the best of the initial designs, but the process itself would build up design and development skills more generally applicable in the future. Ultimately, this is a

potentially important source of *sustainable* development.

Of course, there is an important trade-off to be made here between the immediate relief given by the new technology and the long-run potential for indigenous development. How best to perform this trade-off, and how such a process would interact with other development issues such as education, is the subject of the next chapter.

## 2.5 OTHER METHODS OF MODELLING RESEARCH AND DEVELOPMENT

The primary aim of this chapter has been to provide some insights into how simple but realistic models of research and development can be built up from first principles. It may therefore be instructive to offer a brief critique of some current approaches to modelling research and development in the light of these results.

The most frequently used method for modelling innovations is the *Poisson* or *memoryless* success rule. Its use began with the ‘patent-race’ literature, starting with Lee and Wilde (1980), but has extended even to models of endogenous growth, such as Grossman and Helpman (1991). No-one could deny its simplicity. If a producer engages in research and development activity of intensity  $z$ , the probability of making an innovation in the time interval  $(t, t + dt)$  is  $zdt$ . In order to model a sequence of innovations with this rule one needs to assume that innovations occur in a sequence of discrete jumps of known size. For example, if  $\omega$  is the number of innovations so far, cost could be given by

$$c(\omega) = \lambda^{\omega-1}, \quad (2.2)$$

where  $\lambda < 1$ .

This rule obviously bears very little affinity to the models of research and development discussed above. While it may to some extent capture uncertainty about *when* innovations occur, it does nothing to capture uncertainty about the *size* of innovations. This could be especially problematic in mod-

els of growth, which may be unduly sensitive to the technological structure imposed by the modeller in relations like that of (2.2). If (2.2) could be deduced from what we know about science and technology, and if  $\lambda$  could be measured, then the use of the Poisson success rule might be justified; however, there is little evidence to support such an assertion.

The simplicity of the Poisson success rule is, of course, a concession to tractability rather than realism. Stochastic systems with ‘memory’ are more difficult to handle—but the task is not impossible, as confirmed by the models of Harris (1988) and Budd, Harris, and Vickers (1993). There are two firms, labeled *A* and *B*. The market share of *A* is given by a point  $s$  in the interval  $[0, 1]$ ; the market share of *B* by  $1 - s$ . *A* and *B* choose efforts (*e.g.* research and development intensities)  $x(s)$  and  $y(s)$  respectively. The motion of  $s$  in  $(0, 1)$  follows a Brownian motion process with drift proportional to the difference in the firms’ effort rates:

$$ds = [x(s) - y(s)]dt + \sigma dw. \quad (2.3)$$

Here,  $w$  represents the Brownian motion process itself:  $dw$  is normally distributed with mean zero and variance  $dt$ ; so  $ds$  is normally distributed with mean  $[x(s) - y(s)]dt$  and variance  $\sigma^2 dt$ . Technological progress is continuous and depends on past as well as present effort rates.

At first sight, this seems closer to the mark. The process in (2.3) captures uncertainty about the timing *and* magnitude of innovations. However, by taking a market variable as the state variable, the process is not tied at all to the underlying technology. The models discussed in this paper, on the other hand, suggest that the mean and variance of innovations declines as the underlying technology improves (see panels (b) and (d) of Figure 2.3).

## 2.6 CONCLUSION

We have modelled research and development by modelling explicitly the process by which firms solve real engineering design problems. Two processes

have been examined: individual experimental search, based on a problem-solving technique known as *simulated annealing*, and partial imitation, based on a problem-solving technique known as the *genetic algorithm*. Studying these processes has suggested some explanations of why a firm's search of a design space seems to become more focused as it moves 'down the learning curve', and has highlighted an important social aspect to technological learning that is rarely modelled.

Some economic implications of these processes have been explored. They suggest a structured way of thinking about stochastic learning curves. They suggest how patents may be designed to encourage a socially desirable level of partial imitation. Finally, they suggest the importance of 'technodiversity'—a pool of seed technologies to encourage indigenous, sustainable development—in the introduction of new technology to developing countries.

## APPENDIX: THE GENETIC ALGORITHM

As mentioned in the main text, the genetic algorithm is a general problem-solving technique inspired by natural genetic processes.

The basic building block of the genetic algorithm is a string of characters encoding some ‘answer’ to a certain problem. The string is generally written in binary form; a class of strings, or *schemata*, is constructed using the additional wildcard symbol ‘\*’, where the ‘\*’ means ‘either 1 or 0’. For example, ‘\*\*0101\*10\*\*’ is a schemata of length-eleven binary strings with defining length—the distance between the outermost non-wildcard characters—seven.

The algorithm acts on a population  $B(t)$  of  $n$  strings of equal length  $k$ . Suppose that in period  $t$ ,  $B(t) = \{s_1, \dots, s_n\}$ . A new set  $B(t + 1)$  is constructed from the following steps:

1. Evaluate the ‘strengths’ of the strings—*i.e.* how ‘good an answer’ they are. These are given by an objective function  $v(s_j, t)$ .
2. Assign each string a probability weight in proportion to its strength:

$$\left\{ p_1, \dots, p_n \mid p_j = \frac{v(s_j, t)}{\sum_{i=1}^n v(s_i, t)}, \quad j = 1, \dots, n \right\}$$

3. Select  $p$  pairs of strings (where  $2p \leq n$ ) using this probability distribution (this is termed ‘roulette wheel’ selection).
4. Create  $2p$  ‘offspring’. For each chosen pair, choose a random *crossover* position  $l$ ,  $1 < l < k$ ; then exchange the segments to the left of position  $l$  in the two strings. (Other genetic operators, such as *mutation*, may also operate at this stage.)
5. Replace the  $2p$  lowest strength strings in  $B(t)$  with these offspring.
6. Set  $t$  to  $t + 1$ .

Figure 2.8 shows an example of its operation, based on an illustration in Holland, Holyoak, Nesbett, and Thagard (1986). The genetic algorithm ‘evolves’

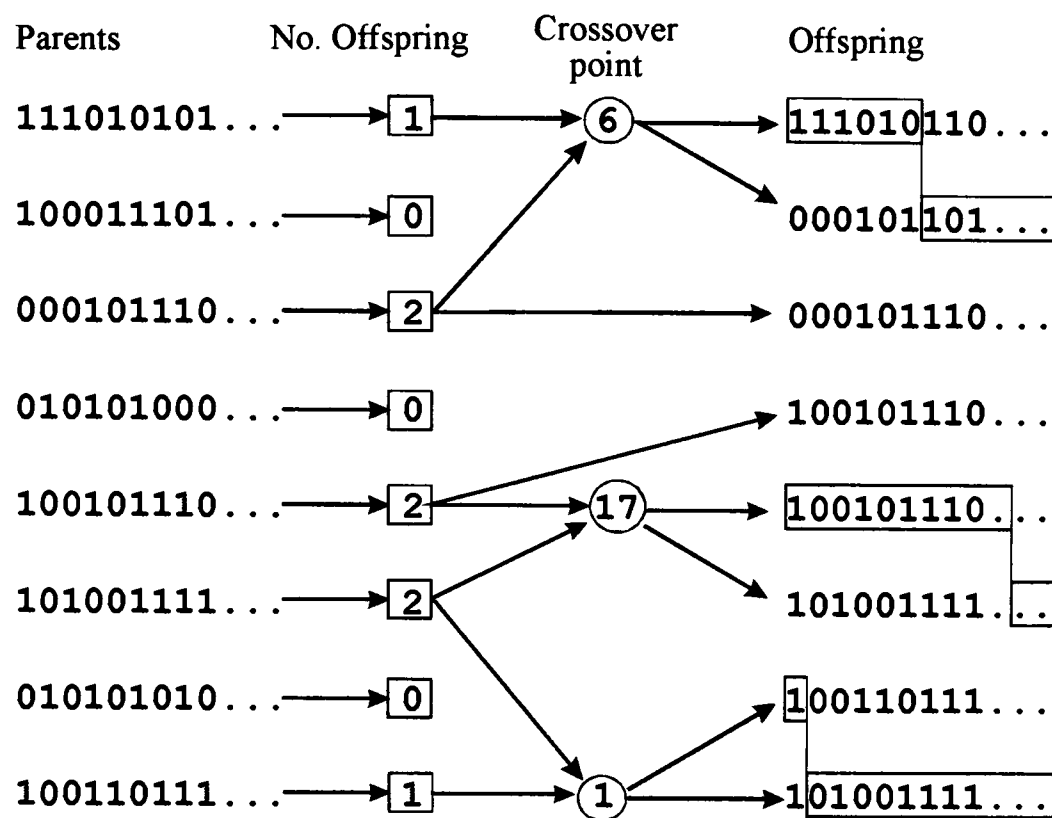


Figure 2.8. An illustration of the genetic algorithm.

better answers to a problem in a way similar to the natural processes that evolve fitter genes for an environment. Holland's *schemata theorem* states that high performance, short defining length schematas receive at least exponentially increasing numbers of trials in successive generations.

The algorithm called 'partial imitation' in the main text is simply a version of the genetic algorithm adapted to allow for:

- an interpretation based on firms independently making design updating decisions (rather than arbitrarily replacing the lowest strength strings, as in the straight genetic algorithm), and
- firms creating new designs from more than two sources.

A useful reference on the genetic algorithm is Goldberg (1989).

## *Chapter 3*

### *Technodiversity*

#### **SUMMARY**

Just as an ecosystem rich in biodiversity has access to a diverse pool of genetic material, an economy rich in *technodiversity* has access to a diverse range of alternative technologies. This chapter presents some models in which innovation is linked to technology diffusion in such a way that technodiversity can be seen to encourage indigenous development.

### 3.1 INTRODUCTION

The new contribution of this chapter is to construct models that link the process of technology diffusion to the process of product or process innovation<sup>1</sup>. In these it is shown that a mix of many different technologies—a situation described as *technodiverse*—may accelerate the adoption of superior technologies and encourage indigenous development. Technodiversity is intended to be analogous to biodiversity. Just as an ecosystem rich in biodiversity has access to a diverse pool of genetic material, an economy rich in technodiversity has access to a diverse range of alternative technologies.

The motivation for studying policies that may encourage indigenous development comes from looking at the problems in the world around us. It is clear to even a casual observer of the world economy that there is a huge disparity in the levels of technology utilised by developed and less developed economies. This is, superficially at least, rather surprising. Why do the less developed economies not simply imitate the technology of more developed economies?—or, if the ‘state-of-the-art’ technology is unsuitable, imitate some older technology?<sup>2</sup> This makes it tempting to conclude that the technological disparity we observe is no more than a short-run phenomenon, a disequilibrium effect. However, there are an increasing number of arguments to suggest that, if left unchecked, this disparity will remain permanent. First, there are the new growth (or ‘endogenous growth’) models. In Romer (1990), for example, levels of human capital below a certain threshold result in stagnation, *i.e.* zero growth; in Redding (1995), coordination failures at intermediate levels of human capital result in a ‘low-level development

---

<sup>1</sup>Throughout this chapter, the word ‘technology’ is used in its most general sense to refer to either a product or a process.

<sup>2</sup>There may even be some advantage in being a technological follower rather than a technological leader. Some eastern European countries, for example, now enjoy better telecommunications networks than their western counterparts because they were able to start from scratch with the latest equipment.

trap', an effect that is exacerbated by imperfect capital markets. Secondly, developing countries may be caught in a 'trade trap' (Cootes 1990); that is, an excessive dependence on the cash commodities needed to, for example, service debt payments—a dependence that diverts resources away from indigenous innovation. Thirdly, we may note that it is agricultural innovation that first frees human capital for indigenous innovation early in the technological development of an economy. However, agricultural innovation is very environment-sensitive, so it may not be possible to redress technological imbalance in this area by straight imitation from more developed economies.

So a strong case can be made for intervention to encourage development. The focus in this chapter will be on microeconomic measures to facilitate innovation. That is not to say that the macroeconomic environment is unimportant—far from it. However, microeconomic measures have received far less attention in the literature and, if the conclusions of this chapter are correct, that is in disproportion to their worth.

In the conventional approach of, for example, Islam (1992), microeconomic measures to facilitate technological progress are divided into two types. First are measures intended to encourage indigenous innovation and development by directly building up human capital. Second are measures to transfer technology from more developed countries (Stewart 1990) which will then, hopefully, spread by some 'diffusion' process to all potential users.

Given that it is insufficient levels of human capital that causes stagnation in many of the new growth models, the first of these measures would seem to commend itself. However, education and training are acutely expensive and the chances of them becoming self-funding depend on the levels of human capital they induce reaching a relatively high level. So this method used alone is likely to lead to frustration. The second method also has its problems. Even if the technology introduced is what Schumacher (1972) would call 'appropriate'—so that it avoids the skills-dependency and damaging concentration often associated with transferring more advanced technology—the transfer does little to directly build up human capital and hence little to en-

courage indigenous innovation<sup>3</sup>.

The claim in this chapter is that this dichotomy of methods into measures to encourage indigenous innovation and methods to transfer or diffuse appropriate technology is a false one. Indeed, when the nature of innovation is properly understood, we find that technological diffusion and technological innovation may go hand in hand. Moreover, this coevolution of diffusion and innovation is best nurtured by an environment in which there is a high level of *technodiversity*; that is, where there is a fairly even mix of a number of alternative technologies designed for the same task or similar tasks.

The models below consider the problem faced by an organisation (*e.g.* a government, international agency or non-government organisation) wishing to transfer some technology to an under-developed region. The question is: which technology should they transfer?—should they even transfer a mix of technologies and, if so, how many and in what proportions? This question is addressed first for the case where the organisation has access to two ‘seed’ technologies, and is addressed for two alternative models of technology diffusion. The possibility of introducing more than two ‘seed’ technologies is then considered. The chapter ends with a discussion of some of the implications of the models, and suggestions for further research.

### 3.2 PARTIAL IMITATION

Before plunging into a description of the models, it may be helpful to consider their central assumption, which is that innovation does not typically originate spontaneously from thin air but is more often the result of intelligently combining technologies that already exist.

This is an idea more extensively studied in the last chapter, where an

---

<sup>3</sup>Although, as mentioned above, transferring *agricultural* technology may release human capital for other purposes. However, finding suitable agricultural technology to transfer is no easy business—it may have to be specially developed for the environment concerned *in situ*.

analogy was drawn between (bisexual) biological evolution and technological development. Just as reproduction in a biological model involves combining the genes of two 'successful' parents, so innovation in a technological model involves combining the designs of two or more 'successful' source designs, a process dubbed *partial imitation*.

Of course, we know by introspection that ideas are very often the product of combining two or more existing ideas, rather than coming to us spontaneously; so it should come as no surprise that technological ideas, *i.e.* designs, might develop in a similar manner. The effect is certainly obvious when we consider large innovations in the history of technology. Combining the power source used for pumping water from mines with primitive rail technology gave us the steam locomotive, for example. Perhaps less obvious is the role that partial imitation plays in the gradual development of a product by a series of small innovations. But we can be sure, for example, that any new car will incorporate design ideas drawn from many of its most recent competitors. Partial imitation is a key feature of innovation because it is so simple and so effective.

Very little has been done to study the possible implications of technology blending for the transfer of technology to developing countries. An exception is Bhalla and James (1986), which suggests that rather than introduce high technology alone, or low ('appropriate') technology alone, why not blend the two? Examples are given where such technology blending has proved successful, including using electronic load controllers to control small-scale hydro-electric projects in Colombia, and using advanced cloning methods to improve palm oil production in Malaysia. However, there is no discussion of exactly what circumstances favour such blending, nor of how blending might subsequently improve indigenous development—these being the subjects of the current chapter.

### 3.3 TWO SEED TECHNOLOGIES

Suppose an organisation introduces two seed technologies,  $a$  and  $b$ , to a large population of users in the proportions  $c$  of technology  $a$  and  $1 - c$  of technology  $b$ , where  $c \in [0, 1]$ . (Equivalently, suppose that an existing population exclusively use a technology  $b$  and we introduce a superior technology  $a$  at time zero, replacing a proportion  $c$  of the inferior technology.) Without loss of generality, suppose  $q_a > q_b$ .

We begin with the case where the ensuing diffusion process updates in discrete steps, separated by  $\tau$  units of time. The agents in the population have ‘Poisson alarm clocks’. That is, at each step a given agent is prompted to update her technology with probability  $\tau$ . An agent prompted to update also has a certain tendency to innovate given by the parameter  $\xi \in (0, \frac{1}{2})$ . With probability  $1 - \xi$  they behave as a ‘non-innovator’ and their updating is based on straight imitation of some technology. But with probability  $\xi$  they behave as an ‘innovator’, constructing some new technology by combining two existing technologies. Here we assume that this combining of technologies operates in a very simple fashion to sometimes give a single new technology  $ab$  according to the following rules:

$$\begin{aligned} a + a &\rightarrow a \\ b + b &\rightarrow b \\ a + b &\rightarrow ab \\ ab + (a \text{ or } b \text{ or } ab) &\rightarrow ab \end{aligned}$$

We assume that the organisation either cannot, or does not want to, introduce the combined technology  $ab$  itself from the start. It may not be feasible, from a cost point of view, for them to design, construct and test the combined technology in a fully-blown pilot scheme. Alternatively, even if the organisation is able to introduce  $ab$  themselves, we shall be able to consider

whether they should do so by comparing the outcome when they do to that when they do not. It is quite possible that the gains from inducing indigenous development and from having  $ab$  designed *in situ* might well outweigh the gains from having  $ab$  available from day one.

Moreover, it seems reasonable to suppose that while the organisation may have some rough idea about the value of  $q_{ab}$  in a known environment, they have a poor idea of the value of  $q_{ab}$  within the target environment. In the models below we assume that the organisation believes  $q_{ab} \sim U(\underline{q}, \bar{q})$ , where  $E(q_{ab}) = \frac{1}{2}(q_a + q_b)$  and  $\bar{q} > q_a$ . Models with different prior beliefs would be similarly straightforward to derive.

The proportion of agents using technology  $a$  at time  $t$  is given by  $x(t)$ , the proportion using technology  $b$  is given by  $y(t)$ ; so the proportion using the new technology  $ab$  is  $1 - x(t) - y(t)$ .

### 3.3.1 A Linear Model

Consider a set of updating rules that result in a linear diffusion process. This is an interesting case not because the rules are the most plausible one could imagine, but because the resulting model is fully solvable for the optimal policy, and it therefore serves as a useful benchmark against which models based on more reasonable rules can be compared. The rules are intended to reflect a situation where communication is good, the two most popular technologies are generally known, and where agents update partly according to quality considerations but are also influenced by relative popularity. They are:

- Let the *dominant technology* at time  $t$  be highest quality of the two most popular technologies in use at that time. In the event of a popularity tie-break, the highest quality technology always wins.
- *Non-innovators* simply adopt the dominant technology.




Figure 3.1. The state space for the linear model.

- *Innovators* take the dominant technology and mix it with their existing design. They compare their new mixed design with the dominant technology and adopt the one with the highest quality.

First suppose that  $q_{ab} > q_a$ . Under these rules, we can separate the interior of the state space into two zones separated by the line  $\max\{1 - 2x(t), \frac{1}{2}(1 - x(t))\}$ , as shown in Figure 3.1. Above this line,  $a$  and  $b$  are the two most popular technologies, so the dominant technology is  $a$ . Non-innovators adopt  $a$ . Innovators adopt  $a$  only if they are currently using  $a$ ; otherwise their mixing of designs produces technology  $ab$ . For  $(x(t), y(t))$  above this line we may thus write:

$$\left. \begin{aligned} E(x(t + \tau)) &= (1 - \tau)x(t) + \tau(1 - \xi + \xi x(t)) \\ E(y(t + \tau)) &= (1 - \tau)y(t) \end{aligned} \right\} q_{ab} > q_a \quad (3.1)$$

On or below this line—apart from the corners  $(0, 1)$ ,  $(0, 0)$  and  $(1, 0)$ , where there is only one technology in use— $ab$  is at least the second most popular technology and is therefore the dominant technology. All those updating will choose  $ab$ . In this lower zone we may therefore write:

$$\left. \begin{aligned} E(x(t + \tau)) &= (1 - \tau)x(t) \\ E(y(t + \tau)) &= (1 - \tau)y(t) \end{aligned} \right\} q_{ab} > q_a \quad (3.2)$$

Taking  $\tau \rightarrow 0$  and writing  $x$  for  $x(t)$  and  $y$  for  $y(t)$ , we obtain  $E\left(\begin{smallmatrix} \dot{x} \\ \dot{y} \end{smallmatrix}\right)_{q_{ab} > q_a} = 0$  at  $(0, 1)$ ,  $(0, 0)$  and  $(1, 0)$ ; otherwise,

$$E\left(\begin{smallmatrix} \dot{x} \\ \dot{y} \end{smallmatrix}\right)_{q_{ab} > q_a} = \left( \begin{array}{l} (1 - \xi)(1 - x) : y > \max\{1 - 2x, \frac{1}{2}(1 - x)\} \\ -x : y \leq \max\{1 - 2x, \frac{1}{2}(1 - x)\} \\ -y \end{array} \right) \quad (3.3)$$

By using either a law of strong numbers combined with the large population size, or following Norman (1972), we may write  $\dot{x}$  for  $E(\dot{x})$  and  $\dot{y}$  for  $E(\dot{y})$ . Neither of these options (taking either the population size or the updating rate to infinity) is entirely plausible. However, we may use an argument of Boylan (1992) to suggest that the deterministic continuous time dynamics thus obtained are a satisfactory approximation of a finite population dynamic updating at a high but finite rate. Equation 3.3 may now be solved to give:

$$\left(\begin{smallmatrix} x \\ y \end{smallmatrix}\right)_{q_{ab} > q_a} = \left( \begin{array}{l} 1 - (1 - c)e^{-(1-\xi)t} : t < \frac{1}{\xi} \log 2 \\ (1 - 2^{\frac{-(1-\xi)}{\xi}}(1 - c))e^{-(t - \frac{1}{\xi} \log 2)} : t \geq \frac{1}{\xi} \log 2 \\ (1 - c)e^{-t} \end{array} \right) \quad (3.4)$$

Going through similar steps for  $q_{ab} \leq q_a$  obtains:

$$\left(\begin{smallmatrix} x \\ y \end{smallmatrix}\right)_{q_{ab} \leq q_a} = \left( \begin{array}{l} 1 - (1 - c)e^{-t} \\ (1 - c)e^{-t} \end{array} \right) \quad (3.5)$$

Suppose we start by defining the optimal policy to be the one maximizing expected discounted total population quality. The task facing the organisation is to find an optimal mix  $c^* \in [0, 1]$  given that  $q_{ab} \sim U(\underline{q}, \bar{q})$ , where




Figure 3.2. Optimal policies in the linear model.

$E(q_{ab}) = \frac{1}{2}(q_a + q_b)$  and  $\bar{q} > q_a$ . If the organisation discounts the future at the rate  $\rho$ , then we can make the following proposition:

**PROPOSITION 3.1** *There exists a  $\hat{q} > q_a$  such that, iff  $\bar{q} > \hat{q}$ , then there exists a  $\hat{\rho} > 0$  such that, iff  $\rho < \hat{\rho}$ , it is (strictly) optimal to set  $c = \varepsilon$ , where  $\varepsilon$  is arbitrarily small. Conversely, if  $\bar{q} \leq \hat{q}$ , or if  $\bar{q} > \hat{q}$  but  $\rho \geq \hat{\rho}$ , then it is (weakly) optimal to set  $c = 1 - \varepsilon$ .*

**PROOF.** The proof is straightforward but rather tedious, and is relegated to the appendix.

That is, there are basically two options open to the organisation, and these are illustrated in Figure 3.2.

First, if the organisation believes that the quality of the combined technology  $ab$  is unlikely to be high, or if they discount the future at a high rate, then it is optimal for it to pursue a ‘conservative’ policy. This involves

introducing almost all of the population to technology  $a$ . If the combined technology  $ab$  transpires to be inferior to  $a$ , then the remaining few without  $a$  quickly adopt it and the system converges quickly to  $B$ . If the combined technology transpires to be superior to  $a$ , then innovation by a very small proportion of the population over  $\frac{1}{\xi} \log 2$  units of time moves the system into the lower zone near  $B$ . From there, it moves by straight imitation of  $ab$  along  $BC$  until everyone has the superior technology.

Secondly, if the organisation believes that the quality of the combined technology  $ab$  is quite likely to be high *and* it discounts the future at a relatively low rate, then it is optimal for it to pursue a ‘radical’ policy. This involves introducing almost all the population to the inferior technology  $b$ . If the combined technology  $ab$  transpires to be inferior to  $a$ , then the system moves by straight imitation of  $a$  along  $AB$ . If the combined technology transpires to be superior to  $a$ , then this policy maximizes the opportunity for innovation. Innovation by a significant proportion of the population over  $\frac{1}{\xi} \log 2$  units of time moves the system to  $D$ , from whence it takes a relatively short time to converge to  $C$ .

The values of  $E(S)$ ,  $\hat{q}$  and  $\hat{\rho}$  can be found from expressions in the appendix, and from these we can obtain the following comparative static results:

$$\left. \frac{\partial E(S)}{\partial \xi} \right|_{\substack{c \rightarrow 0 \\ 0 < \xi < \frac{1}{2}}} > 0, \quad \frac{\partial \hat{q}}{\partial \xi} < 0, \quad \frac{\partial \hat{\rho}}{\partial \xi} > 0$$

That is, by increasing the expected discounted quality under the ‘radical’ policy, increasing the tendency to innovate  $\xi$  reduces both the quality and discounting thresholds above which the ‘radical’ policy is optimal.

### 3.3.2 A Non-Linear Model

Consider now an alternative model with more of an ‘evolutionary’ flavour. Agents prompted to update now do so according to the following rules:

- Agents called upon to update observe the technology used by one other agent drawn from the population at random.




Figure 3.3. The non-linear model,  $q_{ab} > q_a$ .

- *Non-innovators* compare the technology they observe with their current technology and switch iff the quality of the former is strictly greater than the quality of the latter.
- *Innovators* take the technology they observe and mix it with their existing design. They compare their new mixed design with the two source designs, and switch to the best design if the best is strictly better than their current technology.

Following similar steps to those in the section above, we obtain the following dynamical system for  $q_{ab} > q_a$ :

$$\begin{pmatrix} \dot{x} \\ \dot{y} \end{pmatrix}_{q_{ab} > q_a} = \begin{pmatrix} x(x + 2(1 - \xi)y - 1) \\ y(y - 1) \end{pmatrix} \quad (3.6)$$

The basic features, and an example path, of equation (3.6) are shown in Figure 3.3.

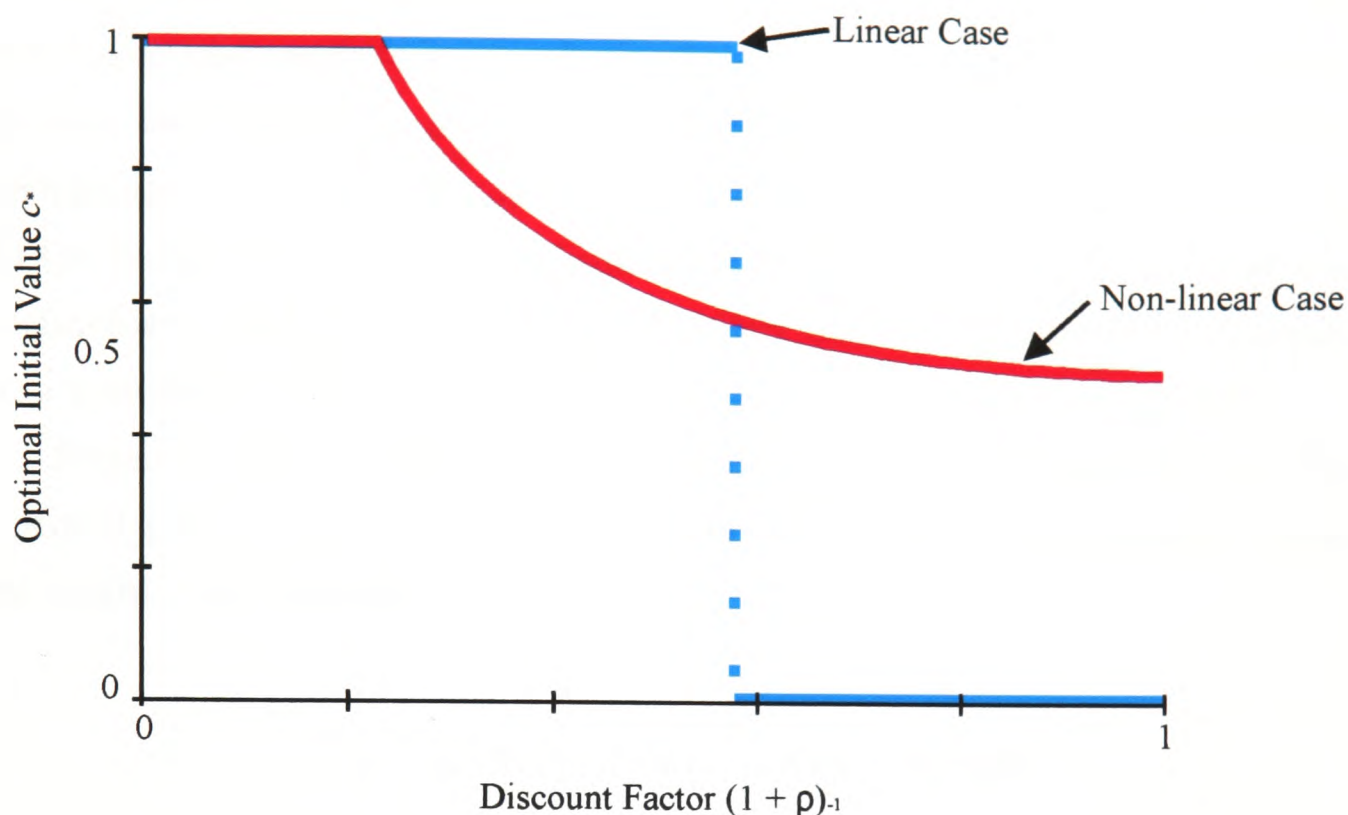


Figure 3.4. Numerical results for the optimal initial value.  $\xi = 0.4$ ,  $q_a = 205$ ,  $q_b = 195$ ,  $\bar{q} = 400$ .

For  $q_{ab} \leq q_a$ , we know that with initial conditions  $(c, 1 - c)$ , the state will change according to  $\dot{x} = x(1 - x)$  and  $\dot{y} = y(y - 1)$  along the line  $AB$ .

This nonlinear model generates time-paths with the characteristic ‘S’-shape of the logistic curves frequently used in econometric studies of diffusion—unlike the linear model, where the time-paths are exponential. For this reason alone, it is to be preferred. Unfortunately, we can no longer solve for the time path explicitly, which makes the sort of welfare analysis conducted above for the linear case rather difficult.

What we can do, however, is numerically simulate the process for given values of  $\xi$ ,  $q_a$ ,  $q_b$  and  $\bar{q}$ . That is, we can simulate the process for a range of discount rates, and, for each discount rate, a range of initial values—in each case noting the initial value that maximizes total discounted quality.

Figure 3.4 shows the results of some simulations of this kind. The non-linear case differs from the linear case in at least two respects.

First, as the discount rate tends to zero, the optimal initial value in the non-linear case tends to an equal mix of technologies  $a$  and  $b$ . This is because, in this case, innovation of  $ab$  requires the innovator to have a different seed technology to the agent she observes, a possibility that is most likely at  $c = \frac{1}{2}$ . In the linear case, by contrast, innovation of  $ab$  occurred when an agent had technology  $b$  and mixed it with the publicly observable superior technology  $a$ , a possibility that is most likely at  $c = \varepsilon$ , where  $\varepsilon$  is arbitrarily small.

Secondly, we no longer get a sharp dichotomy of optimal policies. The optimal policy now typically involves an initial mix of technologies—what we might call *technodiversity*.

### 3.4 MULTIPLE SEED TECHNOLOGIES

There is of course no reason why the organisation should restrict itself to introducing just two ‘seed’ technologies. How many different types of technology, then, should it introduce?

Again, as we try to address this question, it soon becomes apparent that the answer will depend heavily on the particular technology involved. Different technologies originate from very different design spaces and will combine in very different ways. However, there may be one or two things we can say that will apply generally.

Suppose the organisation has access to a range of technologies with qualities  $(q_0, q_1, q_2, \dots)$  such that  $q_i = \lambda q_{i-1}$ , where  $i = (1, 2, 3, \dots)$  and  $\lambda \in [0, 1]$ . The organisation now chooses a number  $N$  such that it introduces the first  $N + 1$  technologies, with qualities  $(q_0, q_1, \dots, q_N)$ . There is an obvious trade-off involved in the choice of  $N$ . First, a high value of  $N$  will depress the initial average quality in the population. Secondly, however, high  $N$  implies a high level of technodiversity, higher levels of innovation and, as we shall see, the possibility of higher long-run average quality.

The first effect is easy to see in this example. If the organisation introduces equal numbers of the  $N + 1$  technologies, then average initial quality

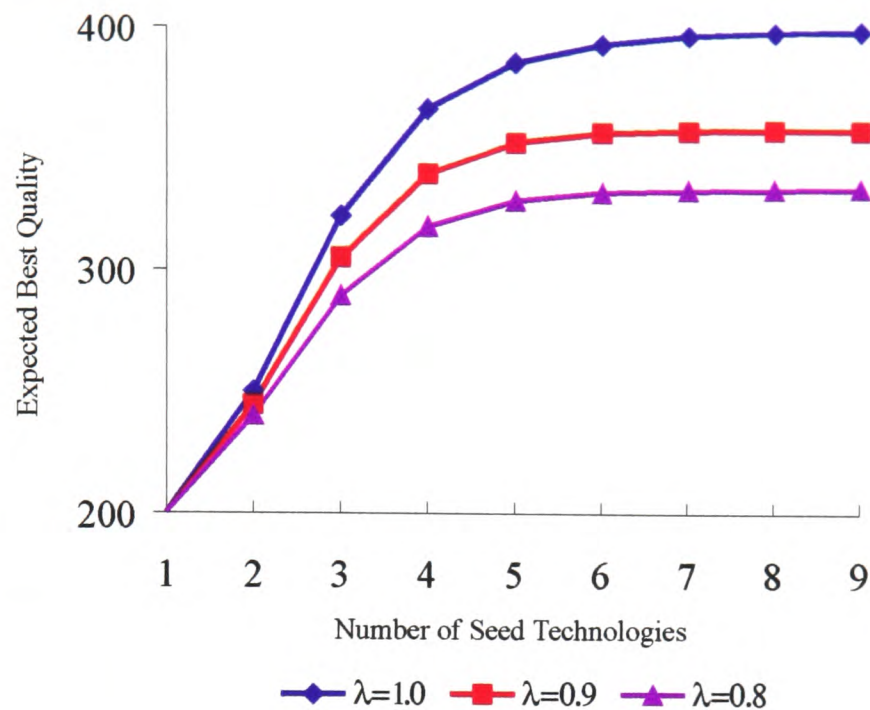


Figure 3.5. Monte Carlo results for expected best quality with multiple seed technologies. (The quality of new technology  $i$  was given by  $\hat{q} + u_i$ , where  $\hat{q}$  was the mean quality of the seed technologies combined to make  $i$  and  $u_i \sim U(-200, 200)$ . The best seed technology had quality  $q_0 = 200$ . Largest 95% confidence region in results was less than 1.3.)

is given by

$$\frac{1}{N+1} \sum_{i=0}^N q_0 \lambda^i = \frac{q_0}{N+1} \left( \frac{1 - \lambda^{N+1}}{1 - \lambda} \right),$$

which is strictly decreasing in  $N$  for  $\lambda < 1$ .

The second effect requires a little more thought. The number of ways of combining the  $N + 1$  technologies is given by:

$$\sum_{i=1}^{N+1} {}_{N+1}C_i = 2^{N+1} - 1$$

In the two seed ( $N = 1$ ) case, there were three technology combinations—*i.e.* just one more than those introduced by the organisation. There, we supposed that the quality of the third technology was distributed with a mean equal to the mean quality of the two seed technologies. As  $N$  increases, the number

of simple technology combinations increases exponentially. If we similarly suppose that new technologies made by combining the seed technologies have a quality distributed with a mean equal to the mean quality of the seed technologies that make it up, then we find that the expected best quality available to the population in the long-run is increasing in  $N$ , as Figure 3.5 makes clear. This is simply because the very large number of combinations associated with high  $N$  give innovators a large number of chances to 'get it right'. With nine seed technologies ( $N = 8$ ), for example, there are 502 new technology combinations for innovators to try. However, as we might expect, the effect is reduced by low values of  $\lambda$ , which lowers the average quality of the seed technologies.

So while we might not be able to say much about the multiple seed case without more information on the structure of the design space, we can say that an organisation interested entirely in short-term relief will want to introduce just one high quality seed, while an organisation more interested in the long-run outcome will want to introduce not just even mixes of seed technologies, but large numbers of *types* of seed technologies.

### 3.5 PLACING MORE WEIGHT ON INDIGENOUS INNOVATION

The comparisons between rival policies so far has been unrelentingly utilitarian in that we have only been interested in the average quality in a population at a given time. This would seem to be a good place to start, but rather glosses over the issue of how much indigenous innovative activity is induced by different policies.

Of course, even from a utilitarian perspective we might be interested in the amount of innovative activity induced by a policy, simply because there is so much here that has been left unmodelled. It could be, for example, that the skills built up by indigenous innovation will change the way the population responds to future influxes of technology. That is, a more radical policy might be favoured in a more complete model because it increases the future

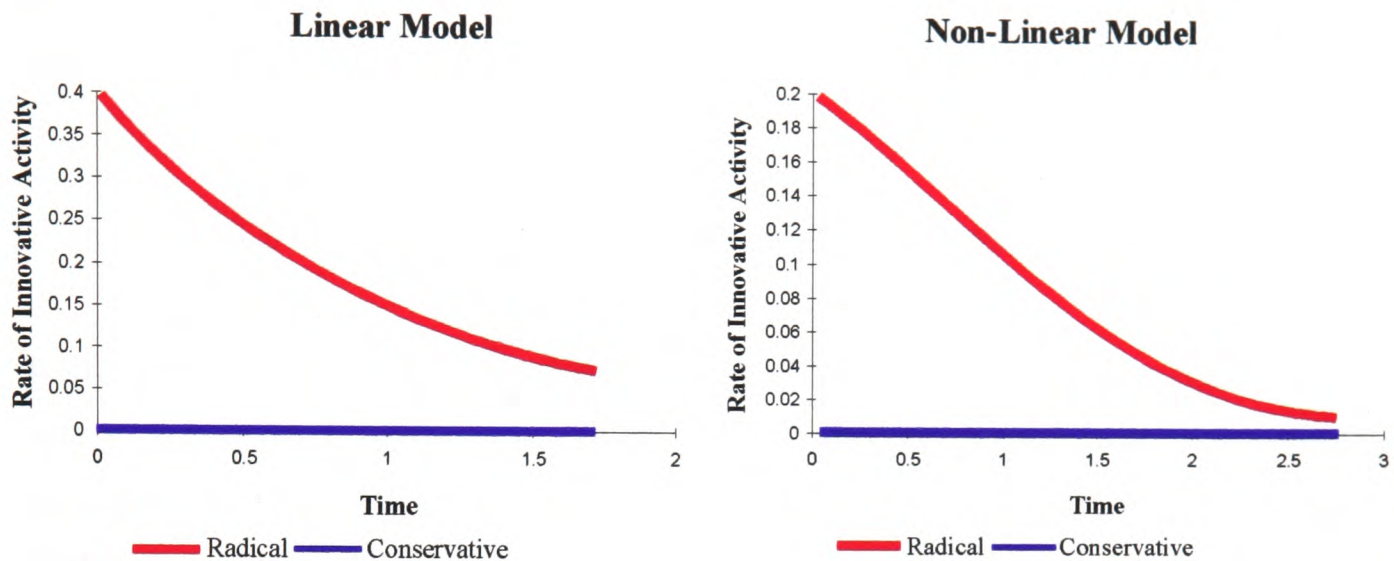


Figure 3.6. Rates of innovation in the two seed models with parameter values as given before. The rate of innovation is defined as the numbers of  $ab$  constructed by combining  $a$  and  $b$  per unit time, as a proportion of the total number of technology units in use. Initial value in the non-linear model were  $c = 0.5$  ('radical') and  $c = 0.99$  ('conservative').

values of  $\xi$ . Or, in the context of one of the 'new growth' models mentioned in the introduction, we might be interested in policies that specifically boost human capital.

But for those who, like the author, find a simplistic utilitarian approach unsatisfying—especially when, as here, it places undue emphasis on ends over means—will find the case for specifically encouraging indigenous innovation in its own right very strong indeed. We know from personal experience how much more satisfying constructive and creative work is compared to the merely routine. Work is not always a disutility—especially if it is creative.

If we are interested in specifically encouraging indigenous innovation, then we should note that the innovative activity induced by 'conservative' policies

is negligible compared to those that are more ‘radical’—as the numerical results in Figure 3.6 for the two-seed models clearly show. For these models, this result means that an organisation capable of introducing the combined technology  $ab$  from the start may not choose to do so.

### 3.6 ENCOURAGING HIGH RATES OF INNOVATION

In the results above, a radical policy can only be better than a conservative policy if  $\xi > 0$ ; moreover, in the linear model we found that increasing  $\xi$  increased the expected discounted average quality under the radical policy. High rates of innovation might also, of course, be desirable in their own right. But how can we encourage higher rates of innovation? The following are a few sketch ideas.

1. *Access.* Even if agents have a high tendency to innovate, they cannot do so unless they are exposed to the seed technologies needed to construct a new design. Policies that encourage communication—*e.g.* improved telecommunication or transport links—might be a help here. The organisation could arrange travelling technology demonstrations or, thinking of the way new innovations are encouraged in the academic community, they could arrange and fund technology conferences intended to share ideas.
2. *Training.* An obvious way to encourage innovation is of course training. The difference here would be that rather than training in basic skills, the training would concentrate on developing creative potential.
3. *Finance.* The costs of both straight imitation and innovation have not been modelled explicitly above. That is appropriate if the technology we have in mind is, for example, some agricultural *technique*—but for, say, light industry, there may be severe constraints facing updating agents. There may be a role here for so-called ‘micro-credit’ schemes to facilitate adoption in the face of capital market imperfections.

4. *Technology Protection.* If, as seems plausible, it is easier or less costly to engage in straight imitation than innovation based on partial imitation, then there may be an in-built bias in favour of the former in the diffusion models considered above. There may even be a free-rider problem in that agents capable of innovating may choose to wait until others have established sufficient amounts of new technology for them to easily imitate directly. There may therefore be a role for some product protection for the seed technologies in the early stages of the diffusion process to encourage innovation, once the superiority of the combined technology has been established.

### 3.7 A CALIBRATED DIFFUSION MODEL

While we may be able to make general conclusions from the sort of analytical models considered above, to find a precise optimal policy in a real-life setting one would need a calibrated diffusion model.

Some of the considerable difficulties in calibrating diffusion models are outlined in Stoneman and Karshenas (1995). These include the low availability of data, the non-linear nature of the process, the *ad hoc* nature of functional forms assumed—all of which make current calibrated diffusion models unable to distinguish between alternative theories of diffusion.

Perhaps the best route forward here is for the organisation involved in a particular technology transfer project to construct such a model. Not only could they collect the relevant data themselves, but they would also have a much better idea of the degree of communication and the tendency to innovate, which would help in choosing suitable functional forms. This would make an interesting area for further research.

### 3.8 CONCLUSION

Technological diffusion and innovation are connected. In the models considered here that connect the two, ‘radical’ transfer policies were those with high levels of technodiversity, while ‘conservative’ policies were those with low levels of technodiversity. If the potential for innovation is high, and if the organisation transferring technology does not discount the future at too high a rate, then radical policies may be superior. This is because, while radical policies depress the initial quality enjoyed by users, they result in faster convergence to new technologies—and perhaps, in the case of multiple seed technologies, to higher long-run qualities. If additional weight is placed on policies that encourage indigenous innovation, then radical policies are even more desirable.

### APPENDIX

PROOF OF PROPOSITION 1. The discounted total population quality along a diffusion path with initial condition  $c \in (0, 1)$ , where  $q_{ab} > q_a$ , is given by:

$$S_{q_{ab} > q_a} = \int_0^\infty e^{-\rho t} (xq_a + yq_b + (1 - x - y)q_{ab}) dt$$

where  $\begin{pmatrix} x \\ y \end{pmatrix}_{q_{ab} > q_a}$  is given by equation (3.4). This integral is easily calculated to give:

$$S_{q_{ab} > q_a} = \frac{q_a}{\rho} + \frac{(q_{ab} - q_a) \frac{\rho(1 + \rho)(1 - c) + (1 + \rho - \xi)2^{-\frac{\rho}{\xi}} - \rho\xi(1 - c)2^{-\frac{1 + \rho - \xi}{\xi}}}{\rho(1 + \rho)(1 + \rho - \xi)} - (q_{ab} - q_b) \frac{1 - c}{1 + \rho}}{\rho(1 + \rho)(1 + \rho - \xi)}$$

The discounted total population quality along a diffusion path where  $q_{ab} \leq q_a$  is given by:

$$S_{q_{ab} \leq q_a} = \int_0^\infty e^{-\rho t} x q_a + (1-x) q_b dt$$

where  $x$  is given by equation (3.5). This gives:

$$S_{q_{ab} \leq q_a} = \frac{q_a}{\rho} - \frac{(q_a - q_b)(1-c)}{1+\rho}$$

If  $q_{ab} \sim U(\underline{q}, \bar{q})$ , where  $E(q_{ab}) = \frac{1}{2}(q_a + q_b)$ , then

$$E(S) = \frac{2\bar{q} + q_a - q_b}{4\bar{q} - q_a - q_b} S_{q_{ab} \leq q_a} + \int_{q_a}^{\bar{q}} \frac{2}{4\bar{q} - q_a - q_b} S_{q_{ab} > q_a} dq_{ab}$$

which gives,

$$\begin{aligned} E(S) = & \frac{1}{(4\bar{q} - q_a - q_b)} \left( (2\bar{q} + q_a - q_b) \left( \frac{q_a}{\rho} + \frac{(q_a - q_b)(1-c)}{1+\rho} \right) + \right. \\ & (\bar{q} - q_a) \left( \frac{2q_a}{\rho} + \frac{(\bar{q} + q_a - 2q_b)(1-c)}{1+\rho} \right) + \\ & \left. (\bar{q} - q_a)^2 \left( \frac{\rho(1+\rho)(1-c) + (1+\rho-\xi)2^{-\frac{\rho}{\xi}} - \rho\xi(1-c)2^{-\frac{1+\rho-\xi}{\xi}}}{\rho(1+\rho)(1+\rho-\xi)} \right) \right) \end{aligned} \quad (3.7)$$

We may ignore  $c = 0$  and  $c = 1$  as optimal policies, as these correspond to fixed points that yield  $\frac{q_b}{r}$  and  $\frac{q_a}{r}$  respectively, both of which are lower than the expected discounted quality given by equation (3.7). As (3.7) is linear in  $c$ , the optimal policy  $c^* \in (0, 1)$  depends on the sign of  $\frac{\partial E(S)}{\partial c}$ .

- Case (a) If  $\frac{\partial E(S)}{\partial c} < 0$ , then  $c^* = \varepsilon$ , where  $\varepsilon$  is arbitrarily small.
- Case (b) If  $\frac{\partial E(S)}{\partial c} = 0$ , then  $c^*$  may take any value from  $(0, 1)$ .
- Case (c) If  $\frac{\partial E(S)}{\partial c} > 0$ , then  $c^* = 1 - \varepsilon$ .

To find the sign of  $\frac{\partial E(S)}{\partial c}$ , note that

$$\frac{\partial E(S)}{\partial c} = -\left(\frac{1 + \rho - \xi 2^{-\frac{1+\rho-\xi}{\xi}}}{1 + \rho - \xi}\right)(\bar{q} - q_a)^2 + (\bar{q} - q_a - 2q_b)(\bar{q} - q_a) + (2\bar{q} + q_a - q_b)(q_a - q_b)$$

That is,  $\frac{\partial E(S)}{\partial c} < 0$  iff

$$\frac{1 + \rho - \xi 2^{-\frac{1+\rho-\xi}{\xi}}}{1 + \rho - \xi} > \frac{(\bar{q} - q_b)^2 + 2\bar{q}(q_a - q_b)}{(\bar{q} - q_a)^2} \quad (3.8)$$

Let  $f(\rho, \xi)$  be defined by the left-hand side of equation (3.8), and  $g(\bar{q}, q_a, q_b)$  be defined by the right-hand side of equation (3.8). Note that since  $\rho \geq 0$ ,  $f_1 < 0$  and since  $\bar{q} > q_a$ ,  $g_1 < 0$ .

Let  $\hat{q}$  be defined by  $f(0, \xi) = g(\hat{q}, q_a, q_b)$  and  $\hat{q} > q_a$ . (As  $g_1 < 0$ , we know this will yield a unique value.)

First consider  $\bar{q} > \hat{q}$ . Let  $\hat{\rho}$  be defined by  $f(\hat{\rho}) = g(\bar{q}, q_a, q_b)$ . (As  $f_1 < 0$ , we know this will yield a unique value.)  $g_1 < 0$ , so  $g(\hat{q}, q_a, q_b) > g(\bar{q}, q_a, q_b)$ . Equivalently, from the definitions of  $\hat{q}$  and  $\hat{\rho}$ , we may write  $f(0, \xi) > f(\hat{\rho}, \xi)$ , which implies that  $\hat{\rho} > 0$ . Thus for  $\bar{q} > \hat{q}$ , we have  $\hat{\rho} > 0$  and  $f(\rho, \xi) > g(\bar{q}, q_a, q_b)$  for  $\rho < \hat{\rho}$  (case (a)),  $f(\rho, \xi) \leq g(\bar{q}, q_a, q_b)$  for  $\rho \geq \hat{\rho}$  (case (b) or (c)).

Secondly, consider  $\bar{q} \leq \hat{q}$ . We may write  $g(\bar{q}, q_a, q_b) \geq g(\hat{q}, q_a, q_b) = f(0, \xi) \geq f(\rho, \xi)$  (case (b) or (c)).

Finally note that:

$$\left. \frac{\partial E(S)}{\partial \xi} \right|_{\substack{c \rightarrow 0 \\ 0 < \xi < \frac{1}{2}}} > 0$$

and

$$\left. \frac{\partial f(\rho, \xi)}{\partial \xi} \right|_{0 < \xi < \frac{1}{2}} > 0 \Rightarrow \frac{\partial \hat{q}}{\partial \xi} < 0, \quad \frac{\partial \hat{\rho}}{\partial \xi} > 0$$

*Q.E.D.*

*Part II*  
*Evolution and Adaptive Economic*  
*Behaviour*

*Chapter 4*  
*Evolutionary Game Theory:*  
*A Constructive Critique*

SUMMARY

The study of evolutionary game theory in economics helps to identify economic contexts where it is a more appropriate modelling method than methods that assume rational choice—and where it is not. However, in contexts where it is an appropriate modelling method, the results obtained with evolutionary game theory are highly sensitive to how one models the learning mechanism driving the results, and how this integrates with the context in question. It is argued that the most common learning mechanisms and equilibrium concepts used by economists applying evolutionary game theory are too highly derivative of biological mechanisms that assume a biological context. Progress in the application of evolutionary game theory to economics therefore requires flexible and creative modelling of context-appropriate mechanisms by theorists, supported by careful empirical and experimental work on the adaptive behaviour of real subjects.

## 4.1 INTRODUCTION

Game theory is the study of well-defined interactive decision problems. *Evolutionary* game theory—as it applies in economic contexts—approaches these interactive decision problems by modelling explicitly the *process* by which the decision-makers (or ‘players’) reach their decisions—be it by individual learning or by some social mechanism.

This is in marked contrast to conventional game theory which, in common with much of microeconomic theory, is not at all interested in this decision-making process *per se*. Instead, the conventional approach assumes that whatever process is involved leads to the same outcome as a situation where each decision-maker makes a choice by solving some well-defined optimization problem. That is, the conventional approach assumes that the players of a game are substantively rational<sup>1</sup>.

The temptation, having contrasted the evolutionary approach to game theory with the rational choice approach, is to set up a false dichotomy that asserts only one is valid. This is extremely unlikely to be true. If the purpose of our inquiries in positive economic theory is to provide a good understanding of real economic phenomena by constructing models which have—in some broad sense—good predictive power, then it seems sensible to use whatever approach works best in the particular circumstance we are interested in.

The purpose of this introduction is to argue that there do exist circumstances in which the evolutionary approach is either more appropriate than the conventional ‘rational choice’ approach or may usefully complement it. The purpose of the remainder of the chapter is, firstly, to argue that in

---

<sup>1</sup>As has been well pointed out by, for example, Binmore (1987), the use of the word ‘rational’ in economics is misleading and may well have stifled research into alternative approaches. We must always remember in economics that when we say ‘This choice is rational’, we do not mean ‘This is the *best* or most reasonable choice’, rather, ‘This is the choice derived from the optimization of a function with certain properties which are admittedly *ad hoc* but which we believe to be plausible’.

circumstances where evolutionary game theory is an appropriate technique, the methods commonly used to implement it too heavily rely on methods designed for a biological context. Secondly, the aim is to outline some examples of models in which a more sensitive approach is taken to economic and social context. These models consider the issues of matching and local interaction, imperfect imitation, and dynamics that are not payoff monotonic.

But first some reasons why the study of evolutionary game theory is worthwhile:

1. “Research on learning and social evolution can help to understand the empirical scope of the rationality hypothesis” (Börgers 1996).

That is, in applications where the predictions of a rational choice model are sometimes, but not always, contradicted by empirical or experimental evidence, some understanding of why this may be so can be provided by an evolutionary model. The examples of section 4.3 below can be interpreted this way. The question remains: why should such an understanding be worthwhile? Well, we have a considerable body of theoretical material based upon the rational choice hypothesis and it would be helpful to have a better understanding of which parts of this body would benefit from reappraisal. Moreover, economic theorists in the future making modelling decisions on which approach to take would benefit from a greater understanding of this issue. It would be helpful to know in advance whether a rational choice model is likely to be entirely inappropriate in a given context. On the other hand, given the complex and unwieldy nature of a typical evolutionary model, it would be helpful to know whether a more concise model based on rational choice may serve the same purpose more elegantly—and perhaps with greater generality.

2. There exist contexts where the predictions of an evolutionary model may be more robust than the predictions of its corresponding rational choice model.

It is not true, of course, that an evolutionary model is superior to a rational choice model *just* because it may more closely mimic the decision-making process being examined. It has been argued at least as far back as Friedman (1953) that there is a real sense in which the ‘realism’ of a model’s structure is irrelevant so long as it predicts well—and it is easy to agree if the phenomenon one is modelling is in a fairly stable environment. Neither is it true to say that an evolutionary model is superior *just* because the calculations necessary in its rational choice counterpart are very complex. Friedman cites the example of leaves on a tree tracking the path of the sun as it moves across the sky. This may be validly modelled as a rational choice model in which the leaves choose a path that maximizes their total exposure to sunlight. Such a model would have good predictive power despite the leaves having no computational ability at all.

However, it remains good modelling practice to attempt to model explicitly those aspects of a process that may be subject to, or influenced by, a change in the environment being modelled. To take Friedman’s example, suppose the tree were placed in an environment where the only light source followed a random walk. A rational choice model would predict the leaves bobbing about like a go-go dancer—but one would be very surprised to see this in practice. A model more closely mimicking the biological processes involved would predict better.

Similarly when modelling interactive decision-making. If one is interested in how behaviour in an application would react to some important structural change (*i.e.* ‘comparative statics’, interpreted broadly), then it may be important to model the decision-making process as closely as possible. Now this does not of course *necessarily* mean using an evolutionary model. In unfamiliar situations where the decision-makers have plenty of time and where the payoff rewards are high, a rational choice model may well be more ‘realistic’—in that it more closely mimics the ‘eductive’ (Binmore 1987) process being used. However, in situations

that are relatively familiar and anonymous and where payoff rewards are minor, an evolutionary model of decision-making would be more plausible<sup>2</sup>.

3. There may be situations in which an evolutionary model is less *ad hoc* than its corresponding rational choice counterpart.

It is not possible to model economic situations involving human behaviour without making some choices that can be criticised as arbitrary or *ad hoc*. This is as much true in rational choice models—which, for example, require assumptions on the form of utility functions or impose strong epistemic requirements on players—as it is in evolutionary models that require assumptions about the mechanisms of learning. To the extent that they actually matter (see above), a careful modeller will try to choose assumptions plausibly supported by empirical evidence, from whatever source. It follows that if, in a particular application, there is more information about learning mechanisms than utility functions, then the evolutionary model may rest on the more solid foundation.

4. When both rational choice *and* an evolutionary approach are appropriate modelling methods, but rational choice models give inconclusive predictions, then evolutionary game theory may provide a structured way of thinking about equilibrium selection.

This is a line of research begun by Binmore (1987) and continued by many authors, including Nöldeke and Samuelson (1993), Binmore and Samuelson (1995) and Nöldeke and Samuelson (1995). It has its roots in Selten's notion of 'trembling hand perfection' (Selten 1975). However, rather than interpreting 'trembles' as merely a technical device, this

---

<sup>2</sup>Of course, we are making informal appeal to some trade-off between payoffs and computational or complexity costs here. However, when we try to think about this trade-off explicitly, we rapidly fall into a paradoxical discussion along the lines of *e.g.* Selten (1991). Moreover, 'ulterior' models, which try to model choices between choice-making procedures, are prone to the difficulty of infinite regress discussed further below.

literature interprets them literally, and attempts to model their origin by explicitly modelling the process by which equilibria are reached. It may follow that ‘trembles’ are not independent or uncorrelated, as the original concept assumed, and this has implications for which equilibria are selected. Compared to the plethora of literature that attempts to ‘refine’ equilibria by some other means, the evolutionary approach has the advantage of transparency. An evolutionary model of selection is a kind of ‘story generating device’—and while one may not wholly agree with the stories it generates, at least they are spelt out in full<sup>3</sup>.

The chapter proceeds as follows. Section 4.2 outlines the methods of evolutionary game theory and considers some of the most important objections that can be made of them. Section 4.3 then suggests some ways in which these objections may be addressed in models that take a more sensitive approach to an economic context. Section 4.4 concludes.

## 4.2 THE METHODS OF EVOLUTIONARY GAME THEORY

A broad but useful taxonomy of the chief learning mechanisms of interest to economists has been suggested by Selten (1991):

1. rote learning,
2. imitation,
3. belief learning.

---

<sup>3</sup>A similar motivation lies behind the extensive literature that attempts to select between strict equilibria by considering the equilibrium state of perturbed dynamical systems in the ultralong-run. However, although these models are conducted against the back-drop of a dynamical system typically based on some ‘evolutionary’ process, the modelling of perturbations remains rather primitive, as Bergin and Lipman (1996) point out. Explicit models of perturbations (or ‘mutations’) that are plausibly context-sensitive remain few and far between. Even then—unless, like Ellison (1995) or Ely (1995), one adopts a specific local interaction structure—the ultralong-run may be *too* long to be of any interest.

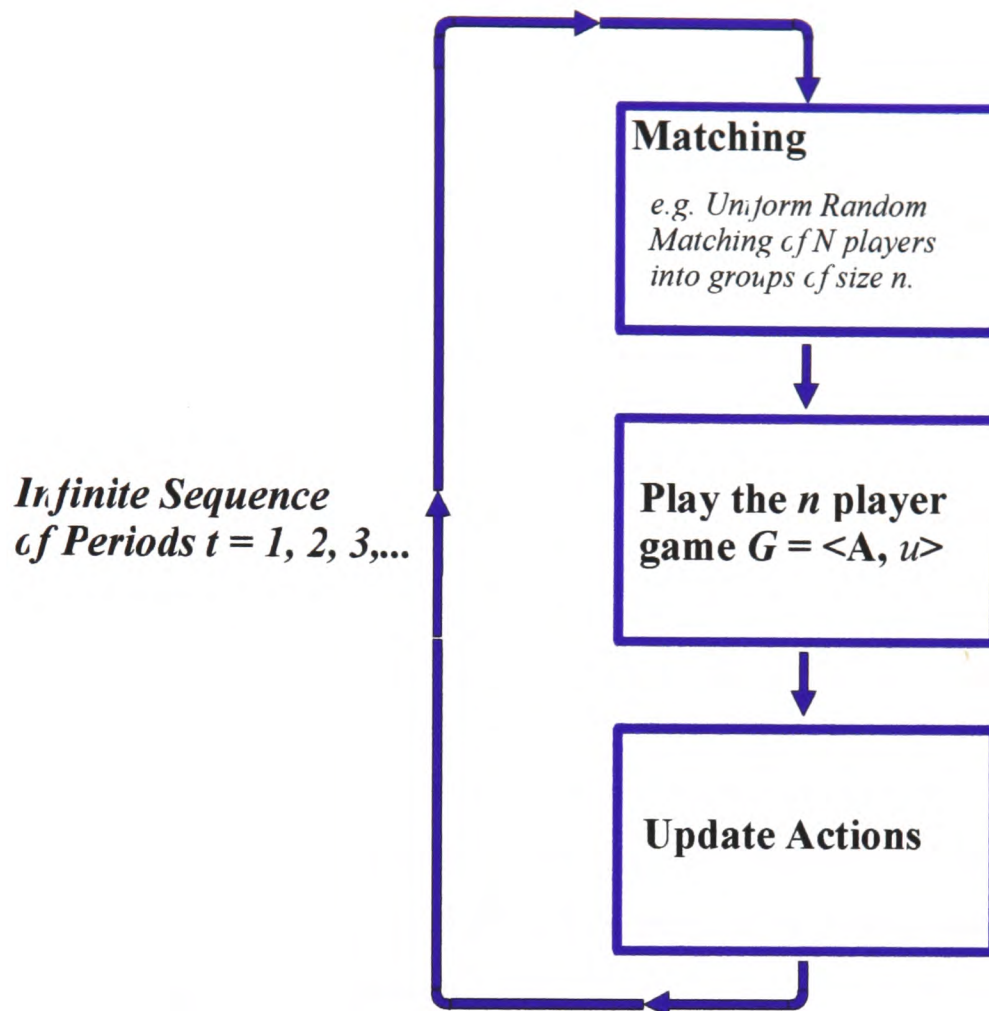


Figure 4.1. 'Evolutionary' dynamics.

However, the great bulk of material that goes under the heading 'evolutionary game theory' has concerned itself with models most closely connected to the second of these. That is, a typical model in evolutionary game theory considers the dynamics of *social* learning in a large population of players who update actions in a game by some process of contagion or imitation.

At least implicit in a model of this kind is a dynamic process that may be expressed schematically (in discrete time) as in Figure 4.1. The players in a large population are *matched* in some fashion to *play* some game  $G$  (which never changes), they receive payoffs according to their current action choice and, finally, at least some of them then *update* their current choice given observations of chosen actions and payoffs. The process repeats itself *ad infinitum* and the modeller is interested in the how the population proportions playing each action of the game change over time—in particular,

their long-run values.

Suppose we restrict our attention to a two-player symmetric game  $G$  with finite action space  $\mathbf{A} = (A_1, \dots, A_m)$  and payoff function  $u$ . (The methods described below extend readily to more general settings.) Let the mixed-strategy set associated with  $\mathbf{A}$  be given by  $\Delta = \{\mathbf{x} \in \mathbf{R}_+^m \mid \sum_{i \in \mathbf{A}} x_i = 1\}$ . The function  $u$  maps  $\Delta^2 \rightarrow \mathbf{R}$ ; so the payoff to a (mixed) strategy  $\mathbf{x} \in \Delta$  against a (mixed) strategy  $\mathbf{y} \in \Delta$  is written  $u(\mathbf{x}, \mathbf{y})$ .

There are basically two approaches to thinking about the long-run equilibrium of such a dynamical system<sup>4</sup>.

The first is to run the following thought experiment. Suppose we have a monomorphous population all playing the strategy  $\mathbf{x} \in \Delta$ . Now consider an *invasion* by a mutant strategy  $\mathbf{y} \in \Delta$ , such that in the post-invasion population a proportion  $\varepsilon$  are playing  $\mathbf{y}$  and a proportion  $1 - \varepsilon$  are playing  $\mathbf{x}$ . If, in this post-invasion population, the expected payoff to the incumbent  $\mathbf{x}$  is greater than the expected payoff to the mutant  $\mathbf{y}$ , then the incumbent is said to *repel* the mutant invasion.

If we assume the *uniform random matching* of players in a very large population, then this condition for repulsion amounts to:

$$u(\mathbf{x}, \varepsilon\mathbf{y} + (1 - \varepsilon)\mathbf{x}) > u(\mathbf{y}, \varepsilon\mathbf{y} + (1 - \varepsilon)\mathbf{x}) \quad (4.1)$$

If  $\mathbf{x}$  can repel all ‘small’ invasions  $\mathbf{y} \neq \mathbf{x}$ , then  $\mathbf{x}$  is called an *Evolutionarily Stable Strategy* (ESS) (Maynard Smith and Price 1973). Formally,

**DEFINITION 4.1**  $\mathbf{x} \in \Delta$  is an ESS if for every strategy  $\mathbf{y} \neq \mathbf{x}$  there exists some  $\bar{\varepsilon}_y \in (0, 1)$  such that (4.1) holds for all  $\varepsilon \in (0, \bar{\varepsilon}_y)$ .

It is straightforward to show (*e.g.* Weibull (1995, p. 37)) that this definition is equivalent to the definition of a mixed-strategy Nash equilibrium *plus* an additional stability requirement. The stability requirement states that if a mutant  $\mathbf{y}$  is an alternative best-reply to  $\mathbf{x}$ , then  $\mathbf{x}$  must be a better reply to  $\mathbf{y}$  than  $\mathbf{y}$  is to itself.

---

<sup>4</sup>The following draws heavily from Weibull (1995).

The second approach to equilibrium is to interpret  $\mathbf{x}$  as a vector of *population proportions*, where these population proportions are subject to some deterministic growth-rate function  $g$ :

$$\frac{\dot{x}_i}{x_i} = g_i(\mathbf{x}), \quad i = 1, \dots, m \quad (4.2)$$

Let  $u(A_i, \mathbf{x})$  denote the payoff to action  $A_i$  under uniform random matching in a large population in state  $\mathbf{x}$ . Attention is generally restricted to growth-rate functions in which relatively high *current* payoffs are associated with relatively high growth rates. In particular, a *payoff monotonic* growth-rate function is one where

$$g_i(\mathbf{x}) = f_i(u(A_i, \mathbf{x}) - u(\mathbf{x}, \mathbf{x})), \quad i = 1, \dots, m,$$

where  $f$  is monotonically increasing. The special case where  $f$  is *linearly* increasing is known as the *replicator dynamics*.

Given such a formally defined dynamical system, one can then find the fixed points of the system and test for stability using the conventional notions of *Lyapunov Stability* and *Asymptotic Stability*.

It is interesting to note how these concepts of long-run equilibrium relate to one another. Let  $\Delta^{ESS}$  be the subset of  $\Delta$  corresponding to the definition of an ESS;  $\Delta^{AS}$  be the subset corresponding to points that are asymptotically stable under payoff monotonic dynamics;  $\Delta^{LS}$  be the set of points that are similarly Lyapunov stable; and  $\Delta^{NE}$  be the set of mixed-strategy Nash equilibria. Then one can show (see, for example, Weibull (1995)):

$$\Delta^{ESS} \subset \Delta^{AS} \subset \Delta^{LS} \subset \Delta^{NE}$$

Now the above is far from an exhaustive account of the methods of evolutionary game theory, but it does cover most of the main areas, and it allows us to make a number of criticisms.

Consider first some criticisms of the methods outlined above from a technical point of view:

1. Taking the first approach to equilibrium, the stability condition built into the concept of an ESS seems excessively strong.

It seems very odd that a population state that is asymptotically stable under payoff monotonic dynamics need not be an ESS. Some progress has been made in plausibly moderating the stability condition in ESS, including the concepts of Neutrally Stable Strategies (in which the strict inequality of equation (4.1) is replaced by a weak inequality) and ‘robustness against equilibrium entrants’ (Swinkels 1992).

2. Taking the second approach to equilibrium, there appears to be some sleight of hand in the way stability is modelled.

Implicit in any statement of deterministic dynamics like that of equation (4.2) is some social learning mechanism from which all stochastic elements have been ‘smoothed out’. This smoothing may result from the assumption of a very large population size. Alternatively, it may result from a discrete-time model of social learning in which both the time between periods and the probability of updating are given by the same parameter  $\tau$ . Börgers and Sarin (1994) explain the conditions under which one may use a theorem of Norman (1972) to obtain a deterministic continuous-time dynamic by taking  $\tau \rightarrow 0$ . Either way, by using conventional stability concepts on the resulting dynamical system, one is in some sense ‘adding back’ small amounts of noise of an arbitrary nature. That is, one is using criteria to distinguish between fixed points of the dynamic that have not been explicitly modelled. If the purpose of the evolutionary model is to address the issue of equilibrium selection, then this may be a serious objection.

In response to this issue (though no doubt for other reasons too), a large number of more recent models have avoided this sleight of hand by, for example, modelling learning mechanisms as Markov processes. To an extent, this trend is to be commended as a laudable attempt to obtain ever greater rigour. However, the models thus generated

are often enormously cumbersome and technically demanding to both write and read. The danger is that an emphasis on rigour has squeezed out economic insight. In some respects it would be refreshing to see a return to more simple ‘low fat’ modelling methods (of which, one might add, the ESS concept is an excellent example).

However, these criticisms are relatively minor compared to the objection that the methods of evolutionary game theory are too highly reliant on the biological context from which they originate. This excessive dependency is revealed in a number of key areas:

1. In an economic context, the assumption of uniform random matching typically makes very little sense.

Actually, random matching does not make a great deal of sense even in a biological context; it is assumed largely for the sake of technical convenience. In an economic context, the assumption does help to justify the assumption of anonymity, making very simple strategic behaviour seem more plausible—but it is hard to think of many applications where something like random matching might actually take place. On the one hand, it seems that the emphasis on *social* learning in the literature may have been misplaced and there has been too little attention to simple learning processes at an individual level (the first of Selten’s categories above)—Börgers and Sarin (1994) and Börgers and Sarin (1997) being notable exceptions. On the other, modellers who are interested in social interaction need to start looking at more realistic matching structures—and there are indeed a number of these ‘local interaction’ models now appearing. Section 4.3.1 below briefly describes the models of Cooper and Wallace (1996) and Cooper and Wallace (1997) in which the interaction structure is to some limited extent endogenised.

2. In an economic context, the assumption that the ‘reproduction’ (which is typically assumed to be either by contagion or imitation) of strategies

from generation to generation is (near) perfect is often invalid.

In particular, in extensive-form games there is no evolutionary pressure on behaviour at unreached information sets: information about behaviour at these sets is not passed on to those next updating their strategies. However, imperfect imitation need not be seen in an entirely negative light. Section 4.3.2 below presents an application of the phenomenon to a simple version of the 'ultimatum' game<sup>5</sup>.

3. In an economic context, the assumption of monotonic dynamics may make less sense than it does in a biological context.

In anything other than the extreme short-run in a biological evolutionary model, the players are dead. Payoff monotonicity, which takes note only of *current* relative performance, makes sense in this context. However, in an socio-economic context we expect our players to last a little longer relative to the process of learning. That is to say, we should not restrict our attention to purely myopic (payoff monotonic) dynamics on the one hand and fully-blown dynamic rationality on the other: there is undoubtedly a rich vein of possibilities in between. Section 4.3.3 below considers the plausibility of dynamics that are not payoff monotonic in the context of coordination games.

4. Unlike the study of biological reproduction, a purely theoretical study of cultural learning mechanisms may be subject to a problem of infinite regress.

The mechanism of transmission in biological evolution (Mendel's laws) are well known and understood. The mechanisms of transmission in cultural evolutionary systems, however, are not; moreover, it seems likely that these transmission mechanisms may vary according to context. If we try to explain why certain mechanisms are used by appeal

---

<sup>5</sup>An application to imperfect imitation in the repeated prisoners' dilemma can be found in the next chapter.

to some ulterior evolutionary process, then we run into the problem of how to explain why *that* process came to be used—and so on<sup>6</sup>.

The only obvious way to cut this Gordian knot is to find out as much as we can about the actual transmission mechanisms used by real subjects—*i.e.* by painstaking empirical and experimental work. However, this is clearly a very long-term research project, so in the meantime it seems sensible to keep our assumptions about transmission mechanisms as flexible and general as possible.

### 4.3 NOTES ON THE TRANSITION FROM A BIOLOGICAL CONTEXT TO AN ECONOMIC CONTEXT

The excessive dependence of popular evolutionary game theory methods on a biological context demands a flexible and creative response from economic theorists. The purpose of the following notes is to outline some contributions to this program of research, addressing the issues of matching, imperfect imitation and payoff monotonicity.

#### 4.3.1 *Beyond Random Matching: Who to Play With?*

It is common in evolutionary game theory to suppose that players are matched randomly each period to play a game. In joint work with Christopher Wallace (Cooper and Wallace 1996), we consider a model where agents have control not only over their actions in a game but also over *with whom* they interact and *from whom* they learn. In particular, agents are able to form

---

<sup>6</sup>Some might argue that the regress stops when we reach strict biological evolution. However, given that the various levels of evolution within such a system would not be stationary relative to each other, but rather co-evolve, the task of constructing such a ‘Theory of Everything Evolutionary’ seems virtually impossible. Moreover, it would probably be a fruitless task, in that such a theory would be unlikely to supply much in the way of falsifiable predictions.




Figure 4.2. Updating in a 'Social Circle'.

*fixed partnerships.* In a partnership situation, the agents' observation sets are the same and their learning is correlated. The partnership outcome picks the best symmetric outcome in a game. In the prisoners' dilemma, for example, we observe the co-evolution of partnerships and cooperation. However, in games where the partnership outcome is inferior to the outcome under random matching, partnerships will not evolve. The following is a brief outline of this model.

Consider  $2N$  players arranged on a *social circle* (Figure 4.2). The two players on a given radius play a  $2 \times 2$  symmetric normal form game. The players are one of four types:

- Type *a* (black circle) play strategy *a* in the game and never move from their current location on the social circle.
- Type *b* (white circle) play strategy *b* in the game and never move from their current location on the social circle.

(Two black circle types on a given radius constitute a *fixed partnership*.)

- Type  $a'$  (white circle) play strategy  $a$  in the game and move to a random spare location on the social circle each period.
- Type  $b'$  (white circle) play strategy  $b$  in the game and move to a random spare location on the social circle each period.

Time, indexed by  $t$  starts at  $t = 0$  and is discrete. Within each time period  $t$ , the sequence of play is divided into 3 stages:

1. The pairs of players around the social circle play the game and receive the payoffs specified by its payoff function.
2. All players then update their type. With probability  $\varepsilon$  they choose at random from the four possible types<sup>7</sup>. With probability  $1 - \varepsilon$  they observe their own pair's types and payoffs and the types and payoffs of pairs close to them. In the social circle the pairs they observe are all those within a distance  $M$  to the left of them and  $M$  to the right of them. Figure 4.2 illustrates this *observation window* for  $M = 1$ . In other words, they observe the types and payoffs of their neighbours. (A player's strategy is observable from stage 1; a player's matching type is observable from her matching behaviour in the previous period.) Updating then takes place. Each player calculates the average payoff received by each of the four types in the observation window, then switches to the type with the highest average payoff.
3. After the players have updated their types they are re-matched. Fixed types (black circles) stay where they are, mobile types (white circles) move to a randomly chosen spare location.

---

<sup>7</sup>This noise can be interpreted in one of several ways. It could simply be a 'trembling hand' error. It could also be the result of an extreme observational error—a player simply not paying any attention to payoffs or behaviour! Or it could be the result of a new player replacing an existing player after payoffs are allocated in stage 1, having no information on which to base a type choice.

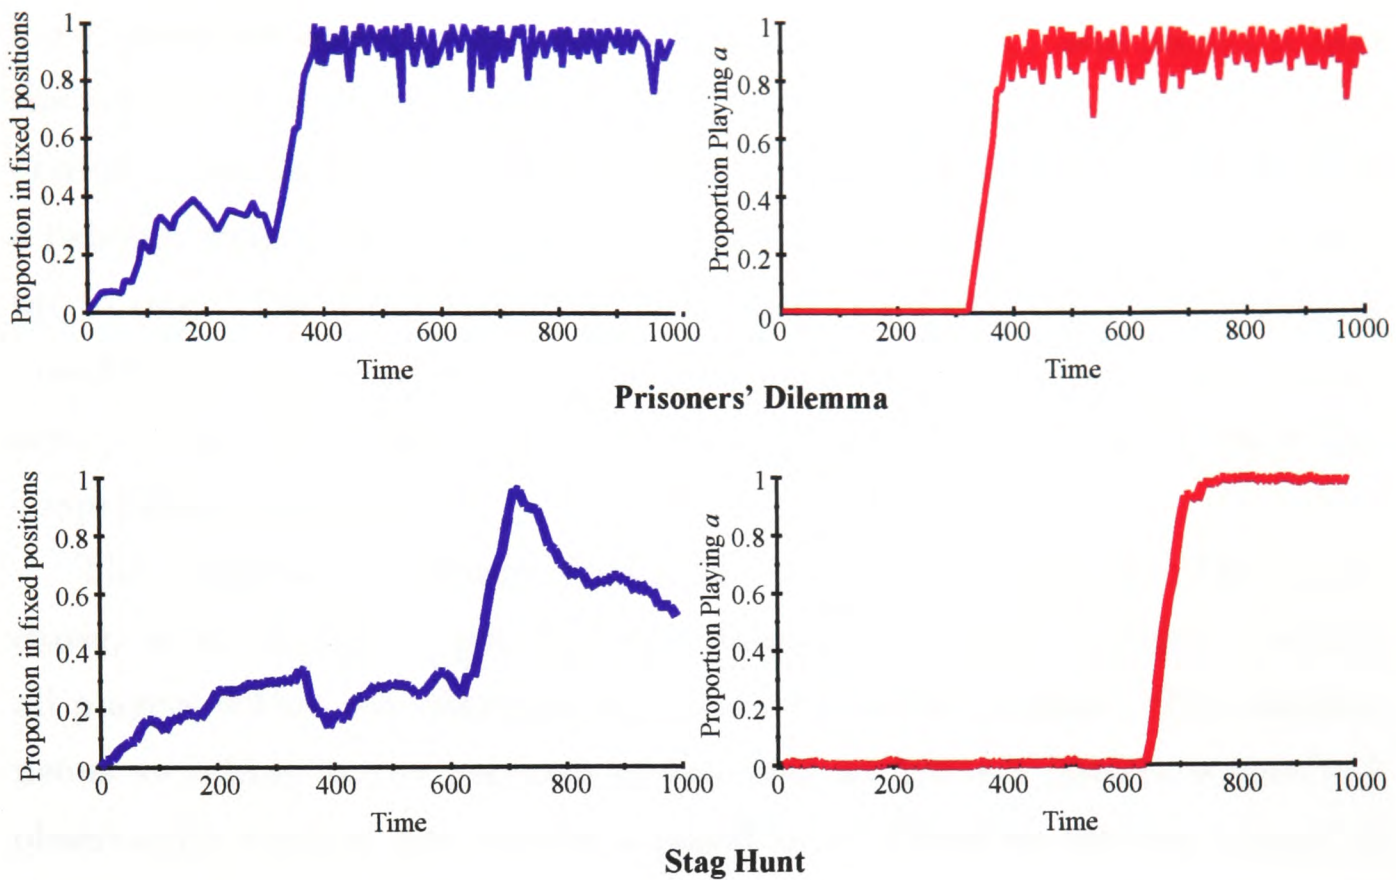


Figure 4.3. Simulation Results.

This dynamical system is quite complex and difficult to analyse using conventional methods. As a preliminary to more formal work, Cooper and Wallace (1996) includes simulation results for a number of example games. We shall consider just two of these games here:

Prisoners' Dilemma		
	<i>a</i>	<i>b</i>
<i>a</i>	(3, 3)	(1, 4)
<i>b</i>	(4, 1)	(2, 2)

Stag Hunt		
	<i>a</i>	<i>b</i>
<i>a</i>	(5, 5)	(0, 4)
<i>b</i>	(4, 0)	(3, 3)

The first example is the prisoners' dilemma. The Nash equilibrium is  $(b, b)$  and is inefficient. If only players could somehow coordinate on  $(a, a)$  they would both be better off. It seems that partnerships are exactly such a mechanism for coordination.

Figure 4.3 shows the evolution over time of the proportions playing in fixed positions and the proportion playing strategy  $a$ . For clarity's sake the simulation is started with everyone randomly matching and playing the dominant but inefficient strategy,  $b$ . The proportion in fixed positions follows a random walk until reaching some threshold value at which point it shoots up to one. The behaviour spreads around the social circle as quickly as possible. At the same time the proportion playing strategy  $a$  moves from zero to one. Both effects clearly reinforce one another. Partnerships and cooperation co-evolve.

The cooperative outcome does seem rather noisy, however. This is because, in an all-type  $a$  population, a single deviation to a 'defect' induces all players within one observation window to switch to defect. The simultaneous switching across partners means that all  $2M + 1$  players within this observation window now receive a payoff of 2. Those on the two 'wings' of the window (*i.e.* all those excluding the middle partnership) observe the superior cooperative payoff of 3, and switch back to  $a$  next period. Finally, the middle partnership switches back to  $a$ .

What we have here is an explanation of cooperation in the prisoners' dilemma that does not depend on the complexities of a repeated game framework and Folk theorem arguments. Players voluntarily submit to an institution in which the incentive to cheat is deliberately reduced by the correlated learning across partners<sup>8</sup>.

In the second example, the Stag Hunt game has two pure Nash Equilibria ( $(a, a)$  and  $(b, b)$ ) and one mixed, when both players play strategy  $a$  with probability  $\frac{3}{4}$  (with the payoffs used here). The second pure strategy Nash

---

<sup>8</sup>We might wonder how a single fully rational patient player would behave in an all type  $a$  population. From the above discussion, it is fairly easy to see that she would be facing a strategy of 'two-tits-for-a-tat'. That is, the maximum limit-of-the-means payoff she could gain from cheating would be  $2\frac{2}{3} < 3$ , and there would therefore be no incentive to deviate from cooperation. This result holds so long as the sum of the temptation to cheat (4 in this case) and twice the punishment payoff is less than three times the cooperative payoff.

Equilibrium is the risk-dominant<sup>9</sup> and the first the payoff-dominant. Figure 4.3 illustrates the outcome of the simulations for this particular game. The simulation is started with all players randomly matching and playing the risk-dominant equilibrium. As time passes the proportion of players in fixed positions follows a random walk. When enough fixed partnerships exist, the payoff dominant outcome in such partnerships becomes attractive to all that observe it and this behaviour is imitated. This imitation spreads the behaviour rapidly around the social circle until all players are playing the payoff dominant equilibrium. At this point, there is no longer any need for fixed partnerships to enforce such an equilibrium since, random matching or not, no-one can gain from deviation. Hence the proportion of the population in fixed positions follows a random walk from that point on. The partnerships are only required to flip the population over from the risk dominant to the payoff dominant equilibrium. The payoff dominant equilibrium is then stable in the long run with any proportion of fixed position players.

Why does this happen? In a world of purely random matching (with window size  $M = 1$ , as here) both the risk-dominant and payoff-dominant equilibria are stable fixed points. In a world of purely fixed partnerships, only the payoff-dominant equilibrium is a stable fixed point. Why? Basically, as soon as random mutations result in at least one partnership playing the payoff-dominant actions, then this partnership is unlikely to observe an alternative action that persuades them to change. Moreover, every player within one observation window of the partnership also copy the payoff-dominant action, and the behaviour quickly spreads around the social circle. So partnerships facilitate the transition to the payoff-dominant equilibrium; however, they are not necessary to sustain it.

---

<sup>9</sup>In a symmetric  $2 \times 2$  game, if both players strictly prefer the same action when there prediction is that their opponent will randomize equally between  $a$  and  $b$ , then the strategy profile where both players play that action is the risk-dominant equilibrium. In such games under monotonic dynamics, population states corresponding to risk-dominant equilibria have the larger basin of attraction.

The aim of Cooper and Wallace (1997) is to construct a formal stability concept to capture some of these ideas. To distinguish this stability concept from those of standard evolutionary game theory it is called *recoverability*.

Let each player in a population be described by a strategy-pair type  $s_i = (x_i, t_i)$ , where  $x_i$  corresponds to a mixed strategy in the symmetric two-player game  $G = \langle \mathbf{A}, u \rangle$  and  $t_i \in \{f, r\}$ , 'f' meaning a fixed partnership type and 'r' meaning a randomly matching type. Each period the players are matched into pairs to play  $G$ .  $f$ -types simply play their partners;  $r$ -types are matched by a uniform matching process to other  $r$ -types.

In a finite population ESS (Vega-Redondo (1996), Schaffer (1988)), one takes one member of a single-strategy incumbent population and replaces it with a mutant, then compares the expected payoff to an incumbent with the expected payoff to the mutant. 'Recoverability' follows from a similar thought experiment. Take a population consisting of players all of strategy-type  $s = (x, t)$ . Consider an 'invasion' by a mutant strategy-type  $s_m = (x_m, t_m)$  such that all pairs observing  $s_m$  copy it, giving an 'invasion zone' consisting of  $W$  pairs of players. Now compare the expected payoff to a player using  $s$  to the expected payoff of a player using  $s_m$ .

Let the probability of meeting a player playing  $y$  given you are playing  $x$  in a pair-type  $t$  be written  $p(y | x, t)$ . Now the expected payoff to a particular player of pair-type  $t$  playing  $x$  in a population using strategies from the strategy set  $Y$  is given by:

$$\sum_{y \in \Delta} u(x, y) p(y | x, t)$$

We are now in a position to define what we mean by 'recoverability'.

**DEFINITION 4.2** *A strategy-type  $s = (x, t)$  is an Evolutionarily Recoverable Strategy-type (ERS) if  $\forall s_m \neq s$ , there exists a  $\bar{W} > 0$  such that for all  $W \leq \bar{W}$ , the following inequality holds:*

$$\sum_{y \in \Delta} u(x, y) p(y | x, t) > \sum_{y \in \Delta} u(x_m, y) p(y | x_m, t_m)$$

This definition gives results in broad accordance with the simulation results.

### 4.3.2 Imperfect Imitation in a 'Mini' Ultimatum Game

The aim of this section is to show how the issue of *imperfect imitation* may help to explain what is apparently 'irrational behaviour' in experimental studies of the ultimatum game.

There may come a time when the *ultimatum game* comes to rival the prisoners' dilemma as a source of discussion and a generator of literature, both theoretical and experimental. It is a two-stage game in which two players divide a unit of money. In the first stage, player I makes an offer  $\lambda \in [0, 1]$ . In the second stage, player II either *Accepts* or *Rejects*. If player II accepts, then payoffs are  $1 - \lambda$  to player I and  $\lambda$  to player II. If player II rejects, then payoffs are zero to both players. Although it is slightly more complex than the prisoners' dilemma, the ultimatum game is relatively uninteresting from the point of view of conventional game theory. There are a large number of Nash equilibria, but only one subgame perfect equilibrium, where  $\lambda = 0$ .

However, the ultimatum game has attracted a great deal of attention because there is now a substantial body of experimental evidence (*e.g.* Thaler (1988), Roth (1995), Bolton and Zwick (1995)) that players tend not to select the subgame perfect equilibrium. Indeed, even in experiments with adequate incentives and sufficient opportunity for learning, the modal offer tends to be  $\lambda = \frac{1}{2}$ —an understandable outcome given that low offers ( $\lambda < \frac{1}{4}$ ) are rejected over half the time. The question is: why does player II leave money 'on the table' when player I makes one of these low offers?

In attempting to construct an answer to this question from the perspective of evolutionary game theory, it seems sensible to follow Binmore, Gale, and Samuelson (1993) and restrict our attention to a 'mini' ultimatum game. The idea is to capture the main strategic aspects of the ultimatum game in a game that may be analysed from an evolutionary point of view within a two-dimensional state-space. Figure 4.4 illustrates the rules of the game and shows its phase diagram under standard replicator dynamics, where  $x$

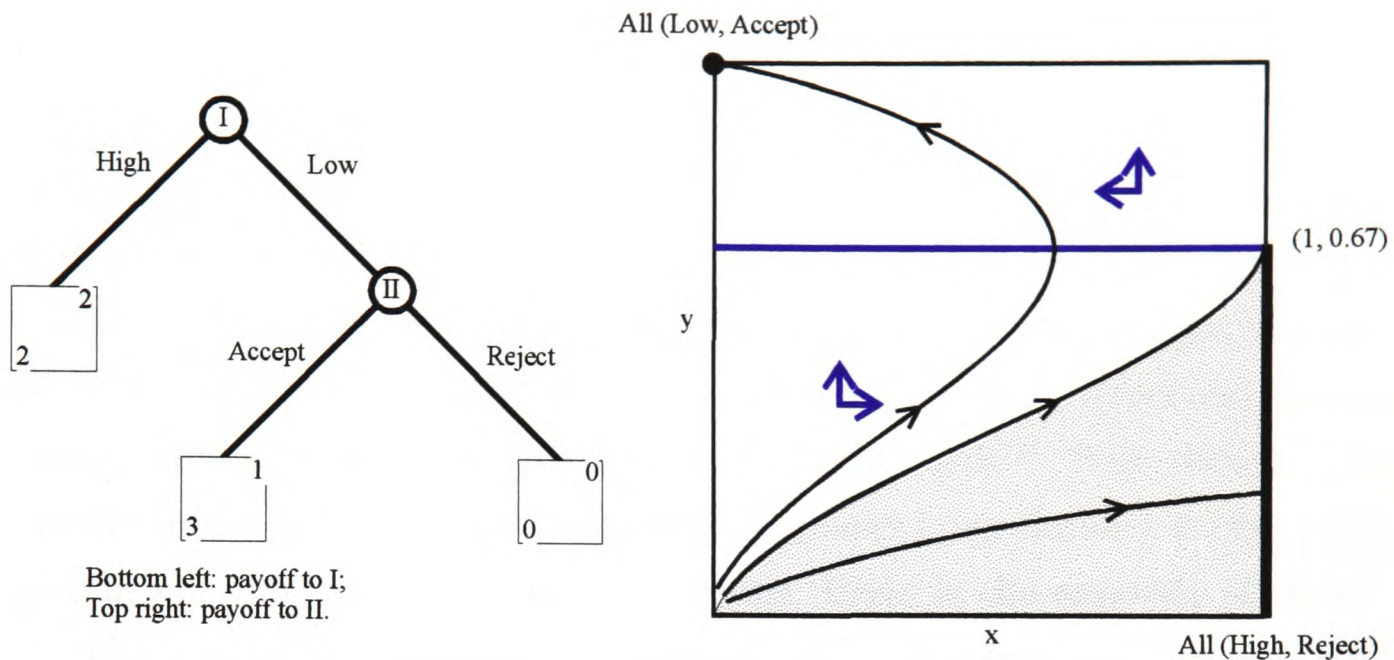


Figure 4.4. The 'mini' ultimatum game with perfect imitation.

represents the proportion in the player I population playing 'High' and  $y$  represents the proportion in the player II population playing 'Reject':

$$\begin{aligned} \dot{x} &= x(1-x)(2-3y) \\ \dot{y} &= y(1-y)(1-x) \end{aligned} \tag{4.3}$$

Under the replicator dynamics, the point  $(0, 1)$ , corresponding to the sub-game-perfect equilibrium of the game, is asymptotically stable, while the component  $(1, [0, \frac{2}{3}])$ , corresponding to the other Nash equilibria of the game, are only *local* attractors. In fact, if we take a point on this component, then a sequence of separate perturbations will gradually nudge the population state up to the point  $(1, \frac{2}{3})$ , from whence it converges rapidly to  $(0, 1)$ . A conventional evolutionary analysis of the game thus fails to provide a good explanation of the experimental results in which player I plays 'High' in the majority of cases.

The approach taken by Binmore, Gale, and Samuelson (1993) is to aug-

ment the standard replicator dynamics by adding an arbitrarily small amount of ‘drift’. Drift is the residue of the noise generated by the learning process that survives the ‘smoothing’ process of applying a strong law of large numbers to make the learning dynamic deterministic. If the drift in the player II population is stronger than the drift in the player I population, then there exists in the component  $(1, (\frac{1}{2}, 2 - \sqrt{2}))$  an asymptotically stable fixed point with a basin of attraction roughly given by the shaded area in Figure 4.4.

This seems to go some way towards explaining why there might be long-run equilibria in which player I chooses ‘High’, but many questions remain unanswered. What is the correct model for the ‘drift’? Why should it be greater in the player II population than the player I population? Is this really a sufficient explanation for player I playing ‘High’ in the *majority* of cases?

The aim of this section is to suggest there is a very good reason why the ‘drift’ in the player II population should actually be very much greater than the drift in the player I population when  $x$  is high and we suppose the underlying learning process is driven by direct imitation. The reason is quite simple: when  $x$  is high, a high proportion of the player I population are playing ‘High’, which means that a player II imitating from the player II population will not, in a large proportion of cases, observe the action taken by player II, whereupon they may have to guess.

Consider the following very simple learning process. Each population member selected to update is given a ‘random aspiration level’  $\Lambda \sim U(0, 3)$ . If her current payoff is at least  $\Lambda$  then she sticks with her current action. If it is less than  $\Lambda$ , then she imitates the action of a population member drawn from her population at random. Binmore, Gale, and Samuelson (1993) show that if this imitation is perfect then this process generates a time-rescaled version of the standard replicator dynamics given in (4.3).

Consider a dissatisfied member of the player II population selecting a population member to imitate. With probability  $x$  this member played against a ‘High’ player I. Suppose that with probability  $p$  she can somehow discern what the member would have played if called upon to play, but with proba-

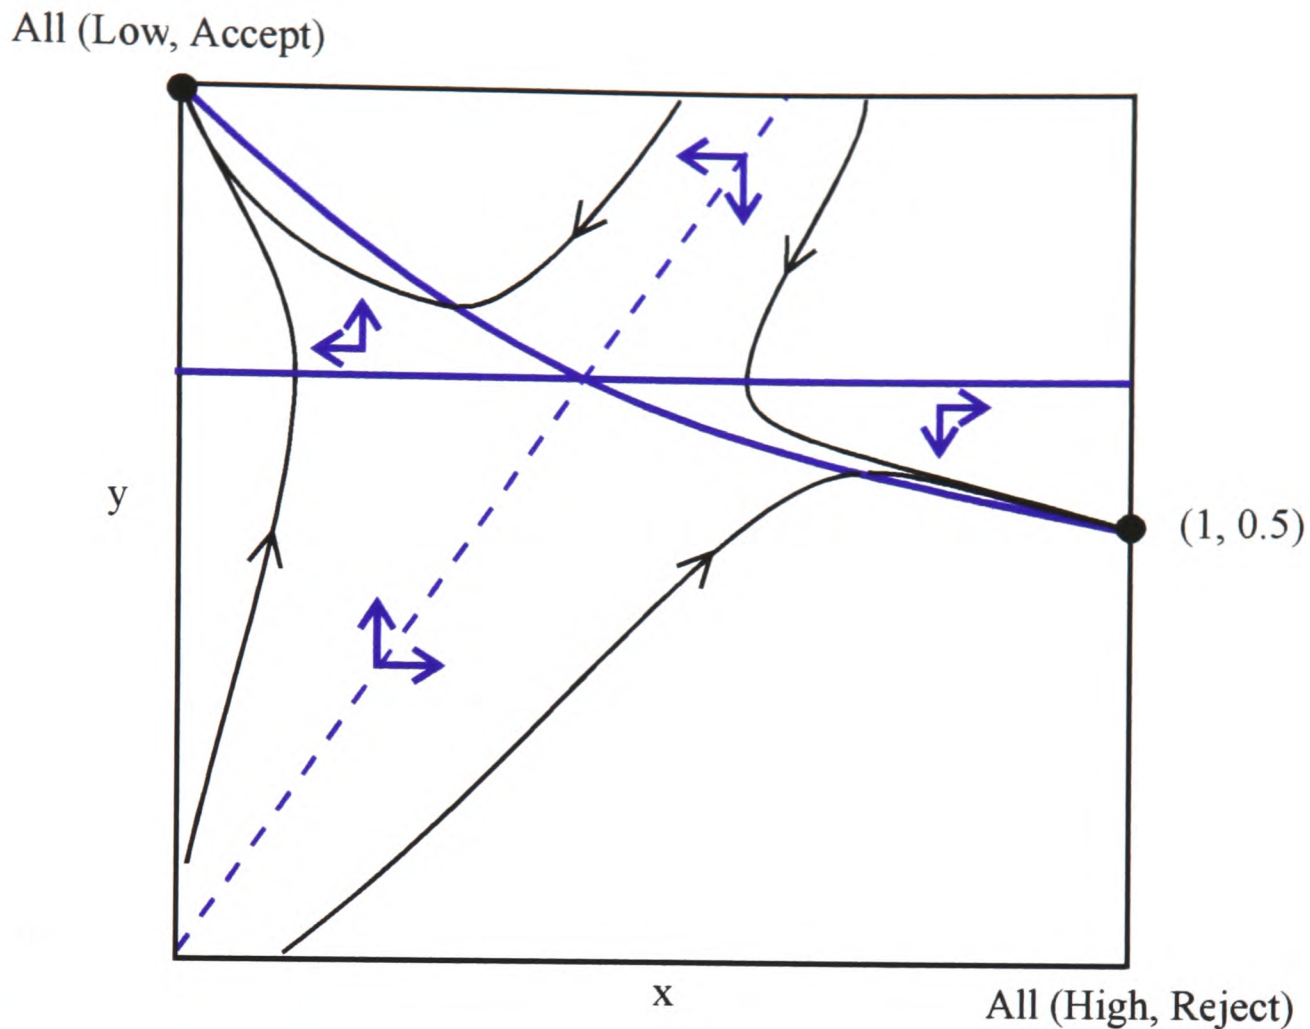


Figure 4.5. The ‘mini’ ultimatum game with imperfect imitation.

bility  $1 - p$  she guesses, placing equal weight on ‘Accept’ and ‘Reject’. The probability of choosing ‘Reject’ from this imitation process is thus:

$$x(py + (1 - p)\frac{1}{2}) + (1 - x)y = (1 - p)x(\frac{1}{2} - y) + y$$

We now obtain a learning dynamic given by:

$$\begin{aligned} \dot{x} &= \frac{1}{3}x(1 - x)(2 - 3y) \\ \dot{y} &= \frac{1}{3} \left( y(x - 2) + \left( (1 - p)x(\frac{1}{2} - y) + y \right) (3 - 2x - y + xy) \right) \end{aligned} \quad (4.4)$$

With  $p = 1$ , this is simply a time-rescaling of (4.3). However, with  $p = 0$ , which seems natural given the structure of the game, we get a very different dynamic illustrated in Figure 4.5. There are now two asymptotically stable fixed points: the subgame-perfect equilibrium at  $(0, 1)$  and a ‘fair’ equilibrium at  $(1, \frac{1}{2})$ . So why would 50% of the player II population leave money ‘on the

table' in the 'fair' equilibrium if called upon to play? Simply because they do not know any better. No player II is actually doing anything, so there is no observable experience from which to learn.

The system with  $p = 0$  differs from the system with  $p = 1 + \text{drift}$  in at least two key respects. First, we now have an explicit explanation of why the system should differ from the standard replicator dynamics. Secondly, the basin of attraction of the 'fair' equilibrium in the case  $p = 0$  is larger than that of the subgame-perfect equilibrium; while in the  $p = 1 + \text{drift}$  case, the reverse is true. We can look at the implications of this in a number of ways. An ultralong-run analysis of discrete-time versions of these dynamics, along the lines of *e.g.* Ellison (1995), would select the 'fair' equilibrium in the former and the subgame-perfect equilibrium in the latter. Alternatively, in the absence of any other model, we might suppose that the initial conditions  $x_0$  and  $y_0$  are independently drawn from  $U(0, 1)$ . This would give us the 'fair' allocation in approximately 30% of cases in the  $p = 1 + \text{drift}$  model, but in 60% of cases in the  $p = 0$  model—a figure more in line with the experimental results.

#### 4.3.3 *The Plausibility of Dynamics that are not Payoff Monotonic*

As was explained above, the growth-rate function of a selection dynamic is *payoff monotonic* if, at any given point in the state simplex, the population proportions of pure strategies with higher expected payoffs at that point grow at a higher rate. Virtually all the selection dynamics discussed in the literature, including the ubiquitous replicator dynamics, are payoff monotonic. The assumption of payoff monotonicity is sometimes justified as 'intuitively obvious'. More often than not, it is assumed in order to provide some continuity between the selection results of evolutionary game theory and the selection results of classical game theory, especially those given by the concept of a Nash equilibrium.

However, if our concern is towards finding *evolutionary* justifications for

proposing one selection dynamic over another, then we need to think more carefully about this assumption. If we imagine some ulterior evolutionary process selecting between learning dynamics<sup>10</sup>, then would such a process select only those that are payoff monotonic? If it were only short-run performance that counted, then the answer would be ‘Yes’ and we could rest easy. This makes sense in a biological context: in anything other than the extreme short-run in a biological evolutionary model the players are dead. However, in a social and economic context players will last longer relative to the progress of learning. It could be the *long-run* performance of a selection dynamic that carries more weight in the ulterior selection between dynamics. If so, then we might expect to see players learning in such a way that they are prepared to forgo short-run payoffs in order to reach payoff dominant long-run outcomes. Such learning would *not* be payoff monotonic.

Pure Coordination		
	<i>a</i>	<i>b</i>
<i>a</i>	(5, 5)	(0, 0)
<i>b</i>	(0, 0)	(3, 3)

Stag Hunt		
	<i>a</i>	<i>b</i>
<i>a</i>	(5, 5)	(0, 4)
<i>b</i>	(4, 0)	(3, 3)

It is in certain coordination games that the assumption of payoff monotonicity seems most counter-intuitive. In the games above, let the proportion of players using *a* at time *t* be given by  $x_t$ . Under payoff monotonic selection dynamics the population will converge to the (3, 3) equilibria for all  $x_0 < \frac{3}{8}$  in the ‘Pure Coordination’ game, and for all  $x_0 < \frac{3}{4}$  in the ‘Stag Hunt’ game. However, the costs of mis-coordination are not especially high<sup>11</sup>, so we might

---

<sup>10</sup>Bearing in mind, of course, the caveat mentioned in section 4.2 that an appeal to ulterior evolutionary processes may logically result in an infinite regress. The argument here is that *any* reasonable ulterior process one step further back than the dynamics being considered may throw into doubt the assumption of payoff monotonicity.

<sup>11</sup>Unlike, for example, the game where players choose which side of the road to drive

expect players to risk mis-coordination in the short-run to get to the (5, 5) equilibria.

Consider a population of  $N$  players and the following class of selection dynamics, similar to that studied above in section 4.3.2. Time is discrete and the players are matched randomly between periods to play a normal-form game. Every period, each player is given a number drawn independently from a uniform distribution with support  $(0, 5)$ . This is their *aspiration level* for the period. If a player's current payoff is less than her aspiration level, then she observes  $n$  'role models' drawn randomly with replacement from the population. She then imitates the action of the role model who receives the highest current payoff.

For  $n = 1$  we may easily derive the dynamics of the population proportion  $x_t$ , dropping the 't' subscript for notational simplicity, for the 'Pure Coordination' game:

$$E(\Delta x) = \frac{1}{5}x(1-x)(8x-3) \quad (4.5)$$

As we noted above, this is simply a constant rescaling of time of the replicator dynamics and is therefore payoff monotonic. The case  $n > 1$  gives a rather more complex expression:

$$E(\Delta x) = \frac{1}{5}x(1-x) \left( 2 + 3x + (1 + 4x)(2^n((1-x)x)^n - 2(1-x^2)^n) \right). \quad (4.6)$$

The equivalent expression for the 'Stag Hunt' game is:

$$E(\Delta x) = \frac{1}{5}x(1-x) \left( 2 - x + 2(1+2x)((1-x)x)^n - (1-x^2)^n \right). \quad (4.7)$$

What the large observation sets do is distort the dynamics as shown in Figure 4.6. For all sizes of observation set, the dynamics of both games have two

---

on. In such a game we would not expect mis-coordinating players to survive to enjoy *any* long-run payoffs!




Figure 4.6. Expected change in population proportions for different sizes of observation set.

stable equilibria: at  $x = 0$  and  $x = 1$ . Large observation sets shrink the basin of attraction of the  $x = 0$  (Pareto inferior) equilibrium. In the ‘Stag Hunt’ game, the  $x = 1$  equilibrium has the smaller basin of attraction for small  $n$  (including  $n = 1$ , the replicator dynamics); but the  $x = 0$  equilibrium has the smaller basin of attraction for large  $n$ .

Moreover, we can ensure that the system converges to the Pareto dominant equilibrium from *any* initial value if we augment the dynamics by allowing updating players to *experiment*. That is, when a dissatisfied player chooses to update, with probability  $\varepsilon$  they experiment, choosing an action at random, but with probability  $1 - \varepsilon$  choose the best action in a peer group of size  $n$ , as above.

Let  $n^*(\varepsilon)$  denote the minimum peer group size needed to ensure that  $x$  converges to 1 from any initial value as  $t \rightarrow \infty$  when the level of experimentation is  $\varepsilon$ . Figure 4.7 shows values of  $n^*$ , calculated numerically, for the two games here.

So how plausible are these non-monotonic dynamics? Learning under a dynamic with  $\varepsilon > 0$  and  $n \geq n^* > 1$  has the following features:

- Updating players choosing new actions are influenced by the perfor-

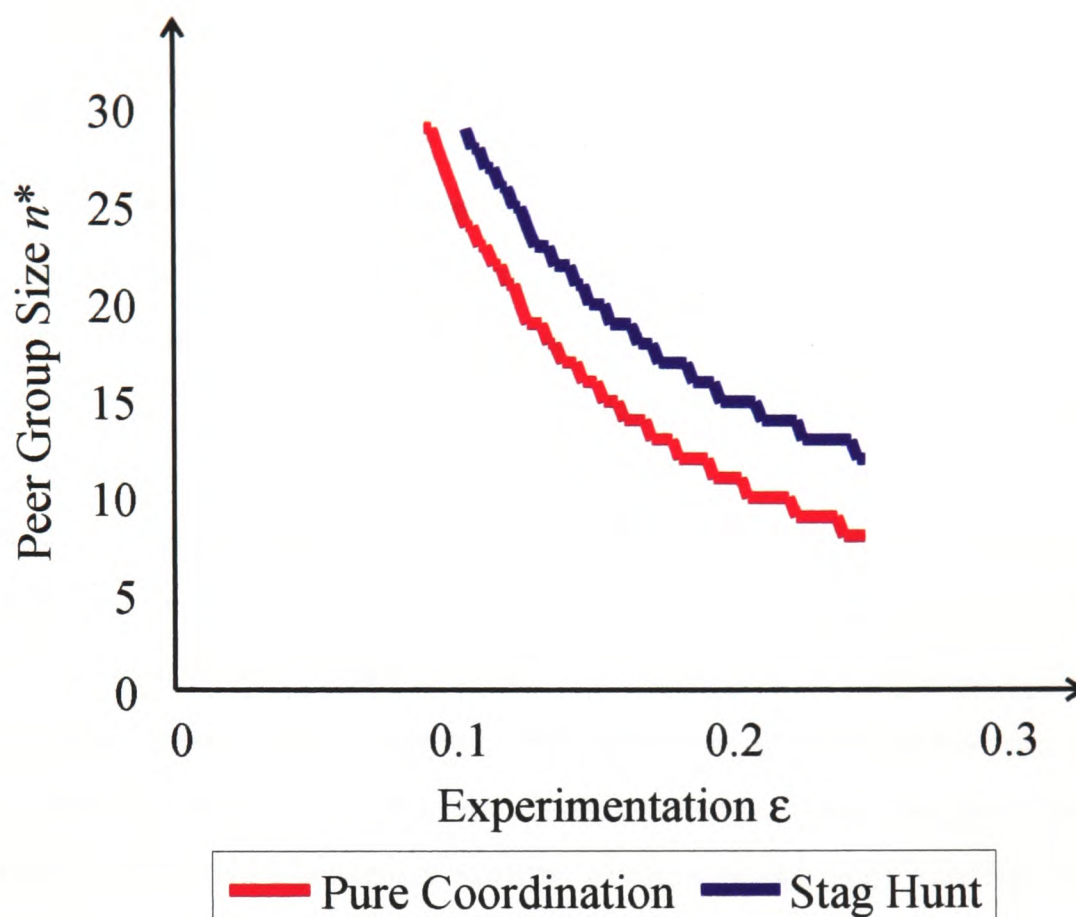


Figure 4.7. Peer group size  $n^*$  required to ensure convergence to the Pareto dominant equilibrium from any initial value when the level of experimentation is  $\epsilon$ .

mance of a wide ‘peer group’ of role models.

- Updating players occasionally experiment.
- If the system starts near an inferior equilibrium of the game, then short-run losses from high initial levels of mis-coordination are compensated by long-run gains from convergence to the superior outcome.

It might be noted that this learning has some of the features that we casually observe in language acquisition. Learning a language is faster when we are ‘immersed’ in a culture speaking that language and can learn from a number of role models. Moreover, one has to be prepared to make mistakes early on if one is going to make any progress. Since, as philosophers including Lewis (1968) have noted, language acquisition can—with a little imagination—be

seen as a giant coordination game in which the ‘players’ attempt to coordinate on common meanings, then this adds plausibility to the dynamics being studied here.

The final verdict on the plausibility of a learning mechanism can really only be, as we noted above, an empirical one. However, before such experimental evidence becomes available, thought experiments like the one we have been considering here do tell us to be cautious about exclusive concentration on one class of learning mechanisms.

Finally, it is important to note that while the learning mechanism studied above performs well in coordination games (in the long-run), there are many games in which it would perform badly<sup>12</sup>. While this could be seen to dampen its plausibility, it could also suggest that there are good reasons why we might learn in different ways in different circumstances. One implication is that experimental studies of learning should always make very careful account of the context in which the learning takes place.

#### 4.4 CONCLUSION

One of the purposes of this chapter has been to argue that the study of evolutionary game theory in economics helps to identify economic contexts where it is a more appropriate modelling method than methods that assume rational choice—and where it is not. This is useful both in assessing the existing theoretical literature and in developing new research programs. Given that it *is* an appropriate approach in many contexts, the methods of evolutionary game theory will become essential tools for the economic theorist.

However, there much to suggest that these tools are currently not fit for the task. In contexts where evolutionary game theory is an appropriate modelling method, the results obtained are highly sensitive to how one models

---

<sup>12</sup> *e.g.* in  $2 \times 2$  symmetric games with just one symmetric (mixed-strategy) Nash equilibrium, these dynamics may perform worse in the long-run than a payoff-monotonic dynamic.

the learning mechanism driving the results, and how this integrates with the context in question. The most commonly used learning mechanisms and equilibrium concepts used by economists applying evolutionary game theory are too highly derivative of biological mechanisms that assume a biological context. The examples of section 4.3 show that this need not be an occasion for depression, but a spur to creative modelling that is linked to specific contexts.

Finally, the long-term acceptance of evolutionary techniques in economics depends on careful empirical and experimental work on the adaptive behaviour of real subjects, for only then will the accusation that evolutionary methods are more *ad hoc* than their rational choice counterparts be properly addressed. In the meantime, evolutionary game theorists must keep as open and general as possible about the learning mechanisms they study.

## *Chapter 5*

# *Copying Fidelity in the Evolution of Finite Automata that Play the Repeated Prisoners' Dilemma*

### SUMMARY

A population of players chooses finite automata to play the infinitely repeated prisoners' dilemma. Players attempt to imitate automata that have proved successful in previous generations in such a way that population proportions follow the replicator dynamics. This imitation may be imperfect, as they do not observe the structure of automata directly, only action sequences. When automata are deterministic, the problem of copying fidelity prevents a utilitarian equilibrium. However, if players occasionally make operating mistakes when using their automata, they may be copied faithfully. Faithful imitation of complex automata is essentially impractical but particularly easy in populations of simple 'clumsy' automata.

## 5.1 INTRODUCTION

As we saw in the last chapter, it has become increasingly common for game-theorists to model the process by which players reach an equilibrium in a game using techniques adapted from evolutionary biology. Several factors have accelerated this tendency. Under the orthodox approach of noncooperative game theory, in which players are assumed to reason carefully, most games of any interest have multiple Nash equilibria and there is a general feeling of dissatisfaction about attempts to select between them by proposing increasingly more elaborate definitions of 'rationality'. Also, evidence from experimental studies of game playing behaviour suggests that players are perhaps not so careful in their reasoning as was once supposed. Finally, seminal work by authors such as Dawkins (1976) and Boyd and Richerson (1985) suggests that *imitation* may be a replicating mechanism that works in a similar manner to genetic replication. The analogy between biological evolution and cultural evolution may thus be close enough for students of the latter, including economists, to learn a great deal from students of the former.

Evolutionary game theory has scored a number of successes, most notably on the issue of equilibrium selection. Attention so far has generally been focussed on very simple games. However, it would certainly be useful to be able to apply evolutionary techniques to more complex situations; in particular, to *repeated* games where, because of the Folk Theorem, the equilibrium selection problem is especially acute. This is the task attempted by the present chapter, with special attention paid to the repeated version of the prisoners' dilemma shown in Table 5.1.

We immediately face a number of difficulties when we attempt to apply evolutionary techniques to repeated games.

First, there is the issue of *timing*. By placing a repeated game in an evolutionary setting, we are forced to imagine a *repeated* repeated game. Would players play a repeated game often enough and frequently enough for

Table 5.1: The Prisoners' Dilemma

	Player 2	
	Cooperate	Defect
Player 1		
Cooperate	$R$	$T$
Defect	$S$	$P$

$$T > R > P > S; 2R > T + P$$

them to learn how to play using the gradual trial-and-error methods that the evolutionary approach supposes? This ceases to be a problem if we note that in a cultural setting a repeated game may be played very quickly. However, the model below is based around supergames: infinitely repeated games with no discounting. Repeating a supergame requires a vivid imagination—but the gain in tractability probably outweighs the effort required to see such games as approximations of relatively long repeated games involving patient players.

Secondly, how should we specify the strategies that players use? In this model, following the approach of many authors interested in boundedly rational behaviour (*e.g.* Rubinstein (1986), Abreu and Rubinstein (1988), Binmore and Samuelson (1992), Piccione (1992)), players use *finite automata* known as ‘Moore machines’ to play a repeated game. These automata are simply a convenient algorithmic way of representing the ‘rules-of-thumb’ that players might use.

Thirdly, what is the appropriate equilibrium concept? It is conventional in evolutionary game theory to find the asymptotically stable state of a population under the replicator dynamics. However, as a population state in a supergame would be defined over an infinite number of possible strategies, this method is essentially impractical. An alternative would be to use a version of the ‘Evolutionarily Stable Strategy’ (ESS) of Maynard Smith (1982),

appropriately modified to account for the possibility of polymorphous equilibrium populations. However, even when thus modified, the stability condition built into the ESS concept is excessively strong and, unlike replicator dynamics, it is difficult to relate directly to some updating mechanism based on imitation.

The approach taken here is to require both *internal stability* and *external stability*. Internal stability requires that the population proportions of automata currently *in use* be stable under the replicator dynamics. External stability requires that the population is not unduly disturbed under replicator dynamics if a proportion of players are given a 'mutant' automaton from outside the set of those currently in use. This is a weaker concept than a stable state in a full-blown dynamical systems analysis, as it does not allow for the possibility of overlapping invasions by polymorphous clusters of mutant automata, but it would seem to be a practical compromise.

Lastly, there is the issue of *copying fidelity*. This issue is of prime importance if one supposes that the 'cultural evolution' in the model is driven by imitation. It is an extension of the problem noted in evolutionary analyses of extensive form games (*e.g.* by Fudenberg and Levine (1993)) that there is no evolutionary pressure on behaviour at unreached information sets. Here, players do not observe supergame strategies directly when they attempt to imitate, they only observe the action sequence the strategies generate in particular matches. They may thus copy a strategy incorrectly. Imitation is imperfect and can introduce new strategies to a population inadvertently.

The chapter proceeds as follows. Section 5.2 describes the model, the equilibrium concept and the process of imitation in more detail. In section 5.3, we consider the infinitely repeated prisoners' dilemma and ask whether a *utilitarian* evolutionary equilibrium is possible—*i.e.* an equilibrium in which the sum of players' payoffs is maximized. The answer given is a qualified 'No'. A population of players receiving the utilitarian payoff is only resistant to invasion if the strategies they use are 'provokable' and punish defection. In a potential equilibrium, this element of a utilitarian strategy is never used

and is thus not copied to the next generation, making the population externally unstable. Binmore and Samuelson (1992) argue that a non-utilitarian evolutionary equilibrium is also impossible. Section 5.4 explains how this argument relies on the overly strong stability condition built into the ESS-type equilibrium concept they employ and why it cannot be used in any model based around the replicator dynamics—as this one is. Section 5.5 discusses how introducing some noise to the model may restore copying fidelity if the automata in a population are very simple. Imitation is particularly easy in populations of stochastic automata with no more than two internal states. In this case, an action choice perfectly reveals which state an automaton is in, making it much easier to reconstruct its transition function from the data provided by an action sequence. Section 5.6 discusses evolutionary equilibrium in populations of two state stochastic automata. We see how a monomorphous population of automata operating a stochastic strategy known as 'PAVLOV' constitute a utilitarian evolutionary equilibrium if the permitted automata are restricted to guarantee copying fidelity. PAVLOV tends to stick with its current action if it receives a payoff of  $R$  or  $T$  in Table 5.1, and to switch actions if it receives anything less.

## 5.2 THE MODEL

### 5.2.1 Basic Elements

Consider a symmetric  $2 \times 2$  strategic-form game  $G = \langle \mathbf{A}, g \rangle$ , where  $\mathbf{A} = A^2 = \{C, D\}^2$  and  $g : \mathbf{A} \rightarrow \mathbf{R}$ . Let  $G^\infty$  denote the supergame of  $G$  (*i.e.* the undiscounted game where  $G$  is repeated infinitely). Most of what follows takes  $G$  to be the prisoners' dilemma of Table 5.1. The realized actions of the players in period  $\tau \in \{0, 1, 2, \dots\}$  of  $G^\infty$  are denoted  $\alpha^\tau \equiv (\alpha_1^\tau, \alpha_2^\tau)$ , and a history is denoted  $h^\tau = (\alpha^0, \dots, \alpha^{\tau-1})$ , where  $h^\tau \in H^\tau = (A)^\tau$ .

In this model, strategies are taken from the subset of strategies that may be played by a particular type of deterministic finite automata known

as a 'Moore machine'. A finite automaton capable of playing, for player  $i$ , the game  $G^\infty$  has  $i$ 's actions in the stage game as its output set and its opponent's actions in the stage game as its input set. An automaton  $a$  is a 4-tuple  $\langle Q_a, q_a^0, \lambda_a, \mu_a \rangle$ , where  $Q_a = \{\omega_0, \dots, \omega_{|a|}\}$  is a finite set of states ( $|a|$  denotes the number of states in  $a$ ),  $q_a^0 \in Q_a$  is  $a$ 's *initial state*,  $\lambda_a : Q_a \rightarrow A$  is  $a$ 's *output function*, and  $\mu_a : Q_a \times A \rightarrow Q_a$  is  $a$ 's *transition function*, which describes how the automaton's state responds to its opponent's last action.

Consider a large but finite population of players each of whom has one Moore machine with which to play  $G^\infty$ . The sequence of play is divided into *generations* (indexed by  $t$ ), each consisting of an infinite number of periods (indexed by  $\tau$ ). In each of these generations, the players are randomly matched to play  $G^\infty$ . Consider one of these matches, involving players with automata  $a = \langle Q_a, q_a^0, \lambda_a, \mu_a \rangle$  and  $b = \langle Q_b, q_b^0, \lambda_b, \mu_b \rangle$ . We get two sequences of states when  $a$  plays  $b$ :

$$\{q_i^\tau\}_{\tau=0}^\infty = (q_i^0, q_i^1, \dots); \quad i = a, b;$$

where

$$q_i^{\tau+1} = \mu_i(q_i^\tau, \lambda_{-i}(q_{-i}^\tau)); \quad i = a, b.$$

The supergame payoff from using  $a$  against  $b$  is:

$$\pi(a, b) = \lim_{\Gamma \rightarrow \infty} \frac{1}{\Gamma} \sum_{\tau=0}^{\Gamma} g(\lambda_a(q_a^\tau), \lambda_b(q_b^\tau)).$$

Alternatively, let  $\mathbf{q}^\tau = (q_a^\tau, q_b^\tau)$  and note that, as the automata are finite, there is a minimal  $\tau_2$  and  $\tau_1 \leq \tau_2$  such that  $\mathbf{q}^{\tau_1} = \mathbf{q}^{\tau_2}$ . Call  $(\mathbf{q}^0, \dots, \mathbf{q}^{\tau_1-1})$  the *initial sequence*, and  $(\mathbf{q}^{\tau_1}, \dots, \mathbf{q}^{\tau_2})$  the *cycle sequence*. Then

$$\pi(a, b) = \frac{1}{\tau_2 - \tau_1 + 1} \sum_{\tau=\tau_1}^{\tau_2} g(\lambda_a(q_a^\tau), \lambda_b(q_b^\tau)).$$

### 5.2.2 Population Dynamics

It would be usual in evolutionary game theory to define the state of the large but finite population of players in this model over the set of all possible

strategies—*i.e.* over the entire set of Moore machines capable of playing  $G^\infty$ , denoted by  $\mathcal{M}$ . However, as this set is infinite, it will be more convenient to work with the subset of  $N(t)$  automata actually *in use* in generation  $t$ , denoted by  $\mathcal{P}(t) = \{m_1, \dots, m_{N(t)}\} \subset \mathcal{M}$ . Let  $x_i(t) > 0$  denote the proportion of players using automaton  $m_i$  at time  $t$ . The population state is now a pair  $(\mathcal{P}, \mathbf{x})$  such that  $\mathbf{x} = (x_1, \dots, x_N)$ , where  $\mathbf{x}$  is in the interior of the unit simplex.

Let  $\pi(a, \mathcal{P})$  denote the expected payoff to a player using  $a$  in a random match in a population where automata  $\mathcal{P} = \{m_1, \dots, m_N\}$  are in use. Clearly,

$$\pi(a, \mathcal{P}) = \sum_{i=1}^N x_i \pi(a, m_i).$$

Finally, let  $\pi(\mathcal{P})$  denote the population average payoff, where

$$\pi(\mathcal{P}) = \sum_{i=1}^N x_i \pi(m_i, \mathcal{P}).$$

Each generation, a proportion of players updates the automata they use. The updating mechanism involves the players imitating automata from the previous generation in such a way that population proportions follow the replicator dynamics. That is, the growth rate of the proportion of players using a strategy is proportional to its relative fitness—*i.e.* the difference between its expected payoff in a random match and the average payoff received in the population. The replicator dynamics do not necessarily imply the genetic conditioning of players. A large number of mechanisms based on imitation have been proposed that result in the replicator dynamics. For example, Binmore, Gale, and Samuelson (1993) describe a mechanism in which players have aspiration levels drawn each generation from a uniform distribution; if a player's current payoff fails to satisfy her aspirations, then she imitates the strategy of a player drawn from the population at random. Section 4.4 of Weibull (1995) describes no less than three updating mechanisms based on imitation that result in the replicator dynamics.

Replicator dynamics implies that, when updating happens continuously (*i.e.* the time between generations tends to zero),

$$\dot{x}_i = x_i(\pi(m_i, \mathcal{P}) - \pi(\mathcal{P})). \quad (5.1)$$

### 5.2.3 Evolutionary Equilibrium

Next, we need to be clear about what we mean by an 'evolutionary equilibrium' in this model.

**DEFINITION 5.1** *A population  $(\mathcal{P}, \mathbf{x})$  is **internally stable** if the population is monomorphous or if  $\mathbf{x}$  is a (Lyapunov) stable fixed point in the interior of the unit simplex under the replicator dynamics of equation (5.1).*

Internal stability is a necessary but not sufficient condition for an evolutionary equilibrium in this model. Suppose we give a proportion  $x_0$  of an internally stable population a 'mutant' automaton  $m_0$  from outside the set of automata currently in use—*i.e.* from  $\mathcal{M} \setminus \mathcal{P}$ . This gives an augmented population  $(\mathcal{P}', \mathbf{x}', x_0)$ , where  $\mathcal{P}' = \{m_0, m_1, \dots, m_N\}$  and  $x'_i = (1 - x_0)x_i$  for all  $i \in \{1, \dots, N\}$ . If such a 'mutant invasion' drives the population proportions of one or more of the existing automata to the boundary of the unit simplex for all possible values of  $x_0$ , then the original population was not externally stable.

**DEFINITION 5.2** *A population  $(\mathcal{P}, \mathbf{x})$  is **externally stable** if there is no augmented population  $(\mathcal{P}', \mathbf{x}', x_0)$  such that, for all  $0 < x_0 < 1$ ,  $x_i \rightarrow 0$  as  $t \rightarrow \infty$  under the replicator dynamics of equation (5.1) for some  $i \in \{1, \dots, N\}$ .*

In this model, a population is in evolutionary equilibrium if it is both internally and externally stable. Despite being a much weaker concept of equilibrium than is usual in evolutionary game theory, we shall still find an equilibrium difficult to obtain with the deterministic automata described in section 5.2.

5.2.4 *Imitation and The Copycat Criterion*

Denote by  $h_{ab}$  the history (*i.e.* infinite sequence of action pairs) generated by automaton  $a$  playing automaton  $b$  in a given play of  $G^\infty$ .

The difficulty facing players who have chosen to imitate is that they do not observe strategies directly, they only observe histories of play. For example, if player has chosen to copy the strategy of a player who played with automaton  $a$  against automaton  $b$  in the previous generation, then she has only the history  $h_{ab}$  to work from. To formalize the process of imitation, first define a *complete copy class* as follows:

DEFINITION 5.3 *The complete copy class of automaton  $a$  given  $h_{ab}$  is given by:*

$$K(a, b) = \{m \in \mathcal{M} \mid \lambda_m(\mu_m(q_m^{\tau-1}, \alpha_b^{\tau-1})) = \alpha_a^\tau \quad \forall \tau = 1, 2, \dots\}$$

In other words,  $K(a, b)$  consists of all finite automata that would have produced the same sequence of action pairs had they been playing in place of automaton  $a$  in a match against automaton  $b$ .

However, most of the members of  $K(a, b)$  will be unnecessarily complex, while it seems reasonable to suppose that a player will seek to choose a simple way of mimicking a strategy. To select the least complex members of  $K(a, b)$ , partition the set as follows. An *automaton-size class* of automaton  $x$  in  $K(a, b)$  is given by:

$$[x] = \{y \in K(a, b) \mid |y| = |x|\}.$$

Next, define a partial ordering over automaton size classes in  $K(a, b)$ :

$$[x] \preceq [y] \Leftrightarrow |\xi| \leq |v|, \quad \text{where } \xi \in [x] \text{ and } v \in [y].$$

DEFINITION 5.4 *The minimal copy class  $C(a, b)$  of automaton  $a$  given  $h_{ab}$  is the minimal automaton-size class of  $K(a, b)$ .*

That is, the minimal copy class is the set of automata with the fewest number of states in the set of automata that would have generated an indistinguishable action sequence had they played in  $a$ 's place.

We can now define precisely what we mean by imitation.

**DEFINITION 5.5** *A player is said to **imitate** an automaton  $a$  used against automaton  $b$  if she chooses at random an automaton from the minimal copy class  $C(a, b)$ .*

When will imitation of this kind not introduce new automata to the population? When can all automata in the population be generated by observing a match between two automata also in the population?

**DEFINITION 5.6** *A population  $(\mathcal{P}, \mathbf{x})$  is a **copycat population** iff (i) all minimal copy classes  $C(a, b)$ ,  $a, b \in \mathcal{P}$ , contain only members of  $\mathcal{P}$ , and (ii) all members of  $\mathcal{P}$  are in some minimal copy class  $C(a, b)$ , where  $a, b \in \mathcal{P}$ .*

**PROPOSITION 5.1** *A population is a copycat population if and only if*

$$\mathcal{P} = \bigcup_{a, b \in \mathcal{P}} C(a, b)$$

**PROOF.** This follows directly from definition 5.6.

### 5.3 UTILITARIAN POPULATIONS AND THE PRISONERS' DILEMMA

Consider the case where  $G$  is the prisoners' dilemma of Table 5.1 and where the population  $(\mathcal{P}, \mathbf{x})$  is a *utilitarian* copycat population. That is, the sum of payoffs across players is maximized and all players receive the cooperative payoff  $R$ . Such a population is clearly internally stable, but is it externally stable?

We shall see that the problem of copying fidelity prevents the perfect propagation of the provokability that makes a utilitarian population resistant to invasion. Imperfect imitation means that a utilitarian population is

progressively corrupted by strategies, which we shall call 'suckers', that fail to punish defection in the long run.

DEFINITION 5.7 A '*sucker*' is a Moore machine with a single cooperative cycle state  $q_c$  and a transition function such that  $\mu(q_c, C) = \mu(q_c, D) = q_c$ .

A trivial example of a 'sucker' is the one-state automaton COOPERATE, but in principle 'suckers' may be quite complex, with very long initial sequences.

PROPOSITION 5.2 Suppose  $(\mathcal{P}, \mathbf{x})$  is a utilitarian copycat population. Then there is a non-empty subset  $\Sigma \subseteq \mathcal{P}$  where all members of  $\Sigma$  are 'suckers'.

PROOF. Take any  $m_k$  in  $\mathcal{P}$ . A member  $a$  of  $C(m_k, m_k) \subseteq \mathcal{P}$  obtains  $R$  against  $m_k$  by cooperating over the entire cycle sequence. As  $C(m_k, m_k)$  is a *minimal* copy class,  $a$  has a single cycle state  $q_c$  with  $\lambda_a(q_c) = C$  and  $\mu_a(q_c, C) = q_c$ . The event  $(C, D)$  is never observed in any symmetric match; so  $\mu_a(q_c, D)$  could map to any state in  $a$ . For  $1/|a|$  of the automata in  $C(m_k, m_k)$ ,  $\mu(q_c, D) = q_c$ , so at least this number are 'suckers'. (In fact, for any  $m_k$  and  $m_j$  in  $\mathcal{P}$ , either  $C(m_k, m_j)$  or  $C(m_j, m_k)$  or both will include 'suckers'.) There is thus a non-empty subset  $\Sigma$  of 'suckers' in this population, as, by proposition 5.1, the population is the union of all minimal copy classes.

*Q.E.D.*

For example, a population consisting of the automata  $m_1$  to  $m_4$  shown in Figure 5.1 is a utilitarian population. It is also a copycat population. Any match yields the action sequence

$$h_{ab} = \begin{pmatrix} D \\ D \end{pmatrix}, \begin{pmatrix} C \\ C \end{pmatrix}, \begin{pmatrix} C \\ C \end{pmatrix}, \begin{pmatrix} C \\ C \end{pmatrix}, \dots \quad a, b \in \{m_1, m_2, m_3, m_4\},$$

and these automata are the complete set of minimal automata capable of generating such a sequence. However, both  $m_1$  and  $m_3$  are 'suckers': in their cycle state they respond to defection by continuing to cooperate.

To show that a utilitarian copycat population is externally unstable, we need to find a mutant automaton capable of taking advantage of the 'suckers' in  $\Sigma$ . The following proposition shows that we will always be able to do this.

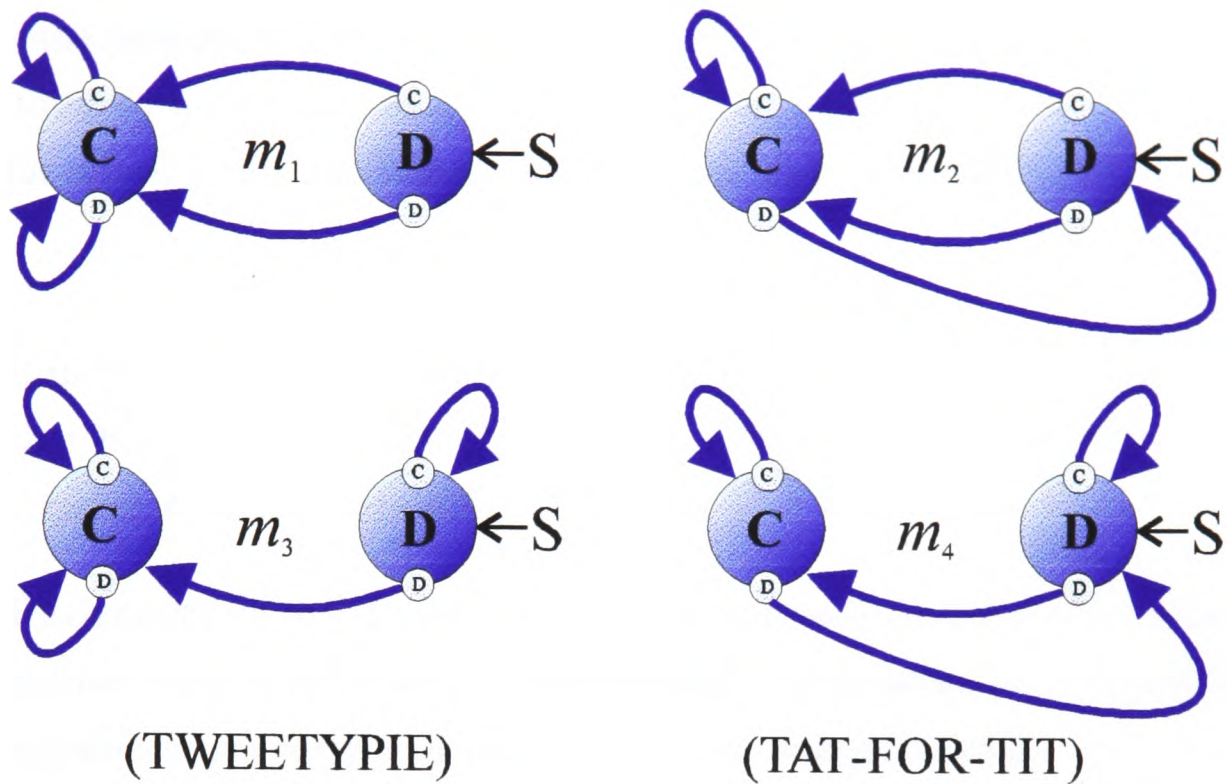


Figure 5.1. A utilitarian copycat population. The larger circles represent the automatas' internal states, each mapping to the action C or D. The arrows represent the automatas' transition functions, a mapping that depends on both the current internal state and the last action of the automatas' opponents, shown in the smaller circles.

PROPOSITION 5.3 When  $G$  is the prisoners' dilemma, for any population of automata  $(\mathcal{P}, \mathbf{x})$  that is utilitarian and copycat, we can find a mutant  $c$  such that  $\pi(c, m_k) = T > R$  for all  $m_k$  in the set of 'suckers'  $\Sigma \subseteq \mathcal{P}$ , but  $\pi(c, m_i) = \pi(m_i, c) = \pi(c, c) = R$  for all  $m_i \in \mathcal{P} \setminus \Sigma$ .

PROOF. Consider a mutant  $c$  that mimics any member of  $\mathcal{P}$  until it reaches a cycle sequence yielding  $(R, R)$ . It then 'tests the water' with a signal. If this is reciprocated, then it knows it is playing itself and it plays C forever more, yielding a supergame payoff of  $R$ . If not, it defects and continues to defect unless it encounters a defection. If its opponent  $m_i$  is a 'sucker', then  $c$  obtains a supergame payoff of  $T > R$ . If not, and it encounters a defection, then it starts a 'recovery sequence'. It experiments until it learns the identity  $m_i$  of its opponent. This leaves  $m_i$  in some known

state  $q$ . As the population is copycat and, by definition 5.6,  $m_i$  is therefore in some minimum copy class  $C(x, m_i)$ ,  $q$  is used against  $m_j$ . The mutant now mimics  $m_j$  and obtains a supergame payoff of  $R$ . We have  $\pi(c, m_k) = T > R$  for all  $m_k \in \Sigma \subseteq \mathcal{P}$ , but  $\pi(c, m_i) = \pi(m_i, c) = \pi(c, c) = R$  for all  $m_i \in \mathcal{P} \setminus \Sigma$ .  
*Q.E.D.*

If imitation post-invasion were perfect then the following proposition shows that the mutant  $c$  would drive out all 'suckers' asymptotically, thus demonstrating external instability.

**PROPOSITION 5.4** *Under perfect imitation, any invasion of an initially utilitarian copycat population by the mutant  $c$  of proposition 5.3 would drive the population proportion of 'suckers' to zero asymptotically.*

**PROOF.** Let  $x_k$  denote the proportion of 'suckers' in the augmented population  $(\mathcal{P}', \mathbf{x}, x_0)$ . The expected payoff of a 'sucker' in a random match,  $\pi(m_k, \mathcal{P}') = x_0 S + (1 - x_0)R$ . The expected payoff to mutant  $c$  in a random match,  $\pi(c, \mathcal{P}') = x_k T + (1 - x_k)R$ . The population average payoff  $\pi(\mathcal{P}') = x_k(x_0 S + (1 - x_0)R) + x_0(x_k T + (1 - x_k)R) + (1 - x_0 - x_k)R$ . The replicator dynamics of equation (5.1) give:

$$\begin{aligned} \dot{x}_k &= x_k x_0 (x_k (2R - S - T) + S - R) < 0 \quad \forall \quad x_k, x_0 > 0 \quad \text{and} \\ \dot{x}_0 &= x_0 x_k (x_0 (2R - S - T) + T - R) > 0 \quad \forall \quad x_k, x_0 > 0. \end{aligned}$$

From any strictly positive initial value,  $x_0$  remains bounded away from zero. Hence  $x_k \rightarrow 0$  as  $t \rightarrow \infty$ .

*Q.E.D.*

Of course, in this model imitation is *imperfect*, so we can only claim that a utilitarian copycat population is *probably* externally unstable. Indeed, the problem of copying fidelity means that the situation after any mutant invasion will be an undignified mess, with many new automata from  $C(c, m)$  and  $C(m, c)$ ,  $m \in \mathcal{P}$ , being introduced to the population. This makes it virtually

impossible to describe the evolution of a population after an invasion. However, if we design the mutant  $c$  so that  $c \in C(c, m)$  for a sufficiently large number of automata  $m$ , then it seems likely that it will drive the 'suckers' in a utilitarian copycat population to extinction.

#### 5.4 NON-UTILITARIAN POPULATIONS

Consider now the case where  $G$  is still the prisoners' dilemma of Table 5.1 but where the population  $(\mathcal{P}, \mathbf{x})$  is a *non-utilitarian* internally stable copycat population. Is such a population externally stable?

Theorem 9.1 of Binmore and Samuelson (1992) suggests that the answer to this question must also be 'No'. They describe a mutant  $b$  that works as follows: Let  $m_l$  denote the automaton in  $\mathcal{P}$  that receives the lowest symmetric payoff  $z = \min\{\pi(m_i, m_i) | m_i \in \mathcal{P}\}$ . The mutant  $b$  mimics  $m_l$  unless a cycle is established yielding a payoff pair of  $(z, z)$ . If this happens, it 'tests the water' with a signal. If this is reciprocated, then it knows it is playing itself and it plays  $C$  forever more, yielding a supergame payoff of  $R$ . If not, it starts a 'recovery sequence'. It experiments until it learns the identity  $m_j$  of its opponent. This leaves  $m_j$  in some known state  $q$ .  $m_j$  must use  $q$  against some  $m_k$  in  $\mathcal{P}$ , so the mutant  $b$  then mimics  $m_k$  thereafter.

However, in order for the mutant  $b$  to ensure itself a higher than average expected payoff post-invasion (whereby it would drive the existing population to extinction) it is necessary that in the pre-invasion population all payoffs are such that  $\pi(m_a, m_b) \geq z$  for all  $m_a, m_b \in \mathcal{P}$ . This guarantees  $b$  at least  $z$  when it mimics  $m_k$  above. This is true under the stability condition of the ESS-type equilibrium concept they use, but it is not in general true for a stable state under replicator dynamics. (*e.g.* Example 3.9 of Weibull (1995) describes a variation of the Rock-Scissors-Paper game with an interior *asymptotically* stable state under replicator dynamics where the minimum payoff is less than the minimum symmetric payoff.)

The existence or otherwise of a non-utilitarian copycat population in evo-

lutionary equilibrium in this model thus remains an open question.

### 5.5 RESTORING COPYING FIDELITY WITH CLUMSY AUTOMATA

The problem of copying fidelity arises because in many matches between finite deterministic automata, some mappings of the transition function never occur. There is thus often no way to faithfully reconstruct a transition function given the information provided by an action sequence. However, suppose we perturb the model above in such a way that each mapping of the transition function is more likely to be used in any given match. For example, we can perturb the transition function of the automata by introducing a uniform tremble:

$$q_i^{\tau+1} = \begin{cases} \mu(q_i^\tau, \alpha_{-i}^\tau) & : \text{ with probability } 1 - \varepsilon \\ q' & : \text{ with probability } \varepsilon \end{cases}$$

where  $q'$  is taken from a uniform distribution over  $Q \setminus \mu(q_i^\tau, \alpha_{-i}^\tau)$ . That is, this 'clumsy' automata generally switches to the state specified by its transition function, but with small probability  $\varepsilon$  switches to some other state. It is clear that in any match between two such automata, all elements of  $Q \times A$  will be reached over an infinite horizon, and hence all mappings of the transition function used.

So why should we expect the automata in a population to be 'clumsy' in this way? After all, clumsiness will often result in lower payoffs: two clumsy utilitarian automata, for example, will always receive at best a slightly lower payoff than their deterministic counterparts. As we saw above, however, deterministic automata will not survive long in an uncorrupted condition. If we take the plausible view that clumsiness is a natural condition for most human players, then the costs of reducing this clumsiness are unlikely to be compensated by slight payoff benefits if the clumsiness also contributes, by improving copying fidelity, to the smooth running of the reproductive process linking generations.

On whether clumsiness may improve copying fidelity, there is, of course, an expanding literature on learning the structure of strategies from action sequences. Kalai and Lehrer (1990), for example, show that Bayesian learning on a sufficiently long action sequence allows players to learn enough about a strategy to predict its future play. For copying fidelity in an evolutionary setting we need the stronger requirement that an observer can learn the full structure of a strategy from an action sequence; but we know that, with trembles like those discussed above, Bayesian learning will *eventually* converge to the real strategy. However, in an evolutionary context, we need to be confident that such learning is feasible for unsophisticated players using simple updating rules. In particular, it seems unlikely that the sort of players we have in mind here will be able to faithfully copy complex strategies; first, because of the length of action sequence necessary and, secondly, because of the number and complexity of the statistical calculations involved.

To address this issue, consider the following numerical simulations of Bayesian learning from action sequences generated by clumsy automata of differing complexity. Let  $\mathcal{M}_\kappa = (m_1, \dots, m_{\nu_\kappa})$  denote the set of Moore machines of size  $\kappa$ , where  $\nu_\kappa = \#\mathcal{M}_\kappa$ : thus  $\mathcal{M}_2$  is the set of all two state Moore machines;  $\mathcal{M}_3$  is the set of all three state Moore machines, and so on. Suppose a player wishes to copy machine  $a$  using the information provided by action sequence  $h_{ab}$ . Fix the complexity of the population at a given  $\kappa$ —*i.e.* the player knows that  $a \in \mathcal{M}_\kappa$ . To learn the identity of  $a$  in a Bayesian fashion, we suppose the player initially has prior beliefs  $p^0$  over  $\mathcal{M}_\kappa$ . Working along the action sequence she updates these beliefs following Bayes' rule,

$$p_i^{\tau+1} = \frac{p_i^\tau Pr(h^\tau | m_i)}{\sum_{s=1}^{\nu_\kappa} p_s^\tau Pr(h^\tau | m_s)}, \quad i = 1, \dots, \nu_\kappa, \quad (5.2)$$

stopping when she is sufficiently confident about the identity of  $a$ .

The difficulty comes in calculating  $Pr(h^\tau | m_s)$  for a given machine  $m_s$ . First note that because of the possibility of a mistake, there may be more than one sequence of internal states  $\{q^u\}_{u=0}^\tau$  consistent with  $m_s$  generating the sequence  $h^\tau$ . Let  $\Theta^\tau$  be the set of all such possible sequences—*i.e.*  $\Theta^\tau =$




Figure 5.2. The automata used in the numerical simulations of Bayesian learning.

$\{\{q^u\}_{u=0}^\tau | \lambda_s(q^w) = \alpha_a^w \forall w = 1 \dots, \tau\}$ . Then,

$$Pr(h^\tau | m_s) = I(m_s) \sum_{\Theta^\tau} \prod_{u=1}^\tau F(m_s, u, \varepsilon), \quad \text{where} \quad (5.3)$$

$$I(m_s) = \begin{cases} 1 & \text{if } \lambda(q_{m_s}^0) = \alpha_a^0 \\ 0 & \text{if } \lambda(q_{m_s}^0) \neq \alpha_a^0 \end{cases}, \quad \text{and}$$

$$F(m_s, u, \varepsilon) = \begin{cases} 1 - \varepsilon & \text{if } \mu(q^{u-1}) = q^u \\ \varepsilon & \text{if } \mu(q^{u-1}) \neq q^u \end{cases}$$

Table 5.2 shows some simulation results of Bayesian of learning for the automata shown in Figure 5.2.

Table 5.2: Simulation Results

Prior:	Uniform		50%	
	Segment Size	Calculations	Segment Size	Calculations
TFT	12	3000	10	2000
GRIM	60	$2 \times 10^5$	50	$1 \times 10^5$
C+	110	$4 \times 10^5$	100	$4 \times 10^5$
TF2T	$> 25$ (0.250) <sup>a</sup>	$3 \times 10^9$	7 <sup>b</sup>	$9 \times 10^5$
NSG	$> 25$ (0.083)	$3.4 \times 10^9$	7	$9 \times 10^5$
TFT+	$\gg 25$ (0.005)	$\approx 4 \times 10^9$	$> 25$ (0.967)	$\approx 4 \times 10^9$

<sup>a</sup>Figure in parenthesis is final belief on true automaton when simulation was terminated.

<sup>b</sup>The fast identification of TF2T and NSG with a '50%' prior is presumably because the large size of  $\mathcal{M}_3$  makes the priors on automata similar to these machines very small.

Of the two-state automata in Figure 5.2, TIT-FOR-TAT is the most easy to identify, as it moves freely between its two states in response to different actions. GRIM is rather harder to identify: once in its defect state it tends (barring mistakes) to stay there, revealing little more about its transition function. C+ is an example of an automaton with a disconnected state: it only defects by mistake. It can thus take a considerable amount of time to identify such an automaton, as one has to wait for a mistake to observe how much of its transaction function works.

Similarly, of the three-state automata, TIT-FOR-TWO-TATS is relatively easy to identify, NOT-SO-GRIM rather harder, and TIT-FOR-TAT-PLUS very hard indeed.

To keep learning times to a minimum, the simulations were conducted with the test automata of Figure 5.2 playing against RANDOM, a stochastic automaton that simply chooses a random action each period. The tremble

for the test machines was set at  $\varepsilon = 0.02$ . Two types of prior beliefs were tried. A 'uniform' prior meant setting equal beliefs over  $\mathcal{M}_2$  for the two-state automata; over  $\mathcal{M}_3$  for the three-state automata. A '50%' prior represented a lucky guess: a prior of 0.5 on the correct automaton, uniform beliefs over the other members of  $\mathcal{M}_2$  or  $\mathcal{M}_3$ . The 'segment size' reported in the table is the average length of action sequence required to place a 0.99 weight on the correct automaton.

To give some indication of the computational complexity of the learning, the number of calculations required to obtain a result was counted, such that calculating  $\prod_{u=1}^{\tau} F()$  in equation (5.3) for one state path represented  $\tau$  'calculations'. A filter was also introduced to set beliefs of automata that dropped below  $1 \times 10^{-6}$  to zero. Although this meant that the learning was not *strictly* Bayesian, and that there was a (very small) possibility of rejecting the correct machine, this was necessary to keep the number of calculations needed for the identification of the three-state machines to an acceptable level.

The results in Table 5.2 are striking. With the more plausible 'uniform' prior, it generally took a longer segment of action sequence to identify (or partially identify) the three-state machines than the two-state machines. More significantly, the computational complexity of identifying a three-state automaton was several orders of magnitude greater than that for the two-state automata, even with the simplifying filter described above. The reason for this is fairly straightforward. First, there are a great deal more members of  $\mathcal{M}_3$  than  $\mathcal{M}_2$ :  $\nu_2 = 32$  but  $\nu_3 = 4374$ . Secondly, for a non-trivial automaton of size  $\kappa$ , there may be up to  $(\kappa - 1)^\tau$  state sequences consistent with an action sequence segment of size  $\tau$ . For a two-state automaton, that is just one possible state sequence; but for a three-state automaton it is up to  $2^\tau$ . For example, for a 25 period segment of action sequence, that is up to 33,554,432 possible state paths.

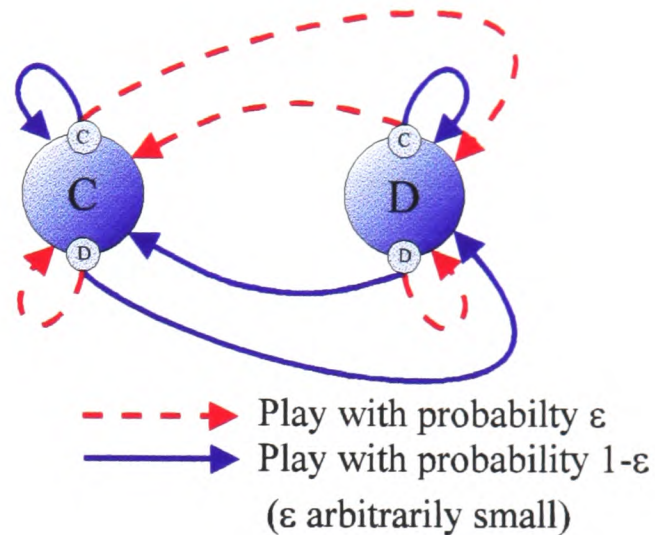
## 5.6 DISCUSSION

To put some of the results of the previous section in perspective, to analyse a 25 period segment of action sequence generated by TFT+ against RANDOM, which involves approximately  $4 \times 10^9$  calculations, takes a fairly modern PC about fourteen hours. A human observer, doing the analysis by hand with calculator and scrap paper, would need a good head for figures and—assuming that she took 0.3 seconds per calculation and worked a forty-eight hour week, fifty weeks a year—about 140 years, by which time she will have long ceased caring about the result. If the observer started with a uniform prior, she would be very little wiser about the identity of the automaton even after this time.

Of course, one can counter with the argument that while human learning does not cover all possibilities as exhaustively as Bayesian learning, it is sufficiently sophisticated to be described as 'Bayes-like'. After all, one could argue, we are rather good at many types of pattern recognition (otherwise, for example, you would not be able to read this text!). Although we do not know exactly how our brains recognise patterns so readily, it seems safe to assume that the learning involved can be approximated as Bayesian, even when learning the identity of quite complex strategies.

Since we do not actually know the 'learning algorithm' of the brain, and what its limitations are, this may be true. However, it seems rather unlikely. While we obviously have the mental hardware to recognise complex visual and aural patterns, it is not so obvious that we have the ability to recognise complex behavioural patterns like those here. Even trained statisticians find the sort of calculations involved difficult and frequently make mistakes—as decades of dubious time-series analysis amply testifies.

This raises the intriguing possibility that we may be able to place an upper bound on the complexity of the strategies that players will successfully copy in this evolutionary context. With the Moore machines we are using here to represent strategies, the natural division suggested by the simulations



*Figure 5.3. The two-state automaton 'Pavlov'. The initial state is irrelevant to payoffs from matches between clumsy automata over an infinite horizon.*

of learning would seem to be to restrict populations to two-state automata. That is, we do not *insist* that all players use two-state automata all the time, just that more complex automata are so difficult to copy that they rapidly die out.

The study of populations of two-state stochastic automata is not new. Simulations conducted by Nowak and Sigmund (1993b) suggest that, at least under replicator dynamics, such populations converge to a population consisting of a single stochastic automaton which has become known in biological circles as 'PAVLOV', following Kraines and Kraines (1988). It is the clumsy counterpart of the deterministic automata that Binmore and Samuelson (1992) call TAT-FOR-TIT and TWEEDLEDUM. PAVLOV tends to stick with its current action if it receives a payoff of  $R$  or  $T$ , and to switch actions if it receives anything less (see Figure 5.3).

However, we might wonder whether in practice people actually use fully stochastic strategies implied by the Nowak and Sigmund simulations. How-

ever, restricting the strategy space to  $\mathcal{M}_2$ , Nowak, Sigmund, and El-Sedy (1995) obtain the same result. That is, PAVLOV is the only Evolutionarily Stable Strategy (ESS) of such an automaton selection game provided that  $2R > T + P$ . (Nowak and Sigmund (1993a) show that for  $2R \leq T + P$  one may get complex cycles or even chaotic behaviour.) A monomorphous population of PAVLOV is thus clearly externally stable under the weaker definition of evolutionary equilibrium in section 5.2.3, so long as mutants are only taken from  $\mathcal{M}_2$ .

To summarize, then. The problem of copying fidelity makes the evolution of populations of deterministic automata extremely complex, but we may effectively rule out the possibility of a utilitarian equilibrium. It therefore seems reasonable to suppose some mechanism that improves copying fidelity. One possibility is to suppose that the automata retain some degree of 'clumsiness'. However, complex automata remain very difficult to identify even if they are clumsy. This suggests that we may impose some upper bound on the complexity of strategies used by players. In one such restriction of the strategy space, a population of players using the PAVLOV strategy stands out as an equilibrium.

However, before claiming that PAVLOV is somehow 'the solution' to the repeated prisoners' dilemma—in a similar manner to the claim made by Axelrod (1984) that TIT-FOR-TAT is the only strategy worthy of attention—some difficulties may be noted:

First, the Nowak, Sigmund and El-Sedy results are interesting but they need to be interpreted with some caution. In calculating the payoffs between clumsy automata, Nowak et al have implicitly taken the discount rate in the repeated game to zero before taking the operating mistake to zero. This means, for example, that they can ignore the automatas' initial states. Given that we are using infinitely repeated, undiscounted games as an approximation of long, finitely repeated games with low discounting, it might be more appropriate to take these limits the other way round, in which case the initial states might matter.

Secondly, we might like to consider whether the automaton selection game would remain valid if one were to endogenise the operating mistakes of clumsy automata—*i.e.* consider some higher-order evolutionary process selecting the value of  $\varepsilon$ . In particular, there is a slight payoff advantage from using TAT-FOR-TIT or TWEEDLEDUM against PAVLOV. On the other hand, TAT-FOR-TIT and TWEEDLEDUM often cannot be copied faithfully. How do these opposing evolutionary forces interact in practice, and what is the resulting value of  $\varepsilon$ ?

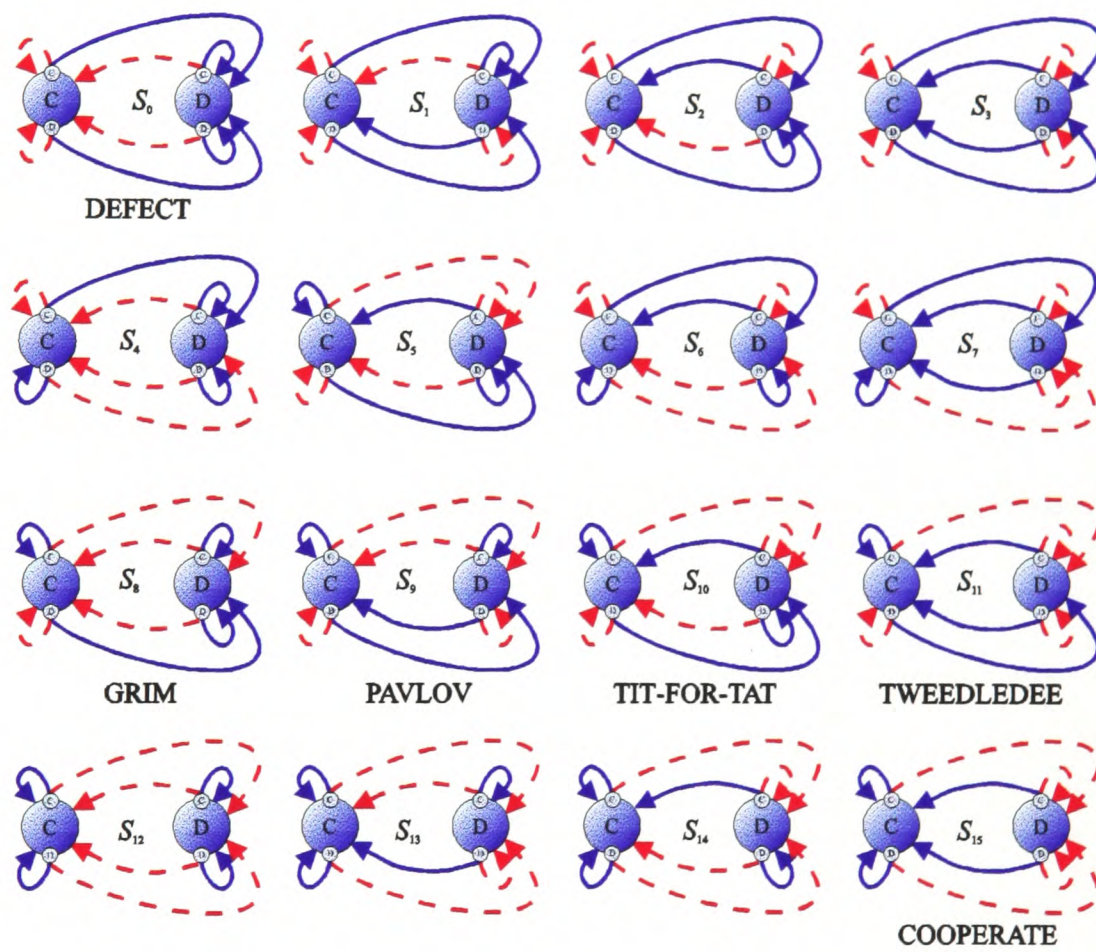
Lastly, we have been considering Moore machines here, but there is obviously more than one way to represent a repeated game strategy algorithmically. Different representations might result in a different restriction of the strategy set to ensure copying fidelity. Even here, some may not be convinced that the upper bound on complexity should be two states. After all, some three-state automata, such as TIT-FOR-TWO-TATS, can be identified from quite short segments of action sequence, even if the computational demands are high. Even amongst the two-state automata, there are differences in the difficulty of identifying an automaton, depending on the structure of its transition function.

There is certainly scope for further work along these lines. Especially interesting would be an experimental study of the ability of real subjects to copy and implement repeated game strategies given only action sequences. This would provide very useful data on how to best represent repeated game strategies algorithmically, and the appropriate complexity restriction to impose.

## APPENDIX

The full set of two-state stochastic automata are shown in Figure 5.4. The method used to calculate supergame payoffs between two-state stochastic automata as  $\varepsilon \rightarrow 0$  is described in Nowak, Sigmund, and El-Sedy (1995). Nowak et al consider the four possible action sequences that may be generated by two purely deterministic two-state Moore machines, depending on their initial states. By considering how perturbations result in movements between these sequences, they are then able to calculate how often the machines will be in each sequence and hence the expected supergame payoff. Taking  $T = 3, R = 2, P = 0$  and  $S = -1$ , this gives the automaton selection game payoff matrix given in the tables below.

The method could also be used to calculate supergame payoffs when players use deterministic machines but make errors when observing their opponent's moves.



*Figure 5.4. The full set of two-state stochastic automata.*

	$s_0$	$s_1$	$s_2$	$s_3$	$s_4$	$s_5$	$s_6$	$s_7$
$s_0$ ('DEFECT')	0	$\frac{3}{2}$	0	$\frac{3}{2}$	1	3	$\frac{3}{2}$	3
$s_1$	$-\frac{1}{2}$	1	$\frac{2}{3}$	1	$\frac{1}{5}$	$1\frac{2}{3}$	3	2
$s_2$	0	$\frac{2}{3}$	$\frac{1}{2}$	1	0	$1\frac{2}{3}$	0	$1\frac{2}{3}$
$s_3$	$-\frac{1}{2}$	1	1	1	$-\frac{1}{2}$	1	1	1
$s_4$	$-\frac{1}{3}$	1	0	$1\frac{1}{2}$	$\frac{1}{2}$	3	1	3
$s_5$	-1	$\frac{1}{3}$	$\frac{1}{3}$	1	-1	1	1	$1\frac{2}{3}$
$s_6$	$-\frac{1}{2}$	-1	0	1	$-\frac{1}{3}$	1	1	$1\frac{2}{3}$
$s_7$	-1	0	$\frac{1}{3}$	1	-1	$\frac{1}{3}$	$\frac{1}{3}$	1
$s_8$ ('GRIM')	0	$1\frac{1}{2}$	0	$1\frac{1}{2}$	1	3	2	3
$s_9$ ('PAVLOV')	$-\frac{1}{2}$	$\frac{1}{3}$	$\frac{2}{3}$	1	$\frac{1}{5}$	1	3	3
$s_{10}$ ('TIT-FOR-TAT')	0	$\frac{2}{3}$	$\frac{2}{3}$	1	0	1	1	$1\frac{1}{3}$
$s_{11}$ ('TWEEDLEDEE')	$-\frac{1}{2}$	$\frac{1}{3}$	1	1	$-\frac{1}{2}$	$\frac{1}{3}$	$1\frac{1}{3}$	$1\frac{1}{3}$
$s_{12}$	$-\frac{1}{2}$	$\frac{1}{4}$	$\frac{1}{4}$	1	0	1	1	$1\frac{3}{4}$
$s_{13}$	-1	-1	$\frac{1}{2}$	$\frac{1}{2}$	-1	-1	1	1
$s_{14}$	$-\frac{2}{3}$	-1	$\frac{2}{5}$	$\frac{1}{2}$	$-\frac{2}{3}$	-1	$\frac{2}{5}$	$\frac{1}{2}$
$s_{15}$ ('COOPERATE')	-1	-1	$\frac{1}{2}$	$\frac{1}{2}$	-1	-1	$\frac{1}{2}$	$\frac{1}{2}$

Table 5.3: Supergame Payoffs to Row against  $s_0$  to  $s_7$  as  $\varepsilon \rightarrow 0$

	$s_8$	$s_9$	$s_{10}$	$s_{11}$	$s_{12}$	$s_{13}$	$s_{14}$	$s_{15}$
$s_0$ ('DEFECT')	0	$1\frac{1}{2}$	0	$1\frac{1}{2}$	$1\frac{1}{2}$	3	2	3
$s_1$	$-\frac{1}{2}$	$1\frac{2}{3}$	$\frac{2}{3}$	$1\frac{2}{3}$	$1\frac{1}{4}$	3	3	3
$s_2$	0	$\frac{2}{3}$	$\frac{2}{3}$	1	$1\frac{1}{4}$	$2\frac{1}{2}$	2	$2\frac{1}{2}$
$s_3$	$-\frac{1}{2}$	1	1	1	1	$2\frac{1}{2}$	$2\frac{1}{2}$	$2\frac{1}{2}$
$s_4$	$-\frac{1}{3}$	1	0	$1\frac{1}{2}$	$1\frac{1}{3}$	3	2	3
$s_5$	-1	1	1	$1\frac{2}{3}$	1	3	3	3
$s_6$	$-\frac{2}{3}$	-1	1	$1\frac{1}{3}$	1	$1\frac{4}{5}$	2	$2\frac{1}{2}$
$s_7$	-1	-1	$1\frac{1}{3}$	$1\frac{1}{3}$	$\frac{3}{4}$	$1\frac{4}{5}$	$2\frac{1}{2}$	$2\frac{1}{2}$
$s_8$ ('GRIM')	0	$1\frac{3}{5}$	0	$1\frac{3}{5}$	$1\frac{1}{3}$	$2\frac{2}{3}$	2	$2\frac{2}{3}$
$s_9$ ('PAVLOV')	0	2	1	2	1	$2\frac{1}{3}$	$2\frac{2}{3}$	$2\frac{1}{2}$
$s_{10}$ ('TIT-FOR-TAT')	0	1	1	$1\frac{1}{3}$	1	2	2	2
$s_{11}$ ('TWEEDLEDEE')	0	2	1	$1\frac{1}{3}$	1	2	2	2
$s_{12}$	0	1	1	$1\frac{3}{4}$	1	2	2	$2\frac{1}{2}$
$s_{13}$	0	1	2	2	$\frac{2}{3}$	$1\frac{1}{2}$	$2\frac{1}{3}$	$2\frac{1}{3}$
$s_{14}$	0	0	2	2	$\frac{2}{3}$	1	2	2
$s_{15}$ ('COOPERATE')	0	$\frac{1}{2}$	2	2	$\frac{1}{2}$	1	2	2

Table 5.4: Supergame Payoffs to Row against  $s_8$  to  $s_{15}$  as  $\varepsilon \rightarrow 0$

## *Chapter 6*

### *Conclusion*

The aim of this thesis has been to show that application of the evolutionary analogy *can* add to our understanding of long-run economic processes if it done carefully and selectively.

Part I was concerned with the evolution of technology. There we saw that one way of thinking about research and development is to recognise that firms are trying to solve particular design problems. We often build these design problems into our models, but are forced to oversimplify them in order to make the models solvable. The approach taken there was to acknowledge that design problems are often insoluble using standard techniques and to model instead the process by which firms solve them. Two such processes were simulated in detail. The first, individual experimental search, was based on a problem-solving technique known as simulated annealing. The second, partial imitation, involved learning at a social level and was based on a problem-solving technique known as the genetic algorithm. Some economic implications of these processes were explored, including their application to stochastic learning curves and patent design. A further application was the potential importance of 'technodiversity' in the introduction of new technology to developing countries. Just as an ecosystem rich in biodiversity has access to a diverse pool of genetic material, an economy rich

in technodiversity has access to a diverse range of alternative technologies. Some models were studied in which innovation was linked to technology diffusion in such a way that technodiversity was seen to encourage indigenous development.

Part II was concerned with the evolution of economic behaviour. It was argued that the study of evolutionary game theory in economics helps to identify economic contexts where it is a more appropriate modelling method than methods that assume rational choice—and where it is not. However, in contexts where it is an appropriate modelling method, it was seen that the results obtained with evolutionary game theory are highly sensitive to how one models the learning mechanism driving the results, and how this integrates with the context in question. It was argued that the most commonly used learning mechanisms and equilibrium concepts used by economists applying evolutionary game theory are too highly derivative of biological mechanisms that assume a biological context. Progress in the application of evolutionary game theory to economics therefore requires flexible and creative modelling of context-appropriate mechanisms by theorists, supported by careful empirical and experimental work on the adaptive behaviour of real subjects. This was pursued in several examples dealing with the issues of matching and partnership formation, imperfect imitation in an extensive game, and payoff monotonicity in the context of coordination games. The issue of imperfect imitation was expanded upon in an evolutionary model where a population of players chose finite automata to play the infinitely repeated prisoners' dilemma. The players attempted to imitate automata that have proved successful in previous generations in such a way that population proportions followed the replicator dynamics. This imitation was often imperfect, as they did not observe the structure of automata directly, only action sequences. When the automata were deterministic, the problem of copying fidelity prevented a utilitarian equilibrium. However, if the players occasionally made operating mistakes when using their automata, they were able to be copied faithfully. Faithful imitation of complex automata is essentially impractical

but particularly easy in populations of simple 'clumsy' automata.

There are many ways in which the above research could be expanded. For example, the work on technology diffusion and development could be extended to cover situations where the quality or 'appropriateness' of a new technology was uncertain. An agency introducing the new technology could use the diffusion system as a kind of 'evolutionary filter'—a feedback mechanism that the agency could use to update its own assessment of the technology. Indeed, although a start is made in chapter 3 into the implications of 'technodiveristy' for technology transfer, there are obviously a large number of ways in which the models there could be made more sophisticated or more specific. There is certainly a huge potential for interesting empirical work on the possible relationship between diffusion and innovation.

Extending the work on partnerships discussed above, it should be possible to construct more general models of group formation. As usual, one would consider populations of players learning behaviour by some simple mechanism; now, however, they also learn which groups of other players to interact with and from whom to learn. Ultimately, research along these lines should start to provide insights into why certain socio-economic institutions have evolved rather than others. In particular, depending on the underlying game, we might expect a population of players to endogenously fragment into groups of a certain size. The emphasis on variable group size in models like this would provide an ideal opportunity to move away from the rather simplistic  $2 \times 2$  games most frequently considered in evolutionary game theory to consider games with a more general structure. Because of the very large number of ways of modelling this population fragmentation effect, it may be necessary to link dynamic models of this kind to specific economic applications—no bad thing, of course. In particular, it may be possible to construct models that test the evolutionary stability of certain market structures.

## *Bibliography*

- ABREU, D., AND A. RUBINSTEIN (1988): "The structure of Nash equilibrium in repeated games with finite automata," *Econometrica*, 56, 1259–1282.
- ARIFOVIC, J. (1994): "Genetic Algorithm Learning and the Cobweb Model," *Journal of Economic Dynamics and Control*, 18, 3–28.
- ARROW, K. J. (1962): "The Economic Implications of Learning by Doing," *Review of Economic Studies*, 29, 155–173.
- AXELROD, R. (1984): *The Evolution of Cooperation*. New York, Basic Books.
- (1987): "The Evolution of Strategies in the Iterated Prisoners' Dilemma," in *Genetic Algorithms and Simulated Annealing*, ed. by L. Davis. Morgan Kaufman.
- BERGIN, J., AND B. LIPMAN (1996): "Evolutionary State Dependant Mutations," *Econometrica*, 64, 943–956.
- BHALLA, A. S., AND D. D. JAMES (1986): "Technological Blending: Frontier Technology in Traditional Economic Sectors," *Journal of Economic Issues*, 20, 453–462.

- BINMORE, K. G. (1987): "Modelling Rational Players: Part I," *Economics and Philosophy*, 3, 179–214.
- BINMORE, K. G., J. GALE, AND L. SAMUELSON (1993): "Learning to be Imperfect: The Ultimatum Game," Lecture delivered at the Nobel Symposium, Björkborn, Sweden.
- BINMORE, K. G., AND L. SAMUELSON (1992): "Evolutionary Stability in Repeated Games Played by Finite Automata," *Journal of Economic Theory*, 57, 278–305.
- (1995): "Evolutionary Drift and Equilibrium Selection," SSRI Working Paper 9529, University of Wisconsin-Madison.
- BOLTON, G. E., AND R. ZWICK (1995): "Anonymity versus Punishment in Ultimatum Bargaining," *Games and Economic Behaviour*, 10, 95–121.
- BÖRGERS, T. (1996): "On the Relevance of Learning and Evolution to Economic Theory," mimeo, University College London.
- BÖRGERS, T., AND R. SARIN (1994): "Learning Through Reinforcement and Replicator Dynamics," mimeo, University College London.
- (1997): "Naive Reinforcement Learning with Endogenous Aspirations," mimeo.
- BOYD, R., AND P. J. RICHARSON (1985): *Culture and the Evolutionary Process*. Chicago Press.
- BOYLAN, R. T. (1992): "Laws of Large Numbers for Dynamical Systems with Randomly Matched Individuals," *Journal of Economic Theory*, 57, 473–504.
- BUDD, C., C. J. HARRIS, AND J. S. VICKERS (1993): "A Model of the Evolution of Duopoly: Does the Asymmetry Between Firms Tend to Increase or Decrease?," *Review of Economic Studies*, 60, 543–573.

- BURKE, J. (1978): *Connections*. Macmillan London Limited.
- BUTLER, S. (1872): *Erewhon*. Penguin Classics (1985).
- COOPER, B., AND C. WALLACE (1996): "The Evolution of Partnerships," mimeo, University of Oxford.
- (1997): "Partnerships and Recoverability," mimeo, University of Oxford.
- COOTES, B. (1990): *The Trade Trap*. Oxfam Publications.
- DARWIN, C. (1859): *The Origin of Species by Means of Natural Selection*. Penguin Classics (1985).
- DAWKINS, R. (1976): *The Selfish Gene*. Oxford University Press.
- ELLISON, G. (1995): "Basins of Attraction and Long Run Equilibria," mimeo, Department of Economics, MIT.
- ELY, J. (1995): "Local Conventions," mimeo, University of California at Berkeley.
- FRIEDMAN, M. (1953): *Essays in Positive Economics*. University of Chicago Press.
- FUDENBERG, D., AND D. K. LEVINE (1993): "Steady State Learning and Nash Equilibrium," *Econometrica*, 61, 547–573.
- GILBERT, R., AND C. SHAPIRO (1990): "Optimal Patent Length and Breadth," *RAND Journal of Economics*, 21, 106–112.
- GOLDBERG, D. E. (1989): *Genetic Algorithms in Search Optimization and Machine Learning*. Addison-Wesley.
- GROSSMAN, G. M., AND E. HELPMAN (1991): "Quality Ladders in the Theory of Growth," *Review of Economic Studies*, 58, 43–61.

- HARRIS, C. J. (1988): "Dynamic Competition for Market Share: An Undiscounted Model," Nuffield College Discussion Paper No. 30.
- HOLLAND, J. H., K. J. HOLYOAK, R. E. NESBETT, AND P. R. THAGARD (1986): *Induction: Processes of Inference, Learning and Discovery*. MIT Press.
- ISLAM, R. (1992): "Transfer, Adoption and Diffusion of Technology for Small and Cottage Industries: An Overview.," in *Transfer, Adoption and Diffusion of Technology for Small and Cottage Industries.*, ed. by R. Islam. ILO-Artep.
- KALAI, E., AND E. LEHRER (1990): "Rational Learning Leads to Nash Equilibrium," Discussion Paper No. 895, The Center for Mathematical Studies in Economics and Management Science, Northwestern University.
- KAUFFMAN, S. A. (1988): "The Evolution of Economic Webs," in *The Economy as a Complex Evolving System*, ed. by K. Arrow, P. W. Anderson, and D. Pines, pp. 125–146. SFI Studies in the Sciences of Complexity, Addison-Wesley, N. Y.
- KAUFFMAN, S. A., AND W. MACREADY (1995): "Technological Evolution and Adaptive Organisations," Discussion paper, Santa Fe Institute Working Paper No. 95-02-008.
- KIRKPATRICK, S., C. D. GELATT, AND M. P. VECCHI (1983): "Simulated Annealing," *Science*, 220, 671–680.
- KLEMPERER, P. D. (1990): "How Broad Should the Scope of Patent Protection Be?," *RAND Journal of Economics*, 21, 113–130.
- KRAINES, D., AND V. KRAINES (1988): "Pavlov and the Prisoners' Dilemma," *Theory and Decision*, 26, 47–79.
- LEE, T., AND L. L. WILDE (1980): "Market Structure and Innovation: A Reformulation," *The Quarterly Journal of Economics*, 94, 429–436.

- LEWIS, D. (1968): *Conventions: A Philosophical Study*. Harvard University Press.
- LINDGREN, K. (1990): "Evolutionary Phenomena in Simple Dynamics," in *Artificial Life II*. The Santa Fe Institute.
- MAYNARD SMITH, J. (1982): *Evolution and the Theory of Games*. Cambridge University Press.
- MAYNARD SMITH, J., AND G. R. PRICE (1973): "The logic of animal conflict," *Nature*, 246, 15–18.
- NÖLDEKE, G., AND L. SAMUELSON (1993): "An Evolutionary Analysis of Backward and Forward Induction," *Games and Economic Behaviour*, 5, 425–253.
- (1995): "A Dynamic Model of Equilibrium Selection in Signaling Markets," mimeo.
- NORDHAUS, W. D. (1969): *Invention, Growth and Welfare; A Theoretical Treatment of Technological Change*. Cambridge, Mass.
- NORMAN, M. F. (1972): *Markov Processes and Learning Models*. Academic Press, New York and London.
- NOWAK, M. A., AND K. SIGMUND (1993a): "Chaos and the Evolution of Cooperation," *Proc. Natl. Acad. Sci.*, 90, 5091–5094.
- (1993b): "A Strategy of Win-Stay, Lose-Shift that Outperforms Tit-for-Tat in the Prisoner's Dilemma Game," *Nature*, 364, 56–58.
- NOWAK, M. A., K. SIGMUND, AND E. EL-SEDY (1995): "Automata, Repeated Games and Noise," *Journal of Mathematical Biology*, 33, 703–722.
- PICCIONE, M. (1992): "Finite automata equilibria with discounting," *Journal of Economic Theory*, 56, 180–193.

- PRESS, W. H., W. T. VETTERLING, S. A. TEUKOLSKY, AND B. P. FLANNERY (1994): *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, second edn.
- REDDING, S. (1995): "Economic Miracles and the Low-level Development Trap," mimeo, University of Oxford.
- ROMER, P. M. (1990): "Endogenous Technological Change," *Journal of Political Economy*, 98, S71–S102.
- ROTH, A. E. (1995): "Bargaining Experiments," in *Handbook of Experimental Economics*, ed. by J. Kagel, and A. E. Roth. Princeton University Press.
- RUBINSTEIN, A. (1986): "Finite Automata Play the Repeated Prisoners' Dilemma," *Journal of Economic Theory*, 39, 83–96.
- SCHERER, F. M. (1972): "Nordhaus' Theory of Optimal Patent Life: A Geometric Reinterpretation," *American Economic Review*, 62, 422–427.
- SCHUMACHER, E. F. (1972): *Small is Beautiful*.
- SELTEN, R. (1975): "Re-examination of the perfectness concept for equilibrium points in extensive games," *International Journal of Game Theory*, 4, 25–55.
- (1991): "Evolution, Learning and Economic Behaviour," *Games and Economic Behaviour*, 3, 3–24.
- SIMON, H. A. (1973): "The Structure of Ill-Structured Problems," *Artificial Intelligence*, 4, 181–201.
- STEWART, F. (1990): "Technology Transfer for Development," in *Science and Technology: Lessons for Development Policy*, ed. by R. E. Evenson, and G. Ranis. IT Publications.

STONEMAN, P., AND M. KARSHENAS (1995): "Technological Diffusion," in *Handbook of the Economics of Innovation and Technological Change*, ed. by P. Stoneman. Blackwell.

SWINKELS, J. (1992): "Evolutionary stability with equilibrium entrants," *Journal of Economic Theory*, 57, 306–332.

THALER, R. H. (1988): "The Ultimatum Game," *Journal of Economic Perspectives*, 2, 195–206.

TIROLE, J. (1988): *The Theory of Industrial Organisation*. The MIT Press.

WEIBULL, J. W. (1995): *Evolutionary Game Theory*. The MIT Press.

