ECONOMIC NETWORKS:
COMMUNICATION, COOPERATION & COMPLEXITY

A thesis submitted for the degree of
Doctor of Philosophy

Simon D. Angus

Supervisor: Dr G. Bose
Co-supervisor: A. Prof K. Meagher

School of Economics,
The University of New South Wales.

2006

Authorship Statement

I hereby declare that this submission is my own work and to the best of my knowledge it contains no materials previously published or written by another person, nor material which to a substantial extent has been accepted for the award of any other degree or diploma at UNSW or any other educational institution, except where due acknowledgement is made in the thesis. Any contribution made to the research by others, with whom I have worked at UNSW or elsewhere, is explicitly acknowledged in the thesis.

I also declare that the intellectual content of this thesis is the product of my own work, except to the extent that assistance from others in the project's design and conception or in style, presentation and linguistic expression is acknowledged.

S. D. Angus

Signed: ....

Date: ....

Copyright Statement

S. D. Angus

Signed: ....

Date: ....

Authenticity Statement

I certify that the Library deposit digital copy is a direct equivalent of the final officially approved version of my thesis. No emendation of content has occurred and if there are any minor variations in formatting, they are the result of the conversion to digital format.

S. D. Angus

Signed: ....

Date: ....

Abstract

This thesis is concerned with the analysis of economic network formation. There are three novel sections to this thesis (Chapters 5, 6 and 8). In the first, the non-cooperative communication network formation model of Bala and Goyal (2000) (BG) is re-assessed under conditions of no inertia. It is found that the Strict Nash circle (or wheel) structure is still the equilibrium outcome for $n = 3$ under no inertia. However, a counter-example for $n = 4$ shows that with no inertia infinite cycles are possible, and hence the system does not converge. In fact, cycles are found to quickly dominate outcomes for $n > 4$ and further numerical simulations of conditions approximating no inertia (probability of updating $\gtrsim 0.8 \rightarrow 1$) indicate that cycles account for a dramatic slowing of convergence times.

These results, together with the experimental evidence of Falk and Kosfeld (2003) (FK) motivate the second contribution of this thesis. A novel artificial agent model is constructed that allows for a vast strategy space (including the Best Response) and permits agents to learn from each other as was indicated by the FK results. After calibration, this model replicates many of the FK experimental results and finds that an externality exploiting ratio of benefits and costs (rather than the difference) combined with a simple altruism score is a good proxy for the human objective function. Furthermore, the *inequity aversion* results of FK are found to arise as an *emergent* property of the system.

The third novel section of this thesis turns to the nature of network formation in a trust-based context. A modified Iterated Prisoners' Dilemma (IPD) model is developed which enables agents to play an additional and costly network forming action. Initially, canonical analytical results are obtained despite this modification under uniform (non-local) interactions. However, as agent network decisions are 'turned on' persistent cooperation is observed. Furthermore, in contrast to the vast majority of non-local, or static network models in the literature, it is found that a-periodic, complex dynamics result for the system in the long-run. Subsequent analysis of this regime indicates that the network dynamics have fingerprints of self-organized criticality (SOC). Whilst evidence for SOC is found in many physical systems, such dynamics have been seldom, if ever, reported in the strategic interaction literature.

Acknowledgements

I have waited till the very end to make my thanks known, although they shall now sit appropriately at the beginning of this work. I have delayed not for reasons of confusion or uncertainty as to whom I owe a debt of thanks – as will be shown below, I have been keeping notes! – but only because for me, to write down my thanks is to re-visit in my mind's eye those who have been responsible for so much kindness, wisdom and care towards me over the course of the past several years. And to have taken this turn prematurely would have been to indulge in the sweetmeats before the last pea was had.

So to begin, I believe appropriately, with those I have never met. To the many who developed the Linux-(K)ubuntu operating system; to Donald Knuth the developer of the original plain TeX and thereafter Leslie Lamport and others who contributed to the beautiful, unfailing document preparation language LaTeX 2_ε in which I currently write, to the developers of packages such as ps-tricks and beamer, to the earliest developers of the text editor vi and now vim in which I currently type, and to the whole barely controlled but effervescent open-source software community ... thank you, thank you, thank you! My experience with computers owes so much to your benevolent expertise and the philosophy that drives it.

To the Australian Government's *Australian Post-graduate Award (APA)* scheme and prior to this, to the now defunct but brilliantly inspired independent scholarship program known as *The Southern Cross Foundation Scholarship*, my sincere thanks for this financial and academic support throughout my undergraduate and post-graduate studies. Additionally, the Faculty of Commerce and Economics at the UNSW deserve special thanks as they generously provided a further APA top-up stipend together with many opportunities to teach in the faculty and so gave not only financial aid, but also valuable and enjoyable experiences of the class-room. Moreover, their support through the

PhD program travel grant scheme made several enriching and stimulating national and international trips possible over the course of this research.

Of the other institutional bodies who arranged various conferences that I attended as part of this thesis, I would like to make mention of the Network Theory Working Group at the Commonwealth Science and Industrial Research Organization (CSIRO), Canberra, Australia; the New England Complex Systems Institute (NECSI), MA, USA; and the Santa Fe Institute (SFI), New Mexico, USA. The latter of which were not only extremely generous in their substantial support of my attendance at the Graduate Workshop on Computational Social Sciences and Complexity, but also have contributed greatly to my understanding of the fascinating area of complexity science. John Miller and Scott Page at the SFI in particular provided excellent input, advice and tuition over the stimulating time I spent residing in Santa Fe in 2004. The work contained in Chapters 7 and 8 were conceived and developed at this time.

Closer to home, to all the members of the School of Economics Theory Group, my thanks for their forbearance, comments and guidance of each of the projects presented herein. In particular, I am grateful to Gautam Bose, Bill Schworm, Kieron Meagher, Argyha Gosh, and Hodaka Morita whose continued questions and comments prompted this work to develop in a fruitful direction. Additionally, I would like to make note of correspondence with Sanjeev Goyal (Essex) concerning Chapter 5, Doyne Farmer (SFI) concerning non-linear dynamics and other matters complex, and very recently, Paul Frijters (QUT) whose electronic correspondence thus far foreshadows a vigorous discussion at the forthcoming Economic and Business PhD conference. I shall leave the many other academic correspondents who have shepherded me unnamed, but not forgotten. Further, my colleague PhD students, notably Lorraine Ivancic, Carmit Schwartz, Iqbahl Syed and Haque Nabin I thank for their friendship and often timely conversations, humour, determination and source of procrastination over this time.

It must be said, that for all of this academic help and guidance, all current errors, whether factual, computational, linguistic or other are my own, and the names mentioned above or below bear no such responsibility.

And finally, I turn to those closest to me but no doubt happily somewhat distant from the work itself. Amongst the many who I cannot name due to space but who know my gratefulness, a few I must pause to mention. My thanks for the friendship of Shane Cox and Simon Prendergast who comprised two-thirds of a very special household now disbanded. To the many brothers and sisters in Christ at Unichurch, CBC, SPUBS and the Graduate Bible Fellowship I am so thankful in the Lord for your encouragement, prayers and fellowship over this time. I did not perceive that undertaking PhD studies would require quite so much of me, but it is clear to me now, and should have been from the start, that the Lord works in all things, including the trivial distresses of PhD studies to grow his people in faithfulness and Godliness. To my parents and parents in-law and siblings, my apologies for allowing this endeavour to be a source of dis-harmony on so many occasions; stubbornness and fragility have contributed unhappily measure for measure. Nevertheless, my thanks to you all for your wisdom and love throughout.

And to end this lengthy meditation on the falseness of man as an isolate, I thank my closest friend and dearest wife, Susan. Part of me wouldn't urge on others marrying as we did amidst our doctoral studies, but the other part knows that had we not taken this step, and not begun this life so tender and nourishing, neither would have maintained the perspective that has ultimately sustained us. I've no idea how we shall relate after these things are over, but no doubt with His help, we will make do!

S.D.A.
Spring, 2006

I have never doubted the truth of signs, Adso; they are the only things man has with which to orient himself in the world. What I did not understand is the relation among signs... I behaved stubbornly, pursuing a semblance of order, when I should have known well that there is no order in the universe.
William of Baskerville, The Name of the Rose (Eco 1984)

Contents

List of Tables

List of Figures

part I

Introduction

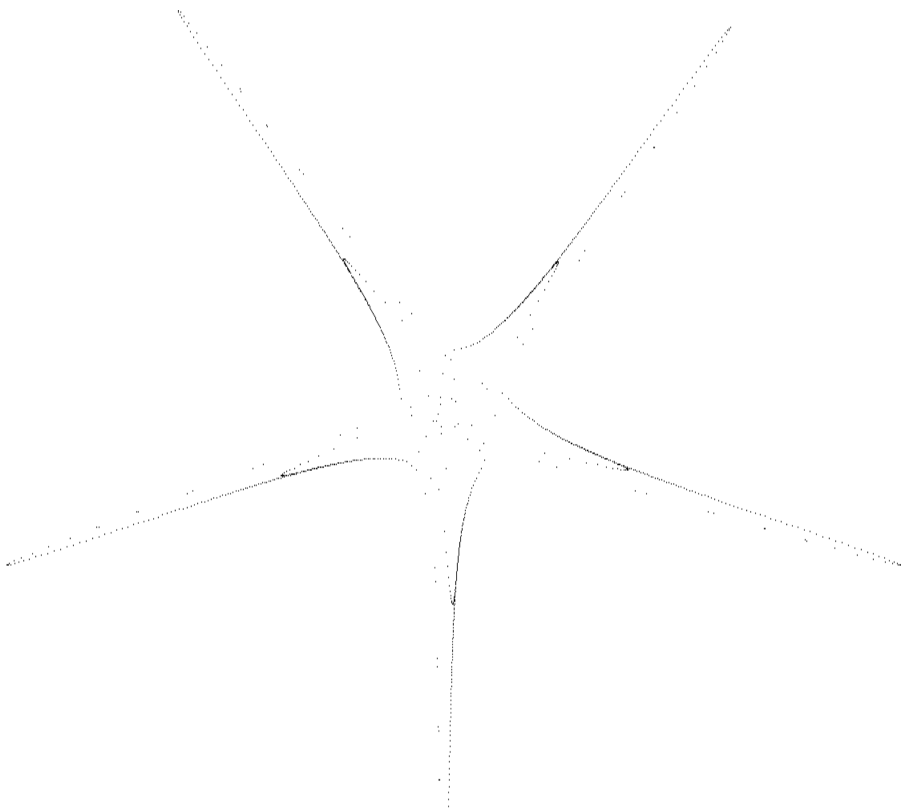

Chapter 1

Introduction to the Thesis

In this chapter, the major themes of this thesis will be introduced. Additionally, specific chapters that deal with particular aspects of economic networks will be refered to. A more complete summary and discussion of Results with an emphasis on future work is given in Chapter 9.

Networks are undoubtedly receiving intensive interest at present both in the scholarly literature and increasingly in the halls of government and the popular media. This new-found interest is due in large part to the realisation that traditional reductionist approaches to inquiry (understand the atom to understand the whole) do not, in fact, provide a rich enough framework of investigation for understanding the system in a variety of contexts. Instead, it has been recognised that it is the *relationships* between the atoms that must be understood in order that the system can be unmasked, modelled, and ultimately predicted. Such relationships, it has been found, can cause feedback that amplify and diminish individual atom activity.

But what of *economic* networks? Do they provide *relationship* driven contexts? And if so, how do social and business relationships develop? What motivates their establishment? How does the inter-connected network of relationships affect individuals or firms, or the whole population? Are there network structures, both local and global, that facilitate one kind of activity better than another? It is the purpose of this thesis to investigate these kinds of questions as they apply to two fundamental economic network contexts, those of *communication* on the one hand and *cooperation* on the other.

Indeed, it is the suggestion of this thesis, that not only are networks prominent in economic contexts, economic networks enforce a further layer of challenge for the inquirer over and above those of the physical sciences due to the unique property of *node-agency*; the nodes think, rationalise, and act. Some reflections on this suggestion and evidences for the importance and fore-going analysis of economic networks is presented in Chapter 2. Moreover, as will be explored in Chapter 3, both structurally and dynamically, economic networks fit well into an emerging field of inquiry known as *complexity science*, or *the science of complexity*. Here, the emphasis is not so much on the 'difficulty' of the problem but rather on the kind of system that the problem presents, such systems being classed as *complex systems*, and are associated with phenomena such as *self-organization*, *emergence* and *autonomy*.

The following three chapters of this thesis (Part II) deal with a model of communication network formation. Chapter 4 introduces a prominent and insightful analytical model of non-cooperative network formation, namely that of Bala and Goyal. This model showed which equilibrium network structures can be expected to form as a result of self-interested, utility maximizing individuals who use a best response rule to unilaterally construct a communication network. Chapter 5 then re-visits this model but considers what might happen if agents do not exhibit *strategic inertia*, as was assumed by Bala and Goyal. Both analytical and numerical simulations indicate that the best-response updating rule leads to non-convergent two-period cycles under these conditions. Furthermore, such cycling is found to be responsible for dramatically increasing the time to convergence of the best-response model when strategic inertia is approximately small. It is this result, together with the experimental work of Falk and Kosfeld that motivates Chapter 6 where the best-response updating rule is relaxed, and in its place *artificial adaptive agents* are constructed to play the non-cooperative network formation game. By utilising a novel cognition and learning architecture, this approach pleasingly replicates many of the results from the human trial, including several features that the Bala and Goyal model did not predict. It also finds that the suggested explanations for non-Nash human behaviour in the two-way information flow case need not be modeled explicitly, but rather arise as *emergent* properties of the dynamic learning-by-imitation environment. Taken together, these chapters suggest useful ways that the non-cooperative communication network framework of Bala and Goyal can be extended and improved.

The following two chapters (Part III), turn to the *cooperation* network context. Firstly, a more detailed survey of how cooperation has been investigated and modeled is given in Chapter 7. This section draws insight from both the biological beginnings of the study of cooperation (or *altruism*) and also from models of cooperation in the economic sciences, especially focusing on models that have attempted to incorporate some kind of interaction network. This chapter suggests that a historical line of inquiry can be drawn through the cooperation literature, beginning with random, or anonymous interactions, and then imposing at first simple (e.g. line, circle, or grid), and then complex (e.g. random, 'small-world' and 'scale-free' graphs), non-uniform interaction structures on agents. This gives rise to the obvious 'next-step' for the analysis of economic networks: the facility for agents *themselves* to construct and dynamically change the interaction setting.

It is the purpose of Chapter 8 to do just this, by making a simple modification to the much studied standard Prisoners' Dilemma set-up, allowing agents to play a costly signal that either (re-)establishes a link with a chosen opponent, or relegates them to a listing of undesirable opponents. This modelling set-up is analytically studied in the simplest case to check that it conforms to well-known results in the literature, and then is analysed computationally to yield interesting results on strategic network formation. Chief amongst these results are: first, that there are minimum criteria that a 'network' context must meet to enable otherwise selected against cooperative types to survive; and second, that for even relatively simple cases, the incorporation of a strategic network formation environment produces un-stable, out-of-equilibrium dynamics that persist in the long-run. Further analysis of this system finds evidence that relates it directly to the foregoing discussion on *emergence*, *self-organization* and *complexity*.

Finally, in Chapter 9, the main results of the thesis are emphasised and discussed, with particular attention given to how these results interact with the wider literature. Recent developments are also discussed and motivate discussions of interesting further work to be pursued as a result of the findings of this thesis.

Chapter 2

An Introduction to Economic Networks

Economic networks are now introduced. Specifically, a review of some analytical results on economic networks is given, tracing the handling of models in *uniform* interaction spaces at first, followed by simple and then complex interaction spaces. Examples of how 'networks' are handled in the economic literature is then given, with a concluding section on methods of analysis.

2.1 Economic Networks: order *sans* design

As crude measures go, networks are a popular topic on the planet. A recent poll of the `Google` internet search engine[1] returned 3.8 billion hits for 'network'.[2] But what is a 'network'? A network is suitably described as an 'interconnected group or system'.[3] Put simply, it expresses a *relationship* that exists between a number (more than one) of objects. Generally speaking, networks fall into one of two categories being either explicit, or implicit. The former are those that can be identified due to their physical presence (e.g. roads, rail, telecommunication lines, power-grids) whilst the latter are those that do not have physical connections, but nonetheless are identifiably 'an interconnected group or system' (e.g. sexual-partner networks, business networks, social networks, food-webs). An important difference between these two kinds of networks is that the former are in the main a product of *design*, and thus are largely fixed in place after construction, whereas the latter develop in an *ad-hoc*, *autonomous* way. Except in the most extreme cases,[4] this second class of networks have no central coordination. Individual nodes in these networks

[1] URL: `http://www.google.com`

[2] Incidentally, to calibrate this check, a search for 'war' gave 1.31 billion hits, whilst 'peace' returned only 578 million hits. Clearly, 'network' is up there in the (digital) world's consciousness.

[3] Source: `http://www.wiktionary.org/wiki/network`.

[4] The friendship relationships of cell-groups in a prison, for example, though even here it might be expected to see sub-groups of social interaction forming.

undertake explicit autonomous actions concerning relationships, or alternatively simply go about their daily behaviours, having only limited interaction with their immediate 'neighbours', oblivious to their rôle in the wider set of dependencies. Nevertheless, over time, these actions aggregate to give rise to enough discernible relationships that a 'network' can be observed and drawn (see §3.6 on *emergence*). At this point, after 'joining the dots' on these relationships, it may be possible to come to some understanding of the directions of influence and dependency that drive the overall system. Indeed, the structure of these relationships may in turn affect the individual behaviours themselves, further contributing to the structure of relationships, and so on. Often such causality leads to discernible system-wide characteristics over time; it is in this sense that autonomous networks can give rise to a sense of 'order *sans* design.'

Economic networks fall largely into this second class described above. In rough terms, an 'economic network' is a system in which a number (more than one) of discernible interactions between entities (e.g. firms, people, institutions), gives rise to the distribution or accumulation of value, character traits, or behaviours amongst them. Examples of economic contexts that fit this kind of description would include communications (information as value), social (satisfaction, fulfillment), political (opinions), firm-to-firm (trade) and cooperation/corruption (information, trade) networks. Furthermore, the way that the literature uses the term 'networks' to describe such systems gives rise to further diversification, from simple conceptions of a sub-set of actors on a 'list' (e.g. Kali (1999); Taylor (2000)), to extremely intricate computational models that attempt to capture every link and node in their description (e.g. Smucker et al. (1994)). However, in each case, these systems are characterised by largely *ad-hoc*, un-planned, and un-coordinated changes in associations and node identities. And indeed, it is these very characteristics that has lead many researchers to the study of economic networks as examples of so-called *complex systems* which I shall return to in detail below (§3.1).

2.2 Why study Economic Networks?

Studying economic networks is important for (at least) two main reasons. The first reason is also the simplest: 'networks matter'. Traditional economic problems conducted on a non-uniform interaction space (i.e. non-uniform probabilities of interaction between agents) can give rise to a reversal of traditional outcomes. This observation will be visited below in more detail, but for now, two examples serve to illustrate the point. In the first, Nowak and May (1992) set agents up on a two-dimensional lattice to play the simple Prisoner's Dilemma game with each of their eight neighbours[5], with scores for a round

[5] In clock-wise divisions of a magnetic compass that is: $\{N, NE, E, ..., W, NW\}$.

being simply the summation of payoffs against these eight neighbours. Using the simplest type of agents, who play either 'always Defect', or 'always Cooperate', and replacing each cell with the type who has the highest score out of the previous owner, and the neighbours, the authors find that in contrast to the canonical result that the fraction of always-defect would reach unity in finite time, they observe patches of cooperative play persisting indefinitely. Secondly, Anderlini and Ianni (1996) consider a model of learning in symmetric, two-player normal form games, situating agents in a similar way to that of Nowak and May, on a two-dimensional lattice, and allow agents to update their strategy based on a comparison between their average payoff from their present strategy and some aspiration level. Again, the traditional result of convergence to a single strategy under non-local interactions is overturned, with various prescriptions giving rise to 'mixed' strategy states under local-learning.

Such results urge a consideration of the way that the overall interaction structure affects the individual agent behaviour and vice versa. Indeed, this interplay between levels of activity: agents on the one hand, and *epi-phenomena* or *emergent* characteristics of the system on the other is the basis of a progression in recent scientific inquiry from pure reductionist analysis, to systems analysis, to the present mixing of the two under *complex systems* analysis (to be followed in detail below §3.2). Indeed, recent reflections on this point, such as that of Heylighen et al. (2004) and Arthur (1999) have noted that many systems, economic ones included, are particularly affected by such multiple scale interplay and interaction.[6]

Second, although incorporating networks (or non-uniform interaction spaces) into economic models is an analytically difficult enterprise, economic networks are ripe for study due to new developments in powerful computational tools.

It is true that progress in purely analytical inquiries of networks as *Graphs* (see below) is not new, especially when the inquirer can estimate graph topology via statistical means (see Erdös and Rènyi (1959)), but the difficulty comes when an attempt is made to place autonomous decision makers in such complex interaction (and/or information) spaces; the 'homogeneous agent' assumption of many economic models necessarily disappears, often taking with it, the power of neat analytical description. Indeed, it has been found that these richer systems regularly exhibit *non-equilibrium* phenomena, with associated dynamics that defy easy description, consequently requiring a literal running of the system from chosen initial conditions to some future point (often many times over) to facilitate

[6]See also Holland (1998), chapter 10.

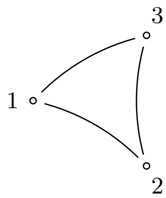

Figure 2.1 Example graph $\Gamma_a(V, E)$.

meaningful characterisation (Markose, 2005). Of course, there has been progress in analytical formulations nonetheless[7] but increasingly, the most interesting technical results in networked agent behaviour are arising in quasi- or full- computational descriptions of these systems.[8] A large part of the explanation for this phenomena is the reduction in costs of computing power and the increasing number of programming languages and development environments amenable to conducting computational, agent-based experiments.[9] Taken together, computational modelling of these previously difficult to study systems is becoming increasingly feasible for any researcher with a desktop computer and a rudimentary technical background in programming it.

2.3 Excursus on Terminology

Some basic terminology will be useful for what follows in the rest of this thesis. Networks are equivalent to *graphs* in the Theory of Graphs.[10] A graph $\Gamma(V, E)$ comprising n *vertexes* and k *edges* is completely defined by a vertex (node) set $V = \{v_1, v_2, \ldots, v_n\}$ and an edge (link) set $E = \{e_1, e_2, \ldots, e_k\}$ where each e_i is an unordered pair (v_i, v_j) representing an edge between the two elements of V (e.g. (a, b) means, 'a is joined to b'). For example, the graph $\Gamma_a(V, E)$ where $V = \{1, 2, 3\}$ and $E = \{(1, 2), (2, 3), (3, 1)\}$ is shown in Fig. 2.1. Two nodes or edges are said to be *adjacent* if they are joined by a shared edge, or incident at a shared node respectively. The number of incident edges at a node is called the *degree* of that node. The graph $\Gamma_a(V, E)$ is also *connected* since there exists an ordered edge vector describing a succession of adjacent edges (a *path*) from each node to every other node (the opposite case being *disjoint*). It is also an example of a *regular* graph since each node has a constant number of adjacent edges. Generally speaking, *regular* graphs are denoted *k-regular* where k is the constant degree of each node (in this case, $\Gamma_a(V, E)$ is *2-regular*). Further, since there exists an edge from each

[7] A selection of such results includes Anderlini and Ianni (1996); Bala and Goyal (1998, 2000); Chwe (2000); Jackson and Watts (2002a,b), returned to in §2.4.

[8] For example, the two papers mentioned above as examples of non-standard results due to non-uniform interactions are both strictly computational in nature. See further examples in the many contributions to the proceedings volume, *The Economy as an Evolving Complex System II* (Arthur et al., 1997).

[9] See the vast listing maintained at `http://www.econ.iastate.edu/tesfatsi/acecode.htm`.

[10] See B. Bollobas, *Modern Graph Theory* (Bollobas, 1998) for a thorough introduction to the area, or, for an example-based, non-technical introduction, see Hayes (2000a,b).

vertex to every other vertex, $\Gamma_a(V,E)$ is a *complete* graph (if this were not true, it would be *incomplete*). Likewise, there is an *Hamiltonian cycle* (or 'circuit') in $\Gamma_a(V,E)$ since we can construct a path $H = \{e_1, \ldots, e_k\}$ with distinct edges of E that travels each edge once (for example, $H = \{(3,1),(1,2),(2,3)\}$ is one such path).

So far we have described an *un-directed* graph, where each edge can be travelled in either direction (for some edge set (a,b) feasible paths are $\{v_a, v_b\}$ *and* $\{v_b, v_a\}$). However, a second class of graphs, known as *directed* graphs, or *di-graphs* are constructed with one-way 'flows' between nodes. Thus E will comprise an *ordered* set of vertex pairs. In short-hand we shall denote a *directed* edge from node a to node b by $a \to b$ and an un-directed edge between the same nodes to be $a \leftrightarrow b$. Similarly, for directed graphs, the degree of a node can be diversified to the *in-* and *out-* degree of a node to describe the count of in-going and out-going edges at a node respectively.

For convenience, I shall assume a direct correspondence between the informal and formal languages of networks and graphs respectively, that is, between: network–graph; node–vertex; and link–edge.

2.4 Some Results from the Analytical Literature

Several workers have attempted to construct dynamic network models analytically, focusing predominately on either the *topology* of the network gained (e.g. star, cycle, empty) and/or, the *kind* of equilibrium (if any) that might persist in the actions of the individuals involved. I will not attempt to make a thorough survey here, rather, some of the influential contributions to this literature will be reviewed to get a feel for the nature of the results.

One of the central questions of the literature has concerned that of *equilibrium selection* in 2-player coordination games, and so I shall consider one general example in some detail to aid our discussion. Suppose $\mathcal{G}_1$ is such a game,

$$\mathcal{G}_1 = \begin{array}{c|cc} & s_1 & s_2 \\ \hline s_1 & a,a & b,c \\ s_2 & c,b & d,d \end{array} \tag{2.1}$$

the strategy set for each player being $S = \{s_1, s_2\}$ with payoffs as given in the table, $\{a, \ldots, d\} \in \mathbb{R}$.[11] Several variations of this game might be specified. For example, if $(a-c)(d-b) > 0$ then there is a dominant strategy; a standard coordination game requires ($a > c$, $d > b$); and a game having a unique mixed strategy equilibrium occurs

[11] The discussion and nomenclature given here derives from Kandori et al. (1993).

when $a < c$, $d < b$. Suppose that n players are randomly paired to play $\mathcal{G}_1$ once each in a period and the count of s_1 players is given by z, then it is straightforward to show that expected (average) payoffs $\pi_1(z)$ and $\pi_2(z)$ for s_1 and s_2 players respectively are given by,

$$\begin{aligned} \pi_1(z) &= \Big(\frac{z-1}{n-1}\Big)a + \Big(\frac{n-z}{n-1}\Big)b, \\ \pi_2(z) &= \Big(\frac{z}{n-1}\Big)c + \Big(\frac{n-z-1}{n-1}\Big)d. \end{aligned}$$

Consider the coordination version of (2.1), and let the two pure strategy equilibria be given by $E_1 = \{s_1, s_1\}$ and $E_2 = \{s_2, s_2\}$. Two related questions are immediately apparent: a) for some given population, *how* do the agents 'know' to coordinate on one of these equilibria? and further, 2) *which* of the two equilibria will be selected in the 'long-run'? The first is a question of convergence to any equilibrium, the second is a deeper question which concerns attempting to predict choice amongst equilibria.

2.4.1 'Trembling' Towards Agreement

In their influential contribution, Kandori et al. (1993) (KMR) considered the second question in detail. They assume that there exists a 'Darwinian' strategy updating function $b(z_t)$ such that,

$$z_{t+1} = b(z_t) \tag{2.2}$$

where,

$$\mathrm{sign}(b(z) - z) = \mathrm{sign}(\pi_1(z) - \pi_2(z)) \quad \text{for} \quad z \neq \{0, n\}. \tag{2.3}$$

That is, (2.3) specifies that the better performing of the two strategies, s_1 and s_2 (given z_t), will occur with higher incidence in the following period. This is a general implementation of the well-known *evolutionary dynamic*, which not surprisingly stemmed from the biological literature (see esp. Maynard Smith (1978) and Maynard Smith and Price (1973)). One such dynamic, which KMR name the 'best reply dynamic' is simply,

$$B(z) = \begin{cases} N, & \text{if} \quad \pi_1(z) > \pi_2(z), \\ z, & \text{if} \quad \pi_1(z) = \pi_2(z), \\ 0, & \text{if} \quad \pi_1(z) < \pi_2(z). \end{cases} \tag{2.4}$$

Further, a mistake-making probability for each agent ϵ, is specified. That is, they allow for agents to 'tremble' in updating their strategy, such that with some small probability, they will make the wrong updating behaviour. In these terms, the 'long-run' is an asymptotic analysis of the situation where $\epsilon \to 0$. It is worth noting, that players here respond *myopically* in the sense that they do not take into account the actions of others (who may be updating strategies just as they are) when considering whether or not to update. Updating functions such as (2.4) are deterministic but *backward looking* in nature, leaving no room for prediction. Significantly, they obtain several powerful results, of which two were especially taken up by the literature. The first, (THEOREM 3 in their paper) is reproduced as follows:

Theorem 1 *Kandori et al. (1993, p.44) Suppose the stage game is a coordination game and $z^* \neq n/2$. For any population size $n \geq 2$ and any adjustment process satisfying (D)* (2.3)*, the limit distribution puts probability one on n if $z^* < n/2$, and on 0 if $z^* > n/2$.*

where, $z^* \in \mathbb{R}$ satisfies,

$$\text{sign}(\pi_1(z) - \pi_2(z)) = \text{sign}(z - z^*).$$

It can be seen that z^* is equivalent to the mixed strategy equilibrium probability on s_1 (for large enough n). Further, they identify the following corollary,

Corollary 1 *Kandori et al. (1993, p.46) Suppose the stage game is a coordination game and $z^* \neq n/2$. If*

$$N \geq \frac{2(a-d)}{(a-c-d+b)},$$

the unique long run equilibrium is the risk dominant equilibrium.[12]

The focus here is on the number of 'trembles' (stochastic shocks) that are needed to get the population to shift from one equilibrium to another (say E_1 to E_2). The intuition for their results is that where a risk dominant equilibrium exists, it will take more trembles to escape from, than the non-risk-dominant equilibrium. Hence, the long-run *stochastically stable* equilibrium is selected to be the risk dominant one.

Before moving on, several features of the KMR model are worth noting, bringing to the foreground its relevance to our present discussion. First and foremost, it must be stressed that the KMR model is a *uniform* (non-local, or *anonymous*) model. Formally,

[12] Risk dominance is a useful way of describing equilibria, based on their risk with respect to the other payoffs in the game table. For the coordination game, the question of risk dominance comes down to a comparison of $(a-c)$ and $(d-b)$, if $(a-c) > (d-b)$ then E_1 is said to *risk dominate* E_2, see Harsanyi and Selten (1988).

let $p_{i;j} \in [0,1]$ denote the probability that some agent i meets agent j in a given period, both from the population $N = \{1,\ldots,n\}$, then *uniform* interactions would imply,

$$p_{i;j} = \frac{1}{n-1} \quad \forall \quad j \in N. \tag{2.5}$$

The implication here is that two agents are just as likely to 'meet' and play $\mathcal{G}_1$ as any other two. Hence, there is no 'memory' or 'logging' of previous interactions which might otherwise be the stuff of friendship-, or acquaintance-networks: agents meet each other 'anew' every time.[13] Infact, the KMR model *requires* that this be the case, since the dynamics are only relevant when each agent responds to the true strategy distribution in the population.[14] Second, and in some ways a re-statement of the first, the agents gather wide-ranging information from the other $n-1$ players, that is, although boundedly rational (and capable of making mistakes) they are nonetheless un-limited in their ability to observe the other players (they end up obtaining *global* information).[15]

2.4.2 Incorporating Economic Networks

Various authors came to point out these short-comings, in particular Ellison (1993) specifically questioned whether the convergence rate of the evolutionary dynamics specified in KMR would lead to any realisable equilibrium for large populations (as might be considered for its application). Other authors questioned the information requirements, arguing that agents rarely are able to take into account such far-reaching information when forming their decisions. To treat this question, Ellison and others[16] considered models of *localised* interaction, where agents typically come into contact with only those members of the population who are their direct 'neighbours.' For example, instead of (2.5) (and following Ellison (1993)), each agent might interact (on a circle) with $2k$ neighbours ($k \in \mathbb{N}$), then,

$$p_{ij} = \begin{cases} \frac{1}{2k} & \text{if} \quad |i-j| \pmod{n} = 1,2,\ldots,k \\ 0 & \text{otherwise.} \end{cases} \tag{2.6}$$

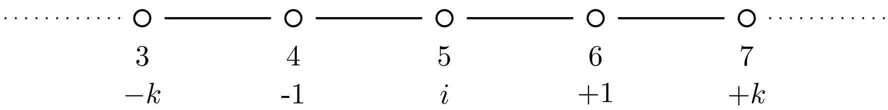

[13] Hence the use of the term 'anonymous' to describe this interaction set-up.

[14] As is pointed out in KMR, this could arise due to tournament-style play, or a very large number of random interactions occurring between agents, such that they come to observe the true strategy distribution.

[15] It should be mentioned that the nature of bounded rationality in learning as an evolutionary process also receives comment. See, for example Borgers (1996).

[16] For example, see Anderlini and Ianni (1996, 1997); Bala and Goyal (1998); Blume (1995); Chwe (2000).

Interestingly, for many of these contributions, the results of the KMR model are upheld – that is, convergence to the risk-dominant equilibrium is obtained for large populations who interact locally (in fact, Ellison concluded that local interactions gave the necessary *acceleration* of the dynamics that would produce plausible convergence to the risk-dominant equilibrium – serving to strengthen further the early KMR result).[17] However, a further strand of the literature investigated models that allowed for the interaction space *itself* to evolve over time.

2.4.3 Allowing the Network to Evolve

Here, two approaches are note-worthy. First, Ely (2002) considered a model in which agents are afforded 'mobility', such that they can periodically choose a new location on a network, and second, Jackson and Watts (2002b) alternatively argue that agents choose their *contacts*, rather than upping and moving to a new location in a social web in one step. I shall consider them in turn.

In Ely's model, players inhabit an incomplete graph[18] and each period, play the coordination game $\mathcal{G}_1$ with each of their neighbours only, receiving an average payoff. It is to be noted that Ely assumes (with KMR) that $\{s_1, s_1\}$ is Pareto efficient, however, in contrast, he sets payoffs such that s_2 is the best reply to an equal mixing, and therefore risk dominant. Then, with probability β each member of the population 'dies' and are replaced by new players who make a best response decision to the current strategy distribution and interaction graph characteristics (they choose an edge set, edges are undirected and costless). Further, a proportion α of the surviving agents revise their strategies, in a best-response manner. In line with KMR, both new agents and (updating) surviving agents will tremble with probability ϵ.[19] Now, suppose Z_k, for $k \in \{1, 2\}$ describes the set of agents in N who play s_k, then the following result is obtained,

Theorem 2 *Ely (2002, p.12) Let* Γ *be any incomplete neighbourhood graph, and let* $\mu^*_{\alpha\beta}$ *be the limit distribution of the general model. Then for sufficiently large* n*, and for every* $\bar{\alpha}, \bar{\beta} \in (0, 1)$

1. $\mu^*_{\alpha\beta}(Z_1) > 0,\ \mu^*_{\alpha\beta}(Z_1 \cup Z_2) = 1$;
2. $\lim_{\alpha \to 0} \mu^*_{\alpha\bar{\beta}}(Z_1) = \mu^*_{0\beta}(Z_1) = 1$; *and*
3. $\lim_{\beta \to 0} \mu^*_{\bar{\alpha}\beta}(Z_1) > 0$.

[17] 'Locally' here is interpreted in a variety of ways determined by the over-all topology of the interaction space. Popular spaces include placing agents on a line, circle (a line with period boundary constraints) or a 2-D grid. See references in previous footnote, for examples.

[18] In his paper, *non-degenerate*. For the following, I shall attempt to be consistent with previously used nomenclature.

[19] As is pointed out in Ely (2002, p.12), in the case of new agents, this can be considered a *mutation*, whereas for existing agents, it is akin to a *mistake*.

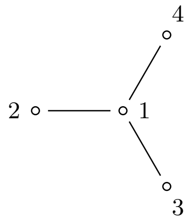

Figure 2.2 Counter example graph Γ_2.

This result suggests that for any strategy-revision (α) and agent-replacement (β) fraction pair the set of agents playing s_1 is non-empty and importantly, as these fractions are pushed towards zero, the efficient strategy is always chosen (s_1), even though it is *not* risk dominant. Ely suggests that through the use of mobility and strategy revision, agents are afforded a mechanism by which they can more successfully identify their opponents, thus favouring the efficient outcome, rather than the risk-dominant one.

Now, as mentioned above, Jackson and Watts (2002b) argue that Ely's conception of location change is too severe. Rather, they contend that, 'in many applications individuals choose with whom they interact in a more discretionary manner, not having to completely uproot to form new relationships.' (p.267) To begin with, they provide a good counter-example to that of the KMR model, showing that even a simple ($n = 4$) network such as Γ_2 (see Fig. 2.2) where $E = \{(1,2),(1,3),(1,4)\}$ gives rise to *two* stochastically stable states (either all play s_1 (let this be state A) or all play s_2 (state B)).

This is due to the simple observation that a single tremble on behalf of the centre player, 1, will cause a population in state A to switch to state B (since the rest only care about player 1's strategy), whilst if in state B, then a tremble by any player will lead to a switch back to A.[20]

In contrast to Ely, the authors consider a single edge at a time between two agents i and j. If it currently does not exist, the link is added if the utility of at least one player increases (the other player's remaining fixed), and in the alternate case, is severed, if the resultant configuration would favour at least one player (again, without damaging the the other). Case (iii) from their first proposition is reproduced below,

Proposition 1 *Jackson and Watts (2002b, pp.275–276) Under constant costs of edge maintenance k,*

(iii) If $(b-k) < 0$ and/or $(c-k) < 0$ and $(a-k) > 0$ and $(d-k) > 0$, then if $(d-b)/(a-c+d-b) > 1/(n-1)$, there are two stochastically stable states, a complete network with all players playing s_1 and a complete network with all playing s_2; while if

[20] As Jackson and Watts (2002b, p.272) point out, A is still more likely, since any tremble from B will cause A whereas a *specific* tremble (that of player 1) is needed to go from A to B.

$(d-b)/(a-c+d-b) \leq 1/(n-1)$ *then the complete network with all playing* s_1 *is the stochastically stable state.*

With values for $\{a, \ldots, d\}$ just the same as KMR, the proposition implies that with endogenous interaction space updating, it is possible to have the population converge on the inefficient *and* non risk-dominant strategy equilibria.

2.4.4 Concluding Remarks

It should also be pointed out that, whereas Jackson and Watts consider mutual, undirected edge formation, other authors consider, for example, unilateral edge formation (Goyal and Vega-Redondo, 1999) or quasi-endogenous location models (Droste et al., 2000) of coordination (or conventions) which bring significant changes to the modeling environment, and consequently find at times in favour, or against the results of Jackson and Watts. Furthermore, whilst not covered in detail here, significant contributions (for example, Bala and Goyal (2000); Jackson and Watts (2002a); Slikker and van den Nouweland (2000)) have been made which find equilibrium topologies for the economic network itself, rather than the strategic implications that might arise due to network endogeneity.

To conclude this section, some general comments from the literature mentioned above can be drawn:

1. The presence of partner selection due to an economic network in theoretical models has the capacity to *reverse* previous results where such mechanisms are not considered;
2. When economic networks are modeled, the outcome for both equilibrium topologies and strategic outcomes are sensitive to assumptions including (but not limited to) the type of non-local interaction (static, dynamic); the nature of edge sponsorship/establishment (unilateral, mutual); the way that stochastic trembles enter the model; and the type of myopia (or broadly, bounded rationality) agents display.
3. Topological properties of equilibrium networks modeled receive characterisation (only) at the highest level (e.g. connected, un-connected, complete, incomplete) or for very small n (e.g. wheel, star); and
4. Models of equilibrium network configuration rely on agents possessing knowledge of the entire network when constructing 'best-response' strategies with respect to edge sets.

In view of the apparent deficiencies, it must be stressed again, that modeling economic systems within deductive mathematical constraints is of the highest difficulty.[21] As al-

[21] Private discussions with mentioned authors such as A. van den Nouweland and others support this view.

ready mentioned (see §2.2), incorporating the demands of economic activity in models of complex interaction spaces that are subject to constant change pushes the investigator to often make major simplifications (e.g. small n, simple graph-characterisation) in order to make any progress at all.

2.5 Examples of Economic Networks

To this point, several specific examples from the literature of economic networks have been mentioned, presently, a more general discussion of the uses of the term 'economic network' in the literature is initiated. Here, emphasis centres on *approaches* to economic network analysis to show the main themes of the literature. These approaches shall be discussed under three main groupings: 1) 'Lists & Subsets' - the network (simply) as a listing of agents; 2) 'Graph Theoretic' - related to the previous discussion – the detail of network topology is considered; and 3) 'Computational' – the network is again studied in detail, however, exploiting the power of computational methods, often studying larger, more complicated systems.

2.5.1 Lists & Subsets

In common conversation, we would rarely talk in explicit terms of some 'network' that we are part of. Rather, such groupings form because of mutual interests (e.g. sporting clubs) and though possessing an internal social structure that might hold relevance to us, the fact that we are *members* of this structure is what is important. This view of networks – a level of detail between that of anonymous interactions on the one hand, and specific dependency graphs on the other – is a useful one when considering just this kind of behavior, where the network membership and associated benefits conferred to members (rather than the effects of network structure itself) is the focus of inquiry.

In game-theoretic terms, such membership is a clear way to distinguish between *types*. Consequently, models that invoke networks in this way are often concerned with the size of a type grouping versus the population size, and the benefit (if any) that the network provides over and above a non-network environment. For example, in a model of quasi-coordination based on such type differentiation, Taylor (2000) studies an 'old-boy' network with just this foundation. Agents are paired to undertake mutually beneficial projects together if each accepts the pairing (a), although having an option to reject the undertaking (r). Thus, strategies are of the set $\{a, r\}$. Two types inhabit the population, type C - cooperative or competent types, and type D - incompetent or opportunistic types. Membership to the network is only available to agents who display cooperative/competent behaviour when playing with a current network member. If accepted by 'showing their colours' (accepting the project, and displaying type C behaviour), the new member's status becomes known to all in the population. Consequently, as the the 'network' grows, it is possible for networked, type C members, to start rejecting all interactions that are not with fellow (networked) type C members. It is suggested that this leads to discrimination against type C agents who aren't yet members, since they will eventually never be able to join. Therefore, the model suggests that that equilibrium network size is always too small to be efficient, and in fact, abolishing it would be Pareto improving. Taylor draws a parallel to Akerlof's (1976) "rat-race", in which, 'no one wishes to run, but any individual who stops is left behind in the dust.'[22]

Such networks are often prevalent where institutional mechanisms of type differentiation (e.g. contract law) are either non-existent, or not stable enough to ensure healthy economic activity.[23] An old network system that has arisen in this way, is the Chinese *guanxi* network.[24] Kali (1999) develops a model of the *guanxi* network, where agents gain access to the 'club' by donating a one-off gift during an interaction with a pre-existing member. In a related finding to that of Taylor, considered above, Kali finds that because of the absorption of honest individuals from the general population into the network, it is inefficient unless very large (on the scale of the population itself).

In these examples, the network is no more than a grouping of agents, who form a (proper) subset of the population for the purposes of type-differentiation. It is to be noted that in this approach, only a single 'network' can be considered (agents are either 'in' or 'out'), and subtleties of *search* within the network, the topology of the relationships, or the structural causes for its rise and decline are beyond these models.

2.5.2 Graph Theoretic

Various contributions that treat economic networks with an eye to the actual structural nuances involved in network formation and activity have already been covered (§2.4). Therefore, only a brief return to such literature is made here.

Obviously, one key economic network, often talked of in a very tangible manner, is that of research and development (R & D), or technology networks. Here, the focus is on seeding one firm or agent with some technology and seeing how the technology spreads throughout the population (diffusion). Or, alternatively, the network – as a collection of R & D collaborations – is studied as it responds to changing firm-to-firm

[22] This point is returned to in Chapter 8.

[23] Travellers to developing countries will be all too aware of the need for good *personal* information when procuring various goods and/or services from the indigenous population.

[24] An unofficial business-to-business, or business-to-government transaction mechanism, relying (largely) on pre-existing personal, or family, relationships. See for example, Fan (2002); Standifird and Marshall (2000); Zhang and Li (2003).

collaborations. Two examples of this approach in the theoretical literature are those of Goyal and Moraga-Gonzalez (2001) and Mitchell (1999) who study strategic R & D networks and technology diffusion respectively. In the former, the contributions of spillover effects (caused by the sharing of effort in collaborative relationships) and market competition implications (caused by the convergence of competitive conditions under collaboration) is studied, whilst in the latter, networks are considered as institutional structures, representing private and public cost and benefit flows when investing in new technologies.

Communication is a further important consequence of network analysis – questions here consider the flow of behaviours in a population or the structures that permit effective communication between agents. For example, Morris (2000) studies a model of *contagion*, which focuses on the former question. In a model of binary (0,1) behavioural attributes, he asks what networks would give rise to contagion – where an action played by a finite subset of players spreads to the whole population. He finds that *low neighbourhood growth* (total neighbour count grows less than exponentially for a given number of 'hops' (edge-movements) away from the agent of interest in the network) and *uniformity* (a similarity of structural features) give rise to conditions for contagion. In a related piece of work, Bala and Goyal (2000) suggest Nash equilibrium structures for communication networks, showing that under basic assumptions regarding link sponsorship and best response dynamics, a finite set of strict Nash networks are formed in equilibrium (see Chapters 4, 5 and 6). Alternatively, communication networks can be considered the product of organizational construction, such as in Chwe (1995), who studies networks formed under the principle of 'strategic reliability' (as opposed to 'network reliability'). Here the focus is on the underlying game that an organization faces when constructing reliable communication channels (compare Epstein (2003)).

The above references should be considered together with the previously mentioned network models. Whilst specific results have not been covered in each case, observations of a general form can be made. First, the network models mentioned consider (in the main) small n populations (normally $n < 10$), or they consider easily described graph-theoretic structures such as *complete*, *regular*, *uniform*, *cycle* or *star* structures. Further, models invariably make use of *best-response dynamics* when constructing networks, as this offers a *deterministic* dynamic in which to progress network characteristics whilst requiring only relatively acceptable assumptions (e.g. myopia, inertia).

2.5.3 Computational

A number of contributions have considered economic networks with the aid of computational analysis. Obviously, the possibility to conduct millions of calculations in a short amount of time has provided a rich environment for the study of network and strategic behaviour. Broadly speaking, the literature falls into two main approaches. First, simulations of agent characteristics are conducted on a *fixed* interaction space (network), or the interaction space itself is allowed to vary.

Examples of the first approach include the statistical mechanics-style treatment of the physics literature (e.g. Arenas et al. (2000); Elgazzar (2001, 2003); Laguna et al. (2003)) where comparisons between *uniform* matching (complete graphs) and so-called *small-world* graphs (see §2.6.1) are popular. The intuition generally concerns attempts to identify network-specific effects by simulating a 'base-line' outcome (under regular, or uniform matching) and comparing it to the 'real-world' network situation (under small-world, or random matching). For example, Elgazzar (2002) constructs a model on a circle where each agent is initially endowed with a technology level $a_i \in \mathbb{R}^{++}$. Provisional payoffs $\pi(a_i|a_j), j \in \{i-1, i+1\}$ contingent on the technology level chosen by their nearest neighbours are then calculated as follows,

$$\pi(a_i|a_j) = \begin{cases} a_i - k_1(1 - e^{-(a_i - a_j)}) & \text{if} \quad a_i \geq a_j \\ a_i - k_2(1 - e^{-(a_j - a_i)}) & \text{if} \quad a_i < a_j \end{cases} \tag{2.7}$$

where k_1 and k_2 are the incompatibility costs resulting from being too advanced or too backward respectively. The payoff to agent i, π_i is then a simple summation across their near-neighbours,

$$\pi_i = \pi(a_i|a_{i-1}) + \pi(a_i|a_{i+1}). \tag{2.8}$$

During a run of the model, a randomly selected agent's technology level is updated by a random variable $\Delta \in (0, 1)$ and during a succession of time-steps, agents compare their current payoff π_i to that of the average payoff for their neighbourhood ($\pi_{av} = (\pi_{i-1} + \pi_i + \pi_{i+1})/3$). Agents alter their technology to be that which optimizes their payoff given the technology levels of $\{a_{i-1}, a_{i+1}\}$ if the comparison shows a payoff less than the average. A 'run' ends when no agent wishes to update their technology level further. These conditions are then compared to a synthetic small-world network model, where with some small probability $\phi \sim 0.05$, connections are re-wired to a different (randomly selected) agent (to make short-cuts).

In general, these authors emphasise the speed at which realistic social networks are able to facilitate attribute convergence or technology adoption compared to more analytically accessible graph structures (such as the line, 2-D grid, or circle).[25] Furthermore, because of the ease of computational data generation, probabilistic curves are produced which display, for instance, the probability distribution function of the total number of technology changes (the *avalanche* size). We return to this kind of analysis below (see §3.5).

Alternatively, authors consider the dynamics of particular strategies on fixed structures as opposed to (say) random interaction spaces. As has been mentioned in the introduction to this chapter (§2.2), key to findings in this area are that normal conclusions concerning strategic homogeneity in coordination games (e.g. Anderlini and Ianni (1996)), or the prevalence of 'selfish' play (through studies of the Prisoner's Dilemma) can be overturned by the imposition of a particular (local) interaction structure (e.g. Burtsev and Turchin (2006); Cohen et al. (2001); Grim (1996); Nowak and May (1992)).

In the second approach, computational flexibility is applied to the generation and ongoing dynamics of the interaction structures themselves. In some cases, this is still based on small n assumptions (Zeggelink, 1995) whilst others consider larger, organizational, or social network structures (Epstein, 2003; Stocker et al., 2002). It is to be noted, that although this literature is growing with considerable speed, the identification of what *constitutes* a network is not constant across contributions. For instance, in some cases, an edge in the network is ascribed to interactions that yield higher than average payoffs (see Smucker et al. (1994, p.10ff), including discussion on this point) or after logging a number of probabilistic choice and refusal rounds (Ashlock et al., 1996). In these models, the 'network' is to be interpreted as a threshold measure of some quantifiable variable, rather than an interaction space resulting purely from individual, agent-based decisions (this point is returned to in Chapter 7).

2.6 Analysing Economic Networks

We have already mentioned some rudimentary terminology that aids network description and analysis, specifically in terms of Graph-Theoretic principles (see §2.3). However, certain other features, in particular the study of the *small-world effect* have come to be synonymous with such research. We therefore devote some detail to this, and to two other forms of economic network analysis frequently encountered in the literature.

[25] See in particular Watts et al. (2002) analysis of search in Milgram's 1977 classic parcel experiment.

2.6.1 The 'Small-world' Phenomenon

The term 'small-world' is now common-place in both our popular conversation and every-day print media. It neatly depicts the apparent binary opposites inherit in our perception of a large, incomprehensible, unknowable *world*, against the frequent experience of the *closeness* of contacts within it – things appear smaller than we might have expected. However, like most popular sayings, the 'small-world' description of our social networks has a technical (though, relatively informal) interpretation in network analysis, and is used to depict a certain kind of graph, in much the same way as *uniform*, *connected*, or *disjoint* might.

The small-world phenomenon was classically studied by Milgram (1977),[26] his research giving rise to the popular phrase 'six degrees of separation.' He sent parcels randomly across the USA with the name and rudimentary details of a single target person on them. Receivers were to add their name to a list and, assuming they did not know the target person personally, re-post the parcel to a personal contact who they believed would be in a better position to speed the parcel on its journey to the target person. It is to be noted that out of 160 parcels, only 42 ($\sim 26\%$) reached their destination.[27] However, of those that reached their destination, the median number of contacts the parcel travelled through was 5.5, giving rise to the popular saying.

Clearly, for economic and social network investigators, being able to (artificially) create realistic networks is a priority due to the significant costs associated with gathering real-life data. In a much cited paper, Watts and Strogatz (1998) suggested that small-world type networks arise due to an underlying ordering of nodes, together with a small number of random short-circuit edges that dramatically reduce the path-length between nodes. A subsequent paper (Watts et al., 2002) proposed a small-world model that could give rise to the Milgram observation. For example, consider an agent set $V_c = \{1, 2, \ldots, n\}$, which is placed, in order, on a 4-*regular* cycle, such that for each agent $i \in V_c$ there exists an edge (i, j) for all $\{j : |i - j| \mod n \in \{1, 2\}\}$. In their procedure, each edge is addressed once, and with some probability $p \in (0, 1]$ the edge is 're-wired' to any $j \in V/\{i\}$ (see Fig. 2.3). Obviously, for $p \to 0$, the *k-regular* graph is simply returned, whilst as $p \to 1$, the graph approaches a random graph. However, in between these two extremes, the authors show that phenomena associated with the so-called *small-world effect* are

[26] For an earlier treatment, see Rapoport and Horvath (1961).

[27] On this point, Milgram relates his work to that of a previous MIT study, conducted in 1961, where data on how many people a random selection of men and women came in contact with over 100 days was collected, the average being 500 persons. Milgram notes that simply saying that people are a median of 5.5 intermediates away from each other should not be taken as a measure of *relational closeness*. (The geometric progression involved means that if people truly do have around 500 acquaintances then the pool of possible people the message could go to is around $500^{5.5}$, or almost 7 with 14 zeros following.)

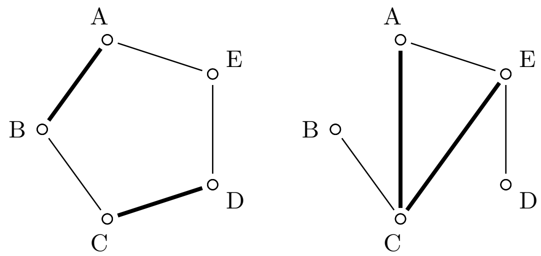

Figure 2.3 Example re-wiring process to form *small-world* graphs as in Watts and Strogatz (1998): two edges, (A, B) and (C, D) in a 2-regular graph (left) are randomly chosen to be re-wired, forming a new graph (right).

observed. First for some undirected graph in this region $\Gamma_{SW}(n, k)$ having n nodes and k edges, versus an equivalent random graph $\Gamma_R(n, k)$, the average minimum path length $L(\Gamma_{SW}) \gtrsim L(\Gamma_R)$; and second, the fraction of feasible *2-regular* triangular cycles in the graph (which they name the 'Clustering Coefficient') $C(\Gamma_{SW}) \gg C(\Gamma_R)$. That is, small-world graphs display a low node-to-node distances, approaching that of random-graphs, but have orders of magnitude higher clustering, or 'cliqueishness' compared to random graphs. This classification test for small-worlds is an easy one for investigators to calculate since the corresponding random-graph properties for large n are given by,

$$L[\Gamma_R(n,k)] \sim \frac{\ln(n)}{\ln(k)}, \text{ and} \tag{2.9}$$

$$C[\Gamma_R(n,k)] \sim \frac{k}{n}. \tag{2.10}$$

Such phenomena has been found to occur in many diverse networks. In Watts and Strogatz (1998) networks constructed from film actor associations,[28] power grids and even the neural network of *C. elegans* display small-world attributes. Likewise, other authors have found the world-wide web (Adamic, 1999; Adamic and Adar, 2003; Clarke et al., 2000), firm-based collaboration networks (Baum et al., 2003), and recent internet dating communities (Holme et al., 2004) all display the properties of a small-world network. It should be mentioned that the Watts and Strogatz rationale for small-world networks is not the only model in the literature. Subsequent contributions include using random networks that permit control over degree distribution and clustering (Volz, 2004) and focusing on small diameter networks with *regular* degree characteristics (Comellas et al., 2000).

2.6.2 Cluster Analysis

Within emerging complex economic networks it is often useful to identify groups of agents that fit one, or both, of two criteria: 1) the group is similar in terms of agent characteristics or attributes; or 2) the group is similar in terms of its topological location (e.g.

[28] Edges between actors are made if the actors have performed in a film together.

density of links etc.). Of course, in many models (such as social networks, collaboration networks) these two descriptions will be highly related. Whilst the first kind of analysis is relatively straightforward, the second offers some difficult computational and philosophical challenges. For instance, in a connected graph component, given that there is an area of high inter-connectivity (i.e. for $A \subset V$ where $|A| = m$ then $|E(A)| \to m(m-1)$), where does the 'cluster' finish? Which of the peripheral nodes should be included? What about nodes that are between clusters – to which cluster should they be assigned?

Consequently, with such inherent subjectivity a variety of approaches have been suggested.[29] We cover here, two recent methods by way of example.

Recursive Neighborhood Means (RNM) (Moody, 2001) In this approach, the intuition is taken from social influence models. Nodes are treated as agents with some randomly endowed vector of opinions. For example, agent i would have opinions $Y^i = \{y_1^i, y_2^i, \ldots, y_m^i\}$, $y \in \mathbb{R}$. In iteration t, the value of each agent's opinion vector is updated to be the mean of their first-degree neighbourhood set from the preceding iteration, say $\mathcal{N}^i(t-1)$ (adjacent nodes):

$$Y^i(t) = \frac{\sum_{j \in \mathcal{N}^i(t-1)} Y^j(t-1)}{|\mathcal{N}^i(t-1)|} \tag{2.11}$$

In this way, agents that are linked, will naturally move closer together in opinions. The analysis can be conducted with any number of dimensions, giving finer and finer detail in terms of clustering.

Of course, there is still 'art' to this process, since the investigator must judge at what point to stop the process and assign grouping within the network. However, a number of methods are available to aid this process (again, see Wasserman and Faust (1994, part III)).

Edge Betweeness (Girvan and Newman, 2001) In this method, a property known as the 'edge betweeness'[30] is calculated for each edge. For some edge e_i, the edge betweeness (say) $B(e_i)$ is defined as the count of all shortest paths between nodes that go through e_i. For somewhere like Sydney, Australia, the Harbour Bridge, or Harbour Tunnel would likely have very high edge betweeness since they connect major commercial hubs on each side of the harbour with no other practical path options available to commuters. The algorithm for identifying clusters proceeds as follows:

1. Calculate $B(e) \, \forall \, e \in E$.

[29] See Wasserman and Faust (1994, part III) for a treatment of earlier work.

[30] Closely related to 'betweeness centrality' of a vertex, first proposed by Freeman, see ref in footnote above.

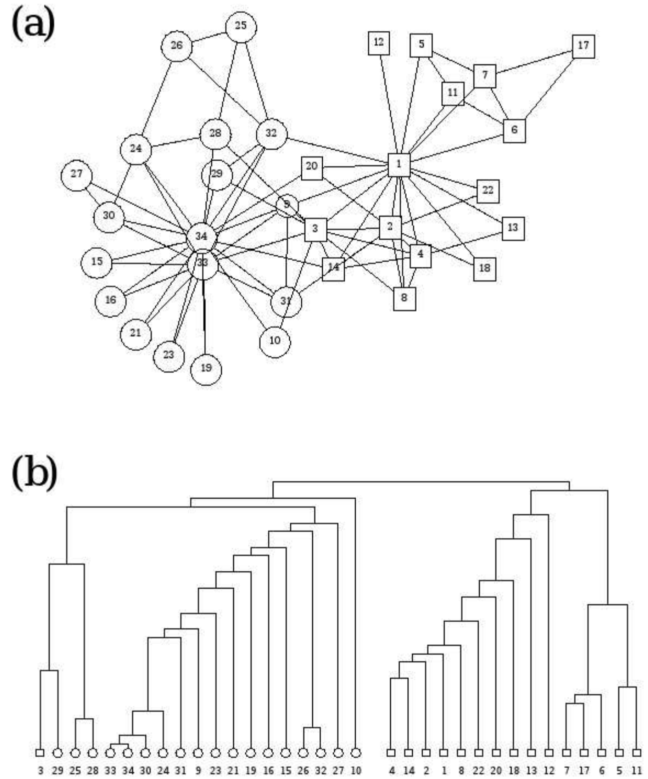

Figure 2.4 Example clustering tree after applying the edge-betweeness algorithm to a sporting club relational graph (a). Each leaf-node in (b) represents a vertex in (a) with lines showing grouping structure (reproduced from Girvan and Newman (2001)).

2. Remove $\{e_m : \max B(e)\}$.
3. Re-calculate $B(e)\,\forall\, e \in E/\{e_m\}$ for only those edges affected by the removal.
4. Repeat from step 2, until no edges remain.

In this way, an hierarchical graph structure can be quickly drawn up, with the first division arising due to the first edge removed, the second sub-division to the second edge removed, and so on (see schematic in Fig. 2.4). The algorithm thus constructs a tree-like hierarchical representation of the graph, with each branching resulting in further sub-branches until no further discrimination remains.[31]

2.6.3 Avalanches & Diffusion

As mentioned previously (§2.5.3), a strand of the computational economic network literature has considered models in which attributes such as technology level spread throughout a collaboration network for the purposes of joint cooperative technology outcomes. Such study comes under a number of banners, namely those of technology *avalanche*, *diffusion*

[31] For further examples of the application of this algorithm, other than the reference given above, see Tyler et al. (2003).

or even *self-organized criticality*. Since the latter is an attempt to describe several seemingly universal attributes of inter-connected, non-linear systems and provides something of an interface between economic network analysis and the science of complexity, we shall consider it separately below (§3.5).

Returning to computational models of technology adoption mentioned above, natural measurement questions follow: How many agents undergo technology updating in a period? What was the total technological improvement in the population for a given initial perturbation? and, How fast did the avalanche travel? In an early computational technology network paper, Arenas et al. (2000) provide useful implementations of such measures. With the technology adoption model as described in (2.7) and (2.8), the count of updating agents in a given period t is easily defined as the count of agents whose technology changed,

$$u(t) = \#\{i : a_i(t) \neq a_i(t-1)\}, \tag{2.12}$$

where $\#\{\ \}$ indicates the cardinality of the set. Likewise, we can simply sum the technology update size for each agent in a period to account for total technological advance $H(t)$,

$$H(t) = \sum_{i=1}^{n} [a_i(t) - a_i(t-1)]. \tag{2.13}$$

With these two measures, quantification of the mean *progress* of technology adoption, given economic network Γ_b and compatibility difference $k = k_1 - k_2$, is given by,

$$\rho = \lim_{T\to\infty} \rho(T) = \lim_{T\to\infty} \frac{\sum_{t=1}^{T} H(t)}{\sum_{t=1}^{T} s(t)}. \tag{2.14}$$

That is, ρ gives an impression of the total update size per updating agent calculated over the entire system perturbation. It can also be seen that ρ can give insight as to the underlying network, or compatibility factors that lead to large or small levels of technology adoption. It is noted by the authors, that for extreme cases, where $k \leq 2$ and $k \to \infty$, the summed technology update H and the number of updaters s are very nearly equal, meaning that technology updating per agent approximates closely the average external perturbation.[32]

[32] Compare Scheinkman and Woodford (1994), an alternative model of system responses to external shocks, studied with a view to criticality.

2.7 Economic Networks: Concluding Remarks

It must be stressed again that this introduction has not attempted to be thorough in depth, rather, it has endeavoured to point to significant considerations in the economic network literature, both of a classical (e.g. analytic) and contemporary (e.g. computational) nature. However, it ought to have become clear that, as with any field of enquiry, the investigator of economic networks must pay careful attention to matching her approach with the question in mind, and, by so doing, embrace the full character of that chosen approach. For instance, an analytic methodology as applied to economic networks necessitates a slender description of individual behaviours (e.g. myopic, best-response) and interaction spaces (e.g. uniform, regular) to achieve progress, whereas, one who utilises a computational model must acknowledge the inherent problems of some-what subjective, statistical description (e.g. cluster analysis, small-world attributes), and lack of generality (e.g. over-specified modeling environments). However, as the examples above attest, each contributes to our understanding of economic networks and will no doubt receive further research effort in the future.

Furthermore, up to this point, we have only hinted at more holistic descriptions of the special behaviour of economic networks. Phenomena such as heterogeneous population attributes at the limit due to localised interactions; the sudden on-set of technology avalanches in spatial models; the small-world nature of economic network topology; and the boundedly rational behaviour that drives realistic endogenous network generation — these are the stuff of much broader classifications of system behaviour, which in many cases, are claimed to have universal application. A few of these classifications have been mentioned already (albeit in passing only) such as *self-organized criticality*, *complexity*, *emergence* and *evolutionary dynamics*, but others such as *diversity*, *chaotic systems* and *non-linearity* could equally be included. Many of these phenomena or descriptions of system behaviour come under the study of *Complex Systems* which broadly speaking, aims to understand how many systems do not, indeed, attain some equilibrious, easily describable state over time, rather, they follow complex, non-equilibrium dynamics, defying prediction and suggesting rich state trajectories; akin to *epochs*, some noisy, others calm.

Traditionally, such research has been the exclusive domain of the physical sciences, in particular that of biology, ecology and mathematical physics. However, as we shall see below, there is a growing interest in such approaches to economic systems, since as mentioned, in many ways, they offer one of the most important, and potentially *complicated* of the many examples of complex systems on Earth! It is to these considerations that we now turn.

Chapter 3

Economic Networks as Complex Systems

Following on from the previous chapter, economic networks are now seen through the framework of *Complexity Science*. For readers who are not aquainted with this area, questions such as What is a 'complex system'? What is meant by 'complexity'? What are the characteristics of a complex system? are covered. Included is a section on the related concept of *self-organized criticality*, which has particular relevance for the analysis of Chapter 8.

3.1 Introduction

Whilst the preceding chapter might have hinted at the appropriateness of considering economic networks as 'complex entities', this is not (only) what I mean by considering economic networks as *complex systems*. Rather, I mean that such problems fit into the broad class of problems that are being increasingly investigated under the 'Science of Complexity' or just 'Complex Systems' research. It should be noted that the breadth of reach of complex systems research is vast, including diverse (and ever-expanding) fields such as geo-morphology, natural ecology, immunology, virology, traffic dynamics, financial market analysis, and so on. Consequently, in what follows, examples that are far removed from economics will be studied, in order to come to a better understanding of general phenomena that might be expected in economic networks.

So what is distinctive about the study of complex systems? In some ways, the answer to this question is itself a work in progress, however, various authors have attempted to described the key factors. For example, Miller and Page (2005) prosaically argue for a consideration of, what they term, 'the interest in between' that these dynamic systems possess,

> Ultimately the study of complex systems illuminates the interest in between the usual scientific boundaries. It is the interest in between various fields, like biology and economics and physics and computer science. ... It is the interest in between the usual extremes we use in modeling. ... It is the interest in between stasis and utter chaos. Much of the world tends not to be completely frozen or random, but rather it exists in between these two states.

Or, in the words of Heylighen et al. (2004) (emphasis retained),

> What distinguishes complexity science is its focus on phenomena that are characterized neither by order – like those studied in Newtonian mechanics and systems science, nor by disorder – like those investigated by statistical mechanics ... but that are situated somewhere in between, in the zone that is commonly (though perhaps misleadingly) called the *edge of chaos*.

So complex systems research inhabits a domain bounded by the twin poles of statistically independent, many-bodied behaviour on the one hand and the dependent, but 'ordered' (and so *predictable*) behaviour of a finite number of bodies on the other. The implications of this 'in betweenness' for the inquirer can be severe; Heylighen et al. capture this point well,

> Order is simple to model, since we can predict everything once we know the initial conditions and the constraints. Disorder too is simple in a sense: while we cannot predict the behavior of individual components, statistical independence means that we can accurately predict their *average* behavior, which for large numbers of components is practically equal to their overall behavior. In a truly complex system, on the other hand, components are to some degree independent, and thus autonomous in their behavior, while undergoing various direct and indirect interactions. This makes the global behavior of the system very difficult to predict, although it is not random.

This 'very difficult' nature of prediction in complex systems is another over-riding theme and ties into fundamental properties expressed by the theory of *chaos* and associated *non-linear* dynamics (returned to in detail below, §3.3).

In the context of the economic literature, Markose (2005) trace a similar difficulty that sprang out of the constraints of Formalistic and Deductive methods to eventually incorporate the *Inductive* Methods and Self Organizing Dynamics of complex systems research. They notably mention the contributions of authors such as Hayek, who introduced the concept of 'the limits of constructivist reason', and so pushed concepts such as formalistic incompleteness into the economic domain. A similar perspective was apparent in the physical sciences long ago, with authors such as Anderson (1972) questioning the assumed implicit connection in so many formal models between the *reductionist* hypothesis (all fundamental elements obey a finite set of laws) and the *constructionist* hypothesis – all happenings in the universe can then be constructed from these fundamental laws. In its place, the gathering results of complexity science would suggest that to know about some systems, the best place to look is not at the component parts (although these cannot be ignored), but at the interplay *between* the components and the whole.

Indeed, this interplay between different *levels* of enquiry is synonymous with complex systems research, authors such as Arthur (1999) point specifically to *increasing-returns* as the locus for complex dynamics in economic systems, providing a means by which the elements (human agents) can (unwittingly) produce a distinct epi-phenomenal patterning, which in turn can constrain or impinge on the elements themselves. Such phenomena provide the conditions for out-of-equilibrium dynamics, in contrast to the standard equilibrium concepts of much formal theory. These principles of *emergence* and *self-organization* prompt deep questions about the fundamental building blocks that together create the observed, and often fascinating, large-scale phenomena that complex systems are known for.

3.2 What is a Complex System?

Leaving a discussion of what we mean by 'complexity' aside for the moment (see §3.3), complex systems are often identified in the literature by several system properties. In the following discussion, I will consider three: 1) system size; 2) nonlinear dynamics; and 3) local interactions. I shall take each in turn.

First, the system will have some (not insignificant) number of interacting, diverse components, normally on a scale that renders traditional analytic approaches inadequate. It is worth pointing out that this does *not* mean that the number of distinct 'rules' in a system must be large, on the contrary, it is widely accepted that even 'simple' system processes can give rise to complex outcomes (see §3.7 below), however, in general, the count of actors (elements, nodes, agents) in a complex system is large.

Second, and importantly, the individual components, or collections of these components, will interact in a way that is often *non-linear*. Formally,[1] consider a function $f : s_t \rightarrow s_{t+1}$, a single-valued function that describes the state of the system at time $t+1$

[1] This discussion follows the very useful introduction to chaos and randomness of Eubank and Farmer (1990).

based on the state in time t. A *linear* function is then described where for two states s_a and s_b (say):

$$f(s_a + s_b) = f(s_a) + f(s_b) \tag{3.1}$$

that is, the principle of *superposition* holds. Conversely, a non-linear system is one in which (3.1) does not hold. One can think of this property in terms of a system perturbation: consider some system described by the function $s_{t+1} = f(s_t)$ (e.g. $f(s_t) = 2s_t$). Now, suppose the system is linear and we perturb s_t by an amount ϵ. Then, in the following period, we would expect,

$$s_{t+1} = f(s_t + \epsilon) = 2s_t + 2\epsilon = f(s_t) + f(\epsilon). \tag{3.2}$$

That is, we simply add the components, with the error term progressing in an highly predictable way. However, consider a simple non-linear example, say, $s_{t+1} = g(s_t) = 2s_t^2$, then if we perturb the system as before we obtain,

$$s_{t+1} = g(s_t + \epsilon) = 2(s_t + \epsilon)^2 = g(s_t) + h(s_t, \epsilon), \tag{3.3}$$

a combination of both the original function $g(s_t)$ and a *new* function $h(s_t, \epsilon)$ (in our example, $h(s_t, \epsilon) = 2(2s_t\epsilon + \epsilon^2)$). It is well known that such a small departure from linearity can have drastic consequences on the trajectory of a system, and hence, the predictability of its dynamics.

Non-linear dynamics relates the study of complex systems to that of *chaotic* systems (§3.3), where the process described above for a small perturbation is characteristically known as *sensitive dependence on initial conditions.*[2] Here, the major insight is that even simple (deterministic) systems can give rise to behaviour, that for all intents and purposes, appears random. Moreover, and causing a re-consideration of previous divisions between the study of *dynamical systems* on the one hand, and *random process* on the other, such complexity of behaviour is generated by the system *itself.* An important consequence of such dynamics is that these systems often display maximal measures of complexity (again, see below), with the most efficient method of analysis being to simply let the system run its course; any other attempt at 'prediction' being frustrated.

Thirdly, as previous discussions have emphasised (§2.4.3), the presence of *localised* interactions between elements creates the conditions for complicated dependencies and associated dynamics. This is perhaps the most difficult aspect of these systems to model

[2] First described by Lorenz (1963).

analytically. A description of each state, and the subsequent transitions between states, would potentially require an huge number of equations, possibly on the scale of the system size itself. This difficulty is particularly apparent in models of asynchronous element-wise updating. Of course, this aspect of complex systems is no doubt one of the key reasons that computational methods are often pursued, since system-wide diverse interactions can be calculated from one step to the next within the large-scale memory environment of modern computers. Not surprisingly, the study of complex systems has become closely associated to that of graph-theoretic, or network analysis (even if the system itself has no underlying spacial need for network approaches), we return to this aspect below (§3.4).

Consequently, authors such as Green (2000) point to two main sources of complexity: *iteration* and *interaction*. In the former, a (small) number of rules is repeated over and over again, giving rise to complex outcomes (e.g. cellular automata), whilst in the latter, agents or elements that are capable of certain behaviours are brought into contact many times over (e.g. ecosystems). Of course, these 'sources' of complexity are quite related, since the line between rules being applied to actors, on the one hand, and actors impacting on other actors by rules, on the other, is rather arbitrary. However, the distinction is an helpful one to keep in mind when approaching examples of complex systems.

Whilst each of the under-girding characteristics mentioned above are to be found in both the physical and social sciences, within the *economic sciences* a *fourth* and distinctive degree of difficulty emerges, that of *autonomy*. As Arthur (1999) notes,

> Unlike ions in a spin glass[3] which always react in a simple way to their local magnetic field, economic 'elements' – human agents – react with strategy and foresight by considering outcomes that might result as a consequence of behaviour they might undertake. This adds a layer of complication to economics not experienced by the natural sciences.

In the same vein, Holland and Miller (1991) comment on traditional analytic methods of analysis (even those that assume some kind of bounded rationality) arguing that,

> Current theoretical constructs, based on optimization principles, often require technically demanding derivations. It is an obvious criticism of these constructs that real agents lack the behavioral sophistication necessary to derive the proposed solutions. This dilemma is resolved if it is postulated that adaptive mechanisms, driven by market forces, lead the agents to act as if they were

[3]For example, a chemical glass, displaying high 'magnetic frustration': as opposed to 'normal' ferromagnetic materials, a spin glass displays multiple resting states, such that by applying heat and magnetic fields in parallel or serial application causes a very path-dependent behaviour. The additive behaviour, however, conforms to a constant for the glass, and each (and *every*) ion within the glass follows a consistent reaction to the applied forces, hence Arthur's comparison above.

> optimizing ... [Artificial Adaptive Agents] explicitly model this link between adaptation and market forces, and can thus be used to analyze the conditions under which optimization behaviour will (not) occur.

Elsewhere Arthur, Durlauf and Lane (1997, Introduction) point to defining features of *economic* systems that make them difficult for the 'traditional mathematics' used in economic analysis, and thus make the case for the complex system approach. They identify six properties of economic systems that distinguish them for the complexity science approach: (i) dispersed interactions; (ii) no global controller; (iii) cross-cutting hierarchical organization; (iv) continual adaptation; (v) perpetual novelty; and (vi) out-of-equilibrium dynamics. Properties (i) and (ii) speak to the autonomous agent phenomenon: many individual agents act in parallel across the system in an autonomous way without central organization. Property (iii) is a consequence of the first two, where emergent organisational *building blocks* form in an *ad-hoc* manner across the economy. The last three properties (iv)-(vi) could be described as different faces of the same trait – the economy is seen to be never resting, its actors continually coming up with new technologies and methods of activity that produce new niches to be filled, old ways to be jettisoned.

3.3 Randomness, Chaos & Measures of Complexity

Up to this point, key terms in complexity have not been elaborated on. It will be useful to ask what words such as *complex*, *complexity*, *complex system*, and even *maximally complex* mean? For example, when is anything really 'complex,' as opposed, perhaps, to 'random,' or 'chaotic?' (or just 'difficult' for that matter?) Such questions have traditionally been helpfully treated in the related fields of computer and information science, and so the discussion below shall draw from this literature.

3.3.1 Random vs. Stochastic

To begin with, it is useful to distinguish between the terms 'random' and 'stochastic,' since this will indirectly uncover one measure of complexity. Consider two streams of (binary) information in the form of a string of 20 zeros and ones (known as 'bits'):[4]

$$\begin{aligned} s_1 &= 01010101010101010101; \text{ and} \\ s_2 &= 10011001010011011111. \end{aligned}$$

It is immediately obvious that the first string, s_1 contains a patterning (a repeated sequence), whereas the second string s_2 may be describable by a single pattern, but if so, the pattern is not clear by inspection alone. A consequence of the regular patterning in

[4]The reader unfamiliar with binary representation, could think of the string as a succession of answers akin to the game 'twenty questions,' where $0 \equiv$ *no* and $1 \equiv$ *yes*.

s_1 is that if I were to convey the content of s_1 to another person, I could *compress* the message (e.g. 'ten zero-one pairs in a row.'). By contrast, s_2 does not easily lend itself to such compression – rather, I might be able to compress parts of it (e.g. I might finish my description, 'five ones'), but other parts would simply have to be read out aloud, in full. For clarity, the term *random* is generally kept to describe incompressible streams of information/data such as s_2, whereas streams that are produced by chance processes (e.g. repeated coin-flips), that may indeed contain sections of compressible information, are called *stochastic*.

3.3.2 Crude Complexity

Now if a description of the strings in terms of *complexity* was sought, it is clear it would be desirable for this description to say that s_2 is more complex than s_1. As a first attempt, a description of each string in words could be written down and then the length of each description measured: the longer the description, the higher the complexity. Clearly, the level of detail that the description gives (e.g. the proportion of zeros and ones, or the exact sequence) will bear on the measurement, such a procedure is termed *coarse graining* in the literature. The measure just described is called (for obvious reasons) *crude complexity*, it is described by Gell-Mann (1994) as follows:

> 'The length of the shortest message that will describe a system, at a given level of coarse graining, to someone at a distance, employing language, knowledge, and understanding that both parties share (and know they share) beforehand.'

3.3.3 Algorithmic Information Content (AIC)

Crude complexity is not a particularly useful measure of complexity, especially considering the presence of inherent subjectivity (despite the above quotation's careful constraints). A far more useful measure, and one that has great prominence in much of the computational literature is to standardise the descriptive engine, indeed, to suppose that the strings are to be described by a *universal computer*. A universal computer $\mathbf{U}$ can carry out any computable program, even therefore, simulating other universal computers. In the present case, a program that would print out s_1 and s_2 and then stop (halt) is required. The length of such a program, defined as the *algorithmic information content* (AIC), or *algorithmic randomness*, or *algorithmic complexity* is defined as follows (following Zurek 1990):

$$K_{\mathbf{U}}(s) \equiv |s^*_{\mathbf{U}}| \tag{3.4}$$

where $s^*_{\mathbf{U}}$ is the minimal (binary) program that displays sequence s, and the vertical lines mean simply the size (length) of the string in bits. It might be noticed that it is with AIC in mind, that the terms random (*maximal AIC*) and stochastic as discussed above are defined.

3.3.4 Effective Complexity

However, as pointed out by Gell-Mann (1994, p.59) AIC does not really accord with what is normally meant by complexity, since a random string will be measured to (by definition) have *maximal AIC*, but be well described by a very short description such as 'complete gibberish of length l.' That is, what is of importance when describing the complexity of an entity, is not a description of every detail, but rather, the length of the description of the entity's *regularities* (Gell-Mann, 1995). Clearly this can be somewhat context dependent,[5] but given some transformation to (say) a bit-string representation, then objective regularities might be described. Gell-Mann and Lloyd (1996) define such a measure as *effective complexity*, and in their prescription, a statistical-mechanics description of the string is employed, such that rather than computing just $K_{\mathbf{U}}(s)$ a set of strings that *probabilistically* describe s is constructed, each string sharing the regularities of s but differing in other respects (the irregularities). That is, a number of alternative strings r that share the regularities of s together with their respective probabilistic contributions P_r is formed. Such as set of strings (states) is equivalent to a (quantum mechanic) *ensemble* $E = \{(r, P_r)\}$ description of s. The effective complexity is then,

$$\mathcal{E} \equiv K_{\mathbf{U}}(E), \tag{3.5}$$

the length of the shortest program that describes the members of the set together with their probabilities. However, in determining $\mathcal{E}$, an obvious concern is which (and how many) elements will be contained in E – clearly, different observers will find different regularities in s and thus an element of subjectivity is unavoidable.

3.3.5 Shannon Algorithmic Information

To treat this concern, the authors propose an additional step, that of utilising the Shannon measure of information (ignorance) I to assess the uncertainty present in the set of strings in E. For a probability distribution p_r over alternatives r, Shannon's measure is given by,

$$I = -k \sum_r p_r \log p_r, \tag{3.6}$$

[5] Consider the changing fortunes of so-called 'junk' DNA — previous explanations have treated this portion of the human genome as irrelevant, however, recent authors have contended that this region may indeed provide the basis for genetic robustness and is therefore well worthy of study.

where k is an arbitrary positive constant. For example, where $k = 1$ and logarithms are taken in base 2, I will be measured in *bits*. As in Gell-Mann and Lloyd (1996, p.45) for a single fair coin flip ($p_h = p_t = \frac{1}{2}$), $I_{pre} = -\{p_h \log p_h + p_t \log p_t\} = 1$bit prior to looking at the coin. However, after the coin has been seen (say heads), $p_h = 1$ and $p_t = 0$ which correspond to $I_{post} = 0$, that is, one bit of ignorance has been lost, and one bit of information has been gained.

Thus, by considering the AIC of each member r in the ensemble, and computing an equivalent Shannon ignorance measure,

$$I \approx \sum_r P_r K_{\mathbf{U}}(r|E), \tag{3.7}$$

where $K_{\mathbf{U}}(r|E)$ is simply the contingent AIC of each member of the ensemble, we have an *average contingent AIC* of the members of E equivalent to (within 1 bit) the Shannon ignorance measure, or entropy of the ensemble. Taken together, the effective complexity $\mathcal{E}$ then gives a measure of regularities present in s, whilst I indicates the information required to specify the residual irregular components. The summation of these properties is defined to be the *total information*,

$$\Sigma = \mathcal{E} + I, \tag{3.8}$$

and since this should be constant between descriptions (in fact, it should be close to minimal, given by $K_{\mathbf{U}}(s)$) a trade-off will ensue. A longer, more detailed (complicated) description of the string's regularities (higher $\mathcal{E}$) will result in a smaller number of potential entities described by the ensemble (lower I). Conversely, a concise description of the string will (normally) result in a broadening of allowable strings possessing those regularities. It is possible to identify concise (low $\mathcal{E}$), and yet *powerful* descriptions of a string (vanishing I), for example, the (seemingly *maximal AIC*) decimal expansion of π might be described in binary form by a simple program to generate the expansion, however, such instances are rare.

A final point is worth making concerning the relationship between complexity and chaos. In the language of Eubank and Farmer (1990), 'chaos provides the link between determinism and randomness.' That is, the major insight of chaos theory is that relatively simple, deterministic processes can give rise to system dynamics that to all intents and purposes are random. No longer is the dichotomy between the study of *dynamical systems* (determinism with precise knowledge of initial states) and *random processes* (uncertainty concerning initial states or laws of motion) necessary. Chaotic processes, with their exclusively non-linear, though often simple, dynamics, give rise to non-smooth state trajectories (akin to a 'shuffling' of states), resulting in *uncorrelated* processes, even over short time-frames. Hence, traditional time-series measures (such as correlation) are inadequate to describe the system. For this reason, whilst not all 'complex systems' satisfy strict definitions of chaotic processes, they do, in general, require non-linear measures of system state, such as those discussed above, or their equivalents.

3.3.6 Complexity and the Economy: relationships within and without

The above discussion focuses on a potted history of measures of complexity, and in particular, on one recent, but powerful approach to this area.[6] Others could be given (e.g. *mutual information* in random grammars (Kauffman, 1993, p.384)) but none seem to deal so well with the concerns of the scientific application of complexity as Gell-Mann and Lloyd mentioned above. Whilst it is possible to apply such measures algorithmically to system states or population characteristics, the real usefulness of these discussions is first, in providing depth to an otherwise ethereal term, and second, in giving insight as to how one might track the complexity or otherwise of a given system (problem). Indeed, in terms of the previous discussion on complex systems in general, parallels with measurements of complexity are apparent. The idea of complexity science operating 'in between' the normal modes of analysis corresponds directly with notions of maximal effective complexity.

Before leaving this section, it should be noted that the direction of the proposed measures of complexity found above is actually very closely aligned with physical measures of order and disorder, captured by the thermodynamic quantity, *entropy*. Entropy is a state variable of physical systems, and is the only law of thermodynamics that is *statistical* in nature. Loosely speaking, it captures the fact that a closed system of entities (e.g. an insulated box of atoms), having some measurable (and constant) level of various macroscopic quantities such as energy, temperature and pressure, can be found to be in a (very large) number of different configurations (e.g. atomic position, velocity), and although these configurations do not alter the macroscopic measures, they will give rise to different levels of *entropy*. That is to say, it is quite (statistically) *unlikely* (although possible) that a large proportion of the atoms would move (reversibly) to a single corner of the box (low entropy), whereas is highly likely that the atoms will move roughly in a manner such that a uniform density of atoms is found throughout the box (high entropy). The second law of thermodynamics expresses this difference of likelihoods, and sets down that systems such as this will never *spontaneously* decrease in entropy, and hence, by observing the two

[6]For a different view of effective complexity, a comparison with Standish (2001) is worthwhile.

states of the box so described, one could immediately determine which existed in time before the other. This is perhaps the deepest conclusion of this area of thermodynamics, since it suggests a physical interpretation for the 'arrow of time'.

The implications of such findings for economic systems is still unclear, but the possible connections are tantalising. Does a similar 'arrow' exist in social systems? What would the objects of order and disorder be – distribution of wealth? technology diffusion? Such questions are left open and are not intended to be pursued by this thesis, it is enough to note that other workers are presently pursuing these ideas.[7]

3.4 Complexity & Networks

Networks do not necessarily have to comprise individual agents or properties as nodes, in fact, a variety of networks can be constructed that describe alternative relationships such as state trajectories, or variable dependencies in dynamic systems. Indeed, although the strong isomorphism between deterministic processes in general (dynamical systems) and directed graphs (di-graphs) is generally accepted, Green (1993) gave formal proof to this notion by providing three related theorems:

Theorem 3 *The patterns of dependencies in matrix models, dynamical systems and cellular automata are all isomorphic to directed graphs;*

Theorem 4 *In any array of deterministic automata with a finite number of states, the state space forms a directed graph; and*

Theorem 5 *Any deterministic automaton, or cellular automaton, with a finite number of states ultimately falls into either a fixed state, or a finite limit cycle.*

The first theorem states that any system with simple or complex dependencies between variables can be represented adequately by a directed graph. The second and third concern the relationship between states in the state space (the system *trajectory*), suggesting that the state space likewise forms a di-graph, and must, therefore, under finite deterministic rules give rise to either an absorbing state (single node) in the graph, or a cycle between a finite number of states (nodes). It is to be noted that the second theorem can be extended to stochastic automata by assigning edges between states to non-zero probability transitions. However, as might be noted, this destroys the final result for such systems, since a trajectory once entered, will seldom be covered again given the probabilistic state-to-state transitions.

[7] This area is pursued under the general title 'Econophysics', see `http://www.ge.infm.it/~ecph/index.php`.

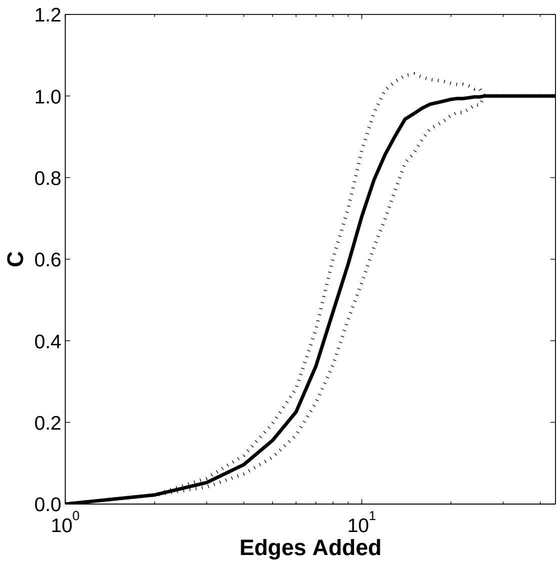

Figure 3.1 Example of the *connectivity avalanche*, undirected edges added randomly to initial 10 node null-graph. Average results from 100 trials (solid) $\pm\sigma$ (dotted).

Whilst in some ways, these results are merely a way of thinking or talking about dynamical systems, they give rise to fascinating outcomes which relate directly to the recent discussions of effective complexity and chaos given above. To see this, consider a second strand of literature which began with the work of Erdös and Rènyi (1959) on undirected graphs, but has subsequently been re-visited on directed graphs by Green (1993) and Seeley (2000). Here, edges are added in some random fashion to a null-graph, say $\mathcal{G}(V, E)$ where $E = \emptyset$ and $|V| = N$, having zero initial connectivity.[8] In both cases, either with undirected or directed graphs, the connectivity starts small, but begins to increase with an ever accelerating rate as edges are added. At this stage, small 'patches' of connected components appear, although each patch is disjoint with respect to the others. However, after time, these patches begin to join (much like the coalescence of oil droplets circulating on water) forming super-patches, and the connectivity begins to rise sharply. After a short time, the addition of just one or two extra edges can cause the super-patches to join, forming a giant-component that encompasses almost all of the nodes ($\gtrsim 85\%$) resulting in a connectivity measure approaching 1 (see Fig. 3.1).

The point at which the system can experience large connectivity jumps is known as the 'double-jump' point, or, as in Seeley (2000), the 'connectivity avalanche.' Whilst various measures undergo their own dynamics with the addition of edges (e.g. the presence of Hamiltonian cycles), the connectivity avalanche appears to be the most striking. At this

[8] Recall, *connectivity* of some graph $c(\mathcal{G})$is defined as the count of node-pairs for which a path exists, normalised by the maximum such count, which for an undirected graph of size N, is given by $\frac{n(n-1)}{2}$ whilst for directed graphs is $n(n-1)$.

point, the addition of a pair of appropriately placed edges has the capacity to cause the system to effectively go from dis-joint to strongly connected. Seeley provides an explanation for the massive rise in connectivity – consider two connected components comprising n_1 and n_2 nodes respectively, then, prior to the connections, their total contribution to connectivity is $n_1(n_1 - 1) + n_2(n_2 - 1)$ whilst after the edges are added, this becomes $(n_1 + n_2)(n_1 + n_2 - 1)$, an increase of $2n_1n_2$. The multiplicative nature of the component sizes thus accounts for the large connectivity jump.

It is to be noted that not every additional edge will have such an impact at this point. Indeed, many additions will have vanishingly small effect on the connectivity, but as mentioned, with small probability, when the additions fall between nodes in each large component, then the connectivity can change dramatically. Clearly, such an observation is just as pertinent for the *subtraction* of edges as it is for their addition, the so-called *connectivity implosion*. For example, consider a firm-based supplier-producer network – the subtraction of one firm, or one partnership will largely have almost no effect on network-wide economic activity, but occasionally such a change will drastically cut off groups of firms, having network-wide ramifications.

Such behaviour is normally addressed under the title *criticality*, to express the fact that the system is at some kind of critical stage between very different states. Since this has a broad application in the sciences, we consider this phenomena in more detail below (§3.5). For now, it is to be noted that in a dynamically evolving network environment, notions of criticality can be a key place to look when attempting to explain certain equilibrium shifts in dynamic systems. In fact, as is discussed presently, some authors argue that naturally evolving interaction networks actually *self-organize* to this region of criticality and thereafter maintain this state.

3.5 Self-Organized Criticality & 'The Edge of Chaos'

In a much cited work on Boolean NK networks Kauffman (1993, ch.5,6) follows just this path. Although initially conceived to study the network effects of genetic activity, the model provides interesting results for the the emergence of equilibria in complex interaction environments. Here, each node (gene) was either 'on' (1) or 'off' (0), and was placed on a random directed graph of size N with an average of K incident edges. The incident edges carry the state of the adjacent nodes, with each node following a rule-based digital transition function to determine its subsequent state. For example, for an $N = 3$, $K = 2$ Boolean network, each node must have a response for $2^K = 4$ possible input states, which in this case are given by the set: $\{00, 01, 10, 11\}$ (an example graph is given

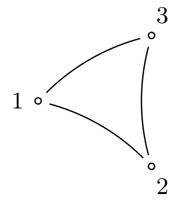

Figure 3.2 Example $N = 3$, $K = 2$ Boolean graph as described in Kauffman (1993, chs 5,6).

in Fig. 3.2). So, for example, the program for node A might be given by (0010), which means, 'return 0 in all cases, except for an input of [10].'

A further control parameter, P was introduced to define the proportion of ones found in each rule (that is, a measure of *bias*). After extensive simulations, tracing the system state (values of each node), it was discovered that either varying K (given N) or varying P (given K and N) resulted in three clear phenomena. In one phase ($K \gtrsim 5$ or $P \to 0.5$), the system displayed chaotic dynamics (attractor sizes scaling exponentially with N), however, as either $K \to 2$ or $P \to 1$, systems quickly establish ordered dynamics, often achieving a limit cycle or fixed state on a time-scale orders of magnitude less than 2^N (the total number of states). Characteristic of this state, was a central *frozen* region in the network (a cluster of nodes with fixed or cycling states), which appeared to force its ordering on the rest of the network via *percolation*. However, in between these regions, a fascinating transient phase was apparent where any frozen sections that did occur, were just as likely to 'melt' amidst seemingly complex signaling regions. That is, order and dis-order existed alongside each other, either in space or time.[9]

The work on random Boolean networks, couched as it was in the context of biological evolution, gave the first hint of the rich behaviour present in the 'middle ground' of many of these systems. In fact, the term 'the edge of chaos' was coined to describe this sub-critical state, where the system could equally display chaotic or ordered dynamics, often returning to the 'edge' only to diverge again. Indeed, Kauffman (1993, p.219) points to the significance of this region as being perhaps *the* characteristic of robust and 'useful' systems, not just with reference to Boolean networks – though they clearly provide a foundational example – but as a *universal* phenomena of complex *adaptive* systems. His description continues,

> "...if a network is deep in the frozen phase, then little computation can occur within it. At best, each small unfrozen, isolated island engages in its own internal dynamics functionally uncoupled from the rest of the system by the frozen component. In the chaotic phase, [the] dynamics [are] too disordered to be useful. Small changes at any point propagate damage to most other

[9] Compare a visual example, based on Class IV cellular automata behaviour in Fig. 3.5 below.

> elements in the system. Coordination of ordered change is excessively difficult. At the boundary between order and chaos, the frozen regime is melting and the functionally isolated unfrozen islands are in tenuous shifting contact with one another. It seems plausible that the most complex, most integrated, and most evolvable behaviour might occur in this boundary region."

Moreover, and receiving great interest of late with respect to self-organized criticality, he makes the following 'bold hypothesis,'

> "Living systems exist in the solid regime near the edge of chaos, and natural selection achieves and sustains such a poised state."(Kauffman, 1993, p.232)

That is, not only do natural systems (including therefore human systems) exhibit 'edge of chaos' dynamics, but these systems naturally evolve *towards* them, even to the extent of returning there after a transient disturbance. In other words, he (and others) would argue that these systems are able to *self-organize* to this state.

Such considerations are clearly related to previous discussions of complexity. Indeed, Gell-Mann (1995) argues that natural systems are those that display high *effective complexity*, having many operating regularities (rules) complimented by some stochastic influence – that is, the transient 'edge of chaos' exists between the two poles of order and chaos. From this view-point, dynamic network systems are themselves examples of complex adaptive systems. The rules or 'schema' they develop over time are used for processing the myriad data transfers within the system (so the Boolean programs of Kauffman's networks). Gell-Mann notes, "a complex adaptive system functions best in a situation intermediate between order and disorder."(Gell-Mann, 1994, p.249) We turn now to some examples of systems that appear to display such phenomena, and in turn, the relative claims for this contentious *universal* phenomena in complex adaptive systems.

In largely theoretical work, Bak et al. (1988) suggested (in their terms) that 'spatially extended dynamical systems' – systems possessing both time- and space- degrees of freedom – would reach, from any starting point, a *critical* state.[10] They consider 'composite' systems (very large systems of a number of non-linearly interactive elements) arguing that these systems 'self-organize' or evolve naturally to a critical state. Once there, as opposed to sub-critical (dampening at short distances) or super-critical (run-away reaction) behaviour, these systems return to the critical state after a collapse. That is, they are always 'at' the critical state. This is surprising, since it was previously believed that such a state could only be achieved through the 'tweaking' of control parameters (such

[10]See also the less-technical treatment in Bak and Chen (1991), and another treatment by Kadanoff et al. (1989).

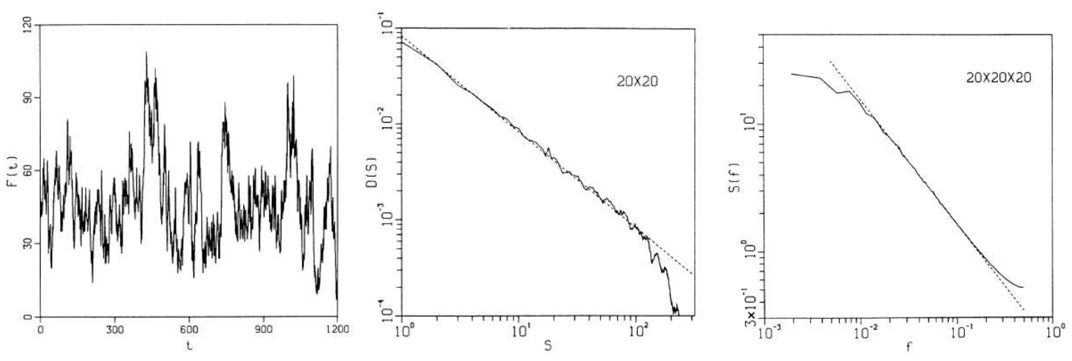

(a) **(b)** **(c)**

Figure 3.3 Example Self-Organized Criticality analysis of a computational simple displacement model as undertaken by Bak et al. (1988) showing power-law scaling in both spatial and temporal dimensions: (a) Example superimposed time series of actual fluctuations of the system; (b) Example frequency distribution of the size of resultant perturbation outcomes; and (c) Example power-spectrum analysis of the time-series.

as the fine control required to maintain the nuclear fission reaction at the critical state). Aleksic (2000, p.113) comments,

> In equilibrium physics criticality is the exceptional case. ... But for the *Game of Life*, and some large dissipative systems the critical scale-free spatial and temporal behaviour naturally self-organize. The critical state of these systems has a large basin of attraction: for a wide range of initial conditions the system will evolve to this critical state. [11]

The 'signatures' of such self-organized criticality appear to be *spatial* self-similarity (fractal scaling) and power-law scaling phenomena, in the spatial- and temporal- dimensions respectively (see Fig. 3.3). The latter being previously referred to as the ubiquitous '$1/f$' noise that had previously been treated as a measurement error, with the power spectrum $S(f)$ scaling as $1/f$ at low frequencies.

Surprisingly, the quintessential example of this behaviour is that of the humble sand-pile. Adding grains of sand to a suitable surface will cause a pile to form, which after time will suffer from a number of slippage events as more sand is added (attaining the angle of 'repose'). Significantly, rather than collapsing completely (a runaway response), or allowing near infinite addition (over-dampening), the edge of the pile always returns to the same angle through a process of many small, but sometimes large, 'avalanches.'

[11] The *Game of Life* is a reference to Conway's (Berlekamp et al., 1982) famous model of an astonishingly simple cellular automaton (covered below) that was capable of displaying highly complex pattern-based behaviour over many, many time intervals.

Indeed, Held et al. (1990) carried out the physical version of this experiment, finding that for small sand-piles, critical mass fluctuations were observed.[12]

Moreover, since such scaling is apparent in a diverse range of physical systems (e.g. earth-quake incidences, light emitted by quasars, current in resistors, sand flow in an hour glass, river flows, stock market fluctuations), the authors generally claim that a *universal* phenomena is being observed. This universality appears to stem from the long-term non-linear coupled interactions that each of these systems possess, the authors note,

> In retrospect, it is hard to see how $1/f$ noise, with long temporal correlations, could possibly occur without long-range spatial correlations, (Kadanoff et al., 1989, p.373)

Clearly, the work by Seeley mentioned above (§3.4) has such thoughts in mind.

It is not difficult to see the relevance for the present study. Many economic systems (e.g. the stock-market, supplier-producer networks, person-to-person interaction networks) display the potential for exactly these long-range spatial interactions. Indeed, elsewhere Bak and Chen (1991, p.33) argue that,

> Conventional models (of economics) assume the existence of a strongly stable equilibrium position for the economy, whereby large aggregate fluctuations can result only from external shocks that simultaneously affect many different sectors in the same way. Yet it is often difficult to identify the reasons for such large-scale fluctuations as the depression of the 1930s. If, on the other hand, the economy is a self-organized critical system, more or less periodic large-scale fluctuations are to be expected even in the absence of any common jolts across sectors.

Indeed, Scheinkman and Woodford (1994) show just such an outcome with a simple model of a supply-demand economy. They argue that economic interactions must display *'local and significantly nonlinear'* (emphasis retained) dependencies for criticality of this sort to emerge.[13] Alternatively, Canning et al. (1998) consider the importance that criticality considerations have for aggregation to the scale of national growth rates, whilst several authors have attempted to point to more esoteric renditions of the world economy (or world *system*) through power-law scaling phenomena. However, in these efforts, it is less clear how the process of aggregation gives rise to the macroscopic outcomes surveyed.[14]

[12]For larger systems, relaxation oscillations were observed instead, however, this could be explained by the necessary change in the experimental drop height.

[13]For an econophysics approach see Arenas et al. (2000).

[14]See for example, Devezas and Modelski (2003) or Matutinović (2002).

To summarize, whether or not self-organized criticality is a phenomenon of all strongly interacting non-linear systems, or just those that fit some subset, it is clear that such considerations cannot be ignored when aggregating from the individual to the whole. In particular, it would seem that the theory's explanatory power is greatest when shocks are able to propagate over extensive inter-temporal scales *via* lagged, or time-intensive spatial interactions. Specifically, when dealing with explicit network-based systems where individual nodes exhibit agency contingent on their connections, scaling phenomena should, at the very least, be unsurprising.

3.6 Emergence & Self-organization

A kind of self-organization related to the processing ability of a system with respect to external shocks has already been mentioned, however, more commonly, self-organization, and its broader partner, *emergence*, are associated with complex systems in describing (often surprising) aggregate behaviour. The property of *emergence* received significant interest in the latter half of the previous century, in no small part due to contributions made by authors in the field of *artificial life*. Here, the emphasis was, and continues to be, on determining the supposed 'simple' rules and processes that drive natural organisms, with a view to exploiting the resultant robust and adaptive properties that these creatures display. In this way, emergence became a criterion of 'good' systems, since it appeared to be synonymous with the way that often rudimentary, autonomous processes – albeit interacting in sometimes complicated ways – could give rise to coordinated or 'intelligent' behaviours.

But what is emergence? At one level, emergence can be thought of as turning on the *level of description* that a system lends itself to. Following Standish (2001), suppose that the fundamental components of a system can be described by $\mathcal{L}_1$, a 'microlanguage,' which, for example, might describe the number of pedestrians on a side-walk, their individual velocities and positions. However, suppose further that because of the myriad interactions of the people on the sidewalk over time, two *flows* are apparent, one traveling in each direction. In this case, a second (presumably shorter) description, $\mathcal{L}_2$ – a 'macrolanguage' – could be used to describe the system. We then have a definition of an emergent phenomenon, Standish summarises,

> An emergent phenomenon is simply one that is described by atomic concepts available in the macrolanguage, but cannot be so described in the microlanguage.

or in the words of Aleksic (2000, p.99),

> The regular patterns that can be discriminated in domains other than the generating domain are examples of emergent phenomena. ... [They] exist as the result of the aggregate behaviour of more elementary parts.

In this view, the property of emergence is closely related to that of self-organization, since to organize is, by definition, to reduce the level of disorder, and hence, to identify regularities that lend themselves to an aggregate description. However, this may not necessarily be the case – a chaotic pattern might indeed result from a very 'simple' system – here, the emergent phenomena would be negatively correlated with organization! Of course, such a discussion is very closely related to previous comments regarding complexity (§3.3).

In many ways, this description of emergence may seem trivial, but the identification of emergent phenomena has had a profound affect on the current understanding of many complex systems. Indeed, some authors would argue that complex systems are worthy of study *because* they display emergent phenomena. For example, Holland (1998) draws lessons from Conway's celebrated *Game of Life*, that: '1) simple rules (absurdly so) produce coherent, emergent phenomena; 2) emergence relies on interactions that are more than the sum of the individual parts (that is, relying on non-linearities); and 3) that persistent emergent phenomena can serve as components of more complex emergent phenomena.' Here, emergence has come to be a meta-description of complex systems in general and that these emergent phenomena can themselves become components for further aggregate behaviour.

Whilst much of the theoretical literature on emergence has studied abstract models such as the *Game of Life* (Berlekamp et al., 1982), Cellular Automata (Wolfram, 1984a), simulated evolution (Ray, 1993) and a variety of Lindenmayer (L-) systems (Lindenmayer, 1971), there is a growing interest in the way that the many decisions made by human actors can aggregate. Concepts such as *herding* in financial systems, or the famous work of Arthur (1994) on inductive reasoning are obvious candidates.[15]

For the economic literature, the most obvious implication of these studies, and of the general phenomena that is described by *emergence*, is that aggregate data need not be described by the combined actions of many self-interested optimizing agents. Rather, through the interactions of many sub-optimal decision making entities, aggregate variables may be generated, sometimes displaying remarkable overall optimizing behaviour, or alternatively, generating seemingly chaotic trajectories. This is not a new realization (see for example, Alchian (1950); Sugden (1989)), but has important ramifications when

[15] See also Surowiecki (2004), *The Wisdom of Crowds*.

Table 3.1 A Fibonacci sequence generated with the two rule L-system (3.9), (3.10).

Iteration	State
0	A
1	B
2	AB
3	BAB
4	ABBAB
5	BABABBAB
6	ABBABBABABBAB
7	BABABBABABBABBABABBAB

coupled with the recent ability to computationally simulate the interactions between large numbers of non-rational agents (Holland and Miller, 1991).

3.7 Modeling Complex Systems

What does a 'complex systems approach' look like (in terms of research method)? The examples that follow can be broadly split into two main approaches: the first, and earlier of the two, aims at uncovering underlying properties of complex systems by investigating abstracted yet transparent models; whilst the second aims at achieving more practical outcomes, either in terms of optimization or explanatory/predictive modeling of real-world systems. I shall take each approach in turn.

3.7.1 Theoretical Approaches

The first style of complex systems modeling is no doubt the most familiar. This owes in part to the surprising insights that this approach originally yielded, and in part to the range and depth of properties that these models have come to be associated with. Examples here include L-systems and Cellular Automata (CA) (including Conway's *Game of Life*, a notable special case of 2-dimensional CAs). L-systems use simple components and rules that together define a *formal grammar* which simulate growth in an algorithmic way. For example, (following Green 2000) an L-system $\mathcal{F}(G)$ on *syntax* G to generate a Fibonacci sequence could be constructed with *variables* $\{A, B\}$, *initial condition* A, and *rules* as given in (3.9) and (3.10).

$$\mathcal{F}_1 : \quad A \longrightarrow B, \text{ and} \tag{3.9}$$

$$\mathcal{F}_2 : \quad B \longrightarrow AB. \tag{3.10}$$

That is, an application of $\mathcal{F}$ for several iterations would yield the sequence in Table 3.1. which can readily be verified to have the structure,

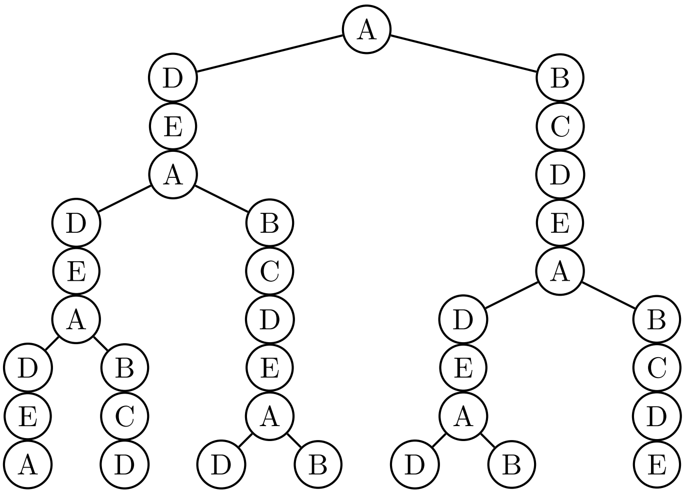

Figure 3.4 Example *Chaetomorpha linum* growth model constructed with L-system $\mathcal{G}$ (see text).

$$|S_n| = |S_{n-2}| + |S_{n-1}|, \ \forall\ n \geq 2, \tag{3.11}$$

where $|S_n|$ denotes the length of state $n \in [0, \infty)$ – the Fibonacci sequence. Alternatively, a model for algal growth (in this case, *Chaetomorpha linum*) $\mathcal{G}$ could be constructed by defining rules:

$$\begin{aligned}
\mathcal{G}_1 : \quad & A && \longrightarrow DB \\
\mathcal{G}_2 : \quad & B && \longrightarrow C \\
\mathcal{G}_3 : \quad & C && \longrightarrow D \\
\mathcal{G}_4 : \quad & D && \longrightarrow E \\
\mathcal{G}_5 : \quad & E && \longrightarrow A\ .
\end{aligned}$$

As can be seen in Fig. 3.4, $\mathcal{G}$ quickly gives rise to a starkly realistic (complex) algal growth pattern. The two examples considered are known as *context-free*, since any rules address only individual elements, rather than groupings (elements in a context). In reality, there is no known bio-chemical basis for the intermediate rules in these grammars (i.e. those that define the timing of branching), and therefore, *context-sensitive* models are more likely a better representation of reality. However, L-systems provide insight as to the generation of self-similar structures since they neatly generate similar characteristics over a range of scales (i.e. fractal properties). The following result (reproduced from Green 2000, THEOREM 2.1) captures this,

Theorem 6 *Let $L(G)$ be a context-free grammar that is generated from the syntax G. If there exists an integer m, such that S_m can be expressed in terms of S_{m-i}, where $i < m$, then this same relationship holds for S_n, for all $n > m$.*

For example, the $\mathcal{F}(G)$ system previously mentioned can be represented,

$$S_n = S_{n-2} \| S_{n-1}, \ \forall \ n \geq 2 \,, \tag{3.12}$$

where $\|$ represents an horizontal concatenation.

CA models, generally attributed to von Neumann (1966) have proved to be a powerful modeling environment in which to understand complex system behaviour. The idea is simple, with a single dimension cellular automaton being a line of sites, $\{a_1, a_2, \ldots, a_n\}$, each capable of k discrete states. Updating of each site follows some deterministic rule (ϕ) contingent on the site's previous value, together with the value of all neighbours up to and including r sites away,

$$a_i^{t+1} = \phi[a_{i-r}^t, a_{i-r+1}^t, \ldots, a_{i+r-1}^t, a_{i+r}^t] \,. \tag{3.13}$$

For example, ϕ might be a *totalistic* rule, such that, a_i^{t+1} is simply defined by,

$$a_i^{t+1} = \hat{\phi} \left[\sum_{j=i-r}^{i+r} a_j^t \right] \,, \ j \in \mathbb{Z} \,, \tag{3.14}$$

where a *look-up* table might be used to define $\hat{\phi}$ for the possible k^{2r+1} summation values (a total of $k^{k^{(2r+1)}}$ such functions exist). Such a two-state ($k = 2$) CA with $r = 2$ was used to generate the selection of sub-figures in Fig. 3.5. Each figure was initialized with a random state seed (top row), with subsequent time steps being printed thereafter.[16]

What is most interesting about even these simple 1-D CAs, is the diverse range of behaviours exhibited. Most of the 'rules' quickly settle to one of the quiescent (Fig. 3.5(a)), periodic (Fig. 3.5(b)) or chaotic states (Fig. 3.5(c)), so-called Class I, II and III behaviour respectively. However, a very small number (measure zero in large sets) display what Wolfram labeled Class IV behaviour (Fig. 3.5(d)), eventually becoming associated with the 'edge of chaos' descriptions already encountered. The CA environment then offers a transparent system to analyse the properties of each dynamic behaviour. Further analysis by Wolfram (1983, 1984b) identified the statistical mechanic properties of these classes (e.g. entropy, information processing) and indicated the possibility that Class IV CAs are

[16] In this figure, black and white pixels indicate cell values 1 and 0 respectively.

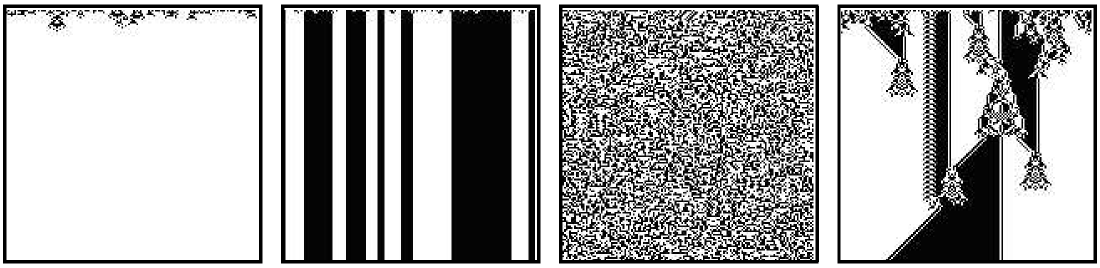

(a) Class I (36) **(b)** Class II (56) **(c)** Class III (10) **(d)** Class IV (52)

Figure 3.5 Example behaviours for Totalistic Cellular Automata: $k = 2$, $r = 2$; (a) homogeneous steady-state; (b) periodic steady-state; (c) chaotic; and (d) transition phase exhibiting complex localized regions, sometimes long-lived. Generating 'rules' given in parenthesis.

able to exhibit *universal computation* (and thus by nature, their state trajectories are inherently *undecidable*). Indeed, the previously mentioned *Game of Life* can be represented as a 2-D CA and similarly has been shown to possess this same property (Berlekamp et al., 1982).

Furthermore, and with clear parallels to Kauffman's work, later investigation by Langton (1992) who defined an order parameter $\lambda \in [0, 1]$ as the percentage of non-zero values in the look-up table, showed that: for $\lambda \to 0$, Class I behavior (quiescent) is observed; Class II behavior (periodic) arises for small, but increasing λ; and as $\lambda \to 1$, Class III behaviour (chaotic) prevails. However, for a small region between Class II and Class III behaviour, a transition state according to the complex patterns of Class IV were observed, suggesting that, as with Kauffman's Boolean nets, a critical value λ_c could be determined to describe the 'edge of chaos.'

Clearly there are many other examples of useful theoretical frameworks for investigating complex system behaviour, we have considered here only two of the most influential. However, from these, several features are apparent, namely: that the 'edge of chaos' is associated with a transient, and quantitative region of the state space; that around this region, emergent phenomena, akin to the frozen/melting sections of random NK nets is apparent; and perhaps most clearly seen in these models, that the *mantra* of many complex systems workers holds true – complex phenomena need not rely on complex generating functions. Gell-Mann (1994, p.313) comments concerning systems such as the CA model that,

> 'The most striking feature of those simulations is the emergence of complex behaviour from simple rules. Those rules imply general regularities, but the working out of an individual case exhibits special regularities in addition.'

3.7.2 Adaptive Approaches

The second, and now ever-popular *adaptive* stream of complex systems research has grown largely from biological metaphors such as *adaptation*, *evolution* and *selection*. Whilst it seems that such an approach is now unavoidable across the academy, the social sciences, and in particular, the economic sciences, have proved eager adopters.

Historically-speaking, evolutionary metaphors that sought to explain relative returns within the economy have been with us for decades (e.g. Alchian 1950). And in particular, with the influential work of Maynard Smith (1978) in mind, the Theory of Games formed a natural substrate for further adoption of (successful) biological metaphors such as random perturbation in equilibrium selection (e.g. Kandori et al. 1993). Moreover, as attributes such as *mistake-making*, *rules-of-thumb*, *conventions* and other forms of non-rational play (e.g. *fairness*) came to form the most surprising results from the experimental/behavioural psychology literature in recent times (e.g. Camerer 1998, 1997), an awareness that reality may in fact be better described by the aggregate behaviour of a number of diverse agents arose within the field (e.g. Arthur 1994). Together with this awareness came the parallel increase in affordability of computing technology such that even standard desk-top PCs now provide a readily accessible platform with which to construct increasingly facile and quantitative models of such complex systems (e.g. Holland and Miller 1991; Tesfatsion 2003).

Subsequently, *adaptive* complex systems approaches generally take one of two related forms, the second being more relevant to the present study. The first harnesses the attractive non-linear optimization properties of adaptive systems to search large-scale multi-dimensional parameter spaces. Here, generations of solutions, typically encoded in binary form, are iteratively selected, combined and 'mutated' to hone in on global optima. The infamous Genetic Algorithm (GA) approach (Holland, 1992) with its diverse applications (e.g. Hillis 1990; Miller 1996a) continues to be the subject of heavy research.

However, it turns out that given the right conditions, such models have appeal in explaining the underlying *behaviour* of agents that gives rise to various aggregate phenomena. For instance, models that employ a GA underpinning often explain generation-to-generation changes in terms of learning (e.g. *via* the *crossover operator* or similar) and innovation, or mistake-making (*via* the *mutation operator*). Alternatively, adaptive models may employ rule-based updating of agent strategy profiles, with (say) some stochastic shock that causes stickiness. Desirable qualities of these approaches are several-fold:

1. Diverse agent behaviours (intelligences) can be implemented, limited only by computational complexity;

2. Agent behaviours can be easily made *dynamic*, allowing for realistic learning/adaptation qualities;
3. Learning rules such as imitation/diffusion and experimentation/mistake-making are afforded within easily tunable parametrisation;
4. Interaction spaces can be varied or endogenised; and
5. Quantification of time-series model attributes such as diversity, self-organization, complexity, similarity etc. are accessible.

As mentioned, the approach was adopted early by workers who studied the evolution of strategic behaviour — games such as the Prisoner's Dilemma was an obvious initial choice, but continues to offer fertile modeling territory (Anderlini and Ianni, 1997; Binmore and Samuelson, 1992; Grim, 1996; Kirchamp, 2000; Lindgren, 1992; Masuda and Aihara, 2003; Miller, 1988, 1996b; Miller et al., 2002; Nowak and May, 1992); as does that of technology transfer/adoption (Elgazzar, 2002), trading (Ashlock et al., 1996; Tesfatsion, 1997), and sector-shock modeling (Scheinkman and Woodford, 1994).

In each of these models, where the representation of firm or agent behaviours and interactions is attempted, several challenges persist. First, because of the necessity to almost exclusively use computational methods in dealing with these many-body modeling environments, the investigator not only has the problem of rendering behaviours in numeric, or even digital form, but also must ensure that any transitions between behaviours should also be 'legal'. That is, that at the end of the operation the agent or firm is still a functioning agent or firm. This is known as the 'encoding problem' and can be a source of great frustration for the researcher. Commonly, a tactic used to overcome this problem is to encode a 'genotype' – a binary representation of each strategy, which after manipulation will still give rise to further (computable) binary strings.

Second, is the challenge of representing realistic learning behaviours by such models. In some ways, this is to be expected, for in attempting to model agents with an higher level of detail than previously implemented, the investigator invites scrutiny as to why some behaviours were included, and others omitted.[17]

Finally, there is the challenge of defining an appropriate 'fitness' criterion. In an optimization setting, adaptive fitness is no more than the focus of the optimization itself. So long as it can be measured, it can be optimized. However, where adaptive agents are used to model real economic and social systems, defining a fitness function is rarely so straightforward. For instance, one of the great insights that these approaches have

[17] An interesting recent comment on this debate is brought to bear by Krakauer (2004), who argues that even the biological story is more complicated than previously thought, suggesting that the traditional simplistic mapping from genotype to phenotype is misplaced, let alone the use of such a mapping in the social sciences.

offered is to show how systems can display both 'intelligent' and sub-optimal behaviour at the aggregate level in the course of a *single* 'run,' a property that traditional theoretic approaches find difficult to capture. Obviously, a *relative* fitness criterion is necessary to facilitate such freedoms, being the basis of adaption itself, but of these, the investigator must be careful – like some views of the market, these models will give back just what the investigator *asks for*, whether or not that is actually what she *wants*.

3.8 The Dangers of 'Complexity'

I shall end this consideration of the complex systems approach with a final reflection on the dangers of complexity. As has been touched on above (§3.7.2), modeling complex economic systems is not without its difficulties. What has not been mentioned so far, however, is the somewhat unique problem of *abundance*. That is, where traditional approaches focus on an extremely rudimentary functional or aggregate level approach, computationally modeling complex systems often requires the investigator to 'drink from a fire-hydrant of information.' The ease with which parameters can be added is astonishing. Very quickly, even the parameter space introduced by the implementation of a model can overwhelm, let alone the actual state space produced by the model itself. Gell-Mann (1994, p.313) – considered by many a founding 'father' of such research – summarises the situation well,

> The trick in designing a manageable simulation is to prune the rules so as to make them even simpler, but in such ways that the most interesting kinds of emergent behaviour remain. The designer of a simulation must then know a good deal about the effects of changes in the rules on behaviour in many different scenarios. ... the design of simple simulations rich in interesting consequences remains more of an art than a science.

For such reasons, the literature is still developing with respect to what constitutes a 'good' contribution in this area, such that these models move beyond the level of simple *metaphor* and begin to provide new and relevant insight. As previous discussions have emphasised (§3.3) this, in part, has resulted in the development of new non-linear measures. However, one suspects that such considerations will be on-going for the near-future at least.

It should not be forgotten that effective execution *is* possible nonetheless, indeed, the survey reported here has covered several contributions that have forced a reconsideration of long-held beliefs about multi-agent interactions and how they aggregate. Indeed, some of the founding contributions have sparked new fields in themselves.[18]

[18] Consider Kauffman and Bak's independent contributions to self-organized criticality, or Wolfram's CA models of complexity and the 'edge of chaos.'

If a 'test' need be applied to the burgeoning field of complex systems in economics in general, or computational economics in particular, then Gell-Mann's further comments, although easily applied to any academic enquiry, seem worthwhile,

> In the end, though, what really matters is the relevance of the simulations to the real-world situations that they imitate. Do the simulations supply valuable intuition about real situations? Do they suggest conjectures about real situations that could be tested by observation? Do they reveal possible behaviours that had not been thought about before? Do they indicate new possible explanations of known phenomena?

3.9 Themes to be Pursued in the Following Chapters

Before moving on, it will be helpful to emphasise a few of the main themes covered in the preceding discussion that will surface again in what follows. The immediately following introduction of the Bala and Goyal analytic model of network formation (Chapter 4), together with the further analysis under no inertia (Chapter 5) in Part II are informed by the discussions above concerning the terminology of networks (as graphs) (§2.3), results covered from the analytical literature (§2.4), and examples of economic networks as handled by the literature (§2.5).

Both the enriched model as presented in Chapter 6 together with the discussions and model concerning cooperation to be found in Chapters 7 and 8, broadly follow the themes of endogenous network formation (§2.4.3) in the analytic sense, but draw on foundations in computational approaches to network modelling (§2.5.3) and graph analysis (§2.6) for their inspiration. Finally, it is in the analysis of state dynamics in Chapter 8 that the more recent discussions concerning complexity (§3.2), measures of complexity (§3.3), and themes of emergence and self-organization (§3.6) and self-organized criticality (§3.5) will be relevant.

part II

Communication Networks

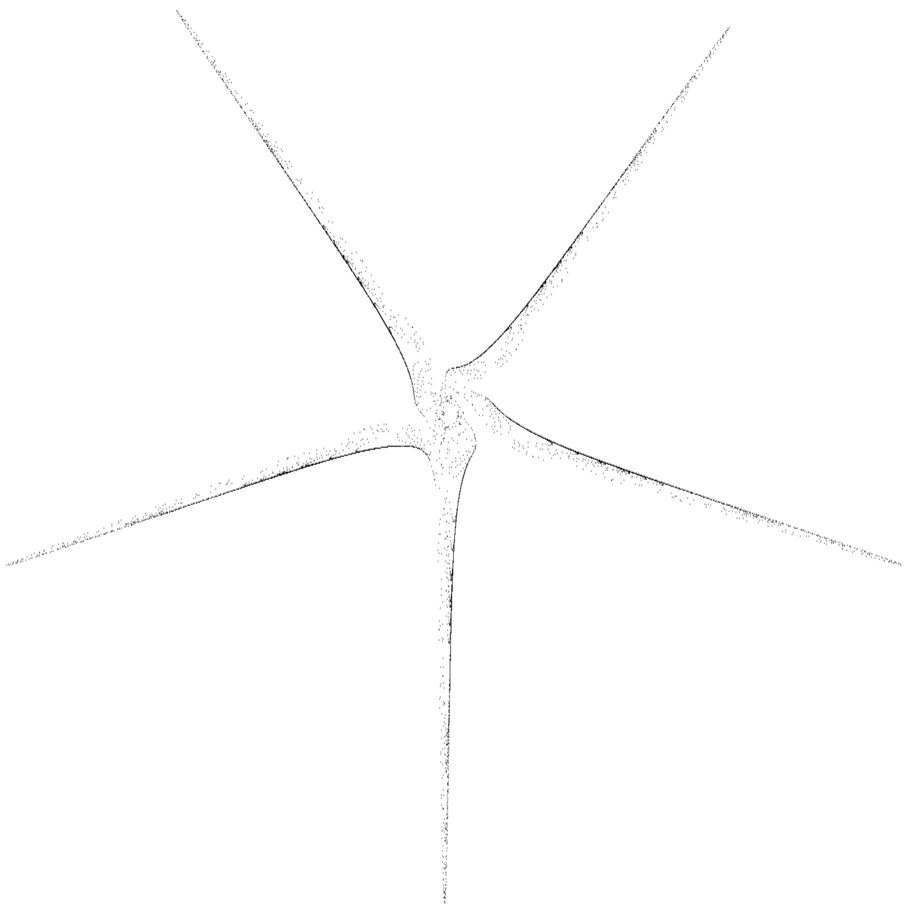

Chapter 4

The Bala & Goyal Model: A Summary

Chapters 5 and 6 to follow speak directly to the non-cooperative communication network formation model of Bala and Goyal. To this end, the present chapter introduces the Bala and Goyal (2000) model, relevant nomenclature, and some important results for the following analysis.

4.1 Introduction

In this part, attention is focused on the non-cooperative network formation model of Bala and Goyal (2000) (hereafter, referred to as BG). The model is a relatively simple implementation of the network formation situation and a summary of its main features is given below. It should also be mentioned that this summary draws on the work of Falk and Kosfeld (2003) (hereafter referred to as FK).

Before going into the details, two features of the BG model that directly impinge on the present work are introduced. First, the BG model assumes that agents update their strategies (in this case, a vector of edge sponsorship decisions to other agents, explained below) based on a myopic *best response* updating rule, i.e. choosing one of their strategies that yield maximum payoffs given the play of the other agents in the preceding period. A point to be returned to presently. Second, agents update this vector with some degree of *inertia*, that is, BG assume that at least one agent decides *not* to respond to the strategies of the other players, and instead, plays their strategy from the previous period again in the present one.

4.2 The BG Model

A set of agents $N = \{1, \ldots, n\}$ play a simple network formation game as follows. Each agent is assumed to hold some information which is valuable not only to himself but to the $n-1$ other players, and importantly, this information can be *observed* by either one of two

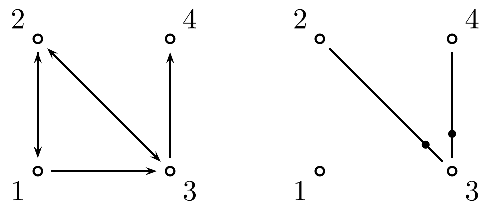

(a) one-way **(b)** two-way

Figure 4.1 Example communication network structures.

network formation processes. In the first case, some representative agent i can sponsor an edge (link) to another agent j and thus gain access to j's information (in addition to their own); such a case is known as the *one-way* information case. In the second instance, an edge sponsored by i to another agent j not only yields access to j's information for agent i, but also, agent j automatically gains access to i's information, in a reciprocal manner; such a case is known as *two-way* information flow.[1] An important implication of the information sharing structure is that each agent obtains access to another agent's complete set of information by linking to them (or being linked to by them, in the case of two-way information flows).

The sponsorship decisions of each agent forms a *strategy* for that agent, which is then an $n-1$ element row vector of the form,

$$g_i = (g_{i,1}, \ldots, g_{i,i-1}, g_{i,i+1}, \ldots, g_{i,n}) \ ,$$

where $g_{i,j} \in \{0,1\}$. The set of all such strategies for agent i is $\mathcal{G}_i$ and has cardinality, $|\mathcal{G}_i| = 2^{n-1}$, consequently, the space of pure strategies for all agents is simply $\mathcal{G} = \mathcal{G}_1 \times \cdots \times \mathcal{G}_n$. It is convenient to depict the strategies of each agent as a graph where the vertex set is simply N and the edge set is an ordered (or un-ordered, in the case of two-way flow) set of vertex pairs such that $e_{i,j} = (i,j) \ \forall \ g_{i,j} = 1$.

Figure 4.1 gives an example of two information networks formed under one-, and two-, way information flows respectively.[2] By way of example, agent 2 in 4.1(a) has information set $\{1,2,3,4\}$, whereas agent 4 has information set $\{4\}$ only. In the two-way case (see 4.1(b)) agent 2 has information set $\{2,3,4\}$ even though she is playing strategy $g_2 = \{0,\ldots,0\}$.

[1]To the Theory of Graphs, these cases are simply *directed* and *undirected* graphs respectively.

[2]Note: following Bala and Goyal, edge sponsorship under two-way flows is indicated by a filled circle on the edge closest to the sponsoring agent.

4.2.1 Payoffs

Nomenclature is now defined for the count of agents that some agent i sponsors a link to, and the number of agents that agent i is explicitly and implicitly connected to (i.e. who they can *observe*). For one-way information flow, define by $N^d(i;g) = \{j \in N | g_{ij} = 1\}$ the set of agents that i maintains a link to, and by $N(i;g)$ the set of agents for whom there exists a *path* from agent i to $k \in N(i;g)$. Note again, that in the one-way flow case, this will imply there exists $g_{i,k} = 1$ or some sequence $g_{i,i_1} = g_{i_1,i_2} = \cdots = g_{i_m,k} = 1$, by convention, $i \in N(i;g)$. The maintenance of edges in the communication network confers costs upon the sponsoring agent together with providing benefits due to information observation in the following way. Define by $\mu_i(g) = |N(i;g)|$ the count of agents that i can observe, and by $\mu_i^d(g) = |N^d(i;g)|$ the number of agents that i sponsors a link to. Then, the agent's payoff function $\Pi_i : \mathcal{G} \to \mathbb{R}$ is given by,

$$\Pi_i(g) = \phi(\mu_i(g), \mu_i^d(g)) \ , \tag{4.1}$$

where ϕ is strictly increasing in $\mu_i(g)$ and strictly decreasing in $\mu_i^d(g)$. In this way, ϕ implements the benefits of information observation accruing to each agent and the costs of edge sponsorship they incur as a result of non-zero elements in their strategy vector g_i.

Under two-way information flow, it is useful to define by $\bar{g} = \mathrm{cl}(g)$ the *closure* of g such that $\bar{g}_{i,j} = \max\{g_{i,j}, g_{j,i}\}$. Consequently, define by $N(i;\bar{g})$ the set of agents that i observes under two-way communication and hence, $\mu_i(\bar{g}) = |N(i;\bar{g})|$ as its cardinality. It can be observed that under two-way information, $N(i;\bar{g})$ implies that a set of adjacent edges must exist, say $\{e_1, \ldots, e_m\}$ such that $i \in e_1$ and $k \in e_m$ for all $k \in N(i;\bar{g})$. For two-way flow the definitions of $N^d(i;g)$ and $\mu_i^d(g)$ are unchanged. Similarly, the payoff function is defined in an analogous way to (4.1) above,

$$\Pi_i(g) = \phi(\mu_i(\bar{g}), \mu_i^d(g)) \ , \tag{4.2}$$

Bala and Goyal analyse a special case of the function ϕ in some detail, that of linear payoffs,[3] e.g.,

$$\Pi_i(g) = \mu_i(g)V - \mu_i^d(g)C \ , \tag{4.3}$$

where V and C are the *benefit* associated with observing one player's information and the *cost* incurred by sponsoring one link respectively.

[3]The special case is defined in an analogous way for the two-way case.

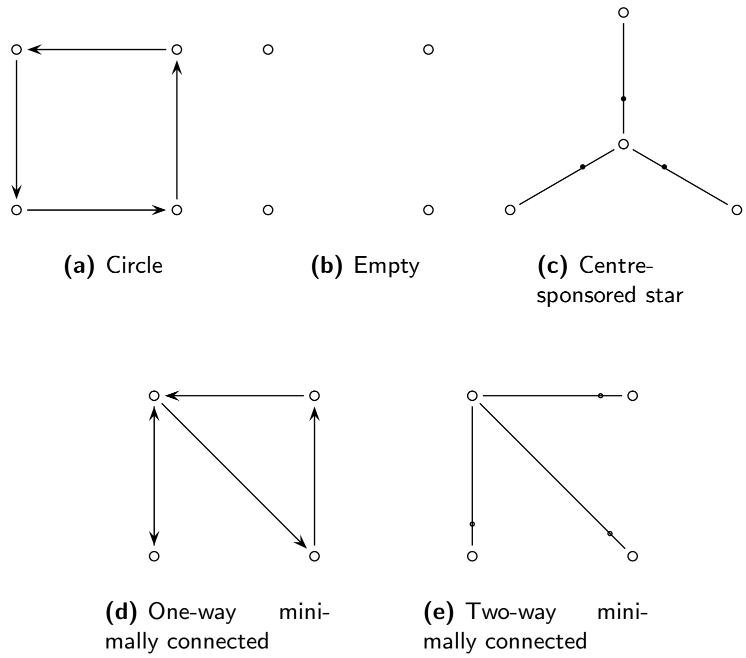

Figure 4.2 (a-c) Nash and Strict Nash equilibrium structures (as referred to in text); and (d-e) Minimally connected structures under one-way and two-way information flows.

4.2.2 Network Structures

It will be useful to introduce key network structures as described by BG, Figure 4.2 shows each of the structures referred to below. The *circle* is related to the *cycle* (of Graph Theory), however, to BG, the path must begin and return at the same node, traversing each node and edge once. The circle structure possesses the useful property that although each agent sponsors only one link, through the coordination successes of the structure, all agents enjoy full node observation through information flows. The *minimally connected* structures demand, in the one-way case, that agents enjoy full observation but the cutting of any one link would result in this not being the case. In the two-way minimally connected structure, this property must hold, and furthermore, there must be no cycles, nor any two agents who *both* sponsor a link to each other.

4.2.3 Nash, Strict Nash, & Efficient Networks

It is useful to define by g_{-i} the network that remains when all of some agent i's links have been removed and shall often be referred to as the *absentee* network for i. Hence, the union of the edges of g_i and g_{-i} will return the full network g. The set of all i's *best responses* to this *absentee network* is denoted by $BR_i(g_{-i})$, where i plays the best response to g_{-i} if,

$$\Pi_i(\{g_i, g_{-i}\}) \geq \Pi_i(\{g'_i, g_{-i}\}) \ , \quad \text{for all} \quad g'_i \in \mathcal{G}_i \ . \tag{4.4}$$

It is to be noted (for our study) that there may be more than one g_i that satisfies the *best response* condition, in which case, it is assumed that the agent chooses one of these strategies in an equiprobable manner. The best response criterion then gives rise to a definition of a *Nash Network* as follows.

Definition 1 Nash Network *The network* $g = \{g_1, \ldots, g_n\}$ *is a Nash Network if, for all agents* $i \in N$*,* $g_i \in BR(g_{-i})$*.*

Furthermore, the refinement of this Nash definition to give a single structure is obtained by the *Strict Nash Network* as follows.

Definition 2 Strict Nash Network *The network* $g = \{g_1, \ldots, g_n\}$ *is a Strict Nash Network if, for all agents* $i \in N$*,* $g_i \in BR(g_{-i})$ *and* $|BR(g_{-i})| = 1$*.*

That is, not only does each agent play a best response strategy, but each has only one such strategy to choose from.

A simple *welfare measure* is defined to be the sum of payoffs accruing to each agent,

$$W(g) = \sum_{i=1}^{n} \Pi_i(g) \ , \tag{4.5}$$

which gives rise to the definition of an *efficient network* as follows.

Definition 3 Efficient Network *A network* $g \in \mathcal{G}$ *is said to be efficient if* $W(g) \geq W(g')$ *for all* $g' \in \mathcal{G}$*.*

4.3 Results

The summary is concluded by reviewing some important results from BG. The first of which is the characterisation of *Nash Networks.*[4]

Proposition 2 (BG, Prop 3.1 and 4.1) *In the 1-way flow model a Nash network is either empty or minimally connected . In the two-way information flow model a Nash network is either empty or minimally two-way connected.*

The second result introduces, for the linear payoffs case, the various structures to be expected when link-sponsorships are low, equal, or high relative to the benefit of observing one agent's information.

Proposition 3 (BG, Prop 3.2 and 4.2) *In the one-way flow model a Strict Nash Network is either the empty network or the circle. In the linear payoffs case, if* $V = 1$ *and* $c < 1$ *the circle is the unique Strict Nash Network. If* $1 < c < n - 1$*, both the empty network and the circle are a Strict Nash Network. If* $c > n - 1$*, the empty network*

[4] The results summary below follows that of FK.

is the unique Strict Nash Network. In the two-way flow model a Strict Nash Network is either the empty network or the center-sponsored star. In the linear payoff case, the center-sponsored star is the unique Strict Nash Network if $c < 1$*, and the empty network is the unique Strict Nash Network if* $c > 1$*.*

Finally, efficient structures are identified, clearly corresponding to Nash play in each case.

Proposition 4 (BG, Prop 3.3) *In the one-way information flow case with linear payoffs and* $V = 1$*, the circle is the unique efficient network if* $c < n - 1$*, while the empty network is the unique efficient network if* $c > n - 1$*. In the two-way information flow case with linear payoffs, if* $c \leq n$*, a network is efficient if and only if it is minimally two-way connected. If* $c > n$ *the empty network is the unique efficient network.*

Chapter 5

Network Formation Under the Best-reply Dynamic: Non-convergence

Having introduced the non-cooperative communication network formation model of Bala and Goyal (2000) in the previous chapter, this chapter develops new results concerning conditions equal to, or approximating, no strategic inertia in agent updating.

5.1 Introduction

The preceding chapter has outlined the main features of one relatively recent, but certainly prominent, analytic network formation model, namely that of Bala and Goyal. This model has many appealing properties, not least of which is its relative simplicity: actions of agents are constrained to a simple yes/no decision vector of sponsorships; this decision vector completely defines an agent's strategy (there is no subsequent strategic intent or interaction on the network); and behaviour is simply described by myopic best-response updating (issues of fore-sight, boundedly rational play, fairness, reciprocity etc. do not interfere). From within this deceptively simple set-up, however, there is a very many ways that the game might be played – the strategy space, even for small n is immense, befitting the difficult realm of *endogenous network formation*.

As has been mentioned, Bala and Goyal derive a number of useful and strong results on the nature of networks formed under various specifications. In the present work however, I wish to investigate in detail, the implications of one of their fundamental modelling assumptions, that of agent-updating as the *best-response* dynamic. Or rather, to be precise, I study the effects of this dynamic when the assumption of *strategic inertia* is relaxed, or is very nearly relaxed.

Strategic inertia for agents describes the modelling prescription that with some non-zero probability, an agent will choose *not* to update their current strategy via the best-response dynamic, but rather, will choose to keep the same strategy as in the previous period. That is, suppose that r is the probability of displaying inertia, then define $p = 1 - r$ to be the probability that an agent will update their strategy (using the best-reply dynamic) in the present period. I shall hereafter refer to p as an agent's *probability of updating*. In their model, Bala and Goyal assume that p is strictly less than unity, and this assumption features prominently in their proof-work for the major network structure propositions. Indeed, this assumption, coupled with the fact that agents are made to *randomize* across optimal strategies if they find more than one best-response to play, ensures that the Markov chain on the space of all network structures never rests on a non-strict Nash network. However, as the authors note, it is possible that 'the Markov chain cycles permanently without converging to a strict Nash network' (see footnote, p.1184).

The question is, how severe is this possibility on the dynamics of the model? In this piece, I wish to investigate this question, and so ascertain the suitability of such dynamics to the network formation problem. I find the following three results: i) that Theorem 3.1(a) (of BG2000) can be extended to the $n = 3$, no inertia case (only); ii) by counter-example, for $n > 3$ and no inertia Theorem 3.1(a) cannot be extended; and iii) as n grows large, the predominant result of the best response dynamics is non-convergence, with a two-period cycle obtained with increasing probability.

5.2 Convergence: $n = 3$

The intuition of the following result rests on observing that the best response decision for any agent with respect to a second agent depends wholly on the strategy decision of the third agent towards the second. That is, in this most simple case, the presence of an exploitable externality completely drives the strategy decisions of each agent. This observation leads to a set of solvable dynamic equations derived directly from the relevant adjacency matrices for two subsequent (periodic) states.

Proposition 5 *Where $p = 1$ and $n = 3$, the only strict-nash equilibrium network arising from one-way information flows ($V > c$) is the wheel network.*

Proof Consider some 3 vertex graph $\Gamma(t)$ having adjacency matrix $G(t)$ with elements g^t_{ij} ($i, j \in \{1, 2, 3\}$; $g_{ii} = 0$). Let each agent's strategy space, S^i be composed of strategy vectors, s^i such that $s^i_j \in \{0, 1\}$ describes the sponsorship decision of agent i towards agent j ($i \neq j$): for example, for $i = 1$, $s^1 = \{s^1_2, s^1_3\}$. Define by $\Pi_i[\hat{s}^i : \Gamma(t)]$ the payoff to agent i given $\Gamma(t)$ and strategy $\hat{s}^i$. Now since V, the value of 'observing' one agent's information is always greater than c, the strategy 'no-link sponsorship', $s^i = \{0, 0\}$, is strictly dominated by strategy 'sponsor-all' $s^i\{1, 1\}$: $\Pi_i[\{1, 1\} : \Gamma(t)] > \Pi_i[\{0, 0\} : \Gamma(t)]$ and hence, all agents must have at least one non-zero element in s^{i*}, their best response strategy. Now further denote $G(t)$ by the matrix:

$$\begin{bmatrix} 0 & A_t & B_t \\ C_t & 0 & D_t \\ E_t & F_t & 0 \end{bmatrix}$$

and, note that for any agent i, the best response strategy $s(t)$ will depend on the value of g_{kj}^{t-1} for $k \neq i, j$:

$$g_{ij}^t = \left\{1 - g_{kj}^{t-1}, \quad \forall\, k \neq i, j\right\} \quad \forall\, i \neq j \tag{5.1}$$

It is the other non-zero element of the jth column of G(t-1) that determines the value of g_{ij}^t.[1] From this observation, several dynamic equations can be written due to BR updating:

$$\begin{aligned} A_t &= 1 - F_{t-1} \\ B_t &= 1 - D_{t-1} \\ C_t &= 1 - E_{t-1} \end{aligned}$$

And in equilibrium, this must be the case for all periods, thus we can omit the time subscripts and write,

$$A = 1 - F \tag{5.2}$$
$$B = 1 - D \tag{5.3}$$
$$C = 1 - E. \tag{5.4}$$

Likewise, knowing that all agents must have at least one non-zero element in their best response strategy we can write three further equations,

$$A + B \geq 1 \tag{5.5}$$
$$C + D \geq 1 \tag{5.6}$$
$$E + F \geq 1. \tag{5.7}$$

Moreover, by summing over (5.2) to (5.4) it is trivial to note that in fact equations (5.5) to (5.7) can be written as equalities. However, since we do not have full specification,

[1] The intuition for this rests on understanding the externality that is created by link formation. If agent i is considering sponsoring a link to agent j, it will *always* be efficient to make the link herself, unless agent k (of three agents) is sponsoring a link to j themselves. In such a case, it is clearly better for i to exploit the externality thus created, and not form a direct link to j.

any choice of one value for some parameter will solve the system. Thus, only two possible solutions arise,

$$G' = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix} \quad G'' = \begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}. \tag{5.8}$$

The clockwise and counter-clockwise 'wheel'.□

Proposition 6 *For $n = 3$, any initial graph subject to best response dynamics as described in BG2000, with $p = 1$ converges in finite time to a wheel network, and consequently Theorem 3.1 can be extended.*

Proof Consider the $n = 3$ adjacency matrix as before with entries g_{ij}, however, let each g_{ij} be continuous Reals on $[0, 1]$. That is, consider them as probabilities that an edge exists from agent i to agent j. Now, we notice, as above, that the full description of an agent's strategy is contingent on the preceding period, that is, $s_i = \{g_{ij}, g_{ik}\}$ where $j, k \in N/\{i\}$, and is given by,

$$g_{ij}^{'} = (1 - g_{kj}) + (\alpha)(g_{jk}g_{kj}) \tag{5.9}$$
$$g_{ik}^{'} = (1 - g_{jk}) + (1 - \alpha)(g_{jk}g_{kj}) \tag{5.10}$$

where $g_{ij}^{'}$ indicates value of g_{ij} in the following period. These equations capture the two cases in which an agent i will sponsor a link to j: i.) if k did not sponsor a link to j; or ii.) if both j and k sponsored a link to each other then i will randomize with probability α of sponsoring to j and the remainder, $(1 - \alpha)$ of sponsoring to k.

Now, dynamics are then given by taking the partial derivatives of (5.9) and (5.10),

$$\frac{\delta g_{ij}^{'}}{\delta g_{kj}} = \alpha g_{jk} - 1 \leq 0 \tag{5.11}$$
$$\frac{\delta g_{ij}^{'}}{\delta g_{jk}} = \alpha g_{kj} \geq 0 \tag{5.12}$$
$$\frac{\delta g_{ik}^{'}}{\delta g_{kj}} = (1 - \alpha)g_{jk} - 1 \leq 0 \tag{5.13}$$
$$\frac{\delta g_{ik}^{'}}{\delta g_{jk}} = (1 - \alpha)g_{kj} \geq 0 \tag{5.14}$$

Which imply that given any starting probabilities of edge sponsorship (that is, any starting network configuration), there will be two sub-sets of edges that within these sets will move in the same direction, whilst between sets will move in opposite directions, namely,

$$\{g_{ij}, g_{jk}, g_{ki}\} \qquad \text{and} \qquad \{g_{ik}, g_{ji}, g_{kj}\}$$

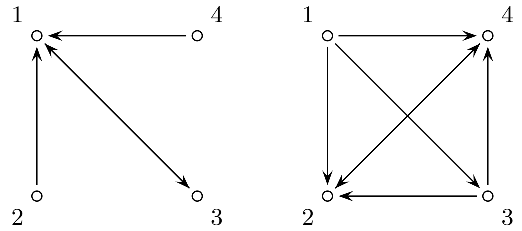

Figure 5.1 Two-period cycle obtained under no-inertia ($p = 1$) and $n = 4$.

or, for clarity, in terms of the alphabetised adjaceny matrix above,

$$\{A, D, E\} \updownarrow \{C, B, F\}$$

Where $\updownarrow$ implies that the sets move in opposite directions. This gives rise two possible absorbing states,

$$G^a = \begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \quad \text{or} \quad G^b = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix}$$

the clockwise and counter-clockwise wheel. □

5.3 Non-Convergence: A Counter-example at $n = 4$

One counter-example at $n = 4$ with best response dynamics and no-inertia is used here (Fig. 5.1) to show that Theorem 3.1(a) cannot be extended past $n = 3$ under these conditions. It is to be noted that this is a *deterministic* cycle, in the sense that all agents have only one BR strategy to the absentee-graph g_{-i} each period (they do not randomise over a number of optimal strategies).

Proposition 7 *Theorem 3.1(a) cannot be extended to the one-way information case without inertia ($p = 1$) for $n > 3$.*

Proof Denote by g^O and g^E the odd (left) and even (right) graphs in Figure 5.1 respectively. Further, we shall use $\{\mu^j(g_i^X) : i, j \in \{1, 2, 3, 4\}; X \in \{O, E\}\}$ to denote the observable information of agent j in either the odd or even graphs, *sans* i's links (by convention, j can always observe themselves). We can then obtain the best response of agent i to g^X and shall use $BR_i(g^X)$ for this purpose. Finally, we also define by $BR(g^X)$ the best response adjacency matrix obtained by this process.

Consider first the odd graph (left). For agent 1: $\mu^2(g^O_{-1}) = \{2\}$; $\mu^3(g^O_{-1}) = \{3\}$; $\mu^4(g^O_{-1}) = \{4\}$; and thus, $BR_1(g^O) = \{2, 3, 4\}$. For agent 2: $\mu^1(g^O_{-2}) = \{1, 3\}$; $\mu^3(g^O_{-2}) = \{1, 3\}$; $\mu^4(g^O_{-2}) = \{1, 3, 4\}$; and thus, $BR_2(g^O) = \{4\}$. For agent 3: $\mu^1(g^O_{-3}) = \{1, 3\}$; $\mu^2(g^O_{-3}) = \{2, 1, 3\}$; $\mu^4(g^O_{-3}) = \{4, 1, 3\}$; and thus, $BR_3(g^O) = \{2, 4\}$. For agent 4: $\mu^1(g^O_{-4}) = \{1, 3\}$;

$\mu^2(g^O_{-4}) = \{1,2,3\}$; $\mu^3(g^O_{-4}) = \{1,3\}$; and thus, $BR_4(g^O) = \{2\}$. Which can be summarised, thus,

$$BR(g^O) = \begin{bmatrix} 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 1 \\ 0 & 1 & 0 & 0 \end{bmatrix},$$

the adjacency matrix of g^E.

Now, consider the even graph (right). For agent 1: $\mu^2(g^E_{-1}) = \{2,4\}$; $\mu^3(g^E_{-1}) = \{2,3,4\}$; $\mu^4(g^E_{-1}) = \{2,4\}$; and Thus, $BR_1(g^E) = \{3\}$. For agent 2: $\mu^1(g^E_{-2}) = \{1,2,3,4\}$; $\mu^3(g^E_{-2}) = \{2,3,4\}$; $\mu^4(g^E_{-2}) = \{2,4\}$; and thus, $BR_2(g^E) = \{1\}$. For agent 3: $\mu^1(g^E_{-3}) = \{1,2,3,4\}$; $\mu^2(g^E_{-3}) = \{2,4\}$; $\mu^4(g^E_{-3}) = \{2,4\}$; and thus, $BR_3(g^E) = \{1\}$. For agent 4: $\mu^1(g^E_{-4}) = \{1,2,3,4\}$; $\mu^2(g^E_{-4}) = \{2,4\}$; $\mu^3(g^E_{-4}) = \{2,4\}$; and thus, $BR_4(g^E) = \{1\}$. Which can be summarised, thus,

$$BR(g^E) = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{bmatrix}.$$

the adjacency matrix of g^O, which completes the proof. □

It is to be noted that computational results suggest that for any $n > 3$, there exists a non-convergent two-period cycle.

5.4 Non-Convergence: Increasing Dominance with n

Obviously, such an observation would be scarcely important if the non-convergent phenomenon was only ever an improbable facet of the dynamics. To test this hypothesis, a computational model[2] was developed to investigate the case in question ($V = 1$, $c \in (0,1)$).

Initially, Bala & Goyal's experiments were replicated to verify the model (see Table 5.1), then, setting $p = 1$ and adding a further 'catch' to the model to identify periodic behaviour, the following convergence table was produced for $n \in \{1,\ldots,8,12\}$ (Table 5.2). A 'periodic cycle' type identification required a double trigger. First, a recurring (twice) two-period graph cycle had to be evident, and second, the count of best response strategy vectors for each agent in all four periods (a double cycle) had to be unity. That is, as per the model presented by Bala and Goyal, if agents determine more than one feasible

[2]All modeling conducted in the MATLAB programming language. Full source code available upon request.

Table 5.1 Model verification, see Bala & Goyal (2000) TABLE 1 (p. 1201), column one.

n	Experiment $p = 0.2$	Bala & Goyal $p = 0.2$
3	14.80 (0.53)	15.29 (0.53)
4	22.86 (0.71)	23.23 (0.68)
5	29.43 (0.89)	28.92 (0.89)
6	37.51 (1.09)	38.08 (1.02)
7	49.90 (1.38)	45.90 (1.30)
8	54.46 (1.63)	57.37 (1.77)

Notes: $c \in (0, 1)$; 500 repeats, standard errors in parenthesis.

best response strategy they will randomise over these strategies, and for the purposes of a recurring cycle, such randomisation is untenable.

Table 5.2 No inertia: $c \in (0, 1)$ and $p = 1$ for $n \in \{1, \dots, 8, 12\}$.

n	Non-conv.	Wheel	Cycle	$log(f_w)$	$log(\hat{f}_w)$
3	0	10,000	0		
4	0	2,044	7,956	−0.690	−0.764
5	0	673	9,327	−1.172	−1.174
6	0	239	9,761	−1.622	−1.584
7	0	91	9,909	−2.041	−1.994
8	0	36	9,964	−2.444	−2.405
12	0	1	9,999	−4.000	−4.045

Notes: f_w and $\hat{f}_w$ are the experimental and estimated wheel convergence frequency respectively, where the basic model $\hat{f}_w = \alpha e^{-\beta n}$, $\alpha = 0.877$ and $\beta = 0.410$ was estimated by a simple $log - linear$ least squared regression.

From Table 5.2, it can be seen that cyclic behaviour is quickly established as the dominant outcome of simulations. Indeed, the wheel graph is obtained as the equilibrium outcome with exponentially diminishing frequency. Thus, rather than the non-convergent behaviour being unimportant or rare, it is the predominant case when $p = 1$.

5.5 Discussion

The present results indicate that the arrival at non-convergent, two-period cyclic behaviour occurs with *greater* speed as n increases ($p = 1$). This stands in stark contrast to Bala & Goyal's reported convergence outcomes with high values of p (periods to convergence to the wheel network approximately scaled exponentially with n).

The existence of cyclic behaviour is not limited to the $p = 1$ case. Indeed, after uncovering the role of cycles in the deterministic case, it was thought that perhaps when

$p \to 1$ cycles might account for the apparent slowing of convergence times at larger n.[3] To test this hypothesis, a number of further numerical experiments were conducted for $p < 1$, with an analytical measure added to count the number of periods spent in a cyclic pattern before the system converged to the wheel structure. Here, as before, a cycle required the double trigger. However, since we are dealing with $p < 1$, it was possible that an agent might not update over the preceding three periods. For the purposes of the investigation, only the updating agents' best response strategies were considered, with non-randomising behaviour again considered as sufficient for a two-period cycle.

Table 5.3 Mean periods to convergence, $p \to 1$.

	p							
n	0.80		0.90		0.95		0.99	
3	6.2	(0.1)	6.6	(0.2)	7.1	(0.2)	7.8	(0.2)
4	13.0	(0.3)	17.3	(0.5)	29.1	(1.0)	98.0	(3.5)
5	29.3	(0.8)	44.6	(1.3)	81.9	(2.7)	350.1	(12.0)
6	54.7	(1.5)	96.9	(3.0)	196.0	(5.8)	858.8	(27.6)
7	111.4	(3.3)	232.5	(7.0)	483.4	(14.8)	2337.8*	(101.0)
8	219.8*	(8.8)	569.3*	(24.4)	1326.9*	(61.1)	6801.8*	(296.1)

Notes: Data represents mean over 1000 independent trials, except where indicated by (*), where 500 trials were conducted; standard errors in parenthesise.

As might be expected, allowing $p \to 1$ gave rise to a further slowing of the convergent dynamics (Table 5.3), but interestingly, for 'moderate' values of $p \in [0.2, 0.8]$, the slowing could not be attributed to cyclic patterns, giving strength to Bala and Goyal's hypothesis that pure mis-coordination is to blame for the slow-down. It is to be noted, that for $p < 0.2$, cyclic behaviour was prominent (and increasingly so as $p \to 0$), with approximately 20% of periods spent in a periodic patterning. However, this is to be expected, since with such a low updating probability, graph sequences over relatively short time frames are likely to be single-state. However, for high, to very-high values of p, two-period cycles did indeed dominate (Table 5.4), rendering an enormous toll on convergence times.

The current findings suggest that the propensity for mis-coordination in a purely best-response dynamic environment is high, with non-sensical cycling of states a feasible and increasingly prevalent outcome under conditions of high strategy revision and 'small' player numbers, and thus, should be borne in mind by practitioners when attempting to model such networks. Indeed, more recent experimental work (Falk and Kosfeld, 2003) suggests that although it is possible for human subjects to attain strict Nash networks in finite time, their decision processes are influenced at least in part by diverse considerations other than the best response. Such thoughts motivate the following chapter.

[3] My thanks to private communications with Sanjeev Goyal for encouraging this line of enquiry.

Table 5.4 Mean fraction of periods spent in cycles, $p \to 1$.

	p			
n	0.80	0.90	0.95	0.99
3	0.000	0.000	0.000	0.000
4	0.003	0.037	0.153	0.507
5	0.003	0.048	0.196	0.650
6	0.002	0.039	0.186	0.685
7	0.001	0.029	0.164	0.683*
8	0.001*	0.020*	0.138*	0.660*

Notes: Data represents mean over 1000 independent trials, except where indicated by (*), where 500 trials were conducted.

Chapter 6

Endogenous Communication Networks with Boundedly Rational Agents

This chapter continues to investigate the non-cooperative communication network formation model of Bala and Goyal (2000). Motivated by the results of the previous chapter, and experimental results of Falk and Kosfeld (2003), a new *artificial adaptive agent* model of agent behaviour is developed and calibrated, finding good agreement with the experimental data.

6.1 Introduction

Communication networks are an apparently ubiquitous feature of many business and interpersonal contexts. In each, depending on the costs of information access and benefits of information content, entities (e.g. firms, individuals) face a strategic problem of who to engage with in mutual or unilateral information sharing partnerships given the actions of other entities that together comprise the information 'landscape'. Although several communication network models have been reported under varying specifications (Chwe, 1995, 2000; Comellas et al., 2000; Jackson and Watts, 2002a; Slikker and van den Nouweland, 2000) this chapter is particularly concerned with the *non-cooperative* communication network formation model of Bala and Goyal (2000)[1] who identified *equilibrium* information network structures under various treatments, including cost of edge sponsorship and direction of subsequent information content flows. Namely, they identify *minimally connected* Nash equilibrium structures in both one- and two- way specifications and using the Strict

[1] To be referred to as BG in this chapter hereafter.

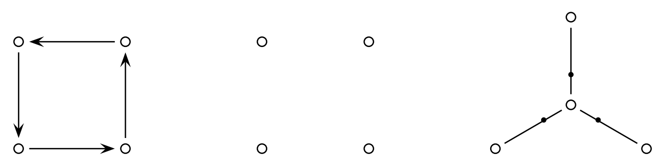

Figure 6.1 Example Strict Nash equilibrium networks, as in BG: (left) Circle (or 'wheel'); (middle) Empty; and (right) Centre-sponsored Star (two-way information flow; smaller dots near centre node indicate sponsorship cost of centre node).

Nash equilibrium refinement[2] obtained specific equilibrium structures, namely: *empty*, *circle*[3] and *centre-sponsored star*, under given treatments (see Fig. 6.1). A more thorough review of this work has been given in Chapter 4 and an investigation of equilibrium convergence was presented in Chapter 5.

In their model, agents are assumed to be able to observe all previous period network sponsorship decisions of their opponents and when given the opportunity to update their strategy, choose the (myopic) *best-response* to this graph. Convergence is guaranteed by in-built strategic *inertia* which describes the condition that at least one agent plays the sponsorship strategy of period $t-1$ in period t (i.e. they choose not to update), a process that searches the space in an incremental fashion, eventually arriving at one of the Nash structures.

However, subsequent human trials of the BG non-cooperative network formation set-up ($n = 4$), conducted by Falk and Kosfeld (2003)[4] find that the BG results receive mixed support in the field. Specifically, in the one-way information flow case, under both low and high costs of edge sponsorship, agents do discover Nash outcomes a majority of the time; especially so towards the end of each treatment. However, in the two-way information case, subjects performed remarkably poorly with respect to finding minimally-connected (Nash) networks as described by BG, barely finding the Strict Nash (centre-sponsored star) outcome *at all*. Such evidence demands further explanation. It is the purpose of this chapter to first construct a richer modelling environment, and second to attempt to uncover reasonable strategic updating assumptions that produce the observed field behaviour.

In their discussion, FK point to various problems for subjects with the two-way Strict Nash equilibrium structure including the serious asymmetry in both a strategic and a

[2]In this case, the Strict Nash refinement implies that not only will each player play a utility maximising strategy in response to each other player's utility maximizing strategy, but each player will only have one such strategy.

[3]That is, arrange all nodes on a circle, connect one to another in a 'daisy-chain' procedure such that each agent's in- and out- degree is unity. NB: BG actually refer to this case as the 'wheel', we shall use the terms interchangeably.

[4]Subsequently referred to as FK in this chapter.

payoff sense, conjecturing that agents dislike such conformations. Further, by use of regression analysis, FK show an apparently strong explanatory correlation between *inertia* (as explained above) and experienced payoff equality in the previous period (controlling for previous best-response play) to support their claim. That is, if other agent payoffs are roughly similar (either low or high) in the previous period, then a subject is more likely to exhibit strategic inertia in the present period.

Such considerations will clearly require a review of the Best Response decision rule to model these contexts. Indeed any form of boundedly rational behaviour due to irrational preferences or other, is not well catered for by the standard best-response decision making rule. It is quite likely that agents are employing (consciously or not) a diverse set of heuristics to determine their strategic play. The problem, of course, is to determine which rules to include in the model. Whilst, for instance, the insightful work of Matros (2004) finds that for generic n-player games, so long as all players are using decision rules from within the union-set of *weakly rational* rules[5] and the best-response rule, then in the short-run the outcomes are 'identical' to that if individuals were constrained to just the best-response rule, giving implicit support to the BG frame-work, it is not clear from FK that such conditions do prevail in reality. Rather, it appears more likely that individuals draw their decision-making rules from outside of such 'minimal curb sets'.[6] Moreover, the FK results show that subjects undergo an incremental improvement processes as they play the game with different opponents. Subjects in the one-way information case found minimally connected structures with greater frequencies in the later stages. Hence, an extension of the BG set-up ought incorporate at least these two elements: one, to allow for a diverse range of feasible decision-making heuristics for each agent such that non-rational play enters in a non-prescriptive fashion; and two, to model some kind of learning process, such that selection between such heuristics occurs in a stage-wise manner.

In the present paper, such considerations are dealt with in a novel way. First, the complexity of decision-making in the network formation game is reduced to a tractable (and implementable) human-decision making process. This is achieved by assuming that individuals are able to recognise and respond to equivalent structures under relabelling as explained below. Second, agent cognition is dealt with by constructing a manageable series of response-rules to each structure under a common reference frame-work but without biasing the relative 'intelligence' of any individual agent. And third, in consequence and in the spirit of Arthur's (1994) artificial bar attendees, since agent cognition architecture

[5] Informally, a set of simple rules is weakly rational if the application of these rules on a sub-set of states (a *minimal curb configuration*) will return one of the member states of this sub-set; the rules are then "consistent with an equilibrium".

[6] Compare the suggested decision-making influences above – those of symmetry and equity.

is common to all agents, between-agent learning and experimentation or mistake-making is afforded and directly observable.

The results of the current approach can be summarised as follows: first, classic payoffs (i.e. benefits *minus* costs, as in BG) appear to be an unsuccessful candidate for the objective criterion used by agents to motivate strategy adoption/revision; second, and in the place of classic payoffs, a *ratio* of benefits and costs, together with a strategic altruism component gives rise to comparable outcomes to that of the experimental data; and third, although equality and best response considerations are found to be significantly associated with strategic inertia as found by FK, these effects appear to be *emergent* in nature, rather than being the driving dynamic in the learning environment as they suppose.

The rest the paper includes a summary of the pertinent results from both BG and FK (§6.2) and a description of the model and its implementation (§6.3). This is followed by computational modelling results and analysis (§6.4) before several concluding comments are drawn (§6.5).

6.2 Background

6.2.1 BG Model Predictions

The key predictions from the BG model are summarised in Table 6.1.

Table 6.1 Predicted structures by BG2000 under each treatment.

Flow	Edge Costs[a]	Structure[b]				
		m1c	circle	empty	m2c	cs-star
One-way	Low	△	▲*			
	High	△	▲*	▲		
Two-way	Low				△*	▲
	High			▲	△*	

Notes: [a] Low $C \leq V$, High $C > V$; [b] structure *m1c* and *m2c* are minimally-connected non-empty graphs in one- and two- way information flow cases respectively. (△) non-empty nash, (▲) strict nash, (*) indicates that the structure is also *efficient* (following FK2003).

Each structure has previously been explained in Chapter 4. Recall that the 'minimally-connected' networks imply that all agents obtain full information, in the sense that they can observe the information of each other agent. However, for the one-way case the minimally-connected property implies that each edge is *necessary*, that is, the removal of any edge will cause some agent(s) to lose their access to the other $(n-1)$ agents' information. In the two-way case, such a criterion results in cycles being ruled out (one edge would be unnecessary), plus any case where two agents mutually sponsor a direct

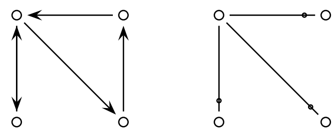

(a) m1c **(b)** m2c

Figure 6.2 Examples of minimally-connected structures (following FK): (a) the one-way information flow case; and (b) the two-way case (periphery-sponsored star).

link to each other (again, one would be unnecessary). Such is the nature of two-way information flows. Examples of each type of minimally-connected structure can be seen in Fig. 6.2.

6.2.2 FK Experimental Findings

The authors of the FK study reproduced the BG network formation context in an extremely faithful manner. For a full description of their study, the reader is referred to the reference. However, to summarise the procedural details, a total of 160 subjects, 32 in each of 5 treatments – 3 under one-way information flows (edge sponsoring costs 5, 15 and 25 points) and 2 under two-way flows (costs 5 and 15 points) – were randomly grouped into 8 mixing groups of 32 subjects to play 5 rounds of a four-person network formation game. After this so-called 'stage', the 32 subjects in each mixing group were shuffled and re-assigned to 8 new groups of 4 (drawn from within the 32 member mixing-group) for a further five rounds. In all, 3 stages were conducted for each treatment. Subjects interacted with a computer screen (after familiarisation) to both observe the sponsorship decisions of their opponents in the preceding round, and to enter their own sponsorship decisions in the current round. Subjects received 10 Swiss Francs for showing-up, and received monetary reward for their play; information observation of one node gained 10 points with all costs (as just mentioned) measured in points with 10 points representing 0.9 Swiss Francs (subjects received around 50 SF on average).

Several FK results are relevant. First, under the four treatments over information flow direction and edge sponsorship costs, the proportions of each graph-type as predicted by the BG results are striking (see Table 6.2). Under one-way information flow, aggregated across all (three) learning stages and all (five) rounds within each stage, a strong tendency for playing Nash more often than not is revealed. Moreover, in a large proportion of cases where a Nash structure was played, the Strict Nash structure (either just the circle-graph, or the circle and the empty graph type) were played. This result is dashed however in the case of two-way flows, where Strict Nash play was non-existent and the Nash structures that did result were played in a significantly lower number of rounds relative to the one-way case, with low edge costs yielding a Nash frequency of around 30% (as opposed to around 50% for one-way), reducing to less than 10% when edge costs were high. Interestingly, average agent edge sponsorship reveals a roughly constant pattern between information flow regimes, indicating that *strategic*, rather than purely edge sponsorship propensity is the likely cause of failure when flows are two-way.

Table 6.2 Summary of selected FK2003 experimental results under relevant treatments. (Compare Table 6.1.)

Flow	Edge Costs^a	Structure					$\langle d_i \rangle$
		m1c^b	circle	empty	m2c^b	cs-star	
One-way	Low (5)	0.48	0.41				1.19
	High (25)	0.59	0.49	0.10			0.76
Two-way	Low (5)				0.31	0.00	0.91
	High (15)			(nr)	0.09		0.75

Notes: ^{a} figures indicated in parenthesise are in 'experimental points' – 10 points equated to 0.9 Swiss Francs (about $US 0.59) and value of one agent's information constant at 10 points; ^{b} minimally connected values include other structures (such as Strict Nash as appropriate); (nr) indicates 'not reported'.

Proportions of Nash play both between and within each learning stage further reveals the problems that subjects experienced with the two-way information regime (see Fig. 6.3). As can be seen, in the one-way case subjects underwent a strong improvement dynamic between stages, and within stages achieved a similar improvement. Within stage improvements, or alternatively, low frequencies in the early rounds of each stage, are attributed by the authors to mis-coordination due to the mixing of subject groups that occurs between learning stages. In the two-way case, the lack of Nash structures in comparison is stark with both the between-stage and within stage improvements not nearly as evident or non-existent.

6.2.3 Background Summary

To summarise, the model introduced in the next section attempts to specifically enrich the standard BG behavioural basis. Agents are initialised with a range of strategic decision rules, which upon revelation through the stage and round-based communication network formation game can be learnt by other agents in the population. Of key interest is whether the dual experimental facts of round-based learning and poor two-way information flow performance can be replicated within this richer framework as a basis for future modelling of human network formation games.

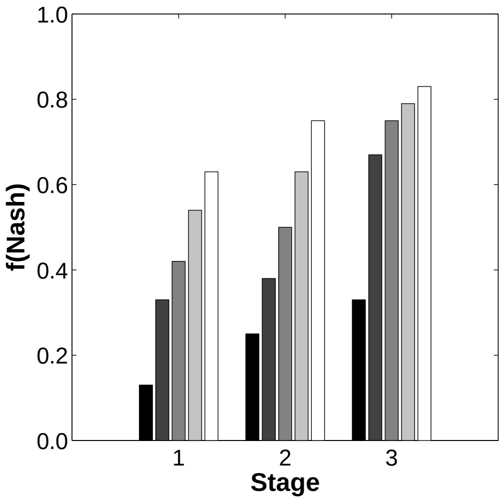

(a) One-way

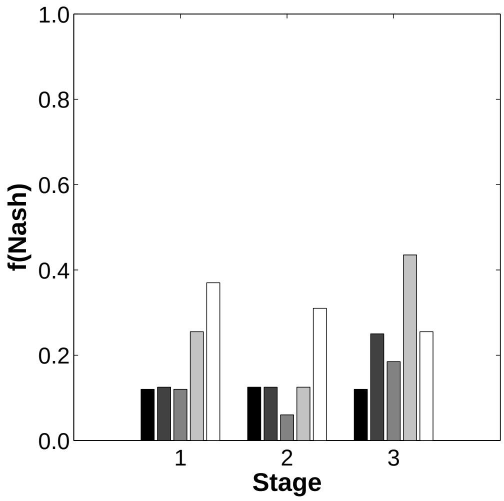

(b) Two-way

Figure 6.3 FK frequency of Nash graph structures under each information flow regime. Bars represent average frequency for all edge costs (5, 15 and 25 points) in the one-way case (a), and over all edge costs (5 and 15 points) in the two-way case (b). Shading represents average frequency in each one of the five rounds per stage across all cost treatments. NB: refer to Table 6.1 for graph structures that are 'Nash' under each information and cost scenario.

6.3 Model

6.3.1 Overview

The present model aims to incorporate reasonable assumptions of human decision-making to investigate and explain the analytic claims and experimental observations of agents in a particular non-cooperative network formation problem. To do this, agents are given a strategy of action that equips them to respond to the network-formation decisions of others through the assumed ability to recognise the type of network before them. As the game progresses, agents who are initially endowed with a mixture of 'good' and 'poor' plays, as measured in action by some objective function, will observe the plays of their counterparts and mimic successful strategies at least in part. Further, between rounds, agents are able to update their own decision-making rules unilaterally through a process of low-level search of the decision space.

6.3.2 The Network Formation Game

Suppose a *playing-group* of agents $N = \{1, \ldots, n\}$ are selected from a population $\mathbf{N}$ to play a communication network formation game.[7] Let the communication network be comprised of vertexes as given by N, and edges as given by the set of edge-sponsorship actions[8]

[7] Nomenclature in this section largely follows that of BG where possible for consistency.

[8] NB: we differ here to the description of BG since g_i gives the *outcome* of an agent's strategy (that is, an action to be taken), rather than comprising the strategy itself, as described by BG.

$g \in \mathbf{g}^n$, $g = \{g_1, \ldots, g_n\}$, where $g_i = (g_{i,1}, \ldots, g_{i,i-1}, g_{i,i+1}, \ldots, g_{i,n-1})$ is an ordered (row) vector of pair-wise sponsorship decisions, and $g_{i,j} \in \{0, 1\},\ \forall\, j \in N/\{i\}$. Correspondingly, let the communication network be represented by the graph $G(N, g) \in \mathbf{G}$.

The communication network confers benefits on agents by gaining them access to the information of other agents. In the terminology of BG, each agent is able to 'observe' the information of an agent they sponsor a link to. Importantly, however, indirect 'flows' of information are allowed. Hence, if $(i \to j) \equiv (g_{i,j} = 1)$ then the graph where $(a \to b)$ *and* $(b \to c)$ is true implies that a can not only observe herself (by convention), they can also observe b (by direct sponsorship) *and* c (by indirect information flow). In the linear benefits specification of BG, followed in both FK and in the present paper, let each agent's information be of value $V \in \mathbb{R}^+$, and let $\{\mu_i(G) \in \mathbb{N} \mid \mu_i > 0\}$ be the number of agents i observes, and further let $C \in \mathbb{R}^+$ be the cost associated with sponsoring one edge, and $\delta_i(G) \in \mathbb{Z}^+$ be the count of edges that i sponsors (that is, the *out-degree* of vertex i). Then define each agent's payoff function $\pi : \mathbf{G} \to \mathbb{R}$ as

$$\pi_i(G) = \mu_i(G)V - \delta_i(G)C \ . \tag{6.1}$$

Further, let a *strategy* (decision-making rule) for an agent i in some stage $t \in \{1, \ldots, T\}$ be a mapping $\mathcal{S}_i^t : \mathbf{G} \to \mathbf{g}$, and denote the set of all such rules $\mathbf{S}$. Within such a stage, each agent will take part in a number of rounds, enumerated by $r \in \{1, \ldots, R\}$. And hence, an agent's strategy $\mathcal{S}_i^t$ will be current for all R rounds of some stage t. Indeed, we can define a communication graph in period r as $G^r(N, \{g_1^r, \ldots, g_n^r\})$, since it will be constructed from each agent's link sponsorship decisions, g_i^r.

The timing of the game is simple in nature. In the first stage, each agent simultaneously reveals their initial sponsorship decisions to form the first round graph $G^1(N, g^1)$. Given the resultant communication network, payoffs are awarded to each agent as in (6.1). After observing the structure of the graph, each agent then applies their decision-making rule, $S_i^1(G^1) = g_i^2$ and then reveals their new sponsorship decisions which forms the new round two communication graph $G^2(N, g^2)$, with payoffs again being awarded accordingly, and so on. This process continues until R rounds have been played, at which point, agents enter a new stage (stage 2) and the process is repeated in full until stage T has been completed. Importantly, however, whilst the retention of an edge sponsorship between rounds is costly (incurring C), the severing of an edge between rounds is not (apart from the potential loss of any information value that will result from the break).

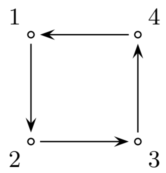

Figure 6.4 Example graph G'.

Example Communication Network

For example, suppose $N = \{1, \ldots, 4\}$, and g' has the form,

$$\begin{aligned} g'_1 &= \{1,0,0\} \ , \\ g'_2 &= \{0,1,0\} \ , \\ g'_3 &= \{0,0,1\} \ , \\ g'_4 &= \{1,0,0\} \ , \end{aligned}$$

then, under one-way information flows, the corresponding graph $G'(N, g')$ would be a cycle, as given in Fig. 6.4.

6.3.3 The Decision-making Rule S

In the present work,[9] $\mathbf{S}$ is implemented so as to permit full inspection of all decision-making processes at any time, and to model possible learning, and/or, decision-making hypotheses. Hence, let $\mathbf{S}$ be of the form (for some agent i), $\mathcal{S}_i = \{s(\mathcal{T}_1), \ldots, s(\mathcal{T}_k)\}$ where $\mathbf{T}(n) = \{\mathcal{T}_1, \ldots, \mathcal{T}_k\}$ is the set of all *minimal absentee graphs (types)* for given n (to be explained below), and $s(\cdot) : \mathbf{T} \to \mathbf{g}$ – the engine room of the decision-making process.

Suppose $\mathcal{G}$ is the set of all n-person graphs, then $\mathbf{T} \subset \mathcal{G}$ is the set of all non-equivalent graphs in $\mathcal{G}$. In fact, all n-person graphs can be constructed from $\mathbf{T}$ under a simple re-labelling of nodes. That is, $\mathbf{T}$ represents the *minimal set of non-equivalent structures*, such that each morphologically similar graph set is represented in $\mathbf{T}$ by a single member of that set. For instance, define a relabelling operator $\mathfrak{R} : G(N, g) \to G''(N'', g)$ where N'' is simply a re-ordering of the elements of N, then

$$\{\mathbf{T} \mid \mathfrak{R} : \mathbf{T} \to \overline{\mathbf{T}}\} \ , \tag{6.3}$$

[9]It is to be noted that for BG, $\mathbf{S}$ has a single member, being the best-response decision making rule, such that g_i^{r+1} (for all $i \in N$) solves the profit maximisation problem,

$$\max_{g_i \in \mathbf{g}} \left[\pi_i(g_i \cap g_j^r)\right] \quad \forall \quad j \in N/\{i\} \tag{6.2}$$

Or, in the case of FK, $\mathbf{S}$ is none other than the decision-making rules that reside in each subject's mind (the locus of our inquiry).

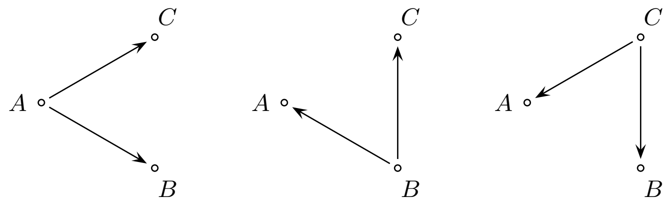

Figure 6.5 A 3-member graph equivalence set.

where $\overline{\mathbf{T}}$ indicates the compliment of $\mathbf{T}$. Then (6.3) implies that any effective relabelling operation applied to some $\mathcal{T}$ in the minimal graph set $\mathbf{T}$ results in a graph outside of the set since by definition $\mathbf{T}$ necessarily cannot contain a second morphologically similar graph. For example, Fig. 6.5 gives the complete set of morphologically equivalent 3-node graphs of the form $\left\{G(N^3, \tilde{g}) \mid \tilde{g} = \{(0,1),(0,1),(0,0)\}\right\}$. Consequently, the set $\mathbf{T}$ (for $n = 3$) must contain exactly one of these (the choice is arbitrary).

The importance of the minimal graph set $\mathbf{T}$ to the present work is the following. First, we shall assume that each agent is able to recognise when two graphs are equivalent under re-labelling, as formulated in A. 1,

A 1 (Type Recognition) *Given k un-identical graphs*

$$\left\{G_1(N_1^n, g), \ldots, G_k(N_k^n, g)\right\}$$

differing only in the ordering of elements in N^n (e.g. $N_1^4 = \{1,2,3,4\}$ and $N_2^4 = \{2,3,1,4\}$), then any agent $i \in N$ will recognise $\{G_1, \ldots, G_k\} \equiv \mathcal{T}_j$, where $\mathcal{T}_j \in \mathbf{T}(n)$.

Such an assumption means that given any graph $G(n)$, an agent will be able to recognise which minimal graph type $\mathcal{T}$ she has in front of her.

Second, as $\mathbf{S}$ has been defined above, the agent must be able to decide on an edge-sponsorship decision g that *applies to the instance*, that is, to the graph G. For this reason, we shall make the second cognitive assumption as below,

A 2 (Context Invariance) *Given any instance of an information network G which corresponds to a minimal graph $\mathcal{T}$, any agent $i \in N$ is able to apply the resultant edge sponsorship decision $s(\mathcal{T})$ to the context, and thus arrive at the g_i that accords to the instance G before her.*

This assumption indicates that an agent, on recognising the basic type that the present graph is drawn from, is able to apply $s(\mathcal{T})$ accurately to the actual graph before them. That is, the agent is assumed to be able to relabel their generic response $s(\mathcal{T})$ correctly to the given circumstance.

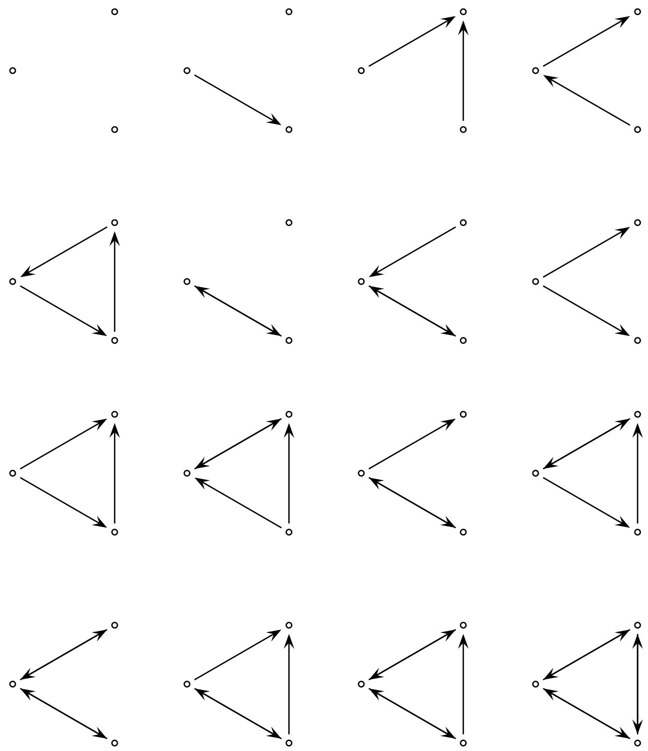

Figure 6.6 Full set of minimal absentee graph types, $\mathbf{T}(3)$, one-way flows.

A final non-trivial simplification of $\mathbf{T}$ is possible for graphs under one-way flows. It can be observed that for some agent $i \in N$, under one-way information flows, the decision making process will only ever need to consider a graph $G(N/\{i\}, g/\{i\})$, that is, the graph G with node i's in- and out- edges removed (by removal of node i). Two related reasons support this observation: first, adjacent inward bound edges – sponsored by another agent to i – do not increase i's observation set, and second, for any decision on some edge $g_{i,j} (j \neq i)$, whether to set $g_{i,j} = (1 \,\text{or}\, 0)$, the state of the link in the previous period is unimportant.[10] Hence, for the n-player game, each agent's strategy vector $\mathcal{S}$ need only prepare responses for the set $\mathbf{T}(n-1)$. For example, if $n = 4$, then the full set of graphs they need to respond to will be comprised of node count $|N/\{i\}| = 3$, and be drawn from the canonical minimal set $\mathbf{T}(3)$ as shown in Fig. 6.6.

For comparison, it is worth noting that the intuition that leads to the reduction of the one-way flow minimal set by one node does not apply in the two-way flow case. Here, since sponsorships result in two way flows, adjacent in-bound edges *are* of importance (they contribute to the current information set of that agent), and so must be considered. Hence, there is not the one-way reduction of important nodes from n to $n-1$. However,

[10]Consider, if $g^r_{i,j} = 0$, then establishing a link will incur cost C, which is the same as if $g^r_{i,j} = 1$, whilst leaving the link un-sponsored has neutral cost impact, which is the same as if $g^r_{i,j} = 1$, as before, since under this specification, there is no added cost for severing a link.

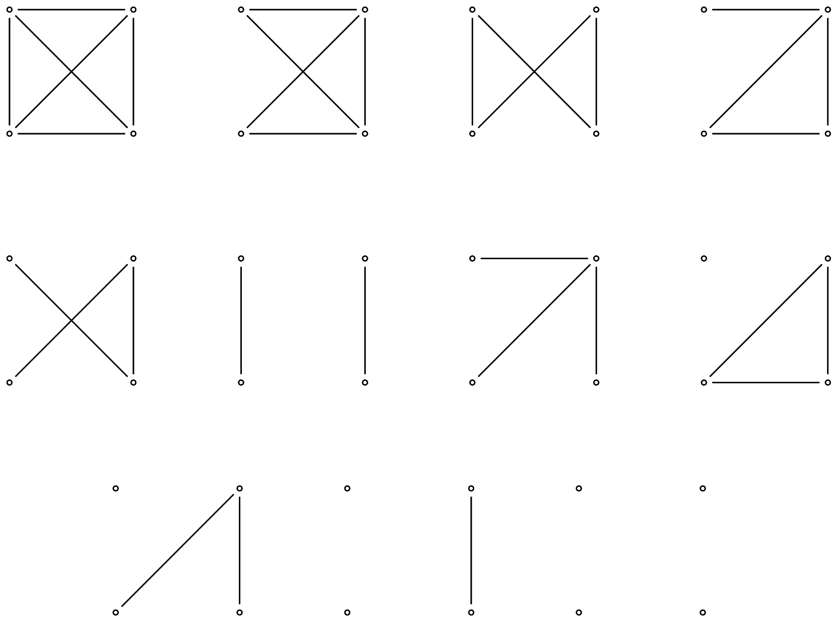

Figure 6.7 Full set of minimal absentee graph types, $\mathbf{T}(3)$, two-way flows.

the situation is made simpler due to the limited number of permutations in the two-way structure. For example the full set of distinct graphs in $\mathbf{T}$ for $n = 4$ has 11 members only and is given in Fig. 6.7.

Example Application of $\mathbf{S}$

Consider a 4-player network formation game under one-way information flows, and suppose that in some stage and round, the current network is of the form $G^r(N, g^r)$, where $g_1^r = (1, 0, 0)$, $g_2^r = (1, 0, 0)$, and $g_3^r = (1, 0, 0)$, whilst $g_4^r = (1, 0, 0)$ (Fig. 6.8(a)). Consider agent 4 (currently sponsoring one link), they must form a decision response to the absentee graph $G/\{4\}$ (Fig. 6.8(b)). Now according to A.1, suppose $G/\{i\}$ is recognised as a member of type $\mathcal{T}_7$, and thus, suppose that she considers her response to such a type (say, g_4^*), and then by A. 2 judges that this response implies the sponsorship decision $g_4^2 = (0, 0, 1)$ in the given context (sponsor a link to 3, as in Fig. 6.8(c)). We might summarise this process as,

$$G^r \longrightarrow G^r/4 \overbrace{\longrightarrow}^{A.\ 1} \mathcal{T}_7 \longrightarrow s(\mathcal{T}_7) \longrightarrow g_4^* \overbrace{\longrightarrow}^{A.\ 2} g_4^{r+1}\ . \tag{6.4}$$

6.3.4 *Properties of the Decision-making Rule* $\mathbf{S}$

Although formal in nature, such a process is directly equivalent to the informal, 'recognise the graph type, and make a response that fits the given instance of that type.' In this

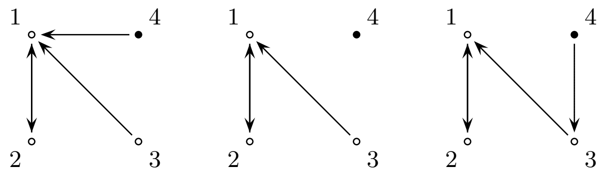

Figure 6.8 Example decision process: (a) The original graph G^r; (b) the absentee graph, or type graph $G^r/\{4\} \equiv \mathcal{T}_7$; and (c) the resultant graph, with 4's decision incorporated $G^r/\{4\} \cup g_4^{r+1}$ (assuming $g_j^{r+1} = g_j^r \ \forall \ j \in N/\{i\}$).

way, the process is intuitive and somewhat obvious, but the reader will realise that such a decision-making process has some pleasing properties.

First, as has been mentioned, the present implementation of $\mathbf{S}$ greatly reduces the cognitive expectation on the artificial agents. It is trivial to show that the total number of distinct graphs that are possible for a given number of agents (without reduction) is $2^{n(n-1)}$ and $2^{\frac{n(n-1)}{2}}$ for directed (one-way) and undirected (two-way) graphs respectively. Which, for example, under one-way information flow and $n = 3$ or $n = 4$ equates to a total of 64 or 4096 graphs respectively. However, in the present formulation, since we focus on the absentee graphs, and then on equivalent *classes* of such graphs only for one-way flows, the total number of distinct structures $|\mathbf{T}(n-1)|$ is 3 and 16 (the count of structures in Fig. 6.6) accordingly; a dramatic reduction in the complexity of the modelling problem.

Thus, we have a strong reason to suppose that the present formulation accords with what an average person could cognitively manage. Obviously, some people will be able to recognise more structures than others, and the numbers given should be treated as the upper-bound of what is required. By way of example, if the full graph set were considered for $n = 5$, a total of $|\mathcal{G}(5)| = 2^{20} \simeq 1 \times 10^6$ graphs would feature, or reducing this to just the absentee set, $|\mathcal{G}(n-1)| = |\mathcal{G}(4)| = 2^{12} = 4096$, whilst by constraining this set to just the *minimal absentee graphs*, that is $|\mathbf{T}(n-1)| = |\mathbf{T}(4)| = 218$, a four order of magnitude reduction is achieved. Figure 6.9 shows these counts for $n = \{2, \ldots, 5\}$ under one-way flows.

Second, as is the intention of this model, the specification of $\mathbf{S}$ in this way clearly allows for all manner of updating strategies. It should be pointed out that the 'best response' function (see footnote above) is accommodated in this framework – the best response is none other than the solution to a profit maximization problem contingent on the absentee graph as described above. However, in the present specification, by constructing $\mathbf{S}$ in the current way, we allow for all kinds of response functions. For example, strategies 'always

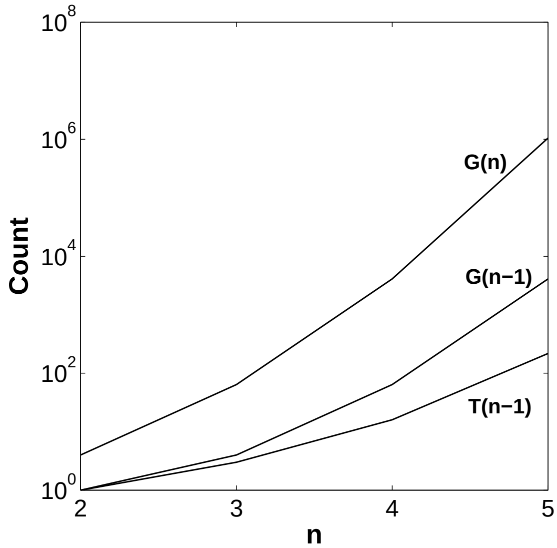

Figure 6.9 Count of distinct graphs in possible response sets under one-way flows: the full graph set $\mathcal{G}(n)$, the absentee full set $\mathcal{G}(n-1)$, and the minimal absentee set $\mathbf{T}(n-1)$.

sponsor-none',

$$s(\mathcal{T}_k) = (0, \ldots, 0) \; \forall \; \mathcal{T}_k \in \mathbf{T}(n) \; ;$$

'always sponsor-all',

$$s(\mathcal{T}_k) = (1, \ldots, 1) \; \forall \; \mathcal{T}_k \in \mathbf{T}(n) \; ;$$

and various strategies in between (e.g. 'uniform random':

$$s(\mathcal{T}_k) = (a, \ldots, a) \; \forall \; \mathcal{T}_k \in \mathbf{T}(n) \; ,$$

where $E[a] = \frac{1}{2}$) are all possible. In fact, the total strategy space $|\mathbf{S}|$ as a function of n under one-way flow is given by,

$$|\mathbf{S}(n)| = 2^{(n-1)|\mathbf{T}(n-1)|}$$

and for n equal to three or four is 512 and 2.8×10^4 respectively! Thus we have a truly rich environment, whereby via inoculation, or some other process (e.g. learning, see below), diverse agents can be modelled and interact.

6.3.5 Learning & the Decision-rule S

It is clear from the experimental work of FK that subjects engaged in this non-cooperative decision problem undergo a process of *improvement* throughout the stages of the game (as indicated by the positive gradient to Nash networks as reported in their paper, see Fig. 6.3). This kind of activity is natural in such a complicated setting since it employs largely foreign terms of reference (subjects would rarely play, or be aware that they are playing, such a game) and so, the game itself provides a forum whereby 'good' plays are revealed to agents over time.

Whilst it is noted that BG (for example) were largely interested in the equilibrium (long-run) outcome of the non-cooperative network formation process, the interest of this paper is on how networks that are realised by human decision making come about. Hence, it is incumbent on the modeler that some attempt at modeling this fact is made.[11]

In the present framework agents are able to develop their strategies by learning from other agents. Such a learning mechanism is equivalent to a (horizontal) *cultural transmission* process (see Cavalli-Sforza and Feldman (1981)). Learning dynamics generally fall into either an R-shaped (limited exponential growth) or S-shaped curve (logistic growth). The former is identified with *environmental* learning factors where agents gain increasing proficiency through practice (see Jovanovic and Nyarko (1995)). The latter is associated with *biased cultural transmission* where the agent imitates (learns) off a subset of other agents, who are often esteemed due to their prestige, or the value of their good decisions. Empirical studies suggest that the majority of human learning is of this second kind (Henrich, 2001).

In the present problem, although the dynamics of learning are ambiguous towards the type of the curve in the empirical data, both the nature of feedback (e.g. full-strategy choices of other agents) and type of problem (i.e. many-agent coordination) point to an n-person coordination problem, and hence, further re-enforce the hypothesis that a social transmission process is at play in the network formation game. It is to be noted, that such *horizontal* transmission, along with being prevalent in human learning, can be a powerful form of problem solving, especially where the problem environment is dynamic (see Acerbi and Parisi (2006)).

In the first stage ($t = 1$), agents are endowed with some decision rule (e.g. random allocation as described in the foregoing paragraph), which they employ for all rounds (for all $r = \{1, \ldots, R\}$) in that stage. At the end of this stage, a measure of the successfulness of each agent strategy is applied (e.g. $\overline{\pi} = \frac{1}{R}\sum_{r=1}^{R} \pi^r$) and used to rank each agent's performance. The highest ranked agent (or agents, if more than one agent shares the

[11] It is to be noted that BG's solution process (best-response updating with inertia) might be construed as a type of 'learning' process, since gradual steps towards Nash structures results. However, in this case, agents are not actually changing anything in their underlying decision-rule – they are playing the Best Response every time. Hence, the BG process was not in fact (nor was it so intentioned) a learning mechanism as defined here.

highest rank) shall then be called ‘teachers’, and the remainder of the agents (assuming a non-empty remainder set) the ‘students’. The learning phase then ensues, with a single student taking on a *public* part of a teacher’s strategy (chosen equiprobably if more than one). That is, since not all graph types will necessarily be observed during a stage, the student only has access to those responses actually employed by the teacher in the previous stage.

It is to be noted that so far, although agents might enter the model with a diversity of decision-making rules, if the model only ever allows for imitation within those rules, then long-run behaviour will be drawn only form the support of the initial rule distribution. Clearly, subjects are prone to discover new ways of solving the network formation problem that are drawn from outside of their counterparts’ rules. Therefore (and as is natural in modelling such *artificial adaptive agents* for precisely this reason), during the learning phase, let $m(\mathcal{S}, e)$ where $m : \mathbf{S} \times [0,1] \to \mathbf{S}$ be a mistake-making or ‘innovation’ filter which is applied such that each agents’ graph-type response decisions are reversed with vanishing probability ($e \simeq 0.01$) (e.g. from ‘sponsor this link’, to ‘do not sponsor this link’).[12]

Taken together, the learning mechanism, whilst related to the Genetic Algorithm (GA) approaches of various authors[13], diverges significantly from the canonical rendition in two main respects. First, agents to do not *combine* their strategic information to give rise to new agent strategies (the ‘cross-over’ operator in standard GAs), but rather, the transmission of strategic information is *one-way*, with the student taking on some of the strategic information of the teacher, without the reverse transmission occurring. Second, rather than agents having access to the *fundamental* strategic information of another agent (that is, the entire strategy encoding, whether revealed or not), the student is only allowed access to the portion of the teacher’s strategy information that has been revealed in the round-based play.[14]

[12] The qualitative naming given to such a filter is somewhat arbitrary. Importantly, what m brings is the capacity for a broader solution-space than was present initially to be searched by the agents, and thus, they are able to realise solutions not previously known to them, or even known *in-part* to a member of their playing group.

[13] See for an introduction Holland (1992) and its application in economics with artificial adaptive agents, Holland and Miller (1991).

[14] Other literature would recognise this as imitation based on phenotypic, rather than genetic information.

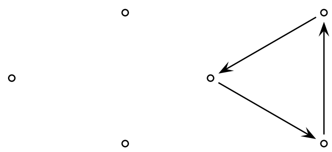

(a) **(b)**

Figure 6.10 Learning example graphs: (a) $\mathcal{T}_1$; and (b) T_5

Example of the Learning Process

Suppose two agents a and b with strategies $\mathcal{S}_a$ and $\mathcal{S}_b$ respectively have played the 4-player network formation game for R periods, and a ranking measure

$$\overline{\pi}_i = \frac{1}{R}\sum_{r=1}^{R} \pi_i^r \;, \quad i = \{a, b\} \tag{6.5}$$

has been applied with outcome $\overline{\pi}_a > \overline{\pi}_b$ .Then further suppose that a has responded to just two minimal absentee graphs in the previous stage (her *public* or *observable* plays), $\mathcal{T}_1$ and $\mathcal{T}_5$ (as shown in Fig 6.10) and in each case, her strategy response has been $s_a(\mathcal{T}_1) = (1, 1, 1)$ and $s_a(\mathcal{T}_5) = (0, 0, 1)$ respectively. Learning then proceeds as follows: first, as a result of the *imitative* part of the learning process,

$$\|s_b(\mathcal{T}_1)\ s_b(\mathcal{T}_5)\| = \|s_a(\mathcal{T}_1)\ s_a(\mathcal{T}_5)\|$$

where $\| \ldots \|$ indicates that a contiguous *sub-vector* of the horizontally concatenated public/observed plays is learnt; and second, the *innovation* filter is applied to the learning agent's decision-making rule, $m(\mathcal{S}_b, 0.01)$.

For example, this might have the outcome for b that

$$\mathcal{S}_b = \big\{(1, 1, \overset{\downarrow}{0}), \ldots, s_b(\mathcal{T}_4), (0, 0, 1), s_b(\mathcal{T}_6), \ldots, s_b(\mathcal{T}_k)\big\} \;;$$

in this example, b has learnt from a by imitation, but has made a single mistake (or change) in this imitation which has caused the generation of an altogether new decision-making rule over the two structures (summarised in Fig. 6.11).

6.3.6 Implementation

In order that the results from the present model could be compared effectively with those of FK, the subject-oriented experimental design was matched as closely as possible (see Algorithm 1). The total population of artificial 'subjects' $\mathbf{N}$ is partitioned into a number

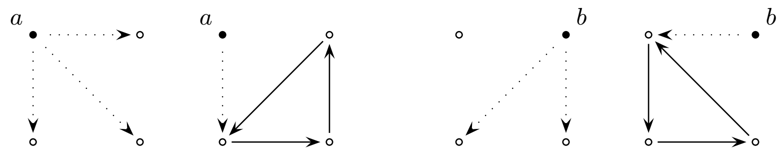

(a) **(b)**

Figure 6.11 Outcome of learning process ([—] absentee graph edges, [$\cdots$] agent responses): agent b (b,left) learns a's (a,left) response to the graph $\mathcal{T}_1$ imperfectly; but agent b (b,right) successfully imitates a's (a,right) response to graph type $\mathcal{T}_5$.

of *mixing groups* $\mathbf{M}$ (line `1.1`), so-called since over the subsequent stages, the *game-playing groups* $\mathcal{N} = N_1, \ldots, N_g$ that an agent will be allocated to (to play the n-player network formation game) will be drawn only from members of her mixing group. Hence, the mixing group will share 'playing information' over time with other members within the group, but not, therefore, with members outside of this group. Figure 6.12 shows this two-stage partitioning method schematically.

Algorithm 1 The main program

1: $\mathbf{M} \leftarrow$ PARTITION POPULATION$(\mathbf{N}, n)$;
2: $\mathcal{S}[0] \leftarrow$ INITIALIZE STRATEGIES;
3: **for** $t \leftarrow 1 \ldots T$ **do** ▷ For all stages
4:     **for all** $\mathbf{M_i} \in \mathbf{M}$ **do**
5:         $\mathcal{N} = N_1 \ldots N_g \leftarrow$ FORM GAME-PLAY GROUPS$(\mathbf{M})$;
6:         **for all** $N_i \in \mathcal{N}$ **do**
7:             $S \leftarrow \mathcal{S}_{N_i}$; ▷ Get strategies
8:             $G \leftarrow S(\mathcal{T}_0)$;
9:             $P[0] \leftarrow G$; ▷ Record plays as public
10:             **for** $r = 1 \ldots R$ **do** ▷ For all rounds
11:                 $G \leftarrow g^* \leftarrow S[T(G)]$; ▷ Update graph
12:                 $O[r] \leftarrow$ OBJECTIVE(G);
13:                 $P[r] \leftarrow G$;
14:             **end for**
15:             $S(N) \leftarrow$ APPLY LEARNING(P, O);
16:         **end for**
17:     **end for**
18: **end for**

At the beginning of a modelling run, each sponsorship decision in the decision-making rule $\mathcal{S}$ is set to an equiprobably (sponsor)/(not-sponsor) element for all players $\{0, 1\}$ (line `1.2`), and at the beginning of each network-formation game (that is, at round $r = 0$), the graph is formed from each agent's first sponsorship action vector $s(\mathcal{T}_0)$. Thereafter, the network is defined by the round-wise application of each agent's decision making rule as described above (line `1.11`).

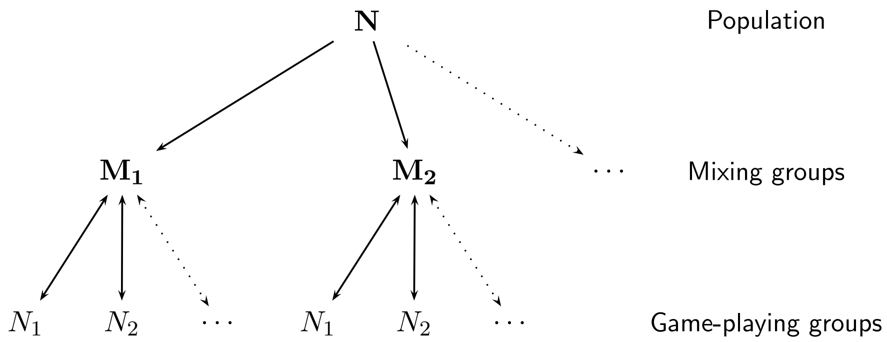

Figure 6.12 Schematic of the two-stage partitioning of the artificial population, following the experimental design of FK. At the start of an 'experiment' the artificial population is partitioned into mixing groups $\mathbf{M}$. As each stage begins, a number of n person playing-groups are formed to undertake the multi-round network formation game. Hence, between stages, agents within the one mixing group have an opportunity to undergo constrained mixing.

After allocation of agents first into their mixing groups and then at the beginning of a stage into their actual game-playing groups, each network formation game is played in turn for a total of R rounds. Here, as described in the model above, agents respond to the current minimal absentee graph ($\mathcal{T}$) as it appears to them from the present full graph (G), subsequently realising some objective measure of each agent's decisions (line `1.12`). Additionally, the sponsorship decisions played by each agent are recorded as 'public' (since they are now revealed) and can then be imitated in the later learning phase (line `1.13`).

After all rounds for a game-playing group have been played, summary objective measures are compiled, and agents are designated by ranking to be either students or teachers, with the students copying something of the public plays of the teachers as revealed through the fore-going rounds, thus forming a new strategy vector for the foregoing stages (line `1.15`). This process continues for all experimental groups, and all stages, ending the modelling procedure.

Agent Learning

The present modelling prescription allows for a relatively simple treatment of learning. As explained above, each agent retains a vector $\mathcal{S}_i$ of sponsorship-decisions towards the $n-1$ other agents for the total number of minimal absentee graph types they might face. Since it is assumed that agents are able to both recognise graph types (from the current instance) and then apply (via manipulation) their chosen response to the instance, inter-agent imitation is quite easily handled by simply reading a teacher's response vector, say $\mathcal{S}_t$ and writing it to the appropriate section of the student's response vector $\mathcal{S}_s$ (see Fig. 6.13). A uniform random process determines how many teacher sections are imitated in each learning process.

$$
\begin{aligned}
\mathcal{S}_t &= \Big(s(\mathcal{T}_1), \ldots, \overbrace{000, 110, 001}^{\substack{\textit{section to be}\\ \textit{imitated}}}, 101, \ldots, s(\mathcal{T}_k)\Big) \\
\mathcal{S}_s &= \Big(s(\mathcal{T}_1), \ldots, 011, 010, 011, 001, \ldots, s(\mathcal{T}_k)\Big) \\
& \qquad\qquad\qquad\qquad \Downarrow \\
\mathcal{S}_s^* &= \Big(s(\mathcal{T}_1), \ldots, 000, 11\underline{1}, 001, 001, \ldots, s(\mathcal{T}_k)\Big)
\end{aligned}
$$

Figure 6.13 Example of a teacher-student learning process for one-way flows and $n = 4$ (thus each $s(\mathcal{T})$ comprises a sponsor-decision over $n - 1 = 3$ edges; $\{0, 1\}$ represent 'not-sponsor' and 'sponsor' respectively; $\underline{1}$ indicates a mistake.

Additional to this learning-by-imitation mechanism is a standard mistake-making operator, applied with vanishing probability (less than one mistake per 100 sponsorship decisions, or $m = 0.0075$) to each learnt decision, allowing agents to innovate, or mistake-make, their way to possible ways of playing not present in the original strategy space.

Further, as can be seen in Fig. 6.13, students learn *sections* at a time from the teacher's revealed strategies. That is, if a particular response strategy requires nine bits (as in the example) to be correctly described, then students will learn those nine-bits, rather than seven bits of this response and two of the next which happen to lie alongside on the string. This methodology is known as 'punctuated learning' and usefully implements the learning of only fully specified responses, rather than part-responses.[15] In all experiments presented below, after ranking, if at least one agent could be ranked below at least one other agent, the single lowest and single highest ranked agents were selected to be the student and teacher respectively. If two or more agents were found to share a ranking as on either level, agents were split equiprobably.

6.4 Results & Analysis

Initial experiments focused on replicating as many of the experimental facts as possible through a process of parameter search and calibration. This process was conducted on the one-way information flow case, since this context yielded the most information in human trials. Once the model achieved a degree of relevance to the one-way case data, two-way information flows were studied to investigate if, and why, a break-down in agent strategies was observed.

6.4.1 Agent Performance

As the foregoing discussion has indicated, the present model focuses attention on how agents improve their performance (find more Nash structures) over time. Supposing

[15] In practice, applying non-punctuated learning does not qualitatively change any results presented below.

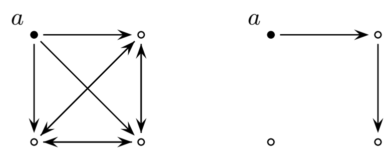

(a) **(b)**

Figure 6.14 Graphs L (a) and R (b) giving rise to equal net payoffs for agent a.

that a process of horizontal transmission between agents coupled with a low-frequency mistake/innovation method can represent the learning process, the key question becomes how to *rank* agent performance for subsequent learning? Who should be 'students' and whom should they learn off (the teachers)? Whilst the possibilities for such a ranking procedure are limitless, three are considered below as a reasonable first pass.

The first and most obvious choice is the *net payoff* to each agent, as conferred by the communication network itself and is simply calculated for a round by (6.1) and for a whole stage by (6.5). This method is attractive for simplicity and basic acceptability in motivating agent behaviour.

However, it is possible that subjects do not rely on individual payoffs alone in determining 'good' structures. An example will illustrate some consequences of a simple payoff regime. Consider the graphs presented in Fig. 6.14, for agent a, it is trivial to calculate π_a under (6.1) (denoting graph (a) L, and (b) R),

$$\begin{aligned}\pi_a(L) &= 4V - 3C;\\ \pi_a(R) &= 3V - C;\end{aligned}$$

which result in *equal* payoffs to agent a for $V = 2C$ (a specific case studied in FK). It is reasonable to assume that the situation on the right (b) displays a better use of existing edge sponsorships for a out of the two, and such an agent would be more likely to receive imitative behaviour. Of course, 'obvious' examples like these do not translate easily into ranking procedures.

Some progress can be made by noting that the non-cooperative network formation game actually gives rise to an *externality* generating environment. Principally due to the assumption of information flows between agents, each agent's edge sponsorships do not just have an effect on their own information observations, but actually provide an information gathering *structure* for other agents.

Thus, a second ranking measure is constructed along these lines: the measure should capture the extent to which an agent has been able to exploit any externalities in the graph structure. Intuitively, a simple ratio of observations (μ) and degree (δ) for each agent

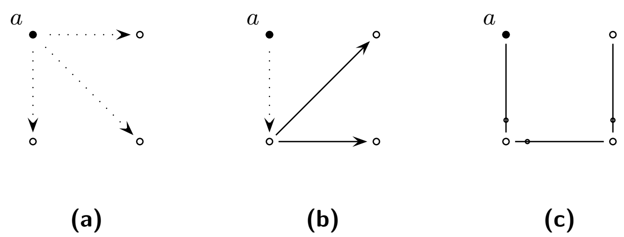

Figure 6.15 Example benefit/cost ratio measure example networks for agent a ($n = 4$): (a) minimum; (b) maximum (one-way flows); and (c) maximum (two-way flows). Agent sponsorship strategy represented by dotted lines. NB: in (c), (following BG), sponsorship is indicated by small dot on edge sponsored by closest agent.

ought provide such a measure. However, the convention of self-observation ($\mu = 1$) with zero link-sponsorship ($\delta = 0$) must be taken into account, hence the following measure is formed,

$$f_i(\mu_i, \delta_i) = \frac{\mu_i V + C}{C(\delta_i + 1)} \ . \tag{6.6}$$

The denominator simply normalizes the $\delta = 0$ case whilst the numerator comprises the benefit to observation plus a further C term as explained presently. It is to be noted, that as desired, $f_\mu > 0$ and $f_\delta < 0$, but moreover, by rearranging, one obtains,

$$f_i(\mu_i, \delta_i) = \left(\frac{1}{\delta + 1}\right) \left[\left(\frac{V}{C}\right) \mu + 1\right] \ ,$$

which indicates that the ratio of observation benefits and edge sponsorship costs V/C weights the relative importance of the number of observations, versus the number of sponsorships. Hence, when V/C is large (i.e. low cost-regime), information observed is favoured in the measure, and numbers of edges sponsored is not, whereas in the alternate case when V/C is small (i.e. high cost-regime), edge-sponsorships become more ‘costly’, having a greater relative effect on benefit/cost ratio. Hence, although the measure captures the main interest by rewarding externality exploiting strategies, it also retains a connection to the actual cost of edge-sponsorship. Without this feature, the measure would be agnostic as to the particular cost-benefit regime they are a part of, which is a key part of both the BG and FK analysis.

By way of example, and to normalize f on $[0, 1]$ the maximum and minimum strategies can be identified by considering the best- and worst- case externality scenario in each of the one- and two- way conditions. In the one-way case, the worst-case scenario is any strategy in which $\mu_i = n$ and $\delta_i = (n-1)$ (sponsor-all); such a case can be seen in Fig. 6.15(a). The best-case result for an agent in the one-way case with low edge sponsorship costs is when they get the highest externality for their strategy, which can be easily verified to be where

$\mu_i = n$ and $\mu_i = 1$; this situation is shown in Fig. 6.15(b). However, it can be shown[16] that in the one-way case, the optimal value of f depends on the relation $C \gtrless V(n-2)$ and that in fact, if $C > V(n-2)$, then the maximal value for an agent is obtained by sponsoring no links at all. That is, the return to sponsoring one link, *even if it returns full information*, is lower than the cost of sponsoring that link. Clearly, this is a natural phenomena and will have important ramifications in modelling.

Table 6.3 Efficiency measure $f(\mu_i, \delta_i)$ for all feasible combinations of μ_i and δ_i with $n = 4$. One- and two- way information flows indicated by $\rightarrow$ and $\leftrightarrow$ respectively.

		$\mathbf{C} = 5$		$\mathbf{C} = 25$	
μ_i	δ_i	$\rightarrow$	$\leftrightarrow$	$\rightarrow$	$\leftrightarrow$
1	0	0.33	0.11	**1.00**	0.38
2	0	-	0.41	-	0.59
3	0	-	0.74	-	0.80
4	0	-	**1.00**	-	**1.00**
2	1	0.11	0.04	0.33	0.13
3	1	0.56	0.19	0.60	0.23
4	1	**1.00**	0.33	0.87	0.33
3	2	0.04	0.01	0.11	0.04
4	2	0.33	0.11	0.29	0.11
4	3	*0.00*	*0.00*	*0.00*	*0.00*

In the two-way case, although the minimum is unchanged, the best-case strategy becomes 'sponsor-none' (as shown in Fig. 6.15(c)), since it is possible to obtain full-information and yet provide no support for the information network. This is akin to a truly non-cooperative ('defect') strategy since all benefits are retained, with costs born by others. The full results for normalized values for f calculated by (6.6) under $n = 4$ and information flow of both types is given in Table 6.3 for comparison. It can be seen that subtle changes in preference ordering do occur within columns due to the cost of link sponsorship; ultimately these will be reflected in learning outcomes when $f(\mu, \delta)$ is the objective ranking measure.

Finally, the third ranking procedure considered was that of a random base-case. This was implemented so as to cause a student to update their strategy by imitating a teacher as often as would have happened under ranking measure f. However, rather than the

[16] We ask, is it possible that sponsoring no links, i.e. gaining $f(1,0)$, could ever be greater than the externality exploiting strategy that yields $f(n,1)$,

$$\begin{aligned} f(1,0) &\geq f(n,1) \ , \\ \frac{V+C}{C} &\geq \frac{nV+C}{2C} \ , \\ C^* &\geq V(n-2) \ . \end{aligned}$$

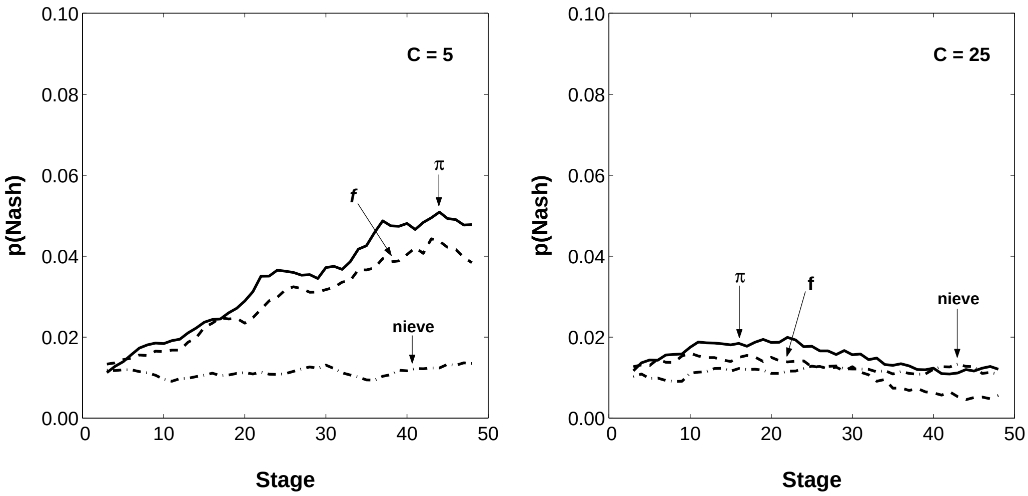

Figure 6.16 Nash (non-empty) structures under one-way information flows: (left) $C = 5$; and (right) $C = 25$, under different objective measures: payoffs (π), benefit/cost ratio (f) and naive (random) learning.

student being identified by ranking due to π or f, the ranking was uniform random, and hence, provides a true test of the trial measures by testing whether pure imitation *alone* can account for observed subject behaviour.

6.4.2 Objective Measures & Learning

To begin, the one-way information flow case was analysed over a range of parameterisations with a relatively short time-scale.[17] Specifically, both high- and low- edge sponsorship costs were considered together with the three ranking procedures measures as described above. Results for these experiments are given in Fig. 6.16.

A primary insight is immediately apparent from these plots, namely that pure imitation alone is not a proxy for human learning in the network formation game. The naive-learning plots both show no significant change through the stages, in contrast to the treatments where ranking was informed by one of the two objective measures, and was therefore non-random.

Second, both non-naive ranking procedures gave reasonable improvement dynamics across stages for low-cost conditions, but the opposite is true of the high-cost regime where these measures are indistinguishable in performance from the naive learning process. To gain a better perspective on the difference between these regimes, average agent degree plots were prepared as reported in Fig. 6.17. It can be seen that in the high cost regime

[17]Results given for 20 independent trials, each comprising of 10 mixing groups of 40 agents each to form the $n = 4$ game-playing groups for each stage play, over 50 stages, playing a 20 round network formation game. Experiment parameters as follows: $V = 10$, $m = 0.0075$ others given in caption to Fig. 6.16. In each line plot, data were smoothed by averaging two points to the left and right to clarify trends, that is a mean of 5 measurements in total.

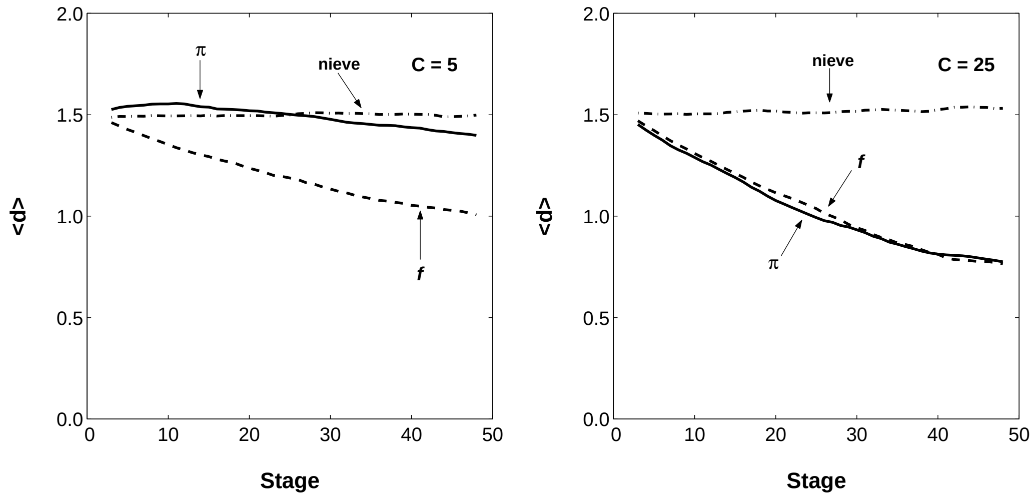

Figure 6.17 Average agent degree under one-way information flows: (left) $C = 5$; and (right) $C = 25$, objective measures as for Fig. 6.16.

(right) the artificial agents quickly chose to sponsor few edges under both non-naive ranking conditions as opposed to the low cost regime where high sponsorship levels, near the naive case, were observed for the π measure, whilst the f measure saw sponsorships decline to around 1 per agent.

The above experiments indicate that the artificial agents were probably under-sponsoring compared to their human counterparts. This is especially so in the high-cost regime, where edge sponsorships continued to trend pass one per agent, and this, regardless of the non-naive objective measure used for ranking. The FK trials show that subjects roughly had a preference for sponsoring about one edge at a time. It can be concluded that the pure non-naive objective measures alone as measures for strategic development are not good proxies for human learning in this context. This is not surprising in the low-cost case since as mentioned above, agent payoffs can carry confusing information about 'good' strategy, and in the high-cost case, payoffs will cause individuals to undersponsor on the whole due to free-rider temptations.

6.4.3 Altruism, Reciprocity, & Learning

With the perspective obtained above, the importance of maintaining average agent sponsorship levels is underlined. Indeed, the game can be considered as a kind of n-player public good game where the benefits of the network are not conferred in a dyadic or triadic relationship, but are available to all players. The difference between the model results and the subject trials indicate that real subjects appeared willing to incur personal costs in sponsoring links to support the public benefit of the network (and subsequent improved personal return). Such behaviour is consistent with empirical evidence which continues to show that humans prefer a kind of altruism often associated with fairness as opposed to strict self-interested play (Camerer, 1995). There are various hypotheses concerning how and why humans exhibit altruistic behaviour, with attendant models such as *reciprocal altruism* relying on repeated interactions (Axelrod and Hamilton, 1981), *indirect reciprocity* through image recognition (Nowak and Sigmund (1998), compare Riolo et al. (2001)), or so-called *strong-reciprocity* involving pro-social rewards *and* anti-social punishments (Bowles and Gintis, 2004).

In the context of the network formation game, the specific form of reciprocity that is at play is unclear. An agent who decides to bear a personal cost to support a link, regardless of the strategic decisions of others, does engage in a kind of reciprocal behaviour, since the benefit they receive is additive to the extent that others imitate this network formation behaviour; the good deed is returned. However, the feedback works in the reverse direction, the agent who does not sponsor a link, if imitated, will find that any free-riding benefit is quickly scotched.

Consequently, an altruism measure was added to the transmission process between stages to augment the objective measures previously described. Specifically it was supposed that for each revealed sponsorship decision within a round, agents are assigned (it is supposed by the observation of the other agents) a basic altruism measure r_i as follows,

$$r_i = \begin{cases} 0 & \text{for} \quad \delta_{in} = \{0,1\}, \quad \delta_{out} = 0\,, \\ 1 & \text{for} \quad \delta_{in} = 1, \quad \delta_{out} \geq 1\,, \\ 2 & \text{for} \quad \delta_{in} = 0, \quad \delta_{out} \geq 1\,, \end{cases} \tag{6.7}$$

where δ_{in} and δ_{out} are an agent's in- and out- degree respectively (the count of edges incident to the agent, and the count of edges the agent sponsors respectively). For each agent over R rounds, the mean value of r is calculated and contributes to the overall ranking measure used to determine 'teachers' and 'students' by simple convex combination,

$$\Omega_i = \alpha\langle r_i\rangle + (1-\alpha)\{\langle \pi_i\rangle, \langle f_i\rangle\} \tag{6.8}$$

where $\alpha \in [0,1]$, and $\langle\cdot\rangle$ indicates arithmetic mean.

To calibrate this new combined ranking measure, the one-way network formation game was again experimentally studied for several values of $\alpha \in \{0.0, 0.2, 0.4, 0.5, 0.6, 0.7, 0.8, 1.0\}$ with the f measure.[18] Results given in Fig. 6.18 indicate that for values of α less than 0.4 little change occurs to the overall down-ward trend in edge sponsorship, however, a non-linear movement for $\alpha \geq 0.4$ yields an increase of the sponsorship levels as desired.

[18] Experiment parameters as per previous experiment.

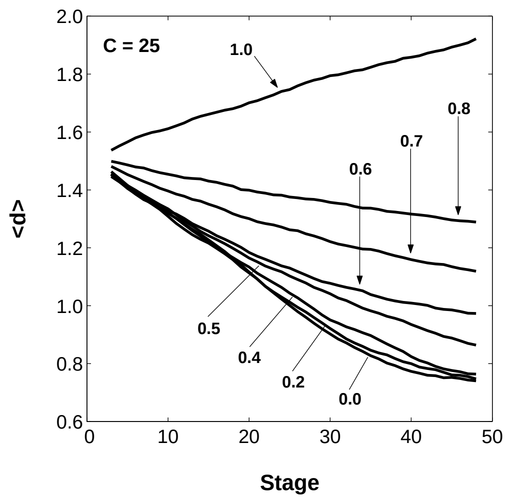

Figure 6.18 Combined (altruism,benefit/cost ratio) objective measure calibration results at different α values.

A value of $\alpha = 0.6$ was therefore chosen since it appeared to provide agreement to the human trial sponsorship levels of about one link per agent. Taking this new *combination* objective measure (setting $\alpha = 0.6$) to a combined low- and high- cost regime study gives rise to more realistic improvement and behavioural outcomes (see Fig. 6.19). In three of the four cases (both measures with $C = 5$, and for f with $C = 25$), agents reduce average edge-sponsorship numbers towards 1.0 or above, with resultant strategic outcomes also displayed in relatively good $p(Nash)$ growth over time. However, the payoff measure, even with the added altruism component is unable to discover Nash structures over time. This result appears to further underscores the low level of information that agents receive via the individual payoff measure. For this reason, the new combined benefit/cost ratio and altruism measure was pursued as an approximate proxy for the ranking measure.

6.4.4 Learning Dynamics over the Long-run: Strategy & Prediction

One-way Flows

To study the dynamics of the learning process for comparison with the subject data, two long-run experiments were conducted. In the first, the conditions as determined above were implemented as is: one-way flow, low- and high- costs of edge sponsorship. In the second, the two-way information flow condition is introduced for the first time, to see if the calibrated conditions chosen above are applicable to replicate the poor subject results of the two-way flow game.

Figure 6.20 gives Nash structure frequencies (left) for the combined benefit/cost and altruism measure over 250 (learning) stages. In this case, the low cost regime does remarkably well, plateauing out at around an average of 16%, whilst the high-cost case, although improved by the altruism measure achieves roughly half this rate. Of further

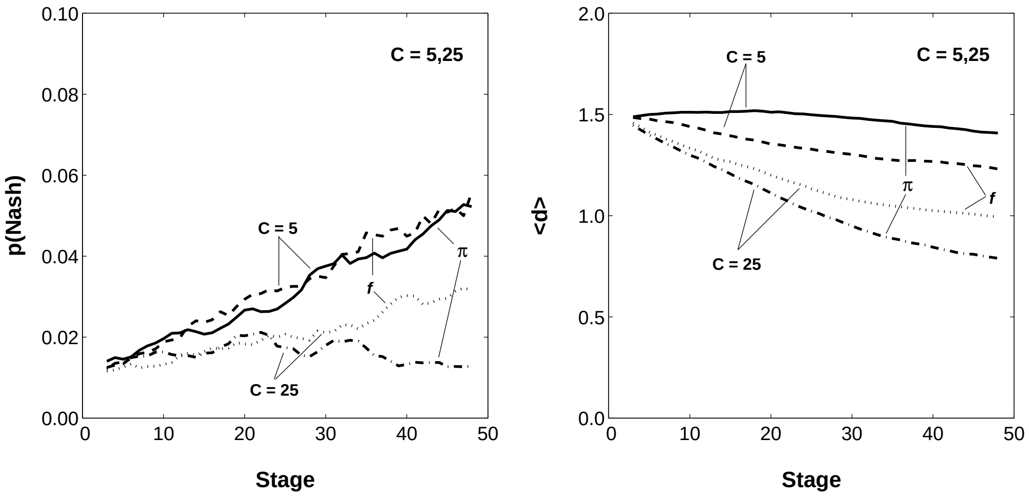

Figure 6.19 Frequency of (non-empty) Nash structures (left) and average agent degree (right) under one-way information flows and *combined* objective measure conditions. In each case, $\alpha = 0.6$, different link costs and objective measure treatments are labelled as in Fig. 6.16.

interest, is the proportion of Nash structures that are actually the *circle* type, the structure seemingly requiring the highest level of coordination. As can be seen in both cost regimes, Nash play is largely accounted for by Strict Nash play. Recall that to play Strict Nash, the artificial agents have individually contributed to forming one of six structures out of a possible $2^{n(n-1)}$, or around 4000 for $n = 4$, structures which is a remarkable result considering the limited cognitive and learning architecture they are endowed with. The naive results presented for the long-run case show conclusively that imitation of others *alone* without the benefit of the combined ranking measure has no effect on frequency of Nash or Strict Nash play. This result resonates with the cultural learning literature that speaks in terms of *biased cultural transmission* rather than simply 'cultural transmission'. The ways of behaving or solving problems are transmitted from a desirable sub-set of the population to the others, rather than just from any (randomly chosen) sub-set (Henrich, 2001).

For ease of comparison with FK, the long-run studies were split into meta-stages so that both play *within* and *between* learning junctions can be compared. To this end, in each bar-chart that follows, the 250 explicit model stages were grouped under 3 meta-stages (that is, $t \sim (1, \ldots, 83), \ldots, (167, \ldots, 250)$) and within each of these, average data for play within specific rounds were split into 5 groups (that is, $r = (1, \ldots, 5), \ldots, (16, 20)$) to represent comparison 'rounds'. The outcome for the one-way case is very interesting (see Fig. 6.21), with improvements clearly visible both within and between phases. This phenomena is exactly as observed in the FK subject trials; the authors hypothesised

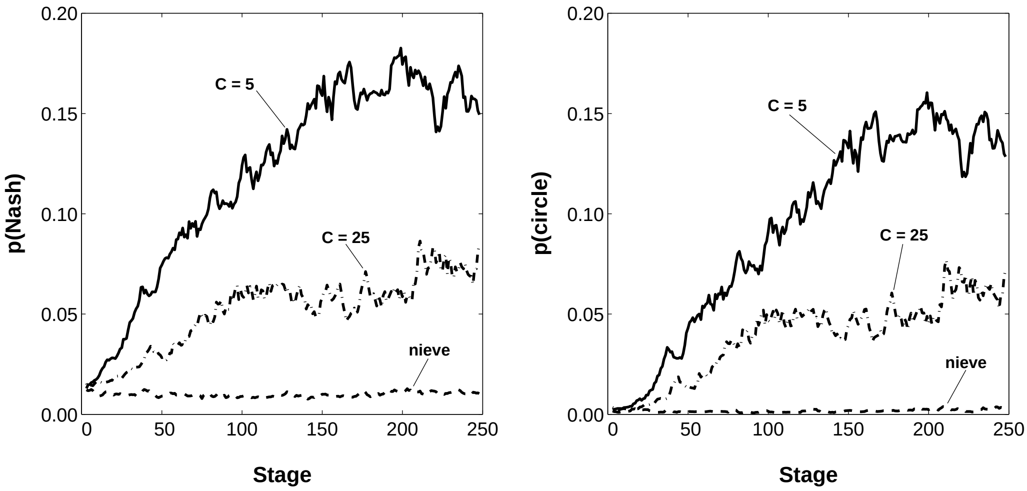

Figure 6.20 Results of long-run study under combined altruism – benefit/cost ratio measure at different costs. Naive learning included as a control. Nash structures (non-empty) are predominantly comprised of the Strict Nash (one-way) circle structure.

that in the early rounds of a stage (recall: groups are re-formed between stages), mis-coordination occurs due to the unfamiliarity with the playing partners, but this goes away over time. Whilst some improvement between stages was to be expected, due to the modelling set-up, that the artificial agents mimic human Nash structure discovery behaviour *within* stages is extremely pleasing.

The calibration work above appears to suggest a seemingly strong correlation between average degree and propensity of groups to find Nash structures, it could be hypothesised that the strong learning results between stages and improvement dynamic within stages are simply the result of random occurrence aided by players lessening the number of edges that they sponsor. To test this hypothesis, the same data treatment is performed on mean agent degree for the one-way case (see right of Fig. 6.21). As can be seen, although there is a slight reduction of average edge sponsorships over the 3 meta-stages which may indeed account for some of the overall learning dynamic between stages, within each stage, there appears to be a vanishingly small difference in agent edge sponsorship numbers.

Whereas the probability of finding the Strict Nash structure (in this case, the circle-graph) grows markedly. Hence, one can conclude that the apparent improvement dynamic within stages of the artificial agents is not by chance alone under lower edge sponsorships, and is therefore due to strategic learning as modelled.

To further investigate the strategic behaviour of agents, a simple measure of prediction was developed, which asks whether an agent's sponsorship decisions would have achieved an higher or lower objective measure against the actions of opponents in the last period,

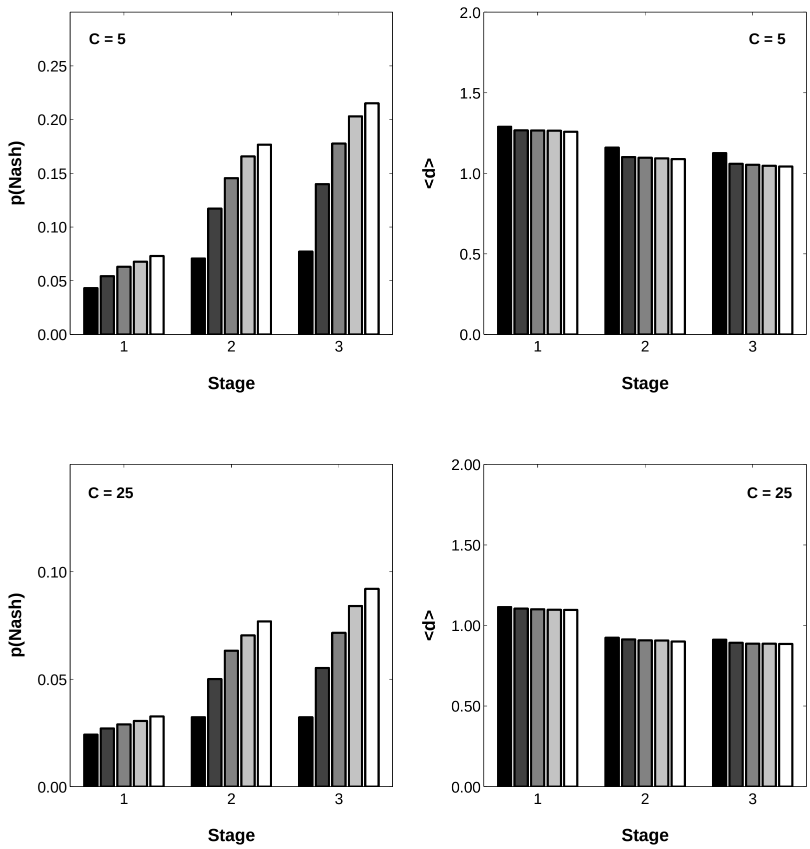

Figure 6.21 Within- and between- stage learning as evidenced by improving (non-empty) Nash structure probability. Average agent degree also shown (right), showing little within-stage variation, despite large equivalent performance variance (left). Data shown is average over all mixing groups and repeats.

or against their actions as revealed in the present period,

$$M_i^r = \operatorname{sign}\Big[f(g_i^r \cap g_{-i}^r) - f(g_i^r \cap g_{-i}^{r-1})\Big] \; . \tag{6.9}$$

Thus the measure will take the value of 0 on average if an agent does no better nor worse by playing their chosen sponsorship vector against the actions of the other players last period, or as revealed this period, whilst a positive value will accord with the sponsorship vector giving better outcomes than not against the future graph, and vice versa. By again taking the average of this measure across the collected rounds and stages, an indication of when agents have been able to play predicatively can be obtained. Figure 6.22 gives these results for the $C = 5$ long-run case as discussed above, with the naive regime also given for comparison.

The plot shows that agents do indeed learn to play predicatively over the stages, with marked increase in this kind of play within round groupings as well when compared to the naive plot. Interestingly, both graphs (naive and non-naive) indicate that the first

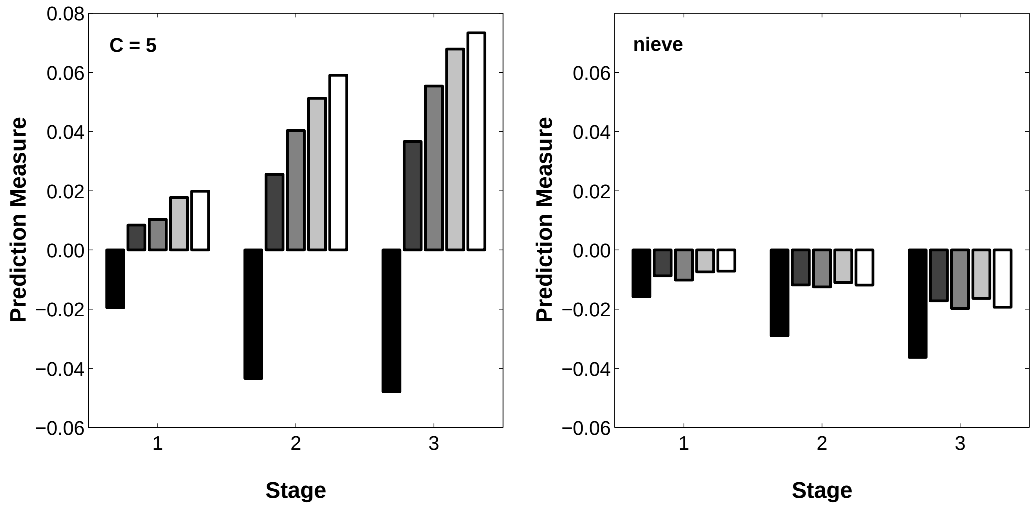

Figure 6.22 Prediction measure results for within- and between- stages for combined and naive learning rules for comparison. A strong correlation with performance is clear (compare Fig.fg:bar:oneway)

few rounds in the non-naive case are played in a *responsive* way, rather than a predictive way. This is not surprising, since the first rounds of a stage are where agents have just been mixed, and will initially mis-coordinate. Recall, that these initial rounds are also when the artificial agents are likely to sponsor more links on average. Putting these two observations together, the artificial agents appear to have learnt to compensate for the lack of early coordination by sponsoring more links.

Two-way Flows

Turning to the two-way information flow case, it is of interest to see whether the present specification of learning calibrated on the one-way outcomes gives rise to problematic play in the two-way case, as was the result in human trials. Figure 6.23 certainly seems to suggest that this is the case. Both in the low- and high- cost regimes, performance overall is poor, with a *regression* in performance observed due to stage-based 'learning', rather than improvement. The results are particularly stark for the Strict Nash structure (cs-star) which finds less than one per cent frequency in the low-cost case. This is in good accordance with the human trials, where the one-way Nash structures were largely due to the Strict Nash cycle structure, whilst in the two-way case, graphs were found to be Nash a smaller number of times, but with only two cs-star structures observed across the entire human trial. Both of these facts are generated in the present model.

Furthermore, Figure 6.24 indicates that sponsorship levels were roughly similar to those of the one-way case, as was found in the human trials, and thus cannot be attributed

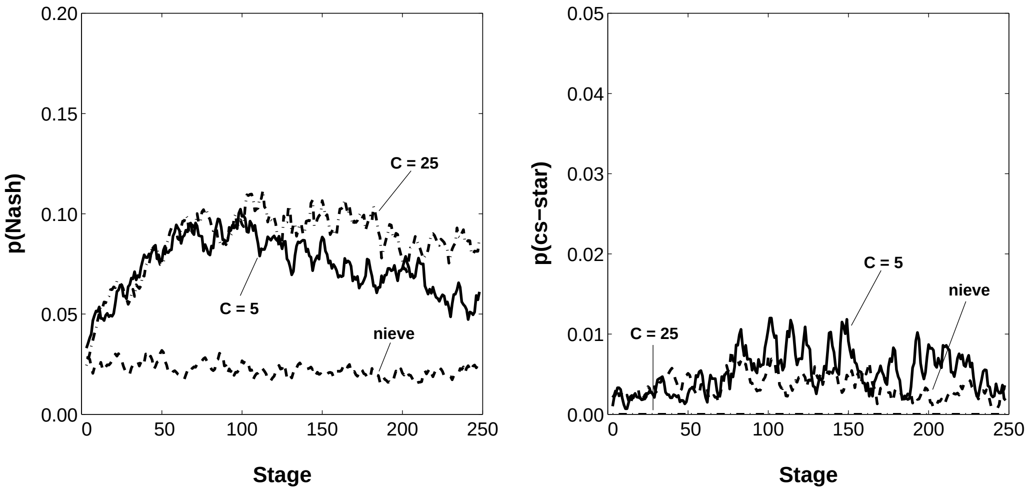

Figure 6.23 Frequency of (non-empty) Nash (left) and CS-Star (Strict Nash) (right) structures under two-way information flows. Data given for combined altruism – benefit/cost ratio under each cost regime and naive learning for comparison. Note change of scale in CS-Star plot.

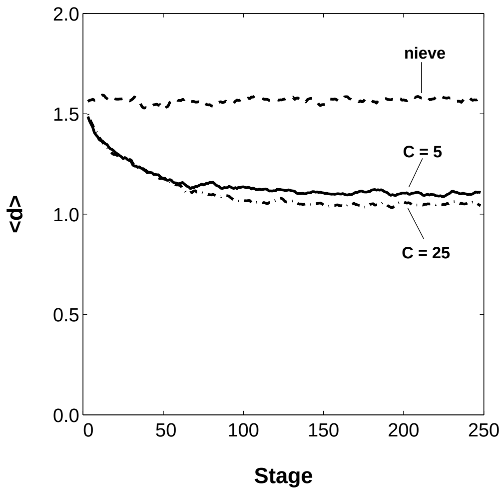

Figure 6.24 Average agent degree under two-way learning, data shown as for Fig. 6.23.

with the main source of error. Indeed, the summary bar-charts shown in Fig. 6.25 indicate that a break-down in any improvement between stages, or within stages is entirely strategic, with the average degree stage and round plot resembling closely that of the more successful one-way case.

6.4.5 Inertia: Inequality & Emergence

Finally, the FK study considered reasons why agents were able to find Strict Nash outcomes so readily in the one-way flow case and not in the two-way case. Their argument centred on two main effects, the first being the natural observation that the Strict Nash structure in the one-way case (the circle) displays a high degree of symmetry, both in

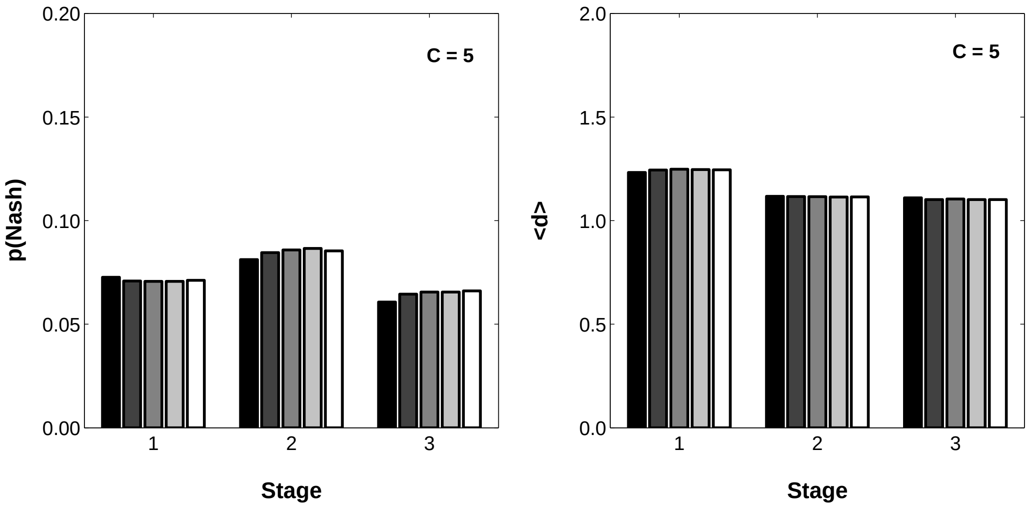

Figure 6.25 Frequency of Nash structures (non-empty) under two-way learning showing limited within- and between- stage learning (compare Fig. 6.21).

payoffs, and in terms of strategies. In contrast, the two-way Strict Nash structure (the cs-star) displays neither of these symmetries. Within the present framework, it is difficult to test this hypothesis.

However, the second argument invoked by FK was a social preference one that suggested that agents certainly dislike inequality of outcomes and probably do not value inefficiency either (see refs in FK for details). To test the social preference hypothesis, FK ran a probit regression of the occurrences when a subject showed *strategic inertia* (played the same sponsorship vector in the present period as in the previous one) and tested for payoff inequality in the previous round. That is, they hypothesised that if an agent experienced higher inequality of payoffs in a round, they were more likely to change their strategy in the next round.[19] They found that inequality was significant and negative, as expected, in explaining agent inertia, even after controlling for whether an agent played a best response (in payoff terms) in the previous round.

To test whether a similar fact was found in the present model, a plot of agent inertia over learning stages was prepared for the one-way flow case (Fig. 6.26). Clearly, both the low- and high- cost cases showed increasing agent inertia over time relative to the naive baseline case. This is very interesting, as it indicates that as well as being strategically predictive, agents are able to choose to play inertia after learning, presumably giving rise to the conditions for Nash play. Secondly, to directly investigate the connection between inequality and inertia, a plot in inertia-inequality space was prepared to show

[19]Note that this is not the same as asking whether or not an agent chose to *update* their strategy or not since they could possibly playing the same strategy as the previous round, and hence would not display strategic inertia.

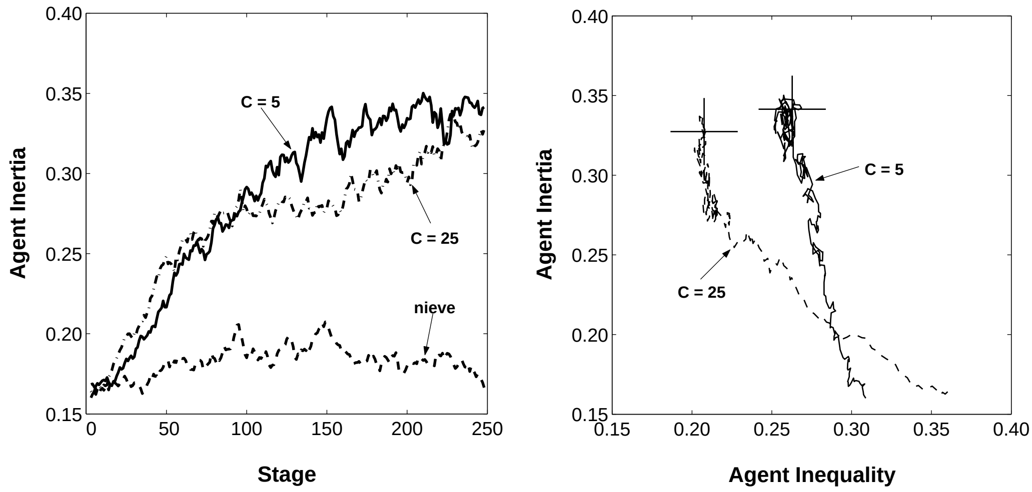

Figure 6.26 Strategic inertia arises endogenously in the model. Agent inertia as a function of learning stage under one-way information flows (left); and the same shown in Inertia-Inequality space showing the relationship between the two, $(+)$ indicate final experiment state (right).

the trajectory of the low- and high- cost cases. Agent inequality (following Fehr and Schmidt (1999), in FK) was calculated as follows,

$$e_i = \sum_j |\pi_i - \pi_j| \ , \tag{6.10}$$

for all $j \neq i$; the sum of absolute differences between an agent's individual payoff and the payoffs of all other agents. As can be seen, a correlation is found in equivalence with the human trial work – lower agent inequality is correlated with higher agent inertia. However, what makes this result interesting is that rather than constructing a model that represented all the observed human preference as an input (as might be a possible course of action), the model constructed above is extremely parsimonious with regards depiction of cognitive structure or process. To find the social preference thesis arise as an *emergent* property of the system further indicates the usefulness of the present approach.

6.5 Conclusions

This work has attempted to construct, calibrate and test a model of human decision making in the richly complicated communication network formation game as proposed by BG and tested in the field by FK. By applying a novel graph-equivalence criterion and mapping technique, it has been shown that the apparently vast and seemingly unfeasible strategy space presented to the artificial agents can be modelled in an intuitive and feasible manner. Whilst this approach is likely useful for $n \lesssim 10$, it is in this region that the problem retains strategic interest given the limits of human functionality.

Further, through a calibration process via human subject data, the limits of a purely payoff driven approach have been shown, with a suggested alternative criterion combining a *ratio* (rather than difference) of benefits and costs together with a simple altruism measure proving a more reliable proxy for human action in this context.

From this basis, key questions in modelling these difficult systems have been answered. Are artificial agents able to play *strategically* from a simple case-based action vector? The steady sponsorship levels between rounds contrasted with improvements in Nash play affirm that the artificial agents are indeed *strategic*. Can artificial agents who have only one round of 'memory' play predicatively? The answer to this appears 'yes', especially so in later rounds of each stage, and later stages of each experiment. This suggests that the deterministic 'observe-recognise-play' method that each player has been endowed with for a round settles to an approximate pattern of play in the graph state space, giving rise to the observed predictive play.

Moreover, whilst (individual) perceptions of altruism were explicitly modelled in the learning cognition for agents, a response to payoff inequality or inefficiency (a social measure) was not. Yet, these features arose as an *emergent* property of the system, suggesting that such preference need not necessarily be incorporated in the model explicitly, but can be derived from it.

The above results suggest that lack of symmetry leading to a coordination breakdown in the case of two-way information flows is the most likely cause of human difficulty in finding the Strict Nash structure. Whilst a large number of facts from the human trial were observed in the present model, the modelling prescription with the combined measure incorporating altruism, paradoxically makes the centre-sponsored star a very unlikely structure to obtain.

It remains to discuss why although many of the features of the subject's Nash play were replicated here, the *level* of play (in the one-way case) was not (falling short by around 40%). Several points can be made. First, a focal structure was not programmed into the agent cognition – there was no explicit ranking of structures that directed the play towards one of the several structures that would appear similar in objective measure to the circle. It seems reasonable to think that humans might actually find the circle an appealing structure to play given the symmetry as discussed above. In effect, the artificial agents were programmed to be fundamentally indifferent to all structures.

Second, within the framework, agents were only endowed with a single round of 'memory'. They observed the structure of the present period, recognised it, and played their response. An obvious extension would be to allow agents to recognise a two- (or more) round sequence of structures and play their response. Clearly, as has been maintained throughout, the network formation problem is fundamentally complex enough for humans to negotiate and thus a focus of the present approach was to keep the cognitive architecture within the bounds of what might be considered humanly capable (e.g. the recognition of *at most* 16 and 11 unique structures in the one- and two- way flow cases respectively). Hence, one suspects the layers of memory cannot honestly be increased much beyond two or three but by doing so agent performance would likely improve.

Third, although raw payoffs alone, or in concert with a measure of altruism appear not to provide enough information to agents to give rise to realistic behaviours, as mentioned above, any number of alternatives could be chosen. The ratio measure above provided a reasonable choice, but others could be investigated (e.g. if memory is expanded, objective measures over *sequences* of plays akin to a *gambit* could be considered, allowing for the adoption of short-run loss/long-term gain strategies).

Nevertheless, this work has indicated a fruitful line of further work in modeling network formation in communication networks in particular, and other economic networks in general. Whilst the BG analysis provided an excellent framework for thinking about non-cooperative network formation with results in the one-way case being largely confirmed in the field, the present computational model has enriched this understanding considerably. Elements such as biased cultural transmission as learning, emergent social preference phenomena, altruism and the exploitation of externalities through modified objective function design have been raised and investigated.

part III

Cooperation Networks: Endogeneity & Complexity

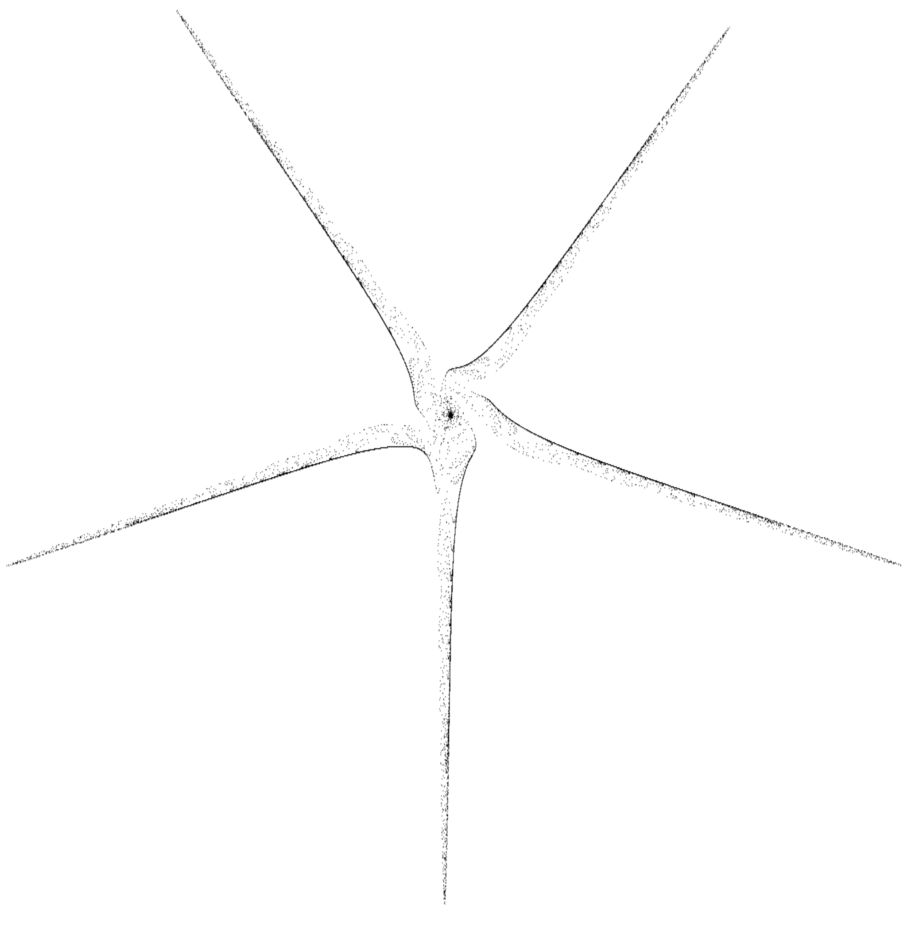

Chapter 7

Cooperation & Networks: towards endogeneity

Attention is now turned to the emergence of *cooperation*. Specifically, this chapter reviews some key contributions to the cooperation literature beginning with biological treatments. The literature is covered within a framework that suggests a progression from uniform (non-local) interactions, to static local interactions, and eventually, to dynamic local interactions. The model to be presented in Chapter 8 relates directly to this trajectory, constructing a truly *endogenous* network formation model of cooperation.

7.1 Introduction

The strategic literature has seen a long-standing interest in the nature of cooperation, with many contributions considering the simple but insightful two-player Prisoner's Dilemma (PD) game,

$$\begin{array}{cc|cc} & & \multicolumn{2}{c}{\mathbf{j}} \\ & & C & D \\ \hline \multirow{2}{*}{\mathbf{i}} & C & (a,a) & (b,c) \\ & D & (c,b) & (d,d) \end{array} \tag{7.1}$$

where $c > a > d > b$ and $a > (b+c)/2$.[1] The rudimentary text-book form of this model is considered as a simple *one-shot*, or *finite repeated game* between two agents. Since the following chapter is predominantly concerned with an evolutionary model of learning and adaptation in a population, the text-book form will be mentioned here in passing only. Moreover, the discussion to follow will not attempt to be exhaustive of this broad topic, but rather, reflecting the model to be presented in the following chapter, will aim

[1] As in Axelrod and Hamilton (1981).

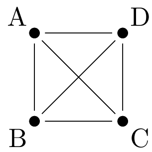

Figure 7.1 Uniform matching, $n = 4$.

to give some background as to how cooperation has been handled from its beginnings in the biological literature to its study in social settings, with a particular emphasis on how agent interactions are handled.

7.1.1 Uniform Interactions

Traditionally, such games were analysed under an uniform interaction specification such that two representative agents i and j from a population $N = 1, \ldots, n$ are selected equiprobably to play a single (or repeated) two-player game such as in (7.1) and then returned to the population pool.[2] It is convenient to think of this kind of matching in graph-theoretic terms as a choice of one edge out of $n(n-1)/2$ edges in a complete graph of size n (Fig. 7.1). This kind of matching seems best suited to situations where agents are essentially the same (homogeneity), have no memory of previous interactions, are essentially un-recognisable to other agents (anonymity), and are not bounded by spatial considerations (non-local). In physical terms, such a matching method is akin to a perfectly mixed, pure atomic gas. As can be seen in what follows, many contributions to the cooperation literature have attempted to enrich this description, bringing in more realistic elements as the domain of enquiry has demanded. Indeed, a progression from the uniform, anonymous interaction case to that of non-uniform and *endogenous* interaction environments as studied in the following chapter can be identified.

7.2 Biological Treatments

In the first (and perhaps oldest-running) case, the biological literature has studied the nature of cooperation (or altruism) in evolutionary game-theoretic models by endowing agents with non-homogeneous identities whilst still assuming uniform interactions. The main questions of interest for these workers is how cooperation or rather, altruistic acts should come about between species or within animals of the same species that ought to care only for themselves (or in more extreme renditions, only for their own genetic lineage[3]).

Early explanations came in different forms and mainly centered on the (continuing) debate between proponents of models motivated by purely *self-interested* but genetically

[2] In the following discussion, we shall refer to this case as 'uniform interaction'.
[3] So Dawkins (1999).

related animals and others who alternatively suggested that cooperation conferred benefits to *groups* of possibly unrelated animals who would be positively selected for by virtue of *group selection*. The former approach forms the basis 'kin-selection' models where animals who are similar enough (genetically) will sacrifice their own reproductive health for the sake of the reproductive health of relatives. The latter approach does not require such relatedness but instead argues that selection does not only apply at the level of individuals and actually may be *dominant* at higher levels, namely groups. A reproductively low-fitness individual may yet be positively selected for if part of a high-fitness group.

To begin with the kin selection camp, Hamilton (1964) classically contributed a now well-known condition to describe why an altruist would sacrificially support a relative,

$$rb - c > 0 \ . \tag{7.2}$$

Hamilton's rule suggests that an altruistic act will occur between two members of the same species if the benefit b in greater related offspring accruing to the recipient is greater than the reproductive cost c incurred by the donor. Significantly, the rule states that this will only occur if the discounting factor r, a measure of genetic relatedness between donor and recipient is high enough. Hence, it came to be the signature model of the 'kin selection' party.

An early (and celebrated) example of such an approach is found in the seminal work of Maynard Smith and Price (1973) (see also Maynard Smith (1964)) who drew inspiration from natural examples of apparent cooperation between related animals such as the mule deer (*Odocoileus hemionus*) whose males fight 'furiously but harmlessly'. Maynard Smith and Price computationally analysed a simple numerical model and showed that more aggressive (and so, energy intensive) strategies such as HAWK (attack to kill) may not be evolutionary superior to more defensive (and so energy conserving) DOVE strategies (mild attack, don't attempt kill). A more recent treatment is given by Riolo et al. (2001) who implement kin selection by modelling agents who match under uniform probabilities but can recognise certain arbitrary characteristics of other agents like themselves (a 'tag') and choose to donate to them. Over time, although the individual concerned gains little directly from their action, members of their 'kin' will prosper through similar behaviours, and so, benefit the species in fitness terms.

Despite the persuasiveness of the kin selection theory, and its subsequent entry as a standard in evolutionary biology texts, the alternative *group selection* perspective, though once virtually banished, is presently making a heavy-weight return. In a 1994 review by two of the major contributors to this debate, Wilson and Sober,[4] were able to list over 200 positive references in support of group selection of various methodological approaches from 1970. Indeed, the title of this review, "Reintroducing group selection to the human behavioral sciences" reflects the 'return from the wilderness' mood amongst its proponents.

Significantly, E.O. Wilson, a major contributor to the early kin selection literature, and founder of the field of *sociobiology*,[5] has recently recanted his previous beliefs and sought to redress the neglect of group selection as a credible theory (Wilson, 2005). Wilson returns to the Hamilton Rule and notes several objections, principle among which is the mounting evidence in biological systems that negates the relatedness assumption required by kin selection (the non-zero value of r in (7.2) above). A signature case is provided by that of the worker castes in social populations of ants and wasps as famously studied for many years by Wilson himself (Wilson, 2001). In contrast to kin selection theory, he finds that these colonies of workers are founded by genetically *unrelated queens*. Consequently, he notes that the previously held idea that agents within cooperative colonies had to be genetically related to each other as a necessary condition for their collective activity, in fact need not be the case. Rather, the apparent relatedness (in the cases that it is found) can be explained as a *product* of the successful colony. He concludes, 'individuals do not form colonies because they are closely related .. they are closely related because they form colonies' (Wilson, 2005). E.O. Wilson suggests a revised version of the Hamilton rule to account for this fact,

$$rb_k + b_e - c > 0 \ , \tag{7.3}$$

where b_k is now the collateral kin selection benefit accruing to the donor based on relatedness to the recipient, and the additional term, b_e is the colony-level benefit independent of relatedness. Even so, a common methodology in group selection models has been to use cohesive groups as the basis for improvements to individual fitness, and so show how group-level organisation can be decisive in retaining members of a population who would otherwise have been selected against (Traulsen and Nowak, 2006; Wilson, 1975, 1976).

[4]In this debate, names are a source of great confusion. Two are unavoidably referred to: E.O. Wilson; and David Sloan Wilson. The former is the eminent insect biologist and founder of *sociobiology* which is based on kin selection, whilst the latter has been a steady (and continuing) proponent of group selection.

[5]Sociobiology sought to find evolutionary underpinnings for all observed social (e.g. colony-level) behaviour in animals on the basis of kin selection motivations, see E.O. Wilson (1975), *Sociobiology: The New Synthesis*, Belknap Press of Harvard University Press: Cambridge.

It should be noted that a popular alternative to kin selection and group selection is that of *reciprocal altruism*. As the name suggests, reciprocal altruism requires cooperative agents to come into contact with other recognisable cooperators and so receive a 'return benefit'. If two agents are required to meet again to obtain their *recompense*, it is the *direct* form of reciprocal altruism, the alternative being indirect reciprocity. In this latter case, it is supposed that such meetings do not have to be repeated between two specific individuals, but merely between two members of the cooperative sub-population. Recent examples include Nowak and Sigmund (1998) and Bowles and Gintis (2004) who suggest reciprocal mechanisms such as 'image scoring' (a kind of community esteem for cooperative individuals) and 'strong reciprocity' (rewards and punishments on normative lines) respectively. Reciprocal altruism differs from kin and group selection in that the cooperating individual them self is the instrument of genetic information transmission, rather than a related individual (so kin selection) or an unrelated, but group-wise helpful individual (so group selection).

The debate in the biological literature will no doubt continue for some time, each side having their battle-hardened protagonists. For the present purposes, the contrasting theories of cooperation are instructive. As will become apparent below, social scientists, and in particular, economic theorists and practitioners have unconsciously incorporated elements of one or more of these biological theories in models of human behaviour. In some ways this is inevitable, and forms the closing remark of Wilson's reflection on the falling stocks of kin selection, commenting,

> "For the present, however, the ongoing shift to group-level selection forced by empirical evidence suggests that it might be profitable to undertake a similar new look at the wellsprings of social evolution in human beings ... where, I believe, surprises also await us." (Wilson, 2005, p.165)

In fact, the incorporation of networks (in different guises) as considered below can be thought of as doing just this. Non-uniform interaction structures, be they set-based, regular grids, lines, circles, or more exotic descriptions, imply the formation of interaction boundaries, or 'groups', and many of the results from the area note how such boundaries facilitate non-uniform equilibrium populations. Whether these models (and future models for economic activity), explicitly attempt to differentiate between the competing kin- and group- (or reciprocal altruism) selection models of the biological literature is less clear. However, that humans engaged in economic activity feel the effects of levels of selective pressure seems obvious (e.g. business units, firms, regions, countries), and hence, one suspects fruitful future interplay between the two literatures.

7.3 Non-uniform Interactions

7.3.1 Networks as Types

As a first step towards incorporating non-uniform interaction pattern formation in human systems, authors have considered cooperation (or corruption) 'networks' as relatively crude typing mechanisms. In these studies, a 'network' is used to describe a (proper) subset of agents in the population who are then distinguished from the majority in some way. For example, in Taylor's 'old-boy network' (Taylor, 2000), agents who are in the network are known to be of a certain type – the qualified/competent type. Membership of this network is conferred upon the individual after 'showing their colours' in an interaction. The mixing of agents is population-wide, and therefore, in this model, the 'network', although giving important *type* information for future transactions, plays no more part in the interaction space, nor does the actual topology of the network matter.[6] Since there is no network exit criterion, nor behavioural dynamic, Taylor finds that networks are rarely socially optimal (as opposed to anonymous transactions) since a bleeding of the 'good' types from the general population ensues. Further, in a model of institutional reliability, Kali (1999) finds the same results, with a comparison between two business regimes for the transaction of an homogeneous good at fixed price being either: *open market* (anonymous transactions); or *network* (buyer and seller operating on given norms). In this case, the network in mind is that of the *guanxi* network in China where a gift given upon joining is a form of type identification, and all transactions in the network are assumed to be honest ones; dishonesty resulting in ejection from the network set.[7] However, these approaches suffer from the constraints imposed by the analytic framework, thus only allowing one (informal) fully-connected component to form with such formation not endogenised. Authors assume that where networks are sustainable they will form.

7.3.2 Structural Imposition

As with some biological examples, social and economic models have increasingly been tied to static non-uniform interaction spaces, especially through games of strategy on these networks. Here, some form of topological imposition, other than the complete graph, such as uni-dimensional play on a line, circle, or higher-dimensional interaction on a regular graph (e.g. a torus) is usually applied, with more recent contributions allowing for richer (statistical) graph environments such as so-called 'small-world' graphs (Elgazzar, 2002; Kirchamp, 2000; Masuda and Aihara, 2003; Stocker et al., 2002).

[6] As is perhaps clear, this is not a network in the sense of a formal graph with an edge set, but can be thought of as a disjoint graph with the 'network' comprising a complete connected component.

[7] For further reading on the *guanxi* network, see Fan (2002); Standifird and Marshall (2000).

For instance, early work by Nowak and May (1992) with simple agents playing a standard PD game with near-neighbours on a regular two-dimensional lattice revealed how cooperative types could co-exist with defection types indefinitely – cooperation with neighbours would be reciprocated directly due to spatial locality (changing the interaction probabilities). Moreover, their work emphasised how even in this simple spatial structuring to simulate local-interactions, richly diverse population distributions with fractal properties could be obtained.[8] In a similar vein Cohen et al. (2001) specifically place agents on a spectrum of regular to irregular interaction networks and again find regions where cooperation can be sustained.

The topological significance of the interacting space has been stressed by these authors as it appears to influence the degree to which cooperation is achievable. Other examples include Anderlini and Ianni (1996) and Anderlini and Ianni (1997) who find that different actions of a pure coordination game survive in the long-run at different locations on the interaction space. Whilst Kirchamp (2000) computationally studies interacting agents on a torus playing the PD and coordination games, with cooperation and non-risk-dominant coordination outcomes observed respectively. Consequently, and reflecting a burgeoning interest in networks of all kinds, much attention has been paid to the study of realistic social networks (Barabasi et al., 2002; Baum et al., 2003), with statistical network characterisation (Watts and Strogatz, 1998; White et al., 2004) and clique analysis similarly receiving interest (Girvan and Newman, 2001; Tyler et al., 2003).

7.4 Non-uniform, quasi-endogenous Interactions

7.4.1 Walking Away

One simple non-uniform interaction specification has been motivated by the observation in both biological and human cases for individuals to 'walk away'. For instance, Joyce et al. (2006) stress that a long-standing feature of animal psychology is that 'animals tend to leave a situation that they find distasteful'. That is, although local-interaction structures promote reciprocal altruist behaviours, they do not allow for defensive types to move away from predators. One very recent contribution (Burtsev and Turchin, 2006) goes some way to admitting this element. They use a novel input-output weighting implementation of species cognition and therefore allow agents to move left and right on a two-dimensional grid (in addition to other actions such as eat, reproduce, attack etc.) in response to environmental stimuli such as predators in near-bye cells. Their main

[8] Interestingly, although Nowak and May's early (1992) work generated striking chaotic patterns in (even) a simple two-dimensional interaction space, the cooperation literature seldom admits such dynamics thereafter. A point to be returned to in the following chapter.

finding is that ecologies of agents with specific niche skills can arise endogenously once the ability to move and exploit multi-agent action is incorporated. Although this work has biological foundations and objectives it is an interesting development, and further motivates the modelling of interactions between the agent and their matching environment as is considered in the following chapter.

On this score, humans appear no different, choosing to leave unhappy interactions (Orbell and Dawes, 1993). Agents respond to good and bad play with their feet. Attempts to capture this modification have concentrated on a simple 'exit-criterion' such that an agent can choose to cease the interaction with a playing partner and be re-assigned to another agent from the 'singles' pool, if some pre-determined trigger is satisfied (e.g. the first defection of an opponent). The effect of such a modification is non-trivial and contrasting.

Computational experiments by Yamagishi and Hayashi (1996) that mimic those of Axelrod and Hamilton's tournament games of the 1980s but add in the exit-option suggest that the strategy OUT-FOR-TAT (exit if partner plays defect) can overcome the famous TIT-FOR-TAT strategy. In this formulation, one would therefore expect that the exit-option will *increase* overall cooperation levels in the population.[9] However, an opposing theory suggests that incorporating an exit-option will *decrease* cooperation by reducing 'perceived dependence on the other'; a 'hit-and-run' type player can flourish in a short-interaction environment without opportunity for retaliation by the victim.

Laboratory experiments reported by Boone and Macy (1999) attempt to tease out the differences in these theories. Their experimental design aimed to discriminate between the countervailing effects of *partner selection* and *relational dependence* (following Pruitt and Kimmel (1977)), respectively capturing the ability to choose a desired partner, and the need for future help from a specific partner without the ability to turn elsewhere. Subjects who were aggressive in nature (favouring retaliation) displayed reduced cooperation in regimes where an exit play was allowed (they had to play with like- outcasts), whereas defensive types increased cooperation. Conversely, where exit was forced on players, defensive types cooperated less (compared to optional-exit), whereas cooperative play by aggressive players was unaffected. The authors conclude that the natural ability of players to exit from an undesirable interaction causes different effects on players for different reasons according to their underlying type, hence producing overall increases or decreases in cooperation accordingly.

[9] This result is replicated in Joyce et al. (2006), although OUT-FOR-TAT is relabelled MOTH (My-Way-Or-The-Highway).

7.4.2 Partner Choice and Refusal

Apart from a few exceptions in the theoretical literature that aim to speak to questions of *optimal* organisation, or communication structures/motifs on a relatively small-scale (Bala and Goyal, 2000; Dutta et al., 1998; Ely, 2002; Jackson and Watts, 2002a; Mitchell, 1999; Slikker and van den Nouweland, 2000), authors have not allowed the interaction environment itself to vary, either (for example) due to some exogenous schedule or as a result of processes endogenous to the model. Such a modelling feature comprises an highly desirable step towards treating realistic economic and social networks. The generally acknowledged *guanxi* business network in China as documented in (Standifird and Marshall, 2000) and (Fan, 2002) appears to be a clear example of a dynamic and strategic network in action. One promising approach that generalises the 'walk-away' models mentioned previously is the preferential partner selection (*choice*) and optional rejection of an offer to interact (*refusal*) literature (or IPD/CR when the game is the IPD). Here, the emphasis is on how the added mechanism of choice and refusal affects the emergence of cooperation in IPD games. As noted above, in discussion of the exit-criterion, such a mechanism is seen as more realistic, from both a biological, and social perspective.[10]

Indeed, by analysing the outcome of interactions between agents in such a set-up, a network of activity can be identified and analysed. For example, Ashlock and coworkers (1996) construct a computational model (see below) to consider the effects on cooperative behaviour with varying levels of preferential selection, finding that most ecologies converge to full cooperative behaviour but that 'wallflower' ecologies are possible if intolerance to defection is high, or costs to social exclusion is low.[11] Such findings are supported to some extent by the experimental work of Hauk and Nagel (2001) who find that cooperative behaviour increases over time under unilateral choice of partners (opponent must accept to play).

A second paper by Smucker et al. (1994) again considers a computational model, but in addition to considering the strategic implications of various levels of choice and refusal, they also perform some characterisation of the evolving network of interactions. In the SSA model, agents are modeled as 16 state Moore machines[12] who are programmed to play the IPD. Importantly, agents keep track of their expected payoffs for interaction

[10]See the introduction to Smucker et al. (1994) on such observations.

[11]See also, Tesfatsion's work on trading games with endogenous partner selection (Tesfatsion, 1997).

[12][As in (Miller, 1996b, p.91), see also Miller (1988)] A Moore machine is defined by the four-tuple $\{Q, q_o, \lambda, \delta\}$ where Q is the set of *internal states*; q_o is the initial state; λ is a mapping from each state to the subsequent action to be played $\lambda : Q \leftrightarrow S_i$, for example, in the PD, $S_i \in \{C, D\}$; and δ is the *transition function* that maps from the current internal state of the machine to the new internal state, contingent on the *opponent's* reported move, $\delta : Q \times S_{\sim i} \leftrightarrow Q$, $S_{\sim i} \in \{C, D\}$ being the opponent's reported move last period (in this case, for the PD).

with all other $n - 1$ agents in the population. For example, player i keeps track of the expected payoff for an interaction with some player j, and if this value is higher than some 'tolerance' τ, then, it is said, i wishes to play with j (choice). Such offers are naturally bounded and subject to the receiver accepting the offer to play ('mutual tolerance'), which forms the 'refusal' mechanism. A single or repeated (IPD) interaction follows, with each player naturally changing its Moore machine state accordingly.

However, the 'network' in the SSA model (as for Ashlock et al. (1996)) is defined by a simple global rule – if the number of interactions between two players is (statistically) significantly larger than the mean interaction count for the whole population, then an edge is assigned between these players. Thus, for SSA, the 'network' is more a record of 'acceptable payoff outcomes' rather than a functional entity which shapes future interactions. This is an acknowledged limitation of the work. Furthermore, in addition to storing payoff outcomes, agents store a unique Moore machine state for every opponent. That is, an agent's Moore machine describes their total behaviour set, but against any given opponent, they choose the 'side' of their behaviour (i.e. they choose an initial state) reflecting their current playing 'history' with that player.

7.5 Towards a Truly Endogenous Strategic Network Model

A developmental thread can be drawn through the literature considered above. From socio-biological beginnings where uniform interactions were assumed and the interactions between individuals enriched, to the imposition of firstly static, and then in some limited way, dynamic networks, the next logical step for this inquiry is to allow agents to determine the interaction structure *themselves*. Although recently considered contributions like that of the SSA model are useful first steps, they again do not permit the agents to define the interaction environment as a result of their own actions and strategies, instead choosing to define a network based on exogenously defined interaction metrics.

The incorporation of network formation as a truly agent-motivated activity suggests two directions of important causality, first with respect to how a changing interaction environment might affect strategic outcomes for agents (the 'topological' effect); and second, how agents *through their own strategic actions* might impact on the very interaction space itself (the 'agency' effect). With the perspective of the biological literature mentioned above, it can be seen that the former effect is a kind of group selection effect: if agents, by virtue of their own specific actions in a connected component, are simply part of a network structure with heavy negative externalities, or poor performance in comparison to another component, they will be selected against. Likewise, the latter effect, that of agency, can be seen as a generalisation of both the 'walk-away' and 'choice-refusal' literature, since agents will mutually choose to construct or break interaction patterns depending on their preferences. It is to be noted, that this will not necessarily mean that an agent is receiving high utility in each interaction, rather, their preference may be to fulfill a kind of mediator role (say) where the emphasis is on the quantity of interactions, not their quality. These dual issues will be jointly tackled in an endogenous strategic network formation model.

With reference to Chapters 2 and 3 of this thesis, it is to be recognised that such a hypothetical system of interacting agents, with heterogeneity in agent behaviour *and* interaction profiles fits well into the so-called 'science of complexity'. As mentioned, this approach seeks to identify and study systems whose components interact in some non-uniform, (and usually) non-linear manner. In particular, due to the inherent unpredictability of such non-equilibrium systems, agent-based computational modeling techniques provide an extremely useful method of enquiry especially where non-rational learning based behavior is also suspected (Arthur, 1994; Holland and Miller, 1991).[13]

The modelling framework that follows in the next chapter reflects such an approach. Specifically, both constraints concerning agent rationality and rigid agent interactions are relaxed within a fundamentally agent-based modelling framework. Moreover, in contrast to the somewhat related approach of SSA mentioned above, agents are given *strategic* abilities to change the interaction space *themselves* (i.e. to change interaction probabilities) during pair-wise game-play. It is in this sense that a 'network' arises in the model, and hence, such a network is said to be a truly *endogenous* feature of the modelling framework; a feature which to this author's knowledge has not been previously handled with boundedly rational agents.

To conclude, the model presented and studied in the following chapter aims to address several of the mentioned shortcomings of the literature. First, by implementing network formation as a *strategic* and therefore truly endogenous process; second, and following on from the first, by allowing for multiple networks to form simultaneously (rather than one connected component only); and third, by implementing agent strategies as finite state automata, both bounded rationality and learning are incorporated in a transparent way.

[13]See also applications of artificial adaptive agents to organizational problems in Choi (2002); Marsili et al. (2004); Stocker et al. (2002).

Chapter 8

A Model of Endogenous Cooperation Network Formation: from simplicity to complexity

Following on from Chapter 7, a model of cooperation in a truly endogenous, dynamic interaction environment is introduced and analysed. Specifically, the model indicates that with a small change to the standard Iterated Prisoners' Dilemma set-up to afford network formation, a variety of strategic networks can result including both cooperation and defection based structures. Furthermore, the long-run dynamics of this model are analysed and relate directly to foregoing discussions of *complexity*, *emergence*, and *self-organization* in Chapter 3.

8.1 Introduction

In the previous chapter a variety of cooperation models was considered with a particular interest in those that have attempted to incorporate a 'network' in some way. As expressed in the preceding chapter, the model to be developed and analysed presently aims to provide a significant step forward in both ways of thinking about, and (surprising) results concerning, the parallel modelling of cooperation and strategic network formation. The approach taken seeks to make a small modification to the standard iterated Prisoners' Dilemma (IPD) set-up, so that in the first case where interactions are uniform in nature, an analytic analysis is tractable and can be directly compared to canonical evolutionary game-theoretic results. Following this, the modelling set-up permits a 'turning-up' of the impact that strategic decisions by agents are having on the interaction structure, that is, a network is allowed to develop. It is from here that the inquiry focuses on how and why networks of various guises might form, and then what happens to these networks over longer time-scales.

The key insights of the present work can be summarised as follows: first, an analytic analysis without network formation reveals that the modification to the standard IPD framework introduced below does not change the canonical behaviour of the system; second, that when network formation is afforded, stable cooperation networks are observed, but only if both a type-selection and enhanced 'activity' benefit of the network are present; third, that the extended system under certain interaction lengths is inherently self-defeating, with both cooperation and defection networks transiently observed in a long-run specification; and fourth, that the network formation process displays self-organized criticality and thus appears to drive the complex dynamics observed in the long-run.

The rest of the paper is organised as follows: first, the model is introduced in overview, and then in detail, paying particular attention to the modelling of agents and the incorporation of network forming behaviour; second, analysis is performed analytically on the basic (non-network forming, or uniform interaction) model before extension to incorporate network formation is performed on both short- and long-run time horizons; and finally, some concluding observations and a discussion of possible extensions is made.

8.2 The Model

8.2.1 Overview

Agents are modeled as finite state automata (FSA) with a maximum number of feasible states. As with normal renditions of these automata, each agent has an initial state which is not contingent on the opponent they are playing, and each state describes both their action for that state, and their state transition contingent on the play of their opponent.

Agents begin with a uniform interaction environment (a null-graph) and within a period undergo at least some minimum number of interactions with other agents to play the IPD. Within each interaction, agents are able to influence the interaction environment by signaling to their opponent that they wish to break the interaction and reveal their positive or negative response to their opponent. If both players play positive signals, an edge is assigned between them, and the two agents will meet each other with higher probability in the future. The exact value of this probability is contingent on how many other agents each has already formed a link with.

In this way, the concept of 'partner-scarcity' is incorporated: though link-formation increases the probability that two agents will meet again, it does not guarantee it. Consequently, successful agents must either protect themselves completely from exploitative players through link formation, or display a depth of complexity in their strategy that can manage playing against undesirable opponents (or a combination of the two).

At the end of a period, total payoffs are determined for each agent, and an 'elite' fraction of the population is retained for the next period, with the remainder being replaced by new agents.[1] New agents are generated from a combination of existing elite behaviours and new behaviours (a type of learning) followed by mistake-making/innovation. Elite agents retain their links between periods (so long as they are to fellow elites) whereas entrants begin with no links, befitting the concepts of incumbency and network dynamism.

In this way, links are established within a period by mutual agreement between two agents. However, links can only be broken when an agent leaves the population after selection, severing all pre-existing links.

8.2.2 Details

Let $\mathbf{N} = \{1, \ldots, n\}$ be a constant population of agents and denote by i and j two representative members of the population. Initially, members of the population are uniformly paired to play the modified IPD game $\mathcal{G}$ described below. When two agents are paired together, they are said to have an *interaction*. Within an interaction, agents play the IPD for up to a maximum of τ iterations, receiving a payoff equal to the sum of the individual payoffs they receive in each iteration of the IPD. An interaction ends prematurely if *either* player plays a 'signal' thus unilaterally stopping the interaction. A strategy for a player s describes a complete plan of action for their play within an interaction, to be explained presently. In addition to the normal moves of cooperate (C) and defect (D), an agent can also play one of two signal actions, $\#_s$ and $\#_w$ respectively. Thus, in any one iteration of the IPD, the action-set for an agents is $\{C, D, \#_s, \#_w\}$. As mentioned above, the playing of a signal by either player leads to the interaction stopping, possibly prior to τ iterations being reached. The playing of a signal can thus serve as an exit move for a player.

The interpretation of the two types of signal is as follows. Although initial pairing probabilities between all players are uniform random, agents can influence these interaction probabilities through the use of the signals. Formally, let some agent i maintain a preference vector,

$$\left\{f^i \,:\, f^i_j \in \{p_s, p_0, p_w\} \;\forall\; j \in \mathbf{N}/i\right\} \tag{8.1}$$

[1] Alternatively, one can think of this as a steady-state *strategy* framework, whereby the stock of agents is constant between periods, but some fraction decide to update their strategies. In what follows we shall continue to think in terms of 'entrants' (new agents), though either interpretation is equally valid.

where f^i_j is the preference status of agent i towards agent j and $p_s > p_0 > p_w$ are Natural and denote *strengthen*, *untried* and *weaken* preferences respectively. Initially all entries are set to p_0 for all $j \in \mathbf{N}/\{i\}$. A probability vector r_i for each agent is constructed from the preference vector by simple normalisation onto the Real line,

$$\left\{ r^i : r^i_j = \frac{f^i_j}{\sum f^i} \quad \forall\, j \in \mathbf{N}/\{i\} \right\} \ , \tag{8.2}$$

such that each opponent occupies a finite, non-zero length on the line $[0,1]$ with arbitrary ordering. Since we study here a model of mutual network/trust formation, preferences can be strengthened only by *mutual* agreement. Specifically, if agents i and j are paired to play the IPD, then when the interaction ends in iteration $t \le \tau$,

$$f^i_j = f^j_i = \begin{cases} p_s & \text{if } \ s^i_t = s^j_t = \#_s \ , \\ p_w & \text{else} \ , \end{cases} \tag{8.3}$$

where s^i_t denotes the play of agent i in iteration t. That is, in all cases other than mutual coordinated agreement, the two agents will lower their relative likelihood of being paired again (though the playing of $\#_w$ might cause the interaction to end prematurely with the same result). Payoffs for each iteration of the PD are given by (8.4) below.

$$\mathcal{G}(\mathbf{I},\mathbf{II}) = \qquad \begin{array}{cc|cccc} & & & \mathbf{II} & & \\ & & \#_s & C & D & \#_w \\ \hline & \#_s & (0,0) & \cdots & & (0,0) \\ & C & \vdots & (3,3) & (0,5) & \vdots \\ \mathbf{I} & D & & (5,0) & (1,1) & \\ & \#_w & (0,0) & \cdots & & (0,0) \end{array} \tag{8.4}$$

The playing of signals, is costly: the instantaneous cost for that period is the foregone payoff from a successful iteration of the IPD.

8.2.3 Example

Let two agents i and j be chosen to play (8.4) in some period and let maximum interaction length $\tau = 3$. Consider the following interaction,

Iteration	P_i	P_j	π_i	π_j
1	C	C	3	3
2	D	C	5	0
3	$\#(s)$	C	0	0
$\sum \pi_x$			8	3

note that i played an unrequited strengthen signal in the third (last) iteration; both players' interaction preference entries would be set to p_w.

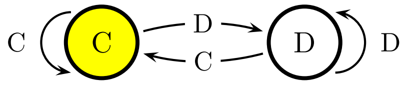

Figure 8.1 Example FSA: TIT-FOR-TAT.

8.2.4 Game Play

In a *period* each agent is addressed once in uniformly random order to undergo m interactions with players drawn from the rest of the population ($\mathbf{N}/\{i\}$). An agent is paired randomly in accordance with their interaction probability vector r^i with replacement after each interaction. Preference and probability vectors are updated after every interaction.

Thus, it is possible that, having previously interacted with all agents, an agent retains only one preferred agent, whilst all others are non-preferred, causing a high proportion (if not all) of their m interactions to be conducted with their preferred partner. However, it is to be noted that the value of m is only a *minimum* number of interactions for an agent in one period, since they will be on the 'receiving end' of other agents' interactions in the same period. In this way, agents who incur an immediate cost of tie strengthening (foregoing iteration payoffs) can gain a long-term benefit through further preferential interactions.

At the end of T periods, the population undergoes selection. A fraction θ of the population is retained (the 'elites'), whilst the remainder $(1 - \theta)$ are replaced by new agents as described below. Selection is based on a ranking by total agent payoffs over the whole period. Where two agents have the same total payoff in a period, the older player remains.[2]

8.2.5 Agent Modeling

Each agent is modeled as a k (maximum) state FSA.[3] Since the interaction will stop immediately after either player plays the signal # each state must include two transition responses only: $R(C)$ and $R(D)$. For example, an agent's first three states might take the form (schematic representation given in Fig. 8.1),

State	P	$R(C)$	$R(D)$
1	C	1	3
2	C	2	1
3	D	1	3

where the first state could be read as,

'play C next, if the opponent plays D, go to state 3; else, stay in state 1.'

[2] Following SSA Smucker et al. (1994).

[3] To facilitate the computational modeling of this environment, agent strategies were encoded into binary format. See Miller et al. (2002) for an analogous description of this method for FSA.

An agent will have k such states as part of their 'strategy'. By convention, the first state is taken as the initial one.

It can be shown[4] that the maximum number of states k possible for an FSA playing some game with count $|R|$ feasible transition responses lasting at most τ rounds is simply,

$$k(R, \tau) = \sum_{t=0}^{\tau-1} |R|^t \ , \tag{8.5}$$

hence, under this regime, the maximum interaction length τ and the number of opponent plays requiring a response defines the maximal FSA length. However, it is to be noted that there is no guarantee that all k states for a given agent will be accessed (consider the example given immediately above, state 2 is present but is strategically redundant). In this way, FSA give a tangible sense of 'strategic complexity' when it comes to individual strategies. An agent who uses all of their $k > 1$ states as part of their strategy will no-doubt display a deeper strategy in action, than an agent who is playing merely (say) ALL-C. Of course, such complexity of strategy may or may not correspond to relative success in the population.[5]

After each period, a fraction θ will stay in the population, with the remaining agents being filled by new entrants. Here, the process of imitation and innovation/mistake-making is implemented via two foundational processes from the *genetic algorithm* (GA) literature.[6] Initially, two agents are randomly selected (with replacement) from the elite population. A one-point crossover operator is applied to each agent, and two new agents are formed. The strategy encoding (bit-strings) of these new agents then undergo point mutations at a pre-determined rate (5 bits per 1000). This process (random selection, crossover and mutation) continues until all the remaining spots are filled. An algorithmic representation of the main procedure is given in Alg. 2.

8.3 Results & Discussion

8.3.1 Uniform interactions

To begin, we study a static uniform interaction space to check any unwanted outcomes due to the modified IPD set-up. In this situation, rather than agents upgrading their

[4]Consider a single-rooted logic tree where each node is a state, and each branch some transition. Without any re-use of states, the tree can be at most $\tau - 1$ nodes deep. Now observe that after the initial node (call this layer $t = 0$) each new layer will produce $|R|^t$ new nodes. The result follows.

[5]For example, it has been shown countless times before that the humble TIT-FOR-TAT (play C until the opponent plays D, then switch to D until the opponent plays C, then switch back to C, and so on.) strategy (and its variants) is often an extremely effective one against all manner of opponents, even though it can be represented by just two states!

[6]See, for example, Goldberg (2002) or for a non-technical introduction, Holland (1992).

Algorithm 2 The main program

$N \leftarrow \{1 \dots n\}$
$s \leftarrow$ INITIALIZE_STRATEGIES(n, τ)
$age \leftarrow 0$ for all N
$f \leftarrow p_0$ for all N ▷ Set all preferences to *unknown*
$r \leftarrow$NORMALIZE(f)
for T periods **do**
    $\phi \leftarrow$ SHUFFLE(N)
    $\Pi \leftarrow 0$ for all N ▷ Initialize period payoffs
    for each $i \in \phi$ **do**
        for m interactions **do**
            $j \leftarrow$ CHOOSE(r^i)
            $[\pi_{i,j}, f^{i,j}] \leftarrow$ PLAY_GAME$(\mathcal{G}, \tau | s_i, s_j)$
            $r \leftarrow$ NORMALIZE(f)
            $\Pi_k \leftarrow \Pi_k + \pi_k$ for $k = \{i, j\}$
        end for
    end for
    $E \leftarrow$ SELECTION(Π, age, θ) ▷ Elites
    $s_F \leftarrow$ GENERATE_STRATEGIES(s_E) ▷ Entrants
    $age \leftarrow 0$ for all F
    $age_i \leftarrow age_i + 1$ for all $i \in E$
    $f \leftarrow p_0$ for all F
    $r \leftarrow$ NORMALIZE(f)
    $N \leftarrow E \cup F$
end for

preference vector after each interaction, the preference vector is uniform and unchanged throughout the model. In this way, the effect of the modification to the standard IPD framework can be analysed. Under such a scenario, the action set for each agent reduces to $\{C, D, \#\}$ since the signal action $\#$ has no interaction space interpretation, but still provides a means of prematurely ending the interaction (thus we may drop the sub-script).

To keep matters simple, we consider a model in which the maximum interaction length $\tau = 2$, which by (8.5) yields a maximum FSA state count of $k = 3$. Under these conditions, a strategy will be composed simply of a first play, and response plays to C and D.

With uniform interaction probabilities, this model can be thought of as a modified evolutionary game theoretic framework. The probability of interacting with a certain agent type is directly equal to the proportional representation in the population of that type. Modification of the standard framework is due to the genetic algorithm approach, the crossover operator providing for imitation in addition to the more standard 'random' mutation operator. However, in terms of evolutionary stable strategies, this modification is insignificant.

In this setting, no evolutionary stable strategy will include $\#$ as a first play, since the payoff for such a strategy with any other agent is 0.[7] This leaves strategies in the form

[7] The interaction would end after the first iteration, and $\mathcal{G}(\#_x|y) = 0$ for all $x \in \{s, w\}$ and $y \in \{C, D, \#(s), \#(w)\}$.

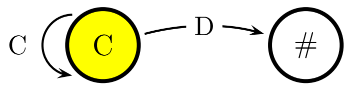

Figure 8.2 Agent strategy $s_{C\#} : \{C, C, \#\}$.

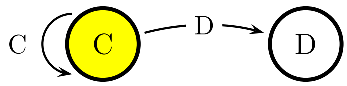

Figure 8.3 Agent strategy $s_{CD} : \{C, C, D\}$

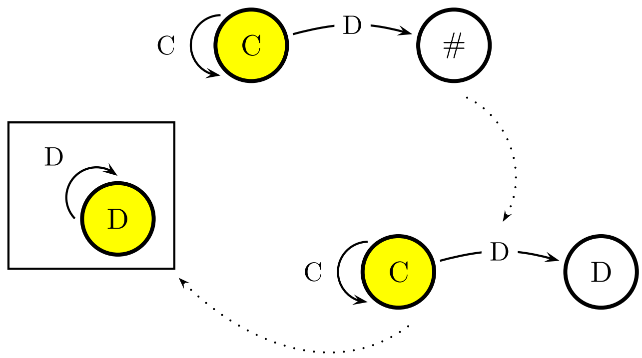

Figure 8.4 (Clock-wise from top) $s_{C\#}$ is stable in the presence of s_D only but is susceptible to attack by s_M which in turn will yield to s_D, the final evolutionary stable strategy component.

of a triplet,

$$s : \{P_1, R(C), R(D)\} ,$$

where $P_1 \in \{C, D\}$ and $R(.)$ indicate subsequent plays in response to either C or D plays by the opponent $R(.) \in \{C, D, \#\}$. In all, 18 unique strategies can be constructed.

It is instructive to consider whether cooperative strategies might be evolutionary stable in this scenario. Clearly, a strategy $s_C : \{C, C, C\}$ will yield strictly worse payoffs than the strategy $s_D : \{D, D, D\}$ in a mixed environment of the two. It can be shown[8] that the strategies $s_{C\#} : \{C, C, \#\}$ and $s_{CD} : \{C, C, D\}$ (Figs. 8.2 and 8.3), constitute the only two ESSs in an environment of $s_{D,D,D}$ only. However, both $s_{C\#}$ and s_{CD} are themselves susceptible to attack by a 'mimic' agent such as $s_M = \{C, D, D\}$, which itself will yield to the familiar s_D.

In this way, even with the added facility of being able to end the interaction prematurely, the only evolutionary stable strategy with respect to the full strategy space is that of $s_{D,D,D}$. Any intermediate resting place for the community will soon falter and move to this end (Fig. 8.4).

---

[8]See Lemma 1 in Appendix.

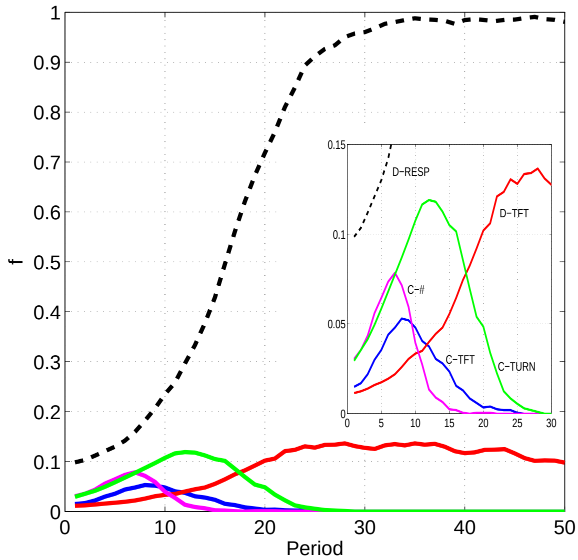

Figure 8.5 Population fraction of strategies under uniform mixing: (D-TFT) 'nasty' TIT-FOR-TAT; (C-TFT) 'nice' TFT; (C-#) plays C, returns D with #; (C-TURN) 'turn-coat', plays C, but D thereafter; (D-RESP) sum of all players who play D and return D with D.

8.3.2 Uniform Interactions: Computational Results

Computational experiments were run under uniform mixing as described above as a method of model validation. Given the maximum interaction length of $\tau = 2$, results for strategies present as a fraction of the total population are given in Figure 8.5.[9] As predicted above, the model shows the clear dominance of s_D under uniform mixing. Additionally, as predicted, the initial 'shake-out' periods ($t < 30$) gave rise to interesting wave-like strategic jostling. Agents playing cooperation first, and replying to D with # were the first to have an early peak, if short-lived, which is not unexpected, since playing the signal is not the best-response to any subsequent play. Thereafter TFT-nice (C) peaked, but were soon overcome by the turn-coat type (who dominates TFT-nice). However, as the stock of C players diminish, 'turn-coat' too, yields to the D-resp type strategies (such as TFT-nasty (D)).

We may conclude then, that the presence of the signal play (#) does little to affect *strategic* outcomes in the standard IPD set-up; defection still reigns supreme in the uniform IPD environment.

8.3.3 Non-uniform Interactions: Network Formation

From the preceding analysis, a natural question arises: under what circumstances, if any, do networks emerge? Whilst it is possible to think of the network formation process as

[9] Modelling notes: each plot-line represents the average result from 20 modelling runs where in each run: $n = 100$ and $m = 20$ with 20 agents replaced after each period, under roulette-wheel type selection of elites (with replacement) and $p(mutation) = 0.5\%$. Results were unchanged in substance for $m \in \{2, 5, 10\}$.

an aid to certain dominated behaviour (e.g. mutual cooperation), it is not obvious to what degree the network must shield certain players from hostile behaviour before their interaction community can be self-supporting.

Two factors that will clearly affect the propensity for cooperative networks to arise are:

1. The 'impact' of the network on the interaction space; to what extent a strengthening signal actually changes mixing probabilities – formally, the values of the preference set $\{p_0, p_s, p_w\}$; and
2. The minimum number of interactions *within* a period over which two agents can exploit a beneficial relationship – formally, the value of m.

These two factors are related, since a low interaction impact may be compensated for by an high absolute number of interactions within a period.

The impact of network formation decisions by agents was parametrised in the computational experiments as follows,

$$p_w = (1-\eta)^2 \quad \text{and} \tag{8.6}$$

$$p_s = (1+\eta)^2 \ , \tag{8.7}$$

where $\eta \in [0, 1)$. The choice of the expression is somewhat arbitrary, however, the current specification retains symmetry about $p_0 = 1$ for all values of η and by taking the squared deviation from 1, the ratio p_s/p_w could be easily varied over a wide range. For example, by choosing $\eta = 0.8$, we yield $p_w = 0.04$ and $p_s = 3.24$, which gives a ratio of preferences in f^i for agent i for two agents with such values of 81. That is, the preferred agent (where an edge is assigned for the purposes of visualisation) will be met around 80 times more regularly than the less preferred agent when i is being addressed. However, actual interaction probabilities are determined by the full vector of preferences. If $n = 100$, and $\eta = 0.8$ and i had only 1 opponent j whom they preferred (the other $n-2$ agents being non-preferred), the matching probability between i and j would be around 0.45.

To determine what conditions are favourable for network formation, a second computational experiment was conducted, this time 'turning up' the *interaction space impact* of any signalling play by the agents. Specifically, the network tuning parameter η was varied in the range $[0.2, 0.95]$ together with the minimum interaction parameter m over $[2, 20]$.

It was found that necessary conditions for sustainable network formation were $\eta \gtrsim 0.8$ and $m \gtrsim 10$. In terms of the population, these accord with a ratio of p_s to p_w (by (8.6) and (8.7)) of around 80 times, and a minimum fraction of interactions per period of around 10% of the population size.

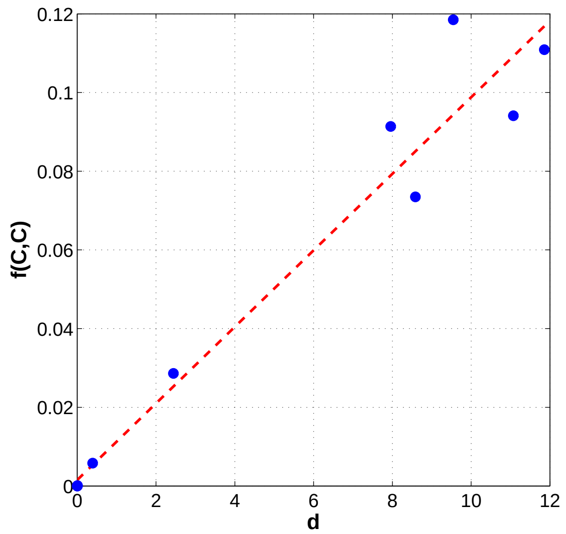

Figure 8.6 Fraction of mutual cooperation per PD vs. average agent degree under each condition listed in Table 8.1. Line (---) indicates simple OLS regression expected values.

Table 8.1 Mean values (over 20 trials) for final period (50) agent degree, $\langle d \rangle$ and fraction of PD plays where mutual cooperation resulted, $f(C, C)$. (see also Fig. 8.6).

	$\langle d \rangle$			$f(C,C)$		
$m \diagdown \eta$	0.80	0.90	0.95	0.80	0.90	0.95
10	0.000	0.000	0.000	0.0000	0.0000	0.0000
14	0.004	0.001	0.391	0.0001	0.0000	0.0058
18	2.441	11.859	8.587	0.0286	0.1109	0.0735
20	7.959	11.073	9.548	0.0914	0.0941	0.1185

Further, the fraction of mutual cooperative plays (of all PD plays) moved in an highly correlated way with degree (see Table 8.1 and Fig. 8.6). It would appear, therefore, that network formation in this model is due to agents who play C first, and $P[R(C)] = \#_s$.[10] A closer look at the dynamics of prevalent strategies under network forming conditions confirms this conclusion (see Fig. 8.7).

We study here an example (m, η) combination at $m = 20$, and $\eta = 0.8$ (see Fig. 8.7. Four agent types are of interest (along with the summed D-responder types): the cooperative network forming type (`C-NET`); the defection network forming type (`D-NET`); and two types which engage in an highly asymmetric relationship – the opportunist (`D-OPP`) and so-called 'sucker' (`C-SCK`) types. Again, the periodic rise and fall of strategy types is

[10] Recall, agents are free to form networks with any kind of behavioural basis.

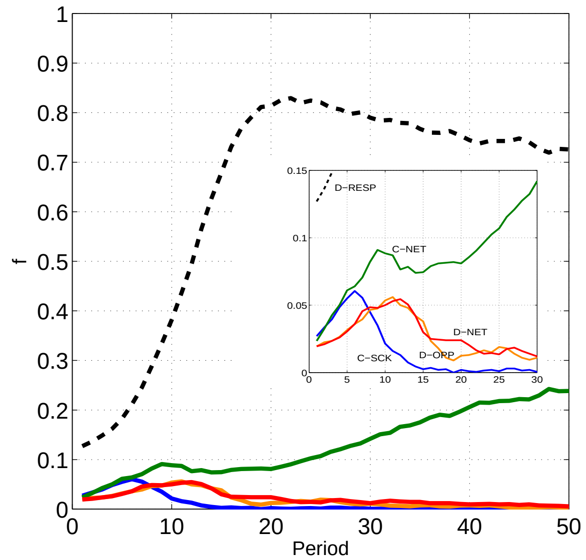

Figure 8.7 Example mean population fraction types playing various strategies (over 20 trials) under necessary network formation conditions ($m = 20, \eta = 0.8$): (C-NET) robust cooperation; (D-NET) robust defection; (D-OPP) opportunist; (C-SCK) 'sucker'; and (D-RESP) mutual defector types.

evident, but importantly, it can be seen that although C-NET, D-NET and D-OPP appear to co-exist for a time, it is only the cooperative network forming type who prevails in the long run.

To better understand these dynamics, a series of network snapshots for one representative network formation trial under the above conditions is shown in Figs. 8.8 and 8.9. Here, at least four distinct phases are discernible.

Phase 1: Amorphous connected (Figs. 8.8(a) and 8.8(b)) The existence of many sucker types leads to a super network with high average degree. In this case, almost all of the cooperative types have formed links to at least one sucker type, whilst the opportunists are largely integrated into the super network, with a range of agent types as adjacent nodes.

Phase 2: Segregated connected (Figs. 8.8(c) and 8.8(d)) The network remains super connected, but clear segregation begins to occur, such that agent-to-agent edges become highly assortative. Fewer sucker types means that opportunists become competitive for activity in the network (e.g. Fig. 8.8(d)). Cooperative and defection communities subsequently establish themselves (higher intra-community connectivity).

Phase 3: Segregated disjoint (Figs. 8.3.3 and 8.3.3) The sucker type disappears, leading to a 'shake-out' in the population – the over-supply of opportunist types is rectified, with only those who were able to integrate with the defection community able to survive. The

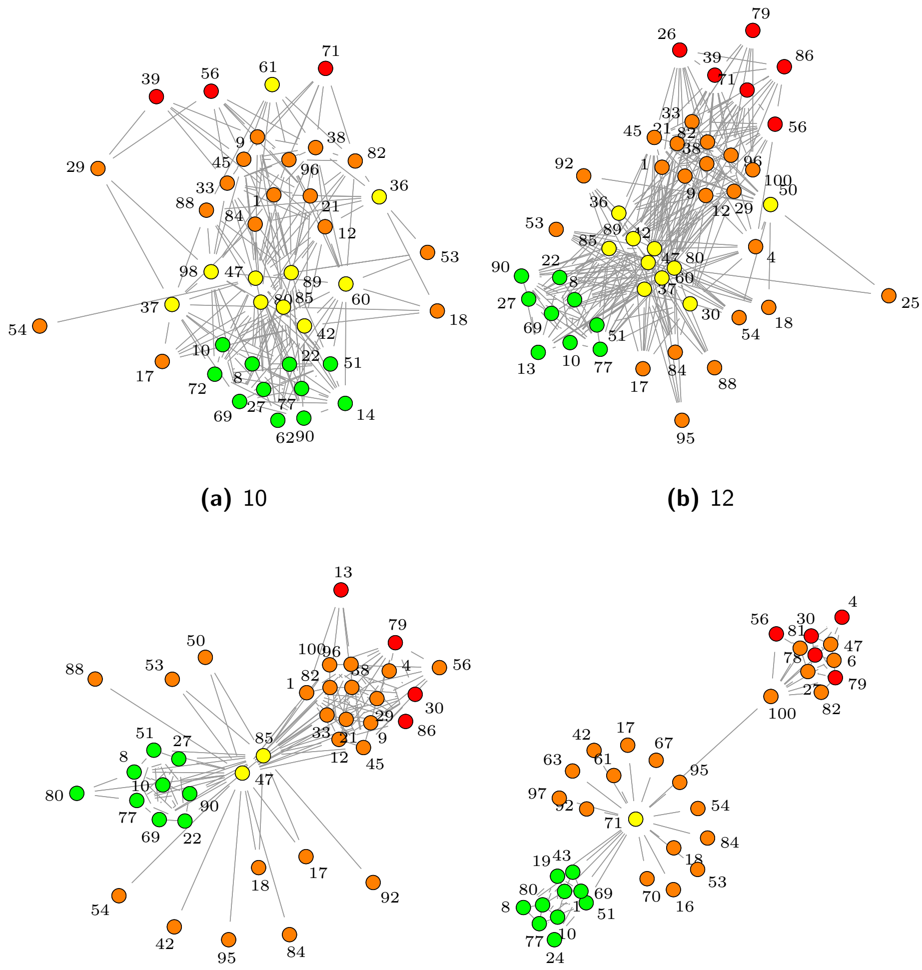

(a) 10 **(b)** 12

(c) 13 **(d)** 16

Figure 8.8 Example network dynamics ($m = 20, \eta = 0.8$): network state at end of indicated period; agent ID show next to each node; agent coloring as follows – (●) robust cooperative; (●) robust defection; (●) opportunist; and (●) 'sucker' (see text for explanation).

network is now dis-joint, with highly defined community characteristics. Further agent survival depends on raw mutual payoff characteristics.

Phase 4: Homogeneous connected (Figs. 8.3.3 and 8.3.3) With the significantly higher intra-community payoffs yielded to the cooperative community, edges here become highly dense, approaching a complete component graph. The defective community disappears, with no possibility of infiltration into the cooperative community (see discussion below). New agents of cooperative network forming type are able to join and be integrated. Some sucker–opportunist relationships arise on margins but are short lived only.

Network Agent Types

A dissection of the prominent strategies that arose in the above experiment was conducted on period 13 (Fig. 8.8(c)). A comparison of the network itself with the agent autopsies

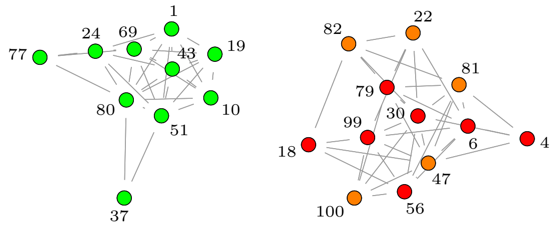

(a) 17

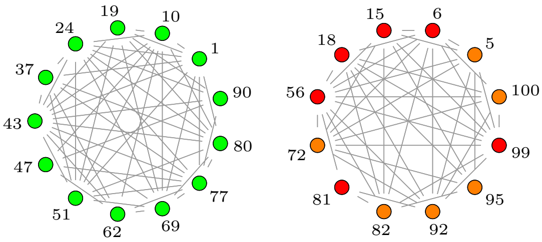

(b) 18

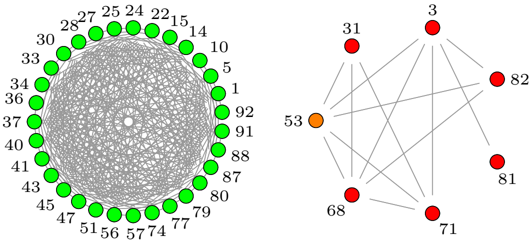

(c) 28

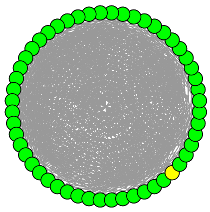

(d) 35

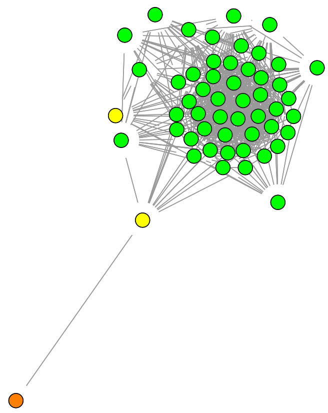

(e) 40

Figure 8.9 Example network dynamics (cont.): network state at end of indicated period (see Fig. 8.8 for explanation). *NB: Agent ID omitted for periods 35 and 40.*

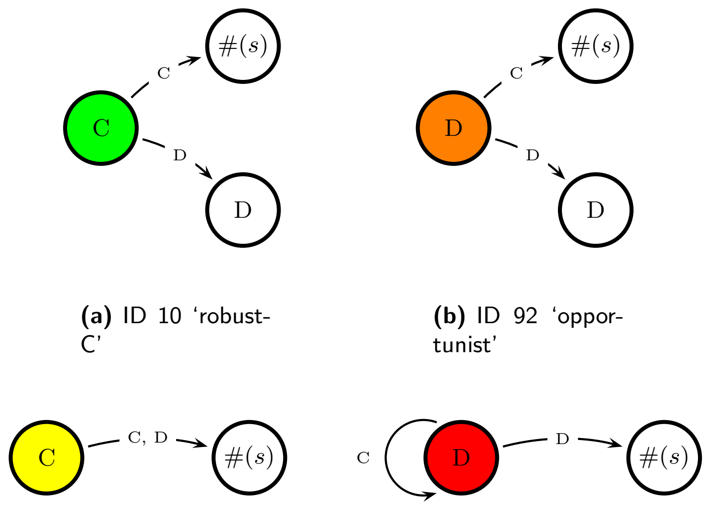

(a) ID 10 'robust-C'

(b) ID 92 'opportunist'

(c) ID 85 'sucker'

(d) ID 79 'robust-D'

Figure 8.10 Example types visible in the network of period 13 above (see Fig. 8.8).

given in Fig. 8.10 makes clear the difference between each agent's activity in the network. Clearly, the interaction of the opportunist and sucker types (Fig. 8.10(b) and 8.10(c) respectively) will lead to tie-strengthening conditions, but with highly asymmetric payoff outcomes.

Importantly, however, the 'robust-C' type (agent 10 in period 13) is immune to this play by responding with D to the opportunist's D opening; a transition that works equally well for agent 10 when facing the robust-D type (agent 79 in period 13). For this reason, as can be seen in the agent networks presented so far, the cooperative types avoid tie-strengthening with either the robust-D or the opportunist types, which in both cases ensures adequate type-selection, but in the latter case, protects the cooperative network forming types from the opportunist shake-out that was inevitable with the decline of the sucker types in periods 13 to 17.

At the statistical level, these interactions are borne out in the periodic struggle of the initial network dynamics (see Fig. 8.7). The initial rise of the sucker types (establishing network ties to any other tie-strengthening agent) provides fertile ground for the opportunist types, who in turn, support the defection network types. However, over time, as each loses its respective 'feed-stock', network dynamics resolve in favour of the cooperative network forming types.

It is important to note that within this boundedly-rational framework, robust network formation is highly dependent on 'purity' of network structures. As can be seen in

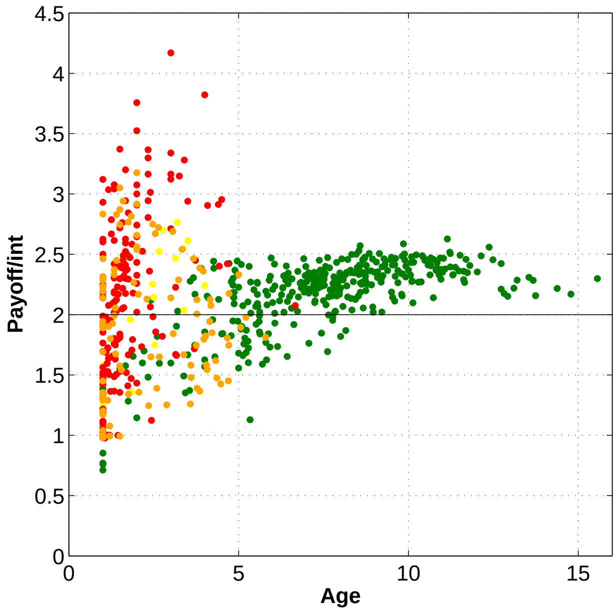

Figure 8.11 Scatter plot, mean Payoff per interaction vs. mean Age for each point which represents a single connected component. Coloring indicates dominant ($> 50\%$) type in each component: (●) 'suckers'; (●) 'opportunists'; (●) defection network builders; and (●) cooperation network builders. NB: line at interaction payoff 2 indicates expected payoff for non-connected D-types. (Other parameters as for Fig. 8.7).

Fig. 8.11, connected components that experience longevity must be able to attain more than the going 'outside' payoff rate of 2 per interaction.[11]

As can be identified, connected components that have a high proportion of sucker or opportunist types will yield large mean payoffs, but are very short-lived (rarely having mean agent ages greater than 5 periods) due to the volatile nature of payoff asymmetries. On the other hand, the cooperative networks who can overcome the short-term heterogeneous phase are very likely to retain higher than 2 average payoffs and so be positively selected for in the end-of-period strategy revision phase. Clearly, ensuring good 'discipline' within a cooperation network must be an high priority for the sustainability for the agents therein.

Interestingly, it appears from the data presented, that although predominantly defection type networks can yield very high payoffs, they will also suffer a type differentiation problem, mixing easily with the opportunist types. In the early stages of population dynamics, this is a positive attribute since it will provide these types with high period payoffs through greater 'activity' (more plays of the IPD), ensuring their individual survival. However, over time, with the propensity for opportunist types to lose valuable

[11] The payoff yield between two ALL-D types (for example) who play a two-iteration IPD game, gaining 1 in each iteration.

payoff opportunities with sucker types, the defection networks yield strictly worse average payoffs than the 'outside' defection population, since they are necessarily sacrificing a unit of payoff every time they re-affirm/establish a link with a fellow defection network type.

8.3.4 Multiple Equilibria & the Long Run

In the previous section, conditions were identified in which stable networks were formed under parsimonious agent specification ($\tau = 2$ implying $k = 3$) to enable correlation with established results in the analytic literature. Here, this constraint is relaxed and instead agents interactions of up to four iterations of the IPD game ($\tau = 4$) are considered and their long-run dynamics studied. Recall, by increasing the length of the IPD game, the maximal FSA state count increases markedly: for $\tau = \{3, 4\}$ maximum state count $k = \{7, 15\}$.

Previous conditions were retained, with $\eta = 0.8$ and $m = 20$, and each trial allowed to run for 1000 periods. Since a full description of the state is not feasible[12] we consider an aggregate description of two fundamental state characteristics, $f(C, C)$ – the fraction of plays in a period where mutual cooperation is observed (strategic behaviour); and $\langle d \rangle$ – mean agent degree (network formation). Results are presented for five long-run trials in Fig. 8.12. Under low interaction length the system moves within 100 steps to one of two stable equilibria – either a stable cooperation network is formed (as was studied in the previous section) or no network arises and a stable defection population sets in. However, as the interaction length increases (and so the associated complexity of behaviour that each agent can display), the dynamics become increasingly erratic, with multiple, apparently stable, equilibria visible in each case, but transient *transitions* between these equilibria observed. This situation is synonymous with that of *complex* system dynamics.

[12] Consider that each time period, a population constitutes $n \times |s|$ bits, where $|s|$ is the length of a string needed to represent each agent's strategy, and the network $n(n-1)/2$ bits; taken together, gives rise to a possible $2^{n(n-1)/2+n|s|}$ states, which for $\tau = 2$ is $2^{9\times10^6}$! (It is possible to reduce this number by conducting automata autopsies, but the problem remains.)

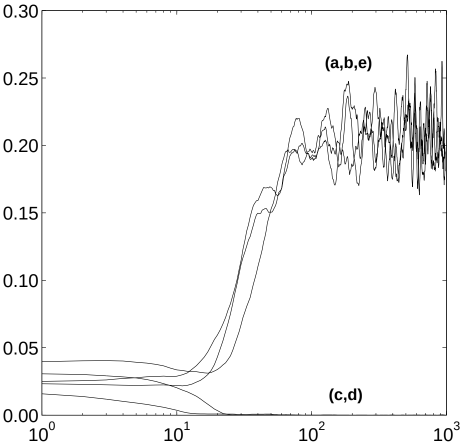

(a) $\tau = 2, \langle d \rangle$

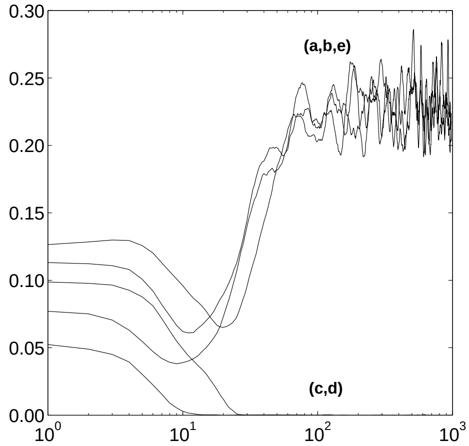

(b) $\tau = 2, f(C, C)$

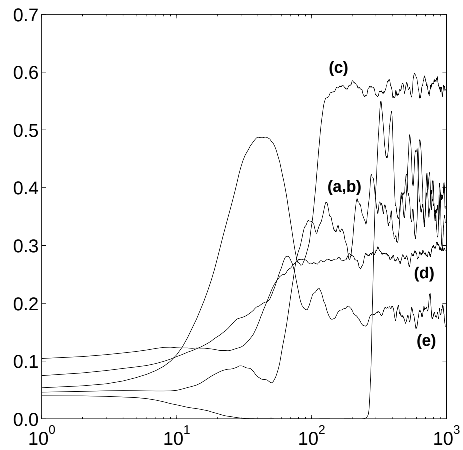

(c) $\tau = 3, \langle d \rangle$

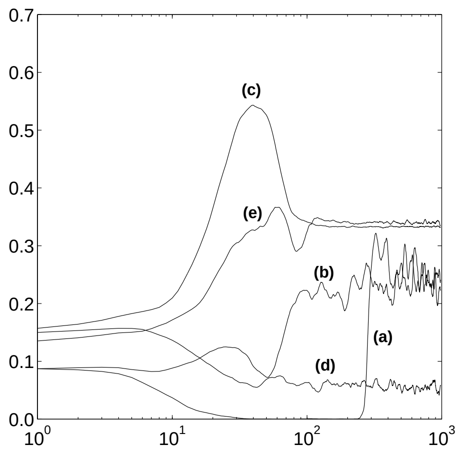

(d) $\tau = 3, f(C, C)$

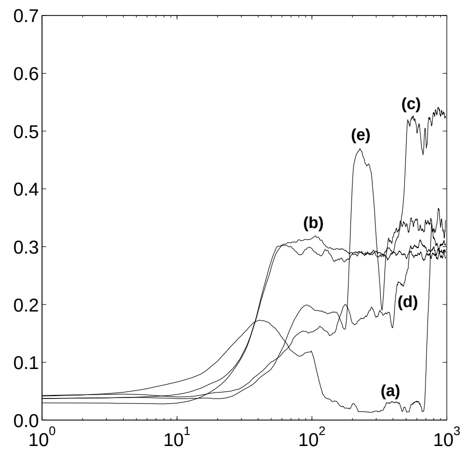

(e) $\tau = 4, \langle d \rangle$

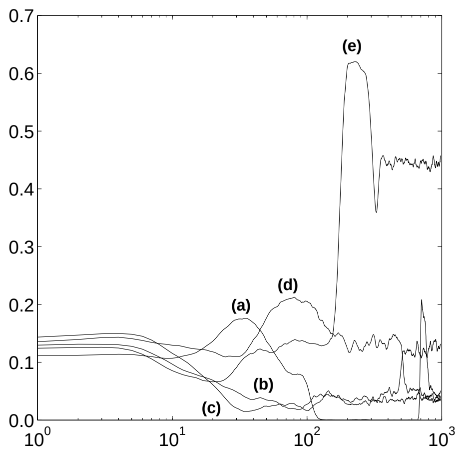

(f) $\tau = 4, f(C, C)$

Figure 8.12 Long-run system dynamics under different maximum interaction lengths indicating increasing complexity; five trials shown at each value of τ (data smoothed over 20 steps).

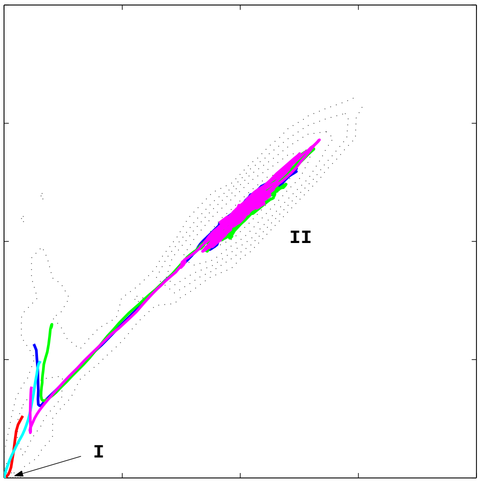

(a) $\tau = 2$

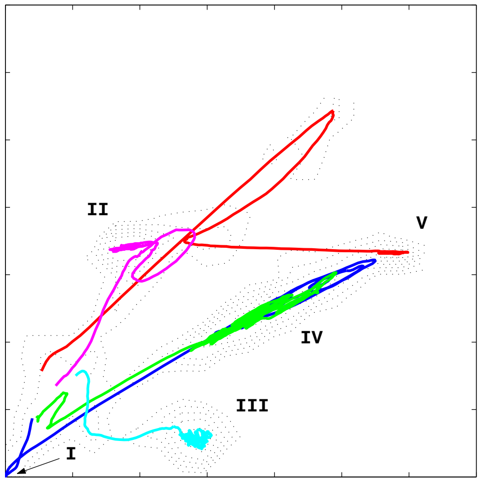

(b) $\tau = 3$

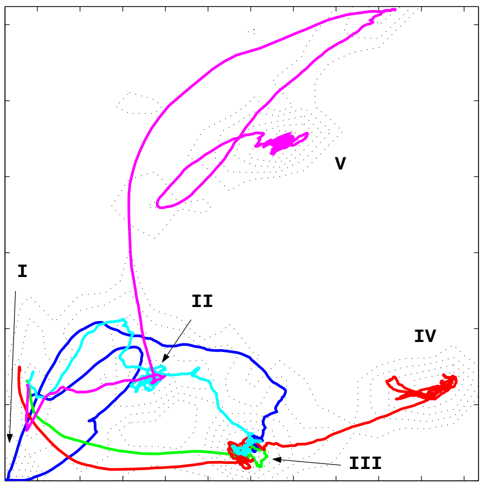

(c) $\tau = 4$

Figure 8.13 State transitions, $f(C, C) - \langle d \rangle$ space. Contours indicate ($\log_{10}$) density of time periods spent at $(f, \langle d \rangle)$ points over all trials; lines represent smoothed trial trajectories (20 steps), coloring shows each trial.

To better see this transition, the locations of the system in $f(C, C) - \langle d \rangle$ state-space were plotted (see Fig. 8.13). Here the transition from relatively well-defined attractors for $\tau = 2$ to complex dynamics at $\tau > 2$ is clear. Indeed, five stationary locations are visible in Fig. 8.13(b) with location I, II and V appearing to be transiently stable, with state trajectories both entering *and* leaving these locations, whilst locations III and IV appear to be absorbing for the system. Interestingly, these absorbing locations give rise to relatively similar average network formation, but different levels of cooperation, being low and moderate respectively.

Similarly, but with greater clarity, the dynamics of $\tau = 4$ shows very erratic behaviour (Fig. 8.13(c)), appearing to have only two absorbing locations, IV and V, whilst each of I, II, and III appear to be transient. In this case, the absorbing locations are very different in character, being an almost complete graph, but similarly defection-based in the first case, or again, with high participation, but markedly cooperative in the second. The former equilibrium suggests an unusual outcome for network formation. In this case, it is not that agents form a network as a type-guarantee (as in say Taylor (2000)) or as protection against unwanted interactions (so Kali (1999)), but rather, the network is being exploited for the extra *activity* that it generates. The message here is that if you want to survive in the heavily defection-based environment, you have to be part of the 'network set'. This has resonance with Akerlof's 1976 reflection on how *indicators* (measures of performance) can give rise to a Rat Race. Here, the private returns for additional interactions via the network outweighs the benefits to cooperation based play.

Surprisingly, such complex dynamics arise in a relatively simple model of network formation. Recall, that the longest that any of the agent interactions can be in these studies was just two, three or four iterations of the modified Prisoner's Dilemma given in (8.4). To be very sure that such dynamics are not a consequence of the encoding of the automata themselves, an identical study was run with $\tau = 4$, but setting $\eta = 0$ such that all interactions would continue to be of uniform probabilities. However, in all cases, the system moved to a zero cooperation regime within the first 100 periods and remained there. Clearly then, we conclude that endogeneity of network formation is driving such complex dynamics as observed above.

8.3.5 Network Formation & Self-Organized Criticality

Next, given that the system displays complex dynamics for given values of τ and that network endogeneity is critical to such dynamics, it is natural to study the dynamics of network formation itself. For these purposes, the size (node count) of the principle (largest) network component that exists at the end of each period is studied. Example time-series for one $\tau = 4$ run are given in Fig. 8.14. In the first figure, the size of the network itself is shown, whilst in the second, the first differences are given (i.e. $S_t - S_{t-1}$). It can be seen from this example, that changes in network size occur both on many time-scales and to various degrees. Such phenomena is synonymous with systems exhibiting critical behaviour (Bak et al., 1988); perturbations to the system cause mostly small, damped outcomes, but can occasionally have dramatic effects, likened to a 'domino-effect' (refer to §3.5 in Chapter 3 for a lengthier treatment of self-organized criticality).

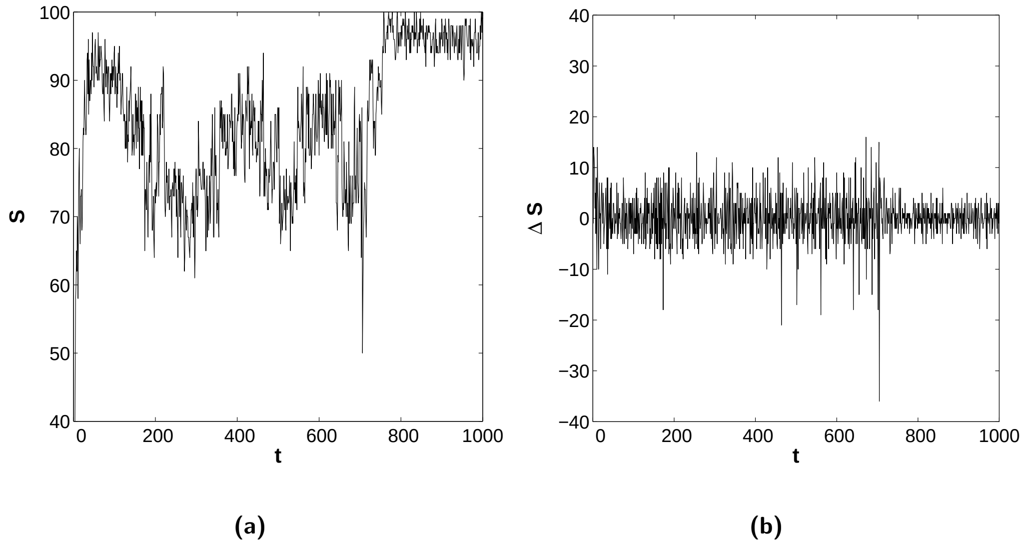

Figure 8.14 Example time series of principle component size (a) and change in size (b) for $\tau = 4$.

To investigate this feature, frequency distributions of average network fluctuation sizes $D(\Delta S)$ were prepared for each interaction length. As can be seen in Fig. 8.15 the distributions appear to follow a power-law behaviour, that is of the form,

$$D(\Delta S) \sim \Delta S^{-\alpha} \; . \tag{8.8}$$

Such a relationship is often termed 'scale-free' since it indicates that the same overall systemic dynamics are operating on all spatial scales; small deviations build up over time and lead to large deviations in the long-run due to connectivity within the system.

Spatial self-similarity is one feature of critical systems, the second is that similar power-law scaling is observed in the temporal domain as well; normally manifesting as so-called '$1/f$' noise, which appears ubiquitous in nature.[13] A power spectrum was therefore prepared of the time-series network size to study this possibility.[14] Fig 8.16 gives the outcome of this analysis, showing clear power-law scaling behaviour. Linear fits were prepared for the first 10 data points[15] with good agreement in all cases. Exponents of

[13]Examples from the introduction to (Bak et al., 1988) include: light from quasars, the intensity of sunspots, the current through resistors, the flow of sand in an hour glass, the flow of the Nile river, and stock exchange price indexes.

[14]Suppose $s(t)$ is the (discrete) times-series of some network size data (as per Fig. 8.14(a)), then using MATLAB a Fast-Fourier-Transform, $F(s)$ was performed with $N = 2^8$ points, followed by the standard power-function, FF'/N, where F' is the complex conjugate of F. Figures show the resultant power spectra without the first constant-shift term $f(0)$, and are cut below $S(f) \leq 5 \times 10^{-4}$.

[15]Fitting power-law models has received some interesting study in recent times due to difficulties in forming goodness-of-fit tests etc. Here we follow Goldstein et al. (2004) in form, fitting the linear specification to only a selection of the primary points, thus avoiding undue bias in the tails (which represent a very small mass of the spectrum).

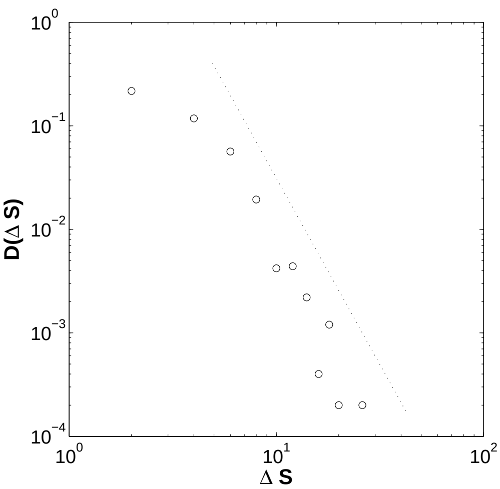

(a) $\tau = 2$

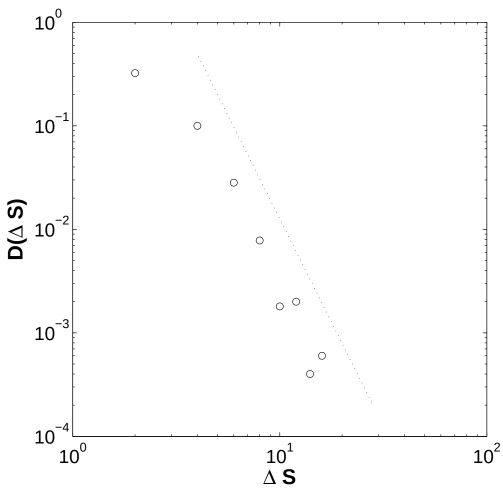

(b) $\tau = 3$

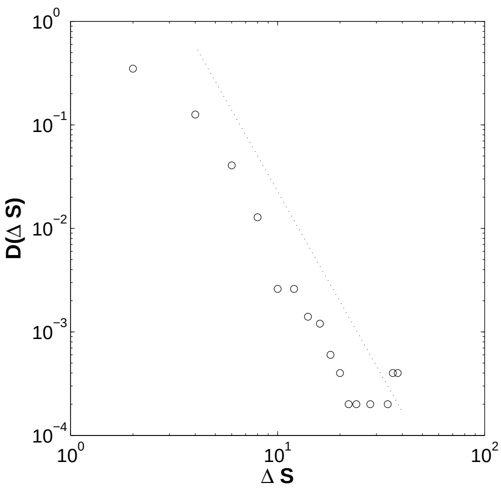

(c) $\tau = 4$

Figure 8.15 Mean principle component network size distribution for each interaction length. Lines given as guide only.

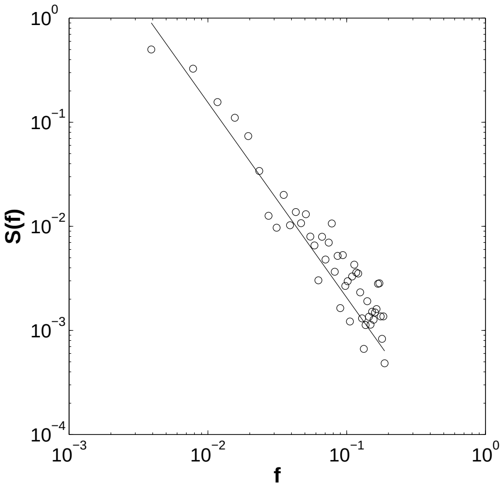

(a) $\tau = 2$

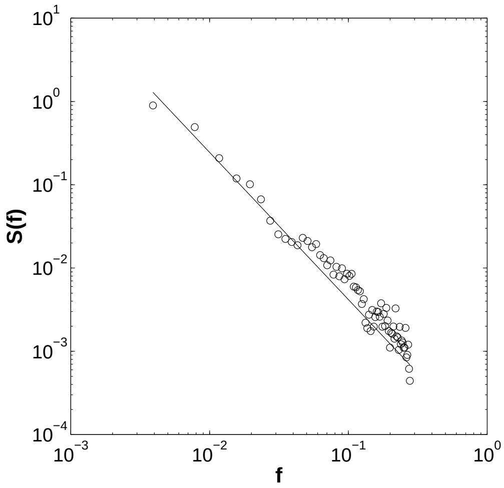

(b) $\tau = 3$

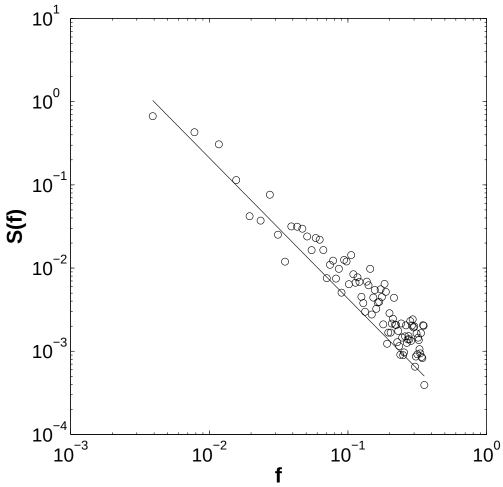

(c) $\tau = 4$

Figure 8.16 Mean power spectra of principle component network size for each interaction length. Lines represent power-law fits to first 10 points with $\alpha = -1.9, -1.8, -1.7$ for $\tau = 2, 3, 4$ respectively.

the relationship,

$$S(f) \sim f^{-\alpha} \ , \tag{8.9}$$

were found all found to be -1.8 ± 0.1.

Taken together, the spatial and temporal fingerprints of criticality observed in the network formation dynamics, indicate that the system is indeed very capable of the kind of complex dynamics observed and discussed above, and that the network formation appears to be a key factor in such behaviour. Furthermore, as has been proposed by various authors, rather than such criticality arising from fine tuning of system parameters such as occurs in designed critical industrial systems (e.g. nuclear fission reactors), the system appears to naturally move towards this critical state, and keep returning to it over time. It is for this property that authors such as Bak et al. (1988) have termed such phenomena 'self-organized criticality'. Indeed, it appears that such phenomena is a strong indicator of complex dynamics, and may indeed be the necessary system state to give rise to the kind of non-equilibrium processes observed in various dissipative systems.[16]

The existence of such dynamics in economic systems has recently received growing interest (see for example, Krugman (1996)). Indeed, power-law behaviour on both a macro (Canning et al., 1998; Devezas and Modelski, 2003) and micro- interactions scale (Arenas et al., 2000; Scheinkman and Woodford, 1994) has been incorporated into both model and empirical evidence, and some assert is fundamental to our understanding and thus modelling of economic systems (Arthur, 1994). This point will be returned to in the following chapter.

8.4 Conclusions

In contrast to previous attempts at capturing the dynamics of strategic network formation, the present model provides a relatively simple foundation, but powerfully rich behavioural and topological environment within which to study the dynamics of strategic network formation. Moreover, in contrast to previous dynamic and strategic network models, by incorporating the network formation decision-process into individual agent strategies, a rich ecology of agent types and consequent network topologies was observed. Significantly, this model suggests that the network formation process must deliver relatively symmetric payoffs to network members. If this is not true, networks formed will likely be heterogeneous in nature, with disruptive edge formation and breaking sequences which can effectively destroy any benefits that the network might have conferred on members (e.g. the opportunist-sucker network volatility of phases I and II mentioned above).

Analytical and subsequent computational components of the present paper indicate that in this simple modified IPD set-up, cooperation is not sustainable without the additional benefits conferred by the type-selection and type-protection network externalities. Specifically, agents require at least some level of repetition of interaction within the current population to gain sufficient incentives to form the network; and secondly, the 'impact' of the network on the interaction space was found to be a necessary condition for cooperative network formation, with network emergence only observed when the network allows for

[16]See for example, Langton (1992) for a discussion on this point.

relatively high (though not complete) discrimination (probabilistically) from the wider population.

Furthermore, the dynamical properties of the present model have been investigated and indicate that even with parsimonious descriptions of boundedly-rational agent strategies, complex dynamics are observed, with multiple and transient stationary locations a feature of the state space. These dynamics increased in complexity with increasing interaction length. Additionally, evidence on the fluctuations in the size of networks over time indicates that the network formation and decay processes themselves are likely the main driving force behind the complex system dynamics, with both spatial and temporal scaling behaviour indicating the existence of so-called 'self-organized criticality'. To this author's knowledge, this is the first strategic network formation model to produce and study such complex dynamics. Such observations clearly raise tantalising avenues for future work; I shall raise a selection in finishing: do realistic cooperative networks display complex dynamics? if not, what mechanism of agency overcomes such instability? if networks can be shown to have such dynamics (admittedly these data are still largely out of reach) what are the implications for supporting cooperative institutions? and finally, given an autonomous, locally interacting world, how should the social planner intervene in such networks to pursue welfare maximizing aims?

Appendix 1: Proof of Lemma 1

Lemma 1 *For a population playing the IPD as given in $\mathcal{G}$ under uniform interaction probabilities and maximum FSA state length $\tau = 2$, the strategy triplet $s_D : \{D, D, D\}$ is the only evolutionary stable strategy (ESS).*

Proof To begin with, consider a population consisting of only two types of agents, where one type is s_D and the other the strategy triplet $s_S : \{P_1^S, R^S(C), R^S(D)\}$ where P_1^S is the agent's play in state one and $R^S(x)$ represents the agent's play in response to opponent's play x in the preceeding iteration. For convenience, call these types **D** and **S** respectively.

Define the total payoff to an agent X undergoing an interaction with agent Y to be $\Pi(X|Y)$ and note that,

$$\begin{aligned}
\Pi(\mathbf{D}|\mathbf{S}) &= \mathcal{G}(D|P_1^S) + \mathcal{G}(D|R^S(D)) && (8.10)\\
\Pi(\mathbf{D}|\mathbf{D}) &= 2 \times \mathcal{G}(D|D) && (8.11)\\
\Pi(\mathbf{S}|\mathbf{D}) &= \mathcal{G}(P_1^S|D) + \mathcal{G}(R^S(D)|D) && (8.12)\\
\Pi(\mathbf{S}|\mathbf{S}) &= \mathcal{G}(P_1^S|P_1^S) + \mathcal{G}(R^S(P_1^S)|R^S(P_1^S)) && (8.13)
\end{aligned}$$

Further, let α be the proportion of type **S** in the population. Then, the expected interaction payoffs for each agent type with uniform mixing is given by,

$$\begin{aligned}
E[\Pi(\mathbf{S})] &= \alpha\Pi(\mathbf{S}|\mathbf{S}) + (1-\alpha)\Pi(\mathbf{S}|\mathbf{D}) && (8.14)\\
E[\Pi(\mathbf{D})] &= (1-\alpha)\Pi(\mathbf{D}|\mathbf{D}) + \alpha\Pi(\mathbf{D}|\mathbf{S})\ . && (8.15)
\end{aligned}$$

Now suppose $\alpha \to 1$, if **S** is to be stable in the presence of **D** then it follows from (8.14) and (8.15) that,

$$\Pi(\mathbf{S}|\mathbf{S}) \geq \Pi(\mathbf{D}|\mathbf{S}), \qquad (8.16)$$

which by (8.10) to (8.13) becomes,

$$\mathcal{G}(P_1^S|P_1^S) + \mathcal{G}(R^S(P_1^S)|R^S(P_1^S)) \geq \mathcal{G}(D|P_1^S) + \mathcal{G}(D|R^S(D))\ . \qquad (8.17)$$

Now suppose that $P_1^S = C$ then by payoffs given in (8.4), (8.17) becomes,

$$\mathcal{G}(R^S(C)|R^S(C)) \geq \mathcal{G}(D|R^S(D)) + 2\ , \qquad (8.18)$$

which implies that $R^S(C)$ must be C and that $R^S(D) \neq C$. Hence, there are only two candidates for **S**, namely $s_{CD} : \{C, C, D\}$ and $s_{C\#} : \{C, C, \#\}$. However, both s_{CD} and $s_{C\#}$ are not stable in the presence of the mimic agent $s_M : \{C, D, D\}$ which itself does not satisfy the condition given in (8.17).

Suppose on the other hand that $P_1^S = D$. This would imply (by substitution of payoffs into (8.17)) that,

$$\mathcal{G}(R^S(D)|R^S(D)) \geq \mathcal{G}(D|R^S(D))$$

which has no solution unless $R^S(D) = \{D, \#\}$.

Now, consider a population where $\alpha \to 0$ and suppose there exists some strategy $\mathbf{S'}$ such that $E[\Pi(\mathbf{S'})] \geq E[\Pi(\mathbf{D})]$ which again yields,

$$\mathcal{G}(P_1^S|D) + \mathcal{G}(R^S(D)|D) \geq 2\ . \tag{8.19}$$

If $P_1^S = C$ (8.19) becomes,

$$\mathcal{G}(R^S(D)|D) \geq 2$$

which has no solution. Likewise, if $P_1^S = D$ we have,

$$\mathcal{G}(R^S(D)|D) \geq 1$$

which has a solution only if $R^S(D) = D$, completing the proof. $\square$

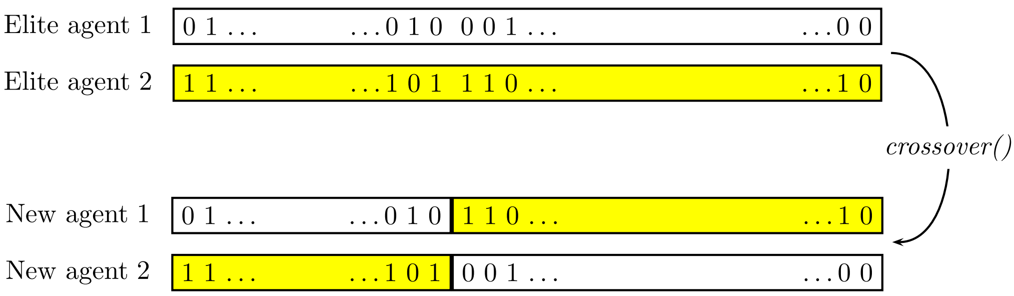

Figure 8.17 Schematic of the crossover operator.

Appendix 2: The Genetic Algorithm

Crossover Operator

In this model, the crossover operator represents learning by imitation. It is assumed that new agents enter the population after having witnessed the behaviour of the incumbent elites, and so enter having copied part or all of these successful strategies.

The crossover operator (Fig. 8.17) is extremely simple to apply: two agent bit-strings from the elite population are selected at random; at some uniform random point along each string, a 'cut' is made[17], and the two trailing pieces are switched, thus forming two 'new' agents (assuming there is at least one difference in these two pieces) from the strategies of the 'old' agents.

It is important to note that the 'cut-point' is selected at random along the string length. In this way, after successive 'generations' of play, successful portions of each strategy will be replicated and amplified in the new population. Such accumulation of strategic insight over time speaks directly to real-life technology adoption or behavioural learning. However, attention is drawn to the fact that in this framework, the new population's technology accumulation content is contingent entirely on the previous period's population only. There is little facility here (except consistency between periods, or chance) for the new population to replicate exactly portions of successful strategies from several or more periods previous to the current one.[18]

[17] This is the so-called 'one-point' crossover operator, other versions include 'two-point', 'three-point' and so on.

[18] One may also care to notice that although the cut-point is uniform random, the probability that a given bit will be 'switched' in the crossover operation is not. Infact, for some bit at position a, with string length L, the probability that a will undergo crossover is $\frac{a}{L}$. Thus, in this formulation, bits to the left (start) of the string are least likely to undergo crossover, whilst those to the right (end) of the string are most likely. Hence, the initial strategy state is more likely to be retained between elite and new agent, with subsequent states more likely to undergo switching.

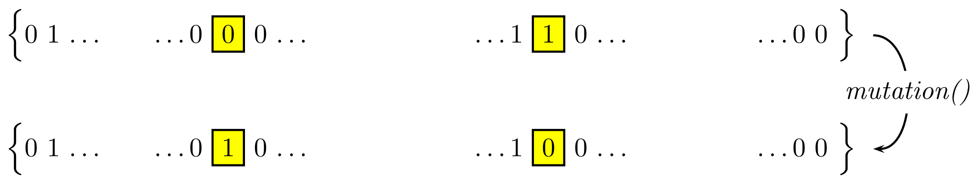

Figure 8.18 Schematic of the mutation operator.

Mutuation Operator

Following application of the crossover operator, each new agent bit string undergoes a mutation operator. Simply put, with low (independent) probability (1 or 5%) each bit in the string is 'flipped': $0 \rightarrow 1$; or $1 \rightarrow 0$ (see Fig. 8.18). For example, at the 1% and 5% levels, a six state strategy (66 bit) would undergo approximately < 1 and 3 bit 'flips' respectively.

Such flipping is akin to the 'mistake making', or 'trembling hand' interpretation elsewhere in the strategic literature. In the current context, it represents either imperfect imitation/learning from successful incumbents, or self-driven innovation by entrants. Regardless, its random strategy generation, taken together with the non-random selection mechanism, provide an opportunity for the population to evolve in a positive manner. Successful components of new strategies are rewarded with more followers in future populations, whilst unsuccessful agents leave, making way for fresh entrants.

Chapter 9

Summary of the Thesis & Directions for Further Work

This thesis has primarily investigated the formation of economic networks. In particular, the two contexts of *communication* and *cooperation* network formation were studied. In this chapter the key findings of this thesis are summarised and placed within the wider body of research. Especially, emphasis is placed on recent related results which indicate further interesting lines of inquiry for the present work.

Chapters 4 to 6 analysed the non-cooperative communication network formation model of Bala and Goyal (BG). After its initial introduction in Chapter 4, several new results were obtained through analysis and numerical simulation in Chapter 5. Primarily, these results suggest that the Best Response updating rule, combined with conditions approximating no strategic inertia can produce extremely time-consuming two-period graph cycles. Indeed, although a result is obtained for $n = 3$ that the Strict Nash graph under one-way flows is still realised (the circle), a counter example is given for $n = 4$ showing that an infinite cycle is possible with no inertia. In some ways, this result is not surprising since similar two-person games can produce such cycles. However, numerical simulations indicate that for conditions *approximating* no inertia, cycles take an enormous toll on convergence times and under exact no inertia conditions cycles are overwhelmingly the predominant case for moderate sizes of n.

These theoretical results, together with the experimental results of Falk and Kosfeld (2003) motivate the development, calibration and analysis of an enriched *artificial adaptive agent* model in Chapter 6. This work incorporated a novel cognition architecture for the agents which enabled the modeling of a biased cultural learning process (learning from observed neighbour plays) between stages. It was found that this process, together with

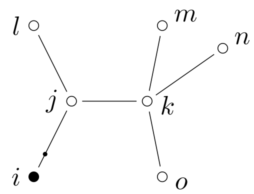

Figure 9.1 In the modified experimental set-up of Berninghaus et al. (2006), although agent i sponsors a two-way link to agent j, they would only receive information from the set $\{i, j, k, l\}$; not $\{i, j, k, l, m, n, o\}$ as in the BG set-up.

an appropriate joint-objective measure based on a ratio of benefits and costs together with an altruism score could produce many of the experimental outcomes with fidelity. Moreover, the suggested *inequality aversion* hypothesis of Falk and Kosfeld to explain the poor performance of human subjects in the two-way flow case was found to be unnecessary as a modeling input, but rather arose as an *emergent* property over learning stages.

Very recently, the BG model has received further experimental investigation. Berninghaus et al. (2006) implemented a human trial of a slightly modified version of the BG model, motivated in particular by the lack of Strict Nash play under two-way flows found in Falk and Kosfeld (2003).[1] Specifically, they modified the BG set-up in two ways: first, rather than assuming information flows along any finite connected path of players, they distinguish between *actively*, *passively* and *indirectly* connected neighbours. Referring to Fig. 9.1, j would be the only active neighbour of agent i, whilst agent i would be a passive neighbour of agent j (j obtains the information of i as an outcome of i's action). Whilst agents l and k would be indirect neighbours of i since they represent 'all actively or passively reached neighbours of all actively reached neighbours of i' (p.240). This leaves the other agents, m, n, o in this example graph not part of i's neighbour set, which is a significant departure from the BG set-up where any finitely connected node to i in the two-way information case would be in i's neighbour set. Essentially, the Berninghaus et al. formulation restricts neighbours to being no more than two hops away. Naturally, this changes the game and so taking this into account, the authors show elsewhere that the Strict Nash network under the new set-up are either empty, or the periphery-sponsored star (ps-star).

Second, and less significantly, Berninghaus et al. set the updating of the agent strategies to being *continuous*, rather than periodic and so simultaneous as in BG. This is of less significance, for the theory of BG, since as mentioned above, and explored in Chapter 5, in the BG proof work, by assuming at least one agent experiences strategic inertia

[1] Indeed, an earlier titling of the Berninghaus et al. work, 'Searching for "Stars": Recent Experimental Results on Network Formation' (Berninghaus et al., 2004) is telling.

(they don't update) each period, the case where only one agent *updates* in a period (so the 'continuous' case) is handled.

Nevertheless, the main findings of their study: that around 30% of the time was spent in a *ps-star* configuration suggest a significant difference to the results of FK where the similarly structured centre-sponsored star (cs-star) was only observed twice in the entire study. In particular, since the ps-star is just as asymmetric as the cs-star, the authors suggest that asymmetry aversion is not a complete story. Further, given that the central agent in the ps-star configuration accumulates a significant 'bonus' over time from actions of the other players, Berninghaus et al. report that indeed they observed a 'sharing' of the central role. For the authors, this suggests that inequity aversion in payoffs is a motivating behaviour for subjects.

Of course, whilst this recent study provides some very interesting further evidence on network formation, the Berninghaus et al. set-up is a different model, and should not be seen in direct comparison to the BG results, nor the FK trials. Nevertheless, the results clearly indicate further avenues of research for the work presented in Chapter 6. For instance, both the neighbourhood and timing modifications could be relatively easily dealt with in the present set-up. The former would require a simple thresholding to the path-length finding procedures of the model, whilst the latter is traditionally handled with a *Poisson* selection process – agent's 'numbers' coming up for updating resulting from the sampling procedure. In fact, although these would bring in more parameters (p in the Poisson distribution and the neighbourhood path-length threshold) they would provide access to very interesting further questions, such as: does the updating probability affect the propensity for Strict Nash play in a smooth manner, or are phase-changes of behaviour apparent? Or, perhaps with larger n, do agents exhibit poor strategic play as their focus becomes more 'local', and hence inefficient (i.e. as the path-length threshold for neighbours goes to zero)? One expects that further work, both theoretical and experimental will continue to provide motivation for the kind of computational inquiry presented in Chapter 6.

Following on from this work on *communication* networks, Chapters 7 and 8 then considered an endogenous network formation model of *cooperation*. The first novel contribution of this work was to model a truly *strategic* form of network formation. This was achieved by implementing Finite State Automata (FSA) agents playing a modified Iterated Prisoner's Dilemma (IPD) game that allowed network formation play as part of the action-set. This approach stands in contrast to previous theoretical models with uniform interactions such as Kandori et al. (1993), or computational contributions on static networks such as Masuda and Aihara (2003), or of the few computational network formation approaches previously reported such as Smucker et al. (1994) or Ashlock et al. (1996). Although many models speak helpfully to the context of network formation, none allows agents *as part of their actual playing strategy* to change the interaction environment. Where the interaction environment does change, it is the result of an externally set tolerance or threshold process.

Secondly, Chapter 8 presents results on the minimum conditions for network formation by utilising a novel tuning parameter. It is found that cooperation is possible due to network formation only where two criteria are met: first, that the network institution has sufficient impact on the interaction probabilities of agents; and second, that there is sufficient basic activity in the population to afford enough opportunities for agents to exploit their costly network linkages. Furthermore, and of most interest, it is found that for relatively small agent strategies, the standard results of long-run equilibrium attainment in cooperation models are not upheld. Instead, a-periodic, complex dynamics in the two-dimensional aggregate measure space of mutual cooperation fraction and average agent degree persist. In accounting for this result, further analysis of the dynamic network/interaction environment suggests the existence of a self-organized critical system.

To place this result, it is worth noting that overwhelmingly, the previous work on the nature of cooperation using computational/numerical-simulation analysis such as that of Axelrod (1981); Axelrod and Hamilton (1981); Maynard Smith and Price (1973); Riolo et al. (2001) with uniform interactions find equilibrium, or simple switching between cooperation and defection regimes in the long-run. The same is apparently true of models which analyse cooperation on non-uniform interaction but static spaces such as Masuda and Aihara (2003) or Cohen et al. (2001), or even models with *random* mobility of agents such as Vainstein et al. (2006). However, in contrast, of the few contributions that do report complex long-run dynamics such as the recent biological model of *intentioned* agent movement of Burtsev and Turchin (2006), or the endogenous local/global interaction model of Choi (2002), or the preferential partner selection IPD model using thresholds as mentioned above of Ashlock et al. (1996), it is the *changing* interaction environment that appears to be the common factor.

Of course, there are exceptions to this rule, for example Nowak and May's (1992) classic (static) spatial chaos model of the Prisoners' Dilemma with evolutionary updating. But even here the long run *average* cooperation fraction is steady; it is the *spatial* distribution of cooperators and defectors which varies chaotically with such remarkable fractal patterns. Similarly, Lindgren's (1992) model of variable memory automata playing a repeated Prisoners' Dilemma style game generated hallmark dynamic complexity in population fractions (see Fig. 9.2) in the long-run (for more than 80,000 generations).

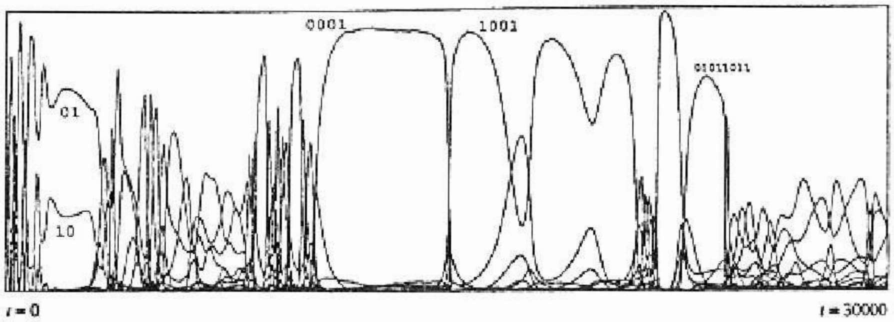

Figure 9.2 Complex dynamics exhibited by population fractions over time in Lindgren's (1992) variable memory, Cellular Automata PD model.

However, in this model, it is to be noted that only around 10% of cases generated the 'interesting' dynamics, with the other 90% getting 'stuck' in an evolutionary stable stasis.

Clearly, it will be some time before a consensus is reached about what might be the ultimate cause of such interesting dynamics as was found in Chapter 8. However, the brief discussion given above suggests that such dynamics accompany situations where the interaction environment is endogenously updated. Certainly, the results presented in Chapter 8 add further weight to this hypothesis, although other explanations would need to be excluded before firm conclusions are drawn. For instance, one would need to run further experiments with different updating methodologies, or encodings. However, given that the complex dynamics disappeared for the uniform interaction case in the present model, all else being equal and that other PD models mentioned above which were run with static but complicated networks (e.g. Masuda and Aihara (2003)) didn't report complex dynamics, it would seem reasonable to maintain the importance of endogenous network dynamics to the phenomena.

Further, evidence of Self-Organized Criticality (SOC) occurring in similar PD systems, has to my knowledge, not been presented in any of the fore-going models. Nevertheless, as the discussion of section §3.5 in Chapter 3 indicated, SOC provides an interesting potential basis for very many well known power-law distributions occurring not just in natural systems, but in financial and economic systems also. Often these arguments stem from the identification of some spatially extended dynamical system which is able to exhibit long-memory effects. Indeed, the link between such under-girding processes as SOC and economic systems that exhibit the ubiquitous power-law (Pareto) distribution is the essential focus of the previously mentioned field of *Econophysics.*[2]

Taking together the results that have been reviewed and discussed above, the interplay between the various conceptual approaches to the study of economic systems can be seen.

[2]The reader is referred to Markose's discussion on this point in (Markose, 2005, §4.2).

Furthermore, as has been argued, especially in Chapters 2 and 3, the difficult area of economic networks will no doubt continue to receive this multi-faceted approach. In fact, a trajectory through this thesis can be drawn from more classical analytical approaches and numerical simulations of rational actors in Chapter 5, to agent-based simulation of non-rational strategic play and learning in a small n model of Chapter 6, to the large n, *complex* environment modeled with artificial adaptive agents of Chapter 8. This trajectory offers a microcosm of a wider movement in several literatures where the interplay between levels of analysis mirrors a growing understanding of the interplay between different scales of systems.

It is the nature of agents who act with strategy, foresight and various shades of reason that in particular distinguish the difficulty associated with the study of economic networks. Nonetheless, the field is very much intensifying in its activity on these problems of late, and no doubt will continue to benefit from the variety of approaches mentioned above. Whether complexity science leads or supports this inquiry is yet to be seen, however, as this thesis has shown, it will at the very least, offer challenging perspectives to the traditional approaches and results. And in the case of computational modeling itself – a ready tool for the study of these deterministic yet unpredictable systems – Kollman et al. (1997, p.485-486) capture well its likely role,

> We are neither advocating nor predicting a dramatic change in the social sciences, where computational techniques replace mathematical techniques. Instead, we foresee the future of computational modeling as a steady development of new methods that extend and occasionally overturn existing theoretical insights.

Indeed, such considerations will no doubt demand a progression in our understanding from purely atomistic conceptions of behaviour, to more holistic *system*-wide understandings and formulations. I shall leave it to one of the earliest workers to conceptualise economic systems as complex systems to conclude on this point, Arthur (1999, p.4) writes,

> After two centuries of studying equilibria – static patterns that call for no further behavioural adjustments – economists are beginning to study the general emergence of structures and unfolding patterns in the economy. Complexity economics is not a temporary adjunct to static economic theory, but theory at a more general, out-of-equilibrium level. ... When viewed in out-of-equilibrium formation, economic patterns sometimes simplify into the simple, homogeneous equilibria of standard economics. More often they are ever-changing, showing perpetually novel behaviour and emergent phenomena. Complexity

therefore portrays the economy not as deterministic, predictable and mechanistic; but as process-dependent, organic and always evolving.

References

Acerbi, A. and Parisi, D. (2006). Cultural transmission between and within generations. *Journal of Artificial Societies and Social Simulation*, 9(1).

Adamic, L. A. (1999). The small world web. In *Proceedings of ECDL '99, LNCS 1696*, pages 443–452.

Adamic, L. A. and Adar, E. (2003). Friends and neighbors on the web. *Social Networks*, 25:211–230.

Akerlof, G. (1976). The economics of caste and of the rat race and other woeful tales. *Quarterly Journal of Economics*, 90:599–617.

Alchian, A. A. (1950). Uncertainty, evolution, and economic theory. *The Journal of Political Economy*, 58(3):211–221.

Aleksic, Z. (2000). Artificial life: growing complex systems. In Bossomaier, T. R. J. and Green, D. G., editors, *Complex Systems*, chapter 4, pages 92–126. Cambridge University Press, Cambridge, 1st edition.

Anderlini, L. and Ianni, A. (1996). Path dependence and learning from neighbors. *Games and Economic Behavior*, 13:141–177.

Anderlini, L. and Ianni, A. (1997). Learning on a torus. In Bicchieri, C., Jeffrey, R., and Skyrms, B., editors, *The Dynamics of Norms*, chapter 5, pages 87–107. Cambridge University Press, New York.

Anderson, P. W. (1972). More is different. *Science*, 177(4047):393–396.

Arenas, A., Diaz-Guilera, A., Perez, C. J., and Vega-Redondo, F. (2000). Self-organized evolution in a socioeconomic environment. *Physical Review E*, 61(4):3466–3469.

Arthur, W. B. (1994). Inductive reasoning and bounded rationality. *American Economic Review, Papers and Proceedings*, 84(2):406–411.

Arthur, W. B. (1999). Complexity and the economy. *Science*, 284:107–109.

Arthur, W. B., Durlauf, S. N., and Lane, D. A., editors (1997). *The Economy as an Evolving Complex System II*, Santa Fe Institute Studies in the Science of Complexity, Santa Fe, NM. Santa Fe Institute, Westview Press.

Ashlock, D., Smucker, M. D., Stanley, E. A., and Tesfatsion, L. (1996). Preferential partner selection in an evolutionary study of prisoner's dilemma. *BioSystems*, 37:99–125.

Axelrod, R. (1981). The emergence of cooperation among egoists. *The American Political Science Review*, 75(2):306–318.

Axelrod, R. and Hamilton, W. D. (1981). The evolution of cooperation. *Science*, 211(4489):1390–1396.

Bak, P. and Chen, K. (1991). Self-organized criticality. *Scientific American*, pages 26–33.

Bak, P., Tang, C., and Wiesenfeld, K. (1988). Self-organized criticality. *Physical Review A*, 38(1):364–374.

Bala, V. and Goyal, S. (1998). Learning from neighbours. *The Review of Economic Studies*, 65(3):595–621.

Bala, V. and Goyal, S. (2000). A noncooperative model of network formation. *Econometrica*, 68(5):1181–1229.

Barabasi, A. L., Jeong, H., Neda, Z., Ravasz, E., Schubert, A., and Vicsek, T. (2002). Evolution of the social network of scientific collaborations. *Physica A: Statistical Mechanics and its Applications*, 311:590–614.

Baum, J. A. C., Shipilov, A. V., and Rowley, T. J. (2003). Where do small worlds come from? *Industrial and Corporate Change*, 12(4):687–725.

Berlekamp, E. R., H, C. J., and Guy, R. K. (1982). *Winning Ways for Your Mathematical Plays*. Academic Press, New York.

Berninghaus, S. K., Ehrhart, K.-M., and Ott, M. (2006). A network experiment in continuous time: The influence of link costs. *Experimental Economics*, 9(3):237–251.

Berninghaus, S. K., Ehrhart, K.-M., Ott, M., and Vogt, B. (2004). Searching for "stars" – recent experimental results on network formation. Working Paper 04-34, Sonder Forschungs Bereich: Rationalitätskonzepte, Entscheidungsverhalten und ökonomische Modellierung, Universitat Mannheim.

Binmore, K. G. and Samuelson, L. (1992). Evolutionary stability in repeated games played by finite automata. *Journal of Economic Theory*, 57:278–305.

Blume, L. E. (1995). Population games. Technical Report 96-04-022, Santa Fe Institute. SFI Working Paper,.

Bollobas, B. (1998). *Modern Graph Theory*. Graduate texts in mathematics. Springer, New York.

Boone, R. T. and Macy, M. W. (1999). Unlocking the doors of the prisoner's dilemma: Dependence, selectivity, and cooperation. *Social Psychology Quarterly*, 62(1):32–52.

Borgers, T. (1996). On the relevance of learning and evolution to economic theory. *The Economic Journal*, 106(438):1374–1385.

Bowles, S. and Gintis, H. (2004). The evolution of strong reciprocity: cooperation in heterogeneous populations. *Theoretical Population Biology*, 65(1):17–28.

Burtsev, M. and Turchin, P. (2006). Evolution of cooperative strategies from first principles. *Nature*, 440(20):1041–1044.

Camerer, C. (1995). *The Handbook of Experimental Economics*, chapter 8, pages 587–703. Princeton University Press, Princeton, New Jersey.

Camerer, C. (1998). Bounded rationality in individual decision making. Social Science Working Paper 1029, California Institute of Technology, Pasadena, California 91125.

Camerer, C. F. (1997). Progress in behavioral game theory. *The Journal of Economic Perspectives*, 11(4):167–188.

Canning, D., Amaral, L. A. N., Lee, Y., Meyer, M., and Stanley, H. E. (1998). Scaling the volatility of gdp growth rates. *Economic Letters*, 60:335–341.

Cavalli-Sforza, L. L. and Feldman, M. W. (1981). *Cultural transmission and evolution: a quantitative approach*. Princeton University Press, Princeton, NJ.

Choi, J.-K. (2002). Play locally, learn globally: The structural basis of cooperation. Working Paper 02-12-066, Santa Fe Institute, Santa Fe, New Mexico, USA.

Chwe, M. S.-Y. (1995). Strategic reliability of communication networks. Technical report, University of Chicago, 1126 East 59th Street, Chicago, Illinois 60637, USA.

Chwe, M. S.-Y. (2000). Communication and coordination in social networks. *Review of Economic Studies*, 67:1–16.

Clarke, I., Sandberg, O., Wiley, B., and Hong, T. W. (2000). Freenet: A distributed anonymous information storage and retrieval system. In *ICSI Workshop on Design Issues in Anonymity and Unobservability*, Berkeley, California.

Cohen, M. D., Riolo, R. L., and Axelrod, R. (2001). The role of social structure in the maintenance of cooperative regimes. *Rationality and Society*, 13(1):5–32.

Comellas, F., Ozon, J., and Peters, J. G. (2000). Deterministic small-world communication networks. *Information Processing Letters*, pages 83–90.

Dawkins, R. (1999). *The selfish gene*. Oxford University Press, Oxford.

Devezas, T. and Modelski, G. (2003). Power law behavior and world system evolution: A millennial learning process. *Technological Forecasting and Social Change*, 70:819–859.

Droste, E., Gilles, R. P., and Johnson, C. (2000). Evolution of conventions in endogenous social networks. Mimeo.

Dutta, B., van den Nouweland, A., and Tijs, S. (1998). Link formation in cooperative situations. *International Journal of Game Theory*, 27:245–256.

Elgazzar, A. S. (2001). Application of the sznajd sociophysics model to small-world networks. *International Journal of Modern Physics C*, 12(10):1537–1544.

Elgazzar, A. S. (2002). A model for the evolution of economic systems in social networks. *Physica A: Statistical Mechanics and its Applications*, 303(3-4):543–551.

Elgazzar, A. S. (2003). Applications of small-world networks to some socio-economic systems. *Physica A*, 324:402–407.

Ellison, G. (1993). Learning, local interaction, and coordination. *Econometrica*, 61(5):1047–1071.

Ely, J. C. (2002). Local conventions. *Advances in Theoretical Economics*, 2(1):Article 1.

Epstein, J. M. (2003). Growing adaptive organisations: an agent-based computational approach. Working Paper 03-05-029, Santa Fe Institute, Santa Fe, New Mexico, USA.

Erdös, P. and Rènyi, A. (1959). On random graphs i. *Publicationes Mathematicae, Debrecen*, 6:290–297.

Eubank, S. and Farmer, D. (1990). An introduction to chaos and randomness. In Jen, E., editor, *1989 Lectures in Complex Systems*, volume II of *Santa Fe Institute Studies in the Sciences of Complexity*, pages 75–190. Addison-Wesley, Redwood City, CA.

Falk, A. and Kosfeld, M. (2003). It's all about connections: Evidence on network formation. Discussion Paper 777, Institute for the Study of Labor (IZA). available at http://ideas.repec.org/p/iza/izadps/dp777.html.

Fan, Y. (2002). Questioning guanxi: definition, classification and implications. *International Business Review*, 11:543–561.

Fehr, E. and Schmidt, K. M. (1999). A theory of fairness, competition and cooperation. *Quarterly Journal of Economics*, 114:817–868.

Fruchterman, T. M. J. and Reingold, E. M. (1991). Graph drawing by force-directed placement. *Software-Practice and Experience*, 21(11):1129–1164.

Gell-Mann, M. (1994). *The Quark and the Jaguar.* W.H. Freeman and Company, New York.

Gell-Mann, M. (1995). What is complexity? *Complexity*, 1(1):16–19.

Gell-Mann, M. and Lloyd, S. (1996). Information measures, effective complexity, and total information. *Complexity*, 2(1):44–52.

Girvan, M. and Newman, M. E. J. (2001). Community structure in social and biological networks. Technical report, Santa Fe Institute and Cornell University, 1399 Hyde Park Road, Santa Fe, NM 87501 and Clark Hall, Ithaca, NY 14853-2501.

Goldberg, D. E. (2002). *The Design of Innovation: lessons from and for competant genetic algorithms*. Kluwer Academic Publishing, Boston.

Goldstein, M. L., Morris, S. A., and Yen, G. G. (2004). Problems with fitting to the power-law distribution. *Cond-mat*, (0402322 v3).

Goyal, S. and Moraga-Gonzalez, J. L. (2001). R & d networks. Technical report, Erasmus University, Rotterdam, The Netherlands.

Goyal, S. and Vega-Redondo, F. (1999). Learning, network formation and coordination. Mimeo.

Green, D. G. (1993). Emergent behaviour in biological systems. In Green, D. G. and Bossomaier, T., editors, *Complex Systems: from Biology to Computation*, pages 24–35. IOS Press, Amsterdam, Netherlands.

Green, D. G. (2000). Self-organisation in complex systems. In Bossomaier, T. R. J. and Green, D. G., editors, *Complex Systems*, chapter 2, pages 12–50. Cambridge University Press, Cambridge, 1st edition.

Grim, P. (1996). Spatialization and greater generosity in the stochastic prisoner's dilemma. *BioSystems*, 37:3–17.

Hamilton, W. D. (1964). The genetical evolution of social behaviour, i and ii. *Journal of Theoretical Biology*, 7:1–52.

Harsanyi, J. C. and Selten, R. (1988). *A General Theory of Equilibrium in Games*. MIT press, Cambridge.

Hauk, E. and Nagel, R. (2001). Choice of partners in multiple two-person prisoner's dilemma games. *The Journal of Conflict Resolution*, 45(6):770–793.

Hayes, B. (2000a). Graph theory in practice: Part i. *American Scientist*, 88(1):9–13.

Hayes, B. (2000b). Graph theory in practice: Part ii. *American Scientist*, 88(2):104–109.

Held, G. A., Solina, D. H., Keane, D. T., Haag, W. J., Horn, P. M., and Grinstein, G. (1990). Experimental study of critical-mass fluctuations in an evolving sandpile. *Physical Review Letters*, 65(9):1120–1123.

Henrich, J. (2001). Cultural transmission and the diffusion of innovations: Adoption dynamics indicate that biased cultural transmission is the predominate force in behavioral change. *American Anthropologist*, 103(4):992–1013.

Heylighen, F., Cilliers, P., and Gershenson, C. (2004). Complexity and philosophy. Technical report, ECCO. working-paper (arXiv: cs/0604072).

Hillis, W. D. (1990). Co-evolving parasites improve simulated evolution as an optimization procedure. *Physica D: nonlinear phenomena*, 42:228–234.

Holland, J. H. (1992). Genetic algorithms. *Scientific American*, pages 44–50.

Holland, J. H. (1998). *Emergence: from Chaos to Order*. Perseus Books, Reading, Massachusetts.

Holland, J. H. and Miller, J. H. (1991). Artificial adaptive agents in economic theory. *American Economic Review, Papers and Proceedings*, 81(2):365–370.

Holme, P., Edling, C. R., and Liljeros, F. (2004). Structure and time evolution of an internet dating community. *Social Networks*. In press.

Jackson, M. O. and Watts, A. (2002a). The evolution of social and economic networks. *Journal of Economic Theory*, 106:265–295.

Jackson, M. O. and Watts, A. (2002b). On the formation of interaction networks in social coordination games. *Games and Economic Behavior*, 41(2):265–291.

Jovanovic, B. and Nyarko, Y. (1995). A bayesian learning model fitted to a variety of empirical learning curves. *Brooking Papers on Economic Activity: microeconomics*, pages 247–305.

Joyce, D., Kennison, J., Densmore, O., Guerin, S., Barr, S., and Thompson, N. S. (2006). My way or the highway: a more naturalistic model of altruism tested in an iterative prisoner's dilemma. *Journal of Artificial Societies and Social Simulation*, 9.

Kadanoff, L. P., Nagel, S. R., Wu, L., and Zhou, S.-m. (1989). Scaling and universality in avalanches. *Physical Review A*, 39(12):6524–6537.

Kali, R. (1999). Endogenous business networks. *The Journal of Law, Economics, & Organization*, 15:615–636.

Kandori, M., Mailath, G. J., and Rob, R. (1993). Learning, mutation, and long run equilibria in games. *Econometrica*, 61(1):29–56.

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

Kirchamp, O. (2000). Spatial evolution of automata in the prisoners' dilemma. *Journal of Economic Behavior & Organization*, 43:239–262.

Kollman, K., Miller, J. H., and Page, S. (1997). Computational political economy. In Arthur, B., Durlauf, S. N., and Lane, D. A., editors, *The Economy as an Evolving Complex System II*, volume XVII of *Santa Fe Institute Studies in the Sciences of Complexity*, pages 461–490, Santa Fe, New Mexico. Santa Fe Institute, Westview Press.

Krakauer, D. (2004). Domains of interaction in evolution: Transmission, construction and selection.

Krugman, P. (1996). *The Self-organizing Economy*. Blackwells, Cambridge, MA.

Laguna, M. F., Abramson, G., and Zanette, D. H. (2003). Vector opinion dynamics in a model for social influence. *Physica A: Statistical Mechanics and its Applications*, 329:459–472.

Langton, C. G. (1992). Life at the edge of chaos. In Langton, C. G., Taylor, C., Farmer, J. D., and Rasmussen, S., editors, *Artificial Life II*, volume 41 of *Santa Fe Institute Studies in the Science of Complexity*, New Mexico. Santa Fe Institute, Addison-Wesley. Proceedings of the Workshop on Artificial Life Held February, 1990 in Santa Fe, New Mexico.

Lindenmayer, A. (1971). Developmental systems without cellular interactions in development, their languages and grammars. *Journal of Theoretical Biology*, 30(3):455–484.

Lindgren, K. (1992). Evolutionary phenomena in simple dynamics. In Langton, C. G., Taylor, C., Farmer, J. D., and Rasmussen, S., editors, *Artificial Life II*, volume X of *Santa Fe Institute Studies in the Science of Complexity*, New Mexico. Santa Fe Institute, Addison-Wesley. Proceedings of the Workshop on Artificial Life Held February, 1990 in Santa Fe, New Mexico.

Lorenz, E. N. (1963). Deterministic nonperiodic flow. *Journal of Atmospheric Science*, 20:130.

Markose, S. M. (2005). Computability and evolutionary complexity: markets as complex adaptive systems (cas). *The Economic Journal*, 115:F159–F192.

Marsili, M., Vega-Redondo, F., and Slanina, F. (2004). The rise and fall of a networked society: a formal model. *PNAS*, 101:1439–1442.

Masuda, N. and Aihara, K. (2003). Spatial prisoner's dilemma optimally played in small-world networks. *Physics Letters A*, 313:55–61.

Matros, A. (2004). Simple rules and evolutionary selection. IZA Bonn research seminar paper, available at http://www.iza.org/.

Matutinović (2002). Organizational patterns of economies: an ecological perspective. *Ecological Economics*, 40(3):421–440.

Maynard Smith, J. (1964). Group selection and kin selection. *Nature*, 201(4924):1145–1147.

Maynard Smith, J. (1978). The evolution of behavior. *Scientific American*, 239(3):136–145.

Maynard Smith, J. and Price, G. R. (1973). The logic of animal conflict. *Nature*, 246(2):15–18.

Milgram, S. (1977). *The Individual in a Social World: Essays and Experiments*, chapter 18 The Small World Problem, pages 281–295. Addison-Wesley Series in Social Psychology. Addison-Wesley, Reading, Massachusetts.

Miller, J. H. (1988). *Two Essays on the Economics of Imperfect Information.* Ph.d. dissertation, University of Michigan.

Miller, J. H. (1996a). Active nonlinear tests (ants) of complex simulation models. Working Paper 96-03-011, Santa Fe Institute, Santa Fe, New Mexico, USA.

Miller, J. H. (1996b). The coevolution of automata in the repeated prisoner's dilemma. *Journal of Economic Behavior & Organization*, 29:87–112.

Miller, J. H., Butts, C. T., and Rode, D. (2002). Communication and cooperation. *Journal of Economic Behavior & Organization*, 47:179–195.

Miller, J. H. and Page, S. E. (2005). The interest in between. Private communication, from Introduction to forthcoming book, 'Complex Adaptive Social Systems'.

Mitchell, G. T. (1999). Evolution of networks and the diffusion of new technology. Technical report, Computational Laboratories Group at the University of California, Santa Barbara, Department of Economics, University of California, Santa Barbara, CA 93106.

Moody, J. (2001). Peer influence groups: identifying dense clusters in large networks. *Social Networks*, 23:261–283.

Morris, S. (2000). Contagion. *The Review of Economic Studies*, 67:57–78.

Nowak, M. A. and May, R. M. (1992). Evolutionary games and spatial chaos. *Nature*, 359:826–829.

Nowak, M. A. and Sigmund, K. (1998). Evolution of indirect reciprocity by image scoring. *Nature*, 393:573–577.

Orbell, J. and Dawes, R. (1993). Social welfare, cooperator's advantage, and the option of not playing the game. *American Sociological Review*, 58:787–800.

Pruitt, D. and Kimmel, M. (1977). Twenty years of experimental gaming: Critique, synthesis, and suggestions for the future'. *Annual Review of Psychology*, 28:363–392.

Rapoport, A. and Horvath, W. J. (1961). A study of large sociogram. *Behavioral Science*, 6:279–291.

Ray, T. S. (1993). How i created life in a virtual universe.

Riolo, R. L., Cohen, M. D., and Axelrod, R. (2001). Evolution of cooperation without reciprocity. *Nature*, 414:441–443.

Scheinkman, J. A. and Woodford, M. (1994). Self-organized criticality and economic fluctuations. *American Economic Review*, 84(2):417–421.

Seeley, D. A. (2000). Network evolution and the emergence of structure. In Bossomaier, T. R. J. and Green, D. G., editors, *Complex Systems*, chapter 3, pages 52–89. Cambridge University Press, Cambridge, 1st edition.

Slikker, M. and van den Nouweland, A. (2000). Network formation models with costs for establishing links. *Review of Economic Design*, 5:333–362.

Smucker, M. D., Stanley, E. A., and Ashlock, D. (1994). Analyzing social network structures in the iterated prisoner's dilemma with choice and refusal. Technical Report

CS-TR-94-1259, University of Wisconsin-Madison, Department of Computer Sciences, 1210 West Dayton Street, Madison, WI.

Standifird, S. S. and Marshall, R. S. (2000). The transaction cost advantage of *guanxi*-based business practices. *Journal of World Business*, 35(1):21–42.

Standish, R. K. (2001). On complexity and emergence. *Complexity International*, 9(Paper ID: standi09).

Stocker, R., Cornforth, D., and Bossomaier, T. R. J. (2002). Network structures and agreement in social network simulations. *Journal of Artificial Societies and Social Simulation*, 5(4).

Sugden, R. (1989). Spontaneous order. *The Journal of Economic Perspectives*, 3(4):85–97.

Surowiecki, J. (2004). *The Wisdom of Crowds*. Doubleday, New York.

Taylor, C. R. (2000). The old-boy network and the young-gun effect. *International Economic Review*, 41(4):871–91.

Tesfatsion, L. (1997). *A Trade Network Game with Endogenous Partner Selection*, volume 6 of *Advances in Computational Economics*. Kluwer Academic Publishers.

Tesfatsion, L. (2003). Agent-based computational economics. Working Paper 1, Iowa State University, Department of Economics, Ames, Iowa.

Traulsen, A. and Nowak, M. A. (2006). Evolution of cooperation by multilevel selection. *Proceedings of the National Academy of Sciences*, 103(29):10952–10955.

Tyler, J. R., Wilkinson, D. M., and Huberman, B. A. (2003). Email as spectroscopy: Automated discovery of community structure within organizations. Technical report, HP labs, 1501 Page Mill Road, Palo Alto, CA, 94304.

Vainstein, M., Silva, A. T. C., and Arenzon, J. J. (2006). Does mobility decrease cooperation. http://arxiv.org/abs/q-bio.PE/0608009.

Volz, E. (2004). Random networks with tunable degree distribution and clustering. *Condensed Matter*.

von Neumann, J. (1966). *The Theory of Self Reproducing Automata*. Univ. of Illinois Press, Urbana.

Wasserman, S. and Faust, K. (1994). *Social Network Analysis: methods and applications*. Cambridge University Press, Cambridge.

Watts, D. J., Dodds, S. P., and Newman, M. E. J. (2002). Identity and search in social networks. *Science*, 296:1302–1305.

Watts, D. J. and Strogatz, S. H. (1998). Collective dynamics of 'small-world' networks. *Nature*, 393:440–442.

White, D. R., Owen-Smith, J., Moody, J., and Powell, W. P. (2004). Networks, fields, and organizations: Micro-dynamics, scale and cohesive embeddings. Working Paper 04-03-009, Santa Fe Institute, Santa Fe, New Mexico, USA.

Wilson, D. S. (1975). A theory of group selection. *Proc. Natl. Acad. Sci. U.S.A.*, 72(1):143–146.

Wilson, D. S. (1976). Evolution on the level of communities. *Science*, 192:1358–1360.

Wilson, D. S. and Sober, E. (1994). Reintroducing group selection to the human behavioral sciences. *Behavioral and Brain Sciences*, 17(4):585–654.

Wilson, E. O. (2001). *Pheidole in the new world: a dominant, hyperdiverse ant genus.* Harvard University Press, Cambridge, M.A.

Wilson, E. O. (2005). Kin selection as the key to altruism: Its rise and fall. *Social Research*, 72(1):159–166.

Wolfram, S. (1983). Statistical mechanics of cellular automata. *Reviews of Modern Physics*, 55(3):601–644.

Wolfram, S. (1984a). Cellular automata as models of complexity. *Nature*, 311:419–424.

Wolfram, S. (1984b). Universality and complexity in cellular automata. *Physica D*, 10:1–35.

Yamagishi, T. and Hayashi, N. (1996). Selective play: Social embeddedness of social dilemmas. In Liebrand, W. and Messick, D., editors, *Frontiers in Social Dilemmas Research*, pages 363–384. Springer, Berlin.

Zeggelink, E. (1995). Evolving friendship networks: an individual-oriented approach implementing similarity. *Social Networks*, 17:83–110.

Zhang, X. and Li, G. (2003). Does *guanxi* matter to nonfarm employment? *Journal of Comparitive Economics*, 31:315–331.

Zurek, W. H. (1990). Algorithmic information content, church-turing thesis, physical entropy, and maxwell's demon. In Jen, E., editor, *1989 Lectures in Complex Systems*, volume II of *Santa Fe Institute Studies in the Sciences of Complexity*, pages 49–65. Addison-Wesley, Redwood City.

Some Technicalities

On the Generation of Pictures for Part Pages

Each picture is the resultant state-space in $\mathbb{R}^2$ of the system of non-linear equtions,

$$\begin{aligned} x_{t+1} &= by_t + F(a, x_t) \ , \quad \text{and} \\ y_{t+1} &= -x_t + F(a, x_{t+1}) \end{aligned}$$

where,

$$F(a, x) = ax + (1-a)\frac{2x^2}{0.5 + x^2} \ .$$

Some will recognise this system as being related to that of Gumowski and Mira. In this case, the parametrs $(a, b) = (-0.809, 0.97)$ were used, based on an interactive evolutionary search process. A program was written in MATLAB to generate a 3×3 grid of state-space pictures with a series of seed parameters, I would then choose the picture I liked the best and whether or not the picture was 'close' to a favourite solution, and the program would create 9 new pictures based on a genetic algorithm updating procedure.

The pictures used in the thesis represent the state space with 1,000, 10,000, and 100,000 points plotted respectively for each successive Part page. In all cases, an initial input of $(x_0, y_0) = (0.5, 0.5)$ was used.

On the Generation of Network & FSA Diagrams

Networks, encoded by their adjacency matrix were optimized for minimum node-overlap by generating random positions in $\mathbb{R}^2$ and applying a modified force-directed-algorithm (see for example, Fruchterman and Reingold (1991)). Two modifications were made, firstly, rather than a 'spring' model, nodes were given masses and inter-atomic forces were applied to the space, a method that avoids problematic 'energy walls' from hindering placement. Second, an annealing schedule was applied to node velocities, in order that cyclic or eliptical orbits were slowed and eventually stopped in the lowest energy state. In this case, programs written in MATLAB were used to automatically generate LaTeX 2$_\varepsilon$code for use with the pst-node package.

Finite State Automata (FSA) diagrams were produced in a simlar way, although in some cases nodes were placed on a circle of set radius rather than doing a free-optimization.

All algorithms used in the production of this thesis are available from the author by request.