Learning in networks: a survey

Sanjeev Goyal*

July 2003

Abstract

This paper presents a survey of research on learning with a special focus on the structure of interaction between individual entities. The structure is formally modelled as a network: the nodes of the network are individuals while the arcs admit a variety of interpretations (ranging from information channels to social and economic ties). I first examine the nature of learning about optimal actions for a given network architecture. I then discuss learning about optimal links and actions in evolving networks.

Forthcoming in, *Group formation in economics: networks, clubs, and coalitions* edited by G. Demange and M. Wooders. Cambridge University Press

*Department of Economics, University of Essex and Tinbergen Institute Amsterdam-Rotterdam. E-mail: sgoyal@essex.ac.uk
I would like to thank Yann Bramoulle, William Brock, Leigh Tesfatsion, Antonella Ianni, Werner Raub, Larry Samuelson, Gerhard Sorger, Brian Skyrms, and Fernando Vega-Redondo for useful comments on an earlier version of the paper.

Contents

1 Introduction

In a wide range of economic situations, individuals make decisions without being fully informed about the rewards from different options. In many of these instances the decision problems are of a recurring nature and it is natural that individuals use their past experience and the experience of others in making current decisions. The experience of others is important for two reasons: one, it may yield information on different actions per se (as in the case of choice of new consumer products, agricultural practices, or medicines prescribed), and two, in many settings the rewards from an action depend on the choices made by others and so there is a direct value to knowing about other's actions (as in the case of which credit card to use, or which language to learn, or whether to buy a fax machine or not). This suggests that the precise way in which individuals interact can influence the generation and dissemination of useful information and that this could shape individual choices and social outcomes. In recent years, these considerations have motivated a substantial body of work on learning in economics, which takes explicit account of the structure of interaction among individual entities. The present paper provides a survey of this research.

I will consider the following simple framework: there is a set of individuals who are located on nodes of a network; the arcs of the network reflect relations between these individuals. At regular intervals, individuals choose an action from a set of alternatives. They are uncertain about the rewards from different actions. They use their own past experience as well as gather information from their neighbours (individuals who are linked to them) and then choose an action that maximizes individual payoffs. I start by studying the influence of network structure on individual and social learning. In this part the network is taken as given. In the second part of the survey I explore learning in a setting where the network itself is evolving due to individual choices on link formation. Here, the focus will be on the way that individual incentives shape the evolution of structure as well as economic performance.[1]

[1]In this survey I will be mostly concerned with learning in models where agents are either fully rational or the departures from full rationality are relatively minor. Moreover, the focus of the survey is entirely on analytical results. I will therefore not be discussing the large literature on agent based modelling and computational economics which studies similar issues. For surveys of this work see Judd and Tesfatsion (2005), Kirman and Zimmermann (2001)).
For surveys on static models of network formation see the accompanying papers by Jackson (2003) and van den Nouweland (2003) in this volume.

The above framework allows for a rich set of interpretations and I provide some examples to illustrate this:

- Adoption of consumer products: Consumers make decisions on brand choice without complete knowledge of the alternatives. They try out different brands and also gather information from market surveys and their friends and acquaintances to make more informed choices.
- Medical innovation: Doctors have to decide on new treatments for ailments without complete knowledge of their efficacy and side-effects; they read professional magazines as well as exchange information with other doctors in order to determine whether to prescribe a new treatment.
- Agricultural practices: Farmers decide on whether to adopt new seeds and farming packages without full knowledge of their suitability for the specific soil and weather conditions they face. They use the experience of neighbouring farms and extension services in making decisions.

In the above examples individuals use their links with others primarily to gather useful information on product characteristics or suitability. In the following examples, there is strategic interaction between individuals and the rewards from an action depend on the actions of others.

- Adoption of new information technology: Individuals decide on whether to adopt fax machines or a new computer operating system without full knowledge of its usefulness. This usefulness is related to the choices of others with whom they would like to communicate.
- Language choice: Individuals choose which language to learn at school as a second language. Here the rewards depend on the choices of others with whom they expect to interact in the future.
- Credit card choice: Individual consumers choose a credit without full knowledge to the benefits of the card, since they do not know how often the card can be used. The benefit in turn depends on the credit cards adopted by shops that they frequent. The shopkeepers face a similar decision problem.

- Social norms on punctuality: Individuals decide whether to be on time for their appointments or to arrive a bit late. The returns from being to time and the costs associated with being late depend on the choices of others whom they are going to meet.

The work discussed in this survey should be seen as part of a larger research programme in economics which examines the role of informal institutions and non-market interaction in shaping economic activity; for general introductions to this area see Goyal (1999), Kirman (1997), and Young (1998). I now briefly mention some closely related strands of research. There is a large and growing body of empirical work which studies the influence of the network/interaction structure on economic outcomes.[2] This work documents the different ways in which individual behaviour is sensitive to the patterns of interaction, and also illustrates how changes in the patterns of interaction lead to changes in individual behaviour and social outcomes. There is also a significant body of experimental research on the effects of networks and non-market interaction on individual behavior and social learning.[3]

The study of learning has been one of the most active fields of research in economics in the last two decades. Different aspects of this research have been surveyed in articles and books; see e.g., Blume and Easley (1995), Fudenberg and Levine (1999) Kandori (1997), Marimon (1997) and Samuelson (1997). The distinctive feature of the present survey is its focus on the role of interaction structures in learning.

I would next like to mention the early work of Coleman (1966) and Schelling (1975) and the large body of work in economic sociology which studies the effects of social structure on economic performance. For an introduction to some of the themes in this this body of work, see Burt (1994), Coleman (1990), Granovetter (1974, 1985), Raub and Weesie (1990), and Smelser and Swedberg (1994). There is also an extensive literature on network formation in mathematical sociology; for recent work in this area see Carley and Banks (1996), Doreian

[2]See e.g., Hagerstrand (1969), Griliches (1957), Ryan and Gross (1943) on diffusion new agricultural practices, Coleman (1966) and Taylor (1979) on diffusion and patterns of medical practices, Young (1998) on spread of traffic norms, Elias (1978) on the history of social customs and manner, Munshi (2003) on migration and social networks, Watkins (1991) on spread of norms in marriage and fertility, Burke and Young (2001) on norms in contracting, and Glaeser, Sacerdote and Scheinkman (2001) on local interaction and criminal activity. For a discussion on the technical issues arising in the measurement of local effects, see Glaeser and Scheinkman (2001), Brock and Durlauf (2001), and Manski (2000).

[3]Kosfeld (2003) provides a survey of the experimental work.

and Stokman (2001), and Snijders (2001). I would like to briefly relate the work in economics and sociology. Traditionally, the relation between the economics research and the sociology research has been seen as follows: In the economics strand, most of the research has a relative simple formulation of the objective function – maximization of a payoff function – but a relatively rich formulation of the strategic possibilities inherent in the network formation process. On the other hand, the sociology work endows individuals with a very rich and varied set of motivations but pays relatively less attention to the strategic aspects of the network formation process. In recent years as the rational choice school has become more prominent in sociology, the research methodology in the two subjects has become more similar. A second difference is the emphasis on the relation between individual incentives and social efficiency in economics, something which seems to be less studied by sociologists.

Finally, I mention the rapidly growing literature in physics on the subject of networks. This work highlights statistical regularities of actual networks such as the World Wide Web, the Internet, the network of co-authors (in different disciplines), network of actors, among others. The empirical work shows that these networks display small world features (the average distance between nodes in the network is relatively short), clustering (high overlap between the connections of connected nodes) and a scale free distribution of links. The research has also developed simple dynamic models of expanding networks which generate these statistical properties. For comprehensive recent surveys of this work in physics see Albert and Barabasi (2002) and Dorogovtsev and Mendes (2002). The distinctive element of the research in economics is the emphasis on individual incentives and strategic interaction.

A brief word on the style of the paper: I believe that the issues addressed here are of general interest and so I have tried to make this survey accessible to readers with different backgrounds. The main results are presented precisely and the intuition behind them is developed in some detail. On occasion, to keep the exposition smooth, I have taken the liberty of omitting some technical assumptions (or qualifications). To make up for this, I have provided complete references for all the results reported, and the enthusiastic reader is encouraged to refer to the originals for the mathematical proofs.

The rest of the paper is organized as follows. In section 2, I introduce networks and present the basic terminology which will be used throughout the survey. I start with a presentation of results on learning within a given network. Section 3 considers learning about optimal actions in non-strategic environments, while section 4 considers learning about optimal actions in strategic environments. I then discuss learning in evolving networks. Section 5 discusses learning about optimal link decisions, while section 6 examines learning about optimal links as well as actions in strategic games. Section 7 concludes.

2 Networks

Let $N = \{1, 2, ..., n\}$ be a finite set of individuals/decision makers, each of whom is located (and identified with) a distinct node of a network. An arc or a link between two individuals i and j is denoted by $g_{i,j}$, where $g_{i,j} \in \{0, 1\}$. Here, $g_{i,j} = 1$ reflects the presence, while $g_{i,j} = 0$ denotes the absence of a link from i to j. In Figure 1, for example, there are 3 players, 1, 2 and 3, and $g_{1,3} = g_{3,1} = g_{1,2} = 1$. We shall denote a network by g and the set of all networks by $\mathcal{G}$. There is a one-to-one correspondence between the set of directed network on n vertices and the set $\mathcal{G}$. We say there is a *path* from j to i in g either if $g_{i,j} = 1$ or there exist distinct players j_1, ..., j_m different from i and j such that $g_{i,j_1} = g_{j_1,j_2} = ... = g_{j_m,j} = 1$. For example, in Figure 1 there is a path from player 2 to player 3. The notation "$j \xrightarrow{g} i$" indicates that there exists a path from j to i in g. A network g is said to be connected if there exists a path between any pair of players i and j in g. The (geodesic) distance from player j to player i in g is the number of links in the shortest path from j to i, and is denoted $d_{i,j}(g)$. We set $d_{i,j}(g) = \infty$ if there is no path from j to i in g.

I now define neighbourhoods of players in a network. Let $N^d(i; g) = \{k \in N | g_{i,k} = 1\}$ be the set of individuals with whom i has a direct link in network g. We shall refer to $N^d(i; g)$ as the set of direct neighbours of i in network g. Let $N(i; g) = \{k \in N | k \xrightarrow{g} i\}$ be the set of individuals whom i can directly or indirectly access in network g. Let $\mu_i^d : \mathcal{G} \to \{0, ..., n-1\}$ be defined as $\mu_i^d(g) \equiv |N^d(i; g)|$. Here, $\mu_i^d(g)$ is the number of individuals with whom i is directly linked in network g. The term $\mu_i(i; g)$ is defined correspondingly. Thus in Figure 1 below, $N(1; g) = \{1, 2, 3\}$, $N(2; g) = \{2\}$, $N(3; g) = \{1, 2, 3\}$, while $N^d(1; g) = \{2, 3\}$, $N^d(2; g) = \phi$, and $N^d(3; g) = \{1\}$

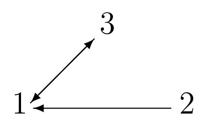

Figure 1

A component of g, is a subset $C \subset N$ and the set of all links between the members of C in g, with the property that for every distinct pair of players i and j in C I have $j \xrightarrow{g} i$, (equivalently, $j \in N(i;g)$) and there is no strict superset C' of C in g for which this is true. A network g is said to be *minimal* if the deletion of any link increases the number of components in g. We can also say that a network g is *connected* if it has a unique component. If the unique component is minimal, g is called *minimally connected*. A network which is not connected is referred to as disconnected. A network is said to be *empty* if $N(i;g) = \{i\}$ and it is called *complete* if $N^d(i;g) = N \backslash \{i\}$ for all $i \in N$. We denote the empty and the complete network by g^e and g^c, respectively. A *star* network has a central player i such that $g_{i,j} = g_{j,i} = 1$ for all $j \in N \backslash \{i\}$ and there are no other links. A *wheel* network is one where the players are arranged as $\{i_1, ..., i_n\}$ with $g_{i_2,i_1} = ... = g_{i_n,i_{n-1}} = g_{i_1,i_n} = 1$ and there are no other links. The wheel network is denoted g^w. Figure 2 presents these networks for a society with 4 people.

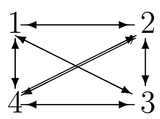

Figure 2a
Complete Network

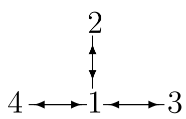

Figure 2b
Star

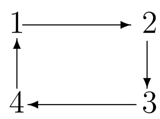

Figure 2c
Wheel

Two networks $g \in \mathcal{G}$ and $g' \in \mathcal{G}$ are equivalent if g' can be obtained by a relabelling of the players in g. For example, if g is the network in Figure 1, and g' is the network where players 1 and 2 are interchanged, then g and g' are equivalent. The equivalence relation partitions $\mathcal{G}$ into classes: each class is referred to as an *architecture*. For example, there are n possible 'star' networks, all of which come under the equivalence class of the star architecture. Likewise, the wheel architecture is the equivalence class of $(n-1)!$ networks consisting of all permutations of n individuals in a circle.

We shall say that a network graph is *regular* if all individuals have the same number of neighbours, $N^d(i;g) = k$, for some $k \in \{0, 1, 2..., n-1\}$. In this case the number of neighbours is referred to as the *degree* of the network.

The above description is for directed networks, i.e., networks in which there is an explicit direction to the arc between two nodes. In particular, in directed networks the presence of a link $g_{i,j} = 1$ says nothing about the status of the link $g_{j,i}$. By contrast, in undirected networks, a link has no orientation/direction and $g_{i,j} = g_{j,i}$. In some parts of the survey, I will discuss undirected networks; the concepts and terminology for these networks can be developed in an analogous manner.

3 Non-strategic interaction

In this section I study learning about optimal actions in a setting where the rewards from an action do not depend on the actions chosen by other individuals.

I consider the following general framework. There are many decision makers, each of whom faces a similar decision problem: to choose an action at regular intervals without knowing the true payoffs from the different actions. The action chosen generates a random reward and also provides information concerning the true payoffs. An individual uses this information as well as the experience of a subset of the society, her *neighbours* to update her prior beliefs. Given the updated beliefs, an individual chooses an action that maximizes one-period expected utility. I study the dynamic process of individuals' beliefs, actions and utilities. Our interest is in the influence of the structure of neighbourhoods on the actions that individuals choose and the transmission of information.

The nature of neighbourhood influence on individual choice has been studied in early papers by Allen (1982) and Ellison and Fudenberg (1992). In the following discussion I draw heavily on more recent papers, by Bala and Goyal (1998, 2001), as these papers fit in more naturally within the general framework of the survey – which involves a finite number of individuals located in a determinstic network structure, finite actions, and myopic best response decision rules – that is used throughout the survey. I will discuss these early papers later in this section.

Decision Problem: Time is assumed to be discrete, and indexed by $t = 1, 2, ...$ There are $n \geq 3$ individuals; an individual i chooses an action from a finite set of alternatives, denoted by S_i. In this section, I assume that $S_i = S_j = A$, for every pair of individuals i and j. I denote by $s_{i,t}$ the action taken by individual i in time period t. The payoffs from an action depends on the state of the world θ, which belongs to a finite set Θ. If θ is the true state and an individual chooses action $a \in A$ then he observes an outcome $y \in Y$ with conditional density $\phi(y, a; \theta)$ and obtains a reward $r(a, y)$. To simplify matters, I will assume that Y is the real line and that the reward function $r(a, .)$ is bounded.

Individuals do not know the true state of the world; their private information is summarized in a prior belief over the set of states. For individual i this prior is denoted by $\mu_{i,1}$. The set of prior beliefs is denoted by $\mathcal{P}(\Theta)$. I assume that priors of all individuals are interior, i.e., $\mu_{i,1}(\theta) > 0$, $\forall\ \theta$, and $\forall\ i \in N$. Given belief μ, an individual's one-period expected utility from action a is given by

$$u(a, \mu) = \sum_{\theta \in \Theta} \mu(\theta) \int_Y r(a, y)\phi(y, a; \theta)dy. \tag{1}$$

I assume that individuals have similar preferences which are reflected in a common reward function $r(., .)$. I will consider the role of heterogeneity later. Let $G : \mathcal{P}(\Theta) \to X$ be the one-period optimality correspondence:

$$G(\mu) = \{a \in A | u(a, \mu) \geq u(a', \mu),\ \forall a' \in A\} \tag{2}$$

Let δ_θ represent point mass belief on the state θ; then $G(\delta_\theta)$ denotes the set of optimal actions if the true state is θ. I now provide an example which is a special case of the framework outlined above.

Example 1: Suppose $A = \{a_1, a_2\}$ and $\Theta = \{\theta_o, \theta_1\}$. In state θ_1, action a_1 yields Bernoulli distributed payoffs with parameter $\pi \in (1/2, 1)$, i.e., it yields 1 with probability π, and 0 with probability $1 - \pi$. In state θ_o, action a_1 yields a payoff of 1 with probability $1 - \pi$, and 0 with probability π. Furthermore, in both states, action a_o yields payoffs which are Bernoulli distributed with probability 1/2. Hence action a_1 is optimal in state θ_1, while action a_o is optimal in state θ_o. The belief of an individual is a number $\mu \in (0, 1)$, which represents the probability that the true state is θ_1. The one period optimality correspondence is given by

$$G(\mu) = \begin{cases} a_1 & \text{if } \mu \geq 1/2 \\ a_o & \text{if } \mu \leq 1/2 \end{cases}$$

Dynamics: For each $i \in N$, let $g_i : \mathcal{P} \to A$, be a selection from the one-period optimality correspondence G. In period 1, each individual chooses $g_i(\mu_{i,1})$ and observes the outcome; individual i also observes the actions and outcomes of each of her neighbours, $j \in N^d(i)$. I assume that individuals use this information to update their prior $\mu_{i,1}$, and then make a decision in period 2 and so on. In particular, I assume that individuals do not infer information about unobserved others from the action choices of their neighbours. There are two reasons for this assumption. The first reason is descriptive realism: I feel that agents either do not have the computational capacity to work through these inferences or do not find the computations involved worthwhile. The second reason is tractability of the model. This assumption as well as the assumption of myopic optimization helps me to simplify the model and allows me to focus on the role of interaction structure directly.[4]

I need some additional notation to descibe the dynamic process. For a fixed θ let $(\Omega, \mathcal{F}, P^\theta)$ define the probability triple where Ω contains all the sample realizations of all the individuals over time and P^θ is the probability measure induced over the sample paths by $\theta \in \Theta$. For a subset $B \subset \Theta$ and H a mesaurable subset of Ω, let $P_i(B \times H)$ be given by

$$P_i(B \times H) = \sum_{\theta \in B} \mu_{i,1}(\theta) P^\theta(H). \tag{3}$$

for each individual $i \in N$. A typical sample path is of the form $\omega = (\theta, \omega')$, where θ is the state of nature and $w' = ((y^a_{i,1})_{a \in A, i \in N}, (y^a_{i,2})_{a \in A, i \in N}, ..., (y^a_{i,t})_{a \in A, i \in N} ...)$, where $y^a_{i,t} \in Y^a_{i,t} = Y$. Let $C_{i,t} = g_i(\mu_{i,t})$ denote the action of individual i in period t, $Z_{i,t}$ the outcome of this action, and let $U_{i,t}(\omega) = u(C_{i,t}, \mu_{i,t})$ be the expected utility of i with respect to her own action at time t. Given this notation I can now write down the posterior beliefs of individual i in period $t+1$ as follows:

$$\mu_{i,t+1}(\theta) = \frac{\prod_{j \in N^d(i) \cup \{i\}} \phi(Z_{j,t}; C_{j,t}, \theta) \mu_{i,t}(\theta)}{\sum_{\theta' \in \Theta} \prod_{j \in N^d(i) \cup \{i\}} \phi(Z_{j,t}; C_{j,t}, \theta) \mu_{i,t}(\theta)}. \tag{4}$$

Our interest is in studying the influence of network structure on the evolution of individual actions, beliefs, and utilities, $(a_{i,t}, \mu_{i,t}, U_{i,t})_{i \in N}$, over time.

[4] See Bramoulle and Kranton (2003) for a model of social learning with fully rational players located in networks.

The following result, due to Bala and Goyal (1998), shows that the beliefs and utilities of individuals converge, in the long run.

Theorem 3.1 *The beliefs and utilities of every individual converge: $lim_{t\to\infty}\mu_{i,t}(\omega) = \mu_{i,\infty}(\omega)$ and $lim_{t\to\infty}U_{i,t}(\omega) = U_{i,\infty}(\omega)$, for every $i \in N$, with probability one.*

The first part of the statement follows as a corollary of the Martingale Convergence Theorem (see e.g., Billingsley, 1985). Let $A_i(\omega)$ be the set of actions that are chosen infinitely often by individual i along sample path ω. It is intuitive that each of these actions must be optimal with respect to limit beliefs and must yield the same utility in each of the states that are in the support of the limit belief $\mu_{i,\infty}(\omega)$. The result on limiting utilities follows from this observation.

I now examine whether all information is communicated efficiently in a connected society. There are different ways of addressing this issue. One way would be to ask if different persons get the same payoffs in the long run. This would suggest that they each possess a similar amount of 'useful' information. The following result, due to Bala and Goyal (1998), is in this spirit.

Theorem 3.2 *Every individual in a connected society gets the same long run utility: $U_{i,\infty} = U_{j,\infty}$ for every $i, j \in N$, with probability one.*

The idea behind the above result is as follows: if i observes the actions and outcomes of j then he must be able to do as well as j in the long run. Next note that this must be true by transitivity for any person k who observes j indirectly. The final step is to note that in a connected society there is an information path from any player i to any player j. This result shows that in a connected society information transmission is good enough to ensure that every person gets the same utility in the long run. The above results lead me to the question: do beliefs converge to the truth and are individuals choosing the optimal action and earning the maximal possible utility in the long run?

We shall assume that θ_1 is the true state of the world. The long run actions of a player i are said to be optimal if $A^i(\omega) \subset G(\delta_{\theta_1})$. Social learning is said to be complete if for all $i \in N$, $A^i(\omega) \subset G(\delta_{\theta_1})$, on a set of sample paths which has probability 1 (with respect to the true state θ_1). The analysis of long run learning rests on the informativeness of actions. An action is said to be fully informative if it can help an individual distinguish between all the states: if for all θ, $\theta' \in \Theta$, with $\theta \neq \theta'$,

$$\int_Y |\phi(y; a, \theta) - \phi(y; a, \theta')| dy > 0. \tag{5}$$

By contrast, an action a is uninformative if $\phi(., a; \theta)$ is independent of θ. In example 1 above, action a_o is uninformative while action a_1 is fully informative.

It is natural that prior beliefs have to be restricted if individuals are to learn the optimal action. For instance, in Example 1 above, if everyone has priors such that the uninformative action is optimal then there is no additional information emerging in the society and nothing is going to change over time. Optimism by itself is, however, not sufficient. To see this consider a society with a finite number of individuals, and suppose that I am in the setting of Example 1 and that everyone has optimistic priors, $\mu_{i,1} > 1/2$, for all $i \in N$. Thus in period 1 everyone will choose action a_1. There is a well defined number of realizations of 0, say T, after which an individual will switch to the uninformative action a_o. Since individual trials are independent, the probability of such a sequence of bad realizations is positive. Given that the number of individuals is finite there is also a positive probability that everyone gets such a poor sequence of realizations in the first T trials. Thus there is a positive probability that everyone chooses the uninformative action a_0 after a finite time T.[5] The above argument also shows that the probability of incomplete learning is strictly positive in any finite society. While this finding is useful, it leaves open two related questions: one, what is the relative attractiveness of different networks for social learning in finite societies and two, what happens to the probability of learning in different networks as the number of players gets large and in the limit goes to infinity? I am not aware of any general results on the first question; Bala and Goyal (1998) develop some results on question two and I report these results now.

Consider therefore a large society where everyone is optimistic, i.e., $\mu_{i,1}(\theta_1) > 1/2$. I explore the role of information networks in this setting. Suppose that the decision problem is as in Example 1 and for concreteness suppose that beliefs satisfy the following condition.

[5]This line of reasoning has been developed and elaborated upon in the Bayesian learning literature; see e.g., Rothschild (1974), Easley and Kiefer (1988), McLennan (1984).

$$\inf_{i\in N}\mu_{i,1} > \frac{1}{2}; \;\; \sup_{i\in N}\mu_{i,1} < \frac{1}{1+x^2} \tag{6}$$

where $x = (1-\pi)/\pi \in (0,1)$. From the optimality correspondence formula, it follows that every person chooses a_1 in period 1. Suppose that individuals are arranged along the integer points of the real line and that the direct neighbourhood of players is as follows: $N^d(i) = \{i-1, i+1\} \cup \{1,2,3,4,5\}$. I shall refer to the commonly observed group of individuals $\{1,2,3,4,5\}$ as the royal family. This structure corresponds to situations in which individuals have access to local as well as some common/public source of information. For example, such a structure arises naturally in the context of agriculture where individual farmers observe their neighbouring farmers but all the farmers observe a few large farms and agricultural laboratories. Similarly, in academic research, individual researchers keep track of developments in their own field of specialization and also try and keep abreast of the work of pioneers/intellectual leaders in their subject more broadly defined.

Suppose that the true state is θ_1 and a_1 is the optimal action. I now argue that there is a strictly positive probability of incomplete learning in this society. The argument is quite simple: suppose that every person in the royal family is unlucky in the first period and gets an outcome of 0. Consider any individual i and note that this person can get at most 3 positive signals from her immediate neighbourhood. Thus any person in this society will have a minimum residual of 2 negative signals on the true state. Given the assumptions on the priors, this negative information is sufficient to push the posteriors below the critical cut-off level of 1/2 and this will induce a collective switch to action a_o in period 2. From then on no further information is generated and society gets locked into the uninformative and sub-optimal action. Notice that this argument does not use the size of the society and thus I have obtained an upper bound which is less than 1, on the probability of learning for all n. This example illustrates how a few common signals can block out and overwhelm a vast amount of locally available information. This example also suggests a possible mechanism for getting out of the problem: looking at networks in which local information is given due weight.

I now present a simple network in which local information is given enough scope and ultimately prevails. Consider a society where for every i, $N^d(i) = \{i-1, i+1\}$. It is possible to show that in this society complete learning obtains. The argument is as follows. First, I fix an individual i and apply the strong law of large numbers to claim that there is a set of sample paths with positive probability on which the experience on the optimal action a_1 always remains positive, on the average. This means that starting with optimistic priors, individual i will persist with action a_1, forever on this set of sample paths, if he were isolated. I then similarly construct a set of sample paths for each the neighbours of player i, on which players $i-1$ and $i+1$, respectively, receive positive information, on average. Exploiting independence of trials across players, I infer that the probability of the three players $i-1$, i, and $i+1$ receiving positive information on average is strictly positive, say $q > 0$. Hence the probability of individual i choosing the sub-optimal action a_o is bounded above by $1-q$. I finally note that along this set of sample paths, the experience of other individuals outside the neighbourhood cannot alter the choices of individual i. Similarly, I can construct a set of sample paths for individual $i+3$, whose information neighbourhood is $\{i+2, i+3, i+4\}$. From the i.i.d nature of the process, I can deduce that the probability of this sample of paths is $q > 0$ as well. Note next that since individuals i and $i+3$ do not share any neighbours, the two events, that neither i nor $i+3$ tries the optimal action in the long run are independent and the probability of this joint event is therefore bounded above by $(1-q)^2$. In a society where $\bar{N}^d(i) = \{i-1, i+1\}$, and given that $q > 0$, it now follows that learning can be made arbitrarily close to 1, by suitably raising the number of individuals.

This example illustrates in a simple way how the architecture of the information network affects the possibilities of social learning. In particular, it shows that the probability of learning can be increased by decreasing the number of information links and thereby restricting flow of information in society. More generally, it helps us identify a structural feature of information networks that facilitate learning, *local independence*. I shall say that two individuals i and j are locally independent if $N^d(i) \cap N^d(j) = \emptyset$. In a society with a royal family the positive information generated on the optimal actions in different parts of the society is overturned by the negative information generated by the royal family. By contrast, in a society with local ties only, negative information does arise over time but it cannot simultaneously overrule the positive information generated in different parts of the society. This allows the positive local information to gain a foothold and eventually spread across the whole society. This insight is fairly general and is summarized in the following result, due to Bala and Goyal (1998). A player i has optimistic prior beliefs if $g_i(\mu_{i,1}) \subset G(\delta_{\theta_1})$.

Theorem 3.3 *Suppose a society is connected. In such a society the probability of learning can be made arbitrarily close to 1, by suitably increasing the number of locally independent optimistic players.*

Conformism vs diversity: Conformism is the state in which everyone chooses the same action, while diversity is a situation in which positive fractions of the population choose each of the two actions. Theorem 3.2 says that in a connected society all individuals will obtain the same utility, in the long run. In a setting with a unique optimal action for every state this implies that all individuals will choose the same action as well. These results are obtained in a setting where all individuals have the same preferences, reflected in their having identical reward functions. In a society with heterogeneous preferences, the analogue of this result would be as follows: in a connected society, all individuals with the same preferences obtain the same utility. In a recent paper, Bala and Goyal (2001) show that this conjecture is false. They present an example to illustrate how preference differences can create information blockages that impede the transmission of useful information and thereby sustain different utility levels for individuals with similar preferences. This leads then to propose a stronger notion of connectedness: *group-wise connectedness*. A society is said to be group-wise connected if for every pair of individuals i and j of the same preference type, either j is a neighbour of i, or there exists a path from j to i with all members of the path being individuals having the same preference as i and j. Theorem 3.2 obtains for members with the same preferences in societies which satisfy this stronger connectedness requirement.

Spatial Patterns: The above framework also allows us to explore the spatial and temporal patterns of diffusion of new technologies. I report some simulations presented by Bala and Goyal (1998). They find that the temporal pattern (percentage of adopters vs. time) is described quite well by the logistic function, and that the rate of adoption is positivey related to the profitability of the new technology. These findings are consistent with the empirical patterns reported in Griliches (1957). The spatial pattern of adoption is as follows: initially small groups of individuals adopt the new technology and it spreads slowly as neighbouring individuals adopt it as well. Eventually these regions of adopters join up and the pace of adoption accelerates. These findings match the empirically patterns reported in Hagerstrand (1969).

I now discuss the papers by Allen (1982) and Ellison and Fudenberg (1992). Allen studies technology adoption by a set of individuals located on nodes of a graph, who are subject to local influences. This is clearly very close in spirit to the motivation behind the framework developed above. She shows that, if every action is chosen with a positive probability, then there exists a unique global (joint) distribution on actions given any local influence structure. These results tell us something about the consistency requirements imposed by local influences, but they leave open the issue of the dynamics of how local influences work their way through to the global level, which is the focus of the present survey. Ellison and Fudenberg (1992) consider a setting with a unit measure of individuals each of whom makes a choice between two actions. The relative profitability of these technologies is unknown. In each period a fraction of the population gets an opportunity to revise their choices. These individuals observe the average payoffs of the two actions in the previous period and pick the action that yields the higher payoff. The authors examine the share of the population adopting different actions over time.

Let f and g be the two technologies and suppose the payoffs are given as follows: $u_t^g - u_t^f = \beta + \epsilon_t$. The value of β is unknown and the ϵ is a random variable with mean 0 and a cumulative distribution H. The distribution of ϵ is such that Probability$[u_t^g - u_t^f \geq 0] = p > 0$. Thus it may happen that action g gets higher payoffs in a period even though the action f is better, i.e., $\beta < 0$.

Let x^t denote the fraction of individuals choosing action g in period t. The first result they obtain says that the time average of x^t converges to its expectation with respect to the unique invariant measure, ν. Moreover, this average corresponds in a simple way to the distribution of the noise in the payoffs function of the two actions: $E_\nu(x) = p$.

It is easy to see that a decision based solely on comparing previous period payoffs does not really allow the superior quality of an action to express itself and as a result the process fluctuates with the noise and does not actually settle down on any one action in the long run. To allow for the history of past experiences to have greater influence, Ellison and Fudenberg (1992) next consider a decision rule that gives some weight to the relative popularity of the different actions. In particular, they use the following decision rule: an individual prefers action g if the observed payoff difference $u_t^g - u_t^f \geq m(1 - 2x^t)$, where m refers to the weight on the relative popularity of the actions. It is clear that given some m, a larger x^t makes it more likely that the inequality will be satisfied and action g will be chosen. For all $x^t \neq 1/2$, a larger m signifies greater weight on the popularity of different actions. Suppose that the distribution of shocks ϵ_t is uniform on $[-\sigma, \sigma]$. Given this assumption, I can state the following result, due to Ellison and Fudenberg (1992), on the behaviour of x^t for different values of m.

Theorem 3.4 *If $m = \sigma$ then x^t converges to* 1 *if g is optimal ($\beta > 0$), and it converges to* 0 *if f is optimal ($\beta < 0$). For $m < \sigma$ the process need not converge, while for $m > \sigma$ it always converges but the limit is sensitive to the initial state x_o.*

This result suggests that there is an optimal weight for popularity in the decision rule: $m < \sigma$ represents underweighting while $m > \sigma$ reflects overweighting of the popularity.

Ellison and Fudenberg also consider a spatial model of learning, where payoffs are sensitive to location. Suppose that the measure of individuals is arranged along a line, and each individual has a window of observation around herself. Moreover suppose that the payoffs are given as follows: $u_t^g(\theta) = \theta + \beta\theta + \epsilon_{gt}$, for technology g at location θ, and $u_t^f(\theta) = \beta\theta + \epsilon_{ft}$, for technology f. With this formulation, it follows that there is a cut-off for the optimal technology at $\theta = 0$, with g being the optimal choice for $\theta > 0$, while f is the optimal choice for $\theta < 0$. In each period, individuals observe the average payoffs of the two technologies in their window of observation. For individual θ, the window is an interval given by $[\theta - w, \theta + w]$. They choose the action which yields a higher average payoff in this window. Suppose there is a boundary point x_0 at the start. Ellison and Fudenberg study the evolution of this boundary over time. From our point of view, their main result pertains to the effect of the size of the window of observation on the long run properties of the system. They find that under some regularity conditions on the distribution of the error terms, the steady state welfare is decreasing in the size of the interval. Thus smaller intervals are better from a long term welfare point of view. However, if w is small then the cut-off point moves slowly over time if the initial state is far from the optimum and this creates a trade-off: increasing w leads to a short-term welfare gain but a long-term welfare loss.

Random networks: So far I have discussed learning in a setting where every player has a fixed set of neighbours and the implicit assumption is that these neighbours constitute a small subset of the whole population. An alternative way to think of information transmission is as follows: in every decision period, an individual gets to observe a sample of other people chosen at random from the population. This individual uses the information obtained from the sample – which could relate to the relative popularity of different actions, or actions and outcomes of different trials – in making her choices. This approach has been explored by Banerjee and Fudenberg (1994), Ellison and Fudenberg (1995), and Smallwood and Conlisk (1979), among others. In this approach the attention has been on the size of the sample, the type of information extracted from the sample, and the nature of the decision rule which maps this information into the choice of individuals (boundedly rational or Bayesian).

It is useful to briefly examine the effects of random sampling in the decision problem framework above. Suppose there are a finite number of players. Then I can employ standard mathematical arguments to show that random sampling implies that every person observes every other person infinitely often (with probability 1). The society will therefore be 'fully connected' in a stochastic sense and I expect that the analogues of Theorems 3.1-3.2 will obtain. The argument on incomplete learning in finite societies can also be extended in a straightforward manner. Thus I will need large , i.e., infinite societies to obtain complete learning within the random network setting. To get an impression of the issues that arise in settings with large populations, I now discuss a model developed by Ellison and Fudenberg (1995).

Consider a unit measure of individuals each of whom makes a choice between two actions. The relative profitability of these technologies is unknown. In each period a fraction of the population gets an opportunity to revise their choices. These individuals observe a random sample of other persons and compare the average payoffs of the two actions in this sample (they also use their own experience in arriving at this average), and pick the action that yields the higher payoff. I examine the share of the population adopting different actions over time.

The two actions are denoted by f and g. Suppose that the payoffs to individual i choosing action f are given by $\bar{f}_t + \epsilon_{ift}$ and the payoffs from action g are given by $\bar{g}_t + \epsilon_{igt}$. The second term reflects the idiosyncratic shocks while the first term refers to the aggregate level of payoffs in that period. The idiosyncratic shocks are assumed to be i.i.d across players and time. They have mean 0 and standard deviation σ. It is assumed that $\beta_t = \bar{g}_t - \bar{f}_t$, and that β_t has a binomial distribution with Probability$(\beta_t = \beta) = p > 0$, while Probability$(\beta_t = -\beta) = 1 - p$.

In each period t a fraction α of the individuals get an opportunity to revise their decisions. Each person faced with this choice gets to observe the actions and payoffs of K randomly chosen other individuals. She supplements this information with her own experience and then arrives at an estimate of the average payoffs from the two actions and she chooses the action with the higher average payoff. In case individual i was choosing action f and does not observe any trials with action g then she is obliged to persist with action f, irrespective of the information on action f. Let x^t denote the fraction of individuals choosing action g in period t. The authors study the behaviour of x^t for different values of the sample size K.

The main results of the paper concern the relative likelihood of social conformism and diversity. The first result, due to Ellison and Fudenberg (1995), is for the case where both actions have the same payoffs.

Theorem 3.5 *Suppose that $p = 1/2$ and the actions are on average equally good. Then x^t converges to an end-point with everyone choosing the same action if the sample size is small. If sample size is large then x^t exhibits diversity, in the long run.*

The intuition behind this result is as follows: If sample sizes are large then individual idiosyncratic noise gets washed out for individuals who are revising their choices, and the process is governed by the aggregate shocks captured in the variable β. Since β is binomial, the process x^t oscillates and a positive fraction of the population chooses each of the two actions, in the long run. On the other hand, if the sample is small then aggregate shocks are mediated by individual shocks and percolate slowly through the system. In such a setting, there is a tendency for popular actions to be reinforced, and the system always converges to an end-point.

The second result, due to Ellison and Fudenberg (1995), covers the case where the two actions have different payoffs.

Theorem 3.6 *Suppose that $p > 1/2$ and action g is, on average, superior. Then x^t converges to an end-point with everyone choosing the same action if the sample size is small. This action can be either f or g and so inefficient conformism can occur. If sample size is moderate then x^t converges to the efficient action g, while if sample size is large then x^t exhibits diversity, in the long run.*

The relationship between sample size and long run learning is due to the combination of aggregate and idiosyncratic noise. To see this I briefly consider the two limit cases: one, with no aggregate uncertainty ($p = 1$) and two, with no idiosyncratic noise. The first case permits a clean comparision with the learning in fixed networks framework presented earlier. In that decision environment, there is a true (stochastic) quality level for each action and individual trials yield independent draws; hence it has no aggregate uncertainty but does have idiosynractic noise. We note that in the above result, the critical sample sizes depend on the value of p, and for the limit case with $p = 1$, it can be shown that efficient conformism obtains for all sample sizes. The intuition for this builds on the arguments presented for Theorem 3.5 above. In case $p = 1$, there is no aggregate uncertainty and as samples get large the influence of idiosynractic noise is less so that the superior technology dominates. Consider next the case where there are aggregate shocks but no idiosyncratic noise. In this case, sample size is only relevant only in so far as it affects the probability of individuals accessing at least one draw of each action.

I would like to conclude by mentioning the literature on herding and informational cascades (Banerjee, 1992; Bikhchandani, Hirshleifer and Welch, 1992). In this literature there is a single sequence of privately informed individuals who take one action each. Before making her choice an individual gets to observe the actions of all the people who have moved earlier. The person moving in a period thus uses the actions of her predecessors as signals for their private information and uses these signals to supplement her own private information. This is quite different from the framework developed above in which individuals can access the entire experience – the actions as well as rewards – of a subset of individuals, their neighbours. Thus there is learning from actions and payoffs, and moreover different individuals have access to different aspects of the social information depending on their location in the society.

4 Strategic interaction

In this section I will study learning of optimal actions in a setting where the rewards from different actions depend on the actions chosen by other individuals. Given our interest in the influence of network structure on individual choice, this leads me to a study of strategic interaction among individuals located in networks. I will consider both games of coordination as well as games of conflict. As before, our interest is in the question: what is the influence of the structure of interaction on individual choice and the behavior of the system as a whole.

4.1 Coordination Games

Suppose there are two players 1 and 2 and they are engaged in the following 2×2 game.

1 \ 2	α	β
α	a, a	d, e
β	e, d	b, b

Figure 3

I shall assume that payoffs satisfy the following restrictions.

$$a > d; \quad b > d; \quad d > e; \ a + d > b + e. \tag{7}$$

These restrictions imply that there are two pure strategy equilibria of the game: (α, α) and (β, β) and that coordinating on either of them is better than not coordinating at all. The assumption that $a + e > b + d$ implies that α is the *risk-dominant* strategy. It is worth noting that α can be risk-dominant even if it is not efficient (that is even if $b > a$). Given the restrictions on the payoffs, these equilibria are strict in the sense that the best response in the equilibrium yields a strictly higher payoff than the other option. It is well known that strict equilibria are robust to standard refinements of Nash equilibrium; thus players engaged in such a game face a coordination problem.

I shall consider a group of players who are engaged in playing this coordination game. The structure of interaction will be modelled in terms of an *undirected* network, whose nodes are the players and an arc between two players signifies that the two players play the game with each other. I start with a discussion of the static problem and then take up the issue of learning.

Suppose there are n players located on vertices of a undirected network with each player being located on a distinct node. To distinguish between a directed and an undirected link, I shall use the notation $\bar{g}_{i,j} \in \{0, 1\}$ to denote a undirected link between players i and j. As before, $\bar{g}_{i,j} = 1$ denotes that the existence of a link, and $\bar{g}_{i,j} = 0$ denotes the absence of a link between players i and j. Let $\bar{g}$ denote an undirected network and let $\bar{\mathcal{G}}$ denote the set of all undirected networks with n nodes. Recall that $N_i^d(g) = \{j \in N | \bar{g}_{i,j} = 1\}$ refers to the set of players with whom i is linked in network $\bar{g}$. I will use s_i to denote a strategy of player i and $S_i = \{\alpha, \beta\}$ to denote the strategy set. I will use $S = \Pi_{i \in N} S_i$ to denote the set of all strategy profiles in the game and s to refer to a typical member of this set. In the above two person game, let $\pi(x, y)$ denote the payoffs to player i when this player chooses action x, while her opponent chooses action y. The payoffs to a player i from a strategy s_i, given that the other players are choosing s_{-i} is:

$$\Pi_i(s_i, s_{-i}) = \sum_{j \in N_i^d(\bar{g})} \pi(s_i, s_j) \tag{8}$$

This formulation reflects the idea that a player i interacts with each of the players in the set $N_i^d(\bar{g})$. A strategy profile $s^* = \{s_1^*, s_2^*, ..., s_n^*\}$ is a Nash equilibrium if $\Pi_i(s_i^*, s_{-i}^*) \geq \Pi_i(s_i, s_{-i}^*)$, for all $s_i \in S_i$, for all players $i \in N$. The equilibrium is strict if the inequalities are strict for every player.

I start by describing the nature of Nash equilibria under different network structures. The *first* point I note is that the strategy profile $s_i = x$, for all $i \in N$, where $x \in \{\alpha, \beta\}$ is a Nash equilibrium given any network structure. This is straightforward to check given the restrictions on the payoffs. Thus the issue is: are there any other equilibria and how is the answer to this question related to the network structure? The *second* point to note then is that if the network is complete, i.e., every pair of players has a link, then the above mentioned outcomes with social conformism are the only equilibria possible. However, if networks are incomplete then a variety of other strategy profiles can arise in a Nash equilibrium. To see this I consider some specific network structures. Consider a society of N players which is divided into two groups, N_1 and N_2, with $N_1 \cup N_2 = N$. Suppose that $\bar{g}_{i,j} = 1$ if and only if $i, j \in N_k$, for $k = 1, 2$. In other words, there exists a link between every pair of players in a group and there are no links across players of the two groups. In this simple network it is an equilibrium for players in group 1 to choose action α, while members of group 2 choose β. (The converse pattern with members of group 1 choosing action β, while members of group 2 choose α is clearly also an equilibrium. These are the only two possible equilibria in this network.) This example exploits the separation of the two groups of players, and leads us to ask: can diversity arise in a connected network? The answer to this question is *yes*. By way of illustration, suppose that the N players are located on a square 2-dimension lattice grid and every player has a neighbourhood of 4 players. (I join the ends of the lattice to ensure that the interaction structure has no asymmetries.) In this network, draw a line through the middle of the graph and consider the strategy configuration in which all players to the right of the line choose α while all players to the left of this line choose β. Any configuration of this form in which at least two adjacent columns (from top to bottom of the lattice) of players choose the same action is a Nash equilibrium. These observations are sumarized as follows.

Theorem 4.1 *The outcome where everyone chooses the same action is a Nash equilibrium in any network. If the network is complete then these are the only equilibria. If the network is incomplete then a variety of other configurations in which players choose different actions can arise in equilibrium.*

This result yields two insights: one, there is a multiplicity of equilibria for networks and two, the possibility and nature of mixed equilibria depend on the network architecture.

These observation lead me to a closer examination of the plausibility of different equilibria and how this is in turn related to the structure of interaction. I shall study plausibility in terms of dynamic stability. I start with a simple decision rule for individuals, and examine the behaviour of the dynamic process generated by this rule.

Suppose that time is discrete and given by $t = 1, 2, 3, ...$ In each period, with probability $p \in (0, 1)$, a player gets an opportunity to revise her strategy. Faced with this opportunity, player i chooses an action which maximizes her payoff, under the assumption that the strategy profile of her neighbours remains the same as in the previous period. If more than one action is optimal then the player persists with the current action. Denote the strategy of a player i in period t by s_i^t. If player i is not active in period t then it follows that $s_i^t = s_i^{t-1}$. This simple best-response strategy revision rule generates a transition probability function $P_{ss'}(\bar{g}) : S \times S \rightarrow [0, 1]$, which governs the evolution of the state of the system $s^t(\bar{g})$. Recall that a strategy profile (or state) is said to be absorbing if the dynamic process cannot escape from the state once it reaches it. Equipped with this notation and terminology, I can now present a general convergence and characterization result.

Theorem 4.2 *Fix some network of interaction $g \in \mathcal{G}$. Starting from any initial strategy configuration, the dynamic process $s^t(g)$ converges to an absorbing strategy profile in finite time, with probability 1. There is an equivalence between the set of absorbing strategy profiles and the set of Nash equilibria of the static social game.*[6]

The equivalence between absorbing states and Nash equilibria of the social game of coordination is easy to see. The arguments underlying the convergence result are as follows: start at some state s_o. Consider the set of players who are not playing a best response. If this set is empty then we are at a Nash equilibrium configuration and this is an absorbing state of the process. Suppose therefore that there are some players who are currently choosing action α but would prefer to choose β. Allow them this choice and let s_1 be the new state of the system (this transition occurs with positive probability, given the above defined process). Now examine the players doing α in state s_1 who would like to switch actions. If there are some such players then have them switch to β and define the new state as s_2. Clearly this process of having α players switch will end in finite time (since there are a finite number of players in the society). Let the state with this property be $\hat{s}$. Either there will be no players left choosing α or there will be some players choosing α in $\hat{s}$. In the former case we are clearly at a Nash equilibrium. Let us take up the second possibility then. Check if there are any players choosing β who would like to switch actions. If there are none then we are at an absorbing state. If there are some β players who would like to switch then follow the process as outlined above to reach a state in which there is no player who wishes to switch from β to α. Let this state be denoted by $\bar{s}$. Next observe that no player who was choosing α (and did not want to switch actions) in $\hat{s}$ would be interested in switching to β. This is true because the game is a coordination game and the set of players choosing α has (weakly) increased in the transition from $\hat{s}$ to $\bar{s}$. Hence we have arrived (with positive probability) at a state in which no player has any incentive to switch actions. This is an absorbing state of the dynamics; since the initial state was arbitrary, and the above transition occurs with positive probability, the theory of Markov chains tells us that the transition to an absorbing state will occur in finite time, with probability 1.

[6] One may wonder if there is any relationship between the Nash equilibria of the social game and the original 2×2 game, I started with. This issue has been studied by Mailath, Samuelson and Shaked (1997) show that the Nash equilibria of the static social game is equivalent to the set of correlated equilibria of the 2×2 game. Ianni (2001) studies convergence to correlated equilibria under myopic best response dynamics.

An early result on convergence of dynamics to Nash equilibrium in regular networks (where every player has the same number of neighbours) is presented in Anderlini and Ianni (1996). In their model a player is randomly matched to play with one other player in her direct neighbourhood. Moreover, every player gets a chance to move in every period. Finally, a player who plans to switch actions can make an error with some probability. They refer to this as noise on the margin. With this decision rule, the dynamic process of choices converges to a Nash equilibrium for a class of regular networks. The result I present here holds for all networks and does not rely on mistakes for convergence. Instead, I rely on inertia and the coordination nature of the game.

The above result shows that the learning process with regard to actions converges in due time. The result also says that every Nash equilibrium (for the given network of interaction) is an absorbing state of the process. Thus one cannot hope to select across the variety of equilibria identified earlier with this dynamic process. This motivates a study of stability with respect to small but repeated perturbations.

This is formally done using the idea of stochastic stability, introduced by Kandori, Mailath and Rob (1993), and Young (1993). I suppose that, occasionally, players make mistakes, experiment, or simply disregard payoff considerations in choosing their strategies. Specifically, I assume that, conditional on receiving a revision opportunity, a player chooses her strategy at random with some small "mutation" probability $\epsilon > 0$. Given a network g, and for any $\epsilon > 0$, the mutation process defines a Markov chain that is aperiodic and irreducible and, therefore, has a unique invariant probability distribution; let us denote this distribution by $\mu_\epsilon(\bar{g})$. I analyze the support of $\mu_\epsilon(\bar{g})$ as the probability of mistakes becomes very small, i.e. as ϵ converges to zero. Define $\lim_{\epsilon \to 0} \mu_\epsilon(\bar{g}) = \hat{\mu}_{\bar{g}}$. I shall say that a state s is stochastically stable if $\hat{\mu}_{\bar{g}}(s) > 0$. This notion of stability identifies states that are relatively more stable with respect to occasional errors or experiments by individuals.

The ideas underlying stochastically stability can be informally described as follows. Suppose that s and s' are the two absorbing states of the best-response dynamics. Given that s is an absorbing state, a movement from s to s' requires an error on the part of one or more of the players. Similarly, the transition from s' to s requires errors on the part of some subset of players. The state s is stochastically stable if it requires relatively more errors to exit than the other state. If it takes the same number of mutations to exit the two states, then the two states are both stochastically stable.

I shall be using the notion of stochastic stability extensively in what follows. It is therefore important to point out some limitations of this approach as a way to select for equilibrium. One limitation is the lack of an explicit model which explains the individual errors and experimentation. A second limitation of this approach is that in most applications the number of mutations needed are of the order of the number of players, and so in large societies, as the probability of errors becomes small, the process moves very slowly and the rate of convergence can be very slow. For an overall account of the concepts, techniques and applications of stochastic stability, see Young (1998).[7].

I start with two examples: interaction with every player in a complete network and interaction with immediate neighbours among players located around a circle. These examples show that the risk-dominant action α is stochastically stable. I first take up the complete network. Suppose that player 1 is deciding on whether to choose action α or β. It is easy to verify that the mimimum number of players choosing α needed to induce player 1 to choose α is given by $k = (n-1)(b-d)/[(a-e)+(b-d)]$. Similarly, the minimum number of players choosing action β needed to induce player to choose action β is given by $l = (n-1)(a-e)/[(a-e)+(b-d)]$. Given our assumption that $a+d > b+e$ it follows that $k < n/2 < l$. It now follows that if we are in a state where everyone is choosing α then it takes l mutations to transit to a state where everyone is choosing action β; likewise, if we are in a state where everyone is choosing β then it takes k mutations to transit to a state where everyone is choosing action α. From the arguments in the above paragraph it now follows that in the complete network, everyone choosing the risk-dominant action α is the unique stochastically stable outcome.

I now turn to interaction with immediate neighbours among players located around a circle. This example is taken from Ellison (1993). Suppose that at the start every one is choosing action β. Now suppose that two adjacent players i and $i+1$ choose action α by way of experimentation. It is now easy to verify that in the next period, the immediate neighbours of i and $i+1$, players $i-1$ and $i+2$ will find it optimal to switch to action α (this is due to the assumption that α is risk-dominant and $a+d > b+e$). Moreover, in the period after

[7]For an explosition of the original mathematical results refer to Freidlin and Wentzell (1984)

that the immediate neighbours of $i-1$ and $i+2$ will have a similar incentive, and so there is a contagion process under way which leads to everyone choosing action α, in finite time. On the other hand, if we were to start in a state with everyone choosing α then it is difficult to generate a similar contagion process. To see why note that a player bases her decision on the actions of immediate neighbours, and so long as at least one of the neighbours is choosing α the optimal action is to do likewise. Hence so long as there are two players choosing action α, the action will revive and take over the whole population. This simple argument suggests that it is relatively easy to perturb the state where everyone is doing β while it is significantly more difficult to perturb the state in which everyone is choosing α. These observations taken along with the earlier remarks on stochastic stabilty show that the everyone choosing the risk-dominant action is the unique stochastically stable action when players are arranged on a circle and interact with their immediate neighbours.

The simplicity of the above arguments may lead one to conjecture that the risk-dominant outcome obtains in all networks. I now present an example (which is taken from Jackson and Watts, 2001b) to show that this conjecture is false. Suppose that players are arranged in a star network. Recall that this is a network in which one player has links with all the other $n-1$ players, while the other players have no links between them. We shall take player 1 to be the central player in our discussion. It is easily verified that in a star network a perturbation which switches the action of the central player is necessary as well as sufficient to get a switch of all the other players. Hence the number of perturbations needed to go from an all α state to an all β state is equal to the number of perturbations needed for the reverse transition. Thus both the states are equally vulnerable and are both stochastically stable as well.

The above arguments illustrate how the network structure can shape the nature of the long run outcome. These examples lead us to the question: are there circumstances under which the stochatically stable states are independent of the network structure? One response to this question is provided by a result in Young (1998). This result proceeds by making the mutations in individual strategy sensitive to payoff losses. To present this result I need to develop some additional notation. I shall say that in every period t, an individual i is drawn at random and chooses an action (say) α according to a probability distribution, $p_i^\gamma(\alpha|s^t)$, where $\gamma > 0$ and s^t is the strategy profile at time t. The probability distribution is obtained via the following formula:

$$p_i^\gamma(\alpha|s^t) = \frac{e^{\gamma.\Pi_i(\alpha,s_{-i}^t)}}{e^{\gamma.\Pi_i(\alpha,s_{-i}^t)} + e^{\gamma.\Pi_i(\beta,s_{-i}^t)}} \tag{9}$$

This is referred to as the log-linear response rule. This rule was first studied in Blume (1993) in the context of games played on lattices.[8] Note that for large values of γ the probability distribution will place most of the probability mass on the best response action. Define $\Delta_i(s) = \Pi_i(\beta, s_{-i}) - \Pi_i(\alpha, s_{-i})$, and say that $\rho = e^{-\gamma}$. Then for large γ I can express the probability of action α as follows:

$$p_i^\gamma(\alpha|s^t) = \frac{e^{-\gamma.\Delta_i(s^t)}}{1 + e^{-\gamma.\Delta_i(s^t)}} \cong e^{-\gamma\Delta_i(s^t)} = \rho^{\Delta_i(s^t)} \tag{10}$$

This expression says that the probability of not choosing the best response is exponentially declining in the payoff loss from the deviation. Equipped with these additional concepts, I am ready to state the following general result on learning to play coordination games in networks, due to Young (1998).

Theorem 4.3 *Let $\bar{g}$ be an arbitrary connected network. Suppose that in each period an individual is picked at random to revise choices. In revising choices this individual uses the log-linear response rule. Then the stochastically stable outcome is one in which every player chooses the risk-dominant action.*

To get some intuition for the result let us briefly discuss the effects of the log-linear decision rule on the dynamic process in the star network. In that example, the simplest way to get a transition is via a switch in the action of the central player. In the standard model, with payoff insensitive mutations, the probability of the central player making a switch from α to β is the same as the other way around. Under the log-linear response rule, matters are different. If there are a large number of players then there is a significant difference in the payoffs losses involved and the probabilities of switching from α to β is significantly smaller as compared to the probability of switching from β to α. This difference is crucial for obtaining the above result.

[8]For a general treatment of the theory of statistical mechanics refer to Liggett (1985).

In the above models the dynamics are Markovian and if there is a unique invariant distribution then standard mathematical results suggest that the rate of convergence is exponential. In other words, there is some number $\rho < 1$ such that the probability distribution of actions at time t, σ^t, approaches the invariant distribution σ^* at a rate approximately given by ρ^t. While this result is helpful, it is easy to see that this property allows a fairly wide range of rates of convergence, depending on the value of ρ. The rate of convergence is important because it clarifies the relative importance of the initial conditions and the evolutionary/dynamic forces, respectively. If ρ is close to 1 then the process is essentially determined by the initial configuration σ_0 for a long period, while if ρ is close to 0 then initial conditions play a less important role and dynamics shape individual choices quickly. The work of Ellison (1993) focused attention on the role of interaction structure in shaping the rate of convergence. He argued that if interaction was random or in a complete network then transition between strict Nash equilibria based on mutations would take a very time in large populations since the number of mutations needed is of the order of the population. By contrast, if interaction takes place between between immediate neighbours who are arranged on a circle then it is easy to see that a couple of mutations (followed by best responses) would be sufficient to initiate a transition to the risk-dominant action. Thus local interaction leads to dramatically faster rates of convergence to the risk-dominant action. In a recent paper, Young (1998) shows that the role of local interaction in speeding up rate of convergence to risk-dominant outcome is quite general: in any society where people are organized in clusters with considerable overlap between neighborhoods the rate of convergence is quite quick and essentially independent of the size of the population.

Related themes: In the last ten years, research on the subject of coordination games has been very active and a significant part of this work deals with interaction models. To conclude this section, I briefly mention some of the other issues that have been explored. I first take up the issue of alternative decision rules (this issue is also related to the issue of how mutations/experiments shape the set of stochastically stable outcomes, which was discussed above.) In this connection, I would like to discuss two issues that have been examined in the literature. The first issue is imitation based decision rules and the second is state-dependent mutations. In the discussion so far, I have assumed that in revising strategies, each player chooses a myopic best response to the current strategy profile. An alternative decision rule would require that an individual compares the realized payoffs from different actions in the previous period and chooses the action that yielded the highest payoffs. I shall refer to this as the imitate the best action rule. Robson and Vega-Redondo (1998) show how this rule taken together with alternative matching rules leads to the efficient action (which is not necessarily risk-dominant) being always stochastically stable. The second issue concerns the modelling of the mutations. Bergin and Lipman (1996) argued that any outcome could be supported as stochastically stable under a suitable mutation structure. This 'anything is possible' result should not come entirely as a surprise given our earlier observations on stochastically stable states when players are located on a star. This result has provoked several responses and I mention two of them here. The first response interprets mutations as errors, and says that these errors can be controlled at some cost. This argument has been developed in van Damme and Weibull (2002). This paper shows that incorporating this cost structure leads us back to the risk-dominant equilibrium. This line of research has been further extended to cover local interaction on general weighted graphs by Baron, Durieu, Haller and Solal (2002). A second response is to argue that risk-dominance obtains for all possible mutation rules, if some additional conditions are satisfied. In this vein, a recent paper by Lee, Szeidl and Valentinyi (2002) argues that if interaction is on a 2-dimensional lattice then the dynamics select for the risk-dominant action for any state-dependent mutation structure, provided the number of players is sufficiently large.

A *second* concern has been the role of random initial configurations. Lee and Valentinyi (2000) study the spread of actions on a 2-dimensional lattice. They suppose that, at the start, every player chooses each of the two actions with positive probability while subsequent decisions are based on a myopic best-response rule. They find that if there are sufficiently large number of players then all players choose the risk-dominant action. This result should be seen in the context of Theorem 4.1 presented above which shows that a variety of mixed equilibria can arise on 2-dimensional lattices. The result of Lee and Valentinyi suggests that such diversity is unlikely to obtain, in a probabilistic sense.

The *third* concern has been the issue of openness of different interaction structures to change. Suppose that there is one-shot introduction of players choosing a new practice into a population who follow a different action. What are the prospects of the new action catching on in the population and how is this likelihood related to the interaction structure? This question has been addressed in Goyal (1996) and Morris (2000). Using a general framework of analysis, Morris shows that diffusion is easier in networks where there is substantial overlap in neighborhoods. In particular, he is able to parameterize the receptiveness of a network to new actions in terms of a contagion threshold and finds that local interaction on a circle is maximally receptive.

The *fourth* issue that has been discussed is the role of interaction structure in shaping behaviour in games where players wish to coordinate on action combinations $\{\alpha, \beta\}$ or $\{\beta, \alpha\}$. In other words, there are two (pure strategy) Nash equilibria but they both involve players choosing different actions. Bramoulle (2000) refers to these as congestion games and finds that network structure affects the static equilibria quite differently as compared to the coordination games I have discussed above. For example, in the coordination games studied above, we observed that the outcomes where everyone chooses the same action are Nash equilibria irrespective of the network of interaction. This is clearly not true when we are dealing with congestion games. More generally, the welfare properties of the networks are quite different as well, since in congestion games, complete bipartite graphs are particularly attractive.

4.2 Games of conflict

I now examine the role of interaction structure in shaping individual behavior in the context of games of conflict. There are a variety of different ways in which conflict in games can be modelled. I study a particularly simple model as it allows me to draw out the role of interaction structure in a straightforward manner and it also points to some interesting open issues in this area.

Suppose that there are a large number of players each of whom has a choice between two actions A and E. I shall think of A as referring to an altruistic or cooperative action and E as referring to an egoistic or 'defect' action. Let $s_i \in S_i = \{A, E\}$ denote the strategy of player i and let $s = \{s_1, s_2, s_3 ... s_n\}$ refer to the strategy profile of the players. I use $n(A, s_{-i})$ to refer to the number of players who choose action A in the strategy profile s_{-i} (of all the players other than player i). Consider first the case where a player interacts with all the other players, i.e., we are in the complete network. In this case, the payoffs to player i from action A given the strategy profile s_{-i} of other players, are as follows:

$$\Pi_i(A, s_{-i}) = n(A, s_{-i}) - c. \tag{11}$$

where $c > 0$ is the cost associated with action A, (or the costs of altruism). Similarly, the payoffs to player i from action E are given by

$$\Pi_i(E, s_{-i}) = n(A, s_{-i}). \tag{12}$$

Since $c > 0$, it follows that action E is strongly dominated by action A. So, if players are payoff optimizers (given the strategies of others) then they will never choose A. Thus it is necessary to have at least some players using alternative decision rules if there is to be any chance of action A being adopted. I follow Eshel, Samuelson and Shaked (1998), and assume that all players use a variant of the imitate the best action rule: each player compares the average payoffs from the two actions and chooses the action that attains the higher payoff. If all players choose the same action, in the current configuration, then a player follows this action.

I first note that if players are in the complete network then, in any mixed configuration, the average payoffs from choosing action E are higher. Thus if players follow the imitate the best action (on average) rule then the outcome where everyone chooses action E will obtain (unless we start with everyone choosing action A, which is an uninteresting case). This negative result on the prospects of action A leads us to consider the role of interaction structure.

By way of illustration, suppose therefore that players are located around a circle and that their payoffs depend on the choices of their immediate neighbours only.[9] Let $\{i-1, i, i+1\}$ be the neighbourhood of player i and suppose that the payoffs to player i are given by $n(A, s_{i-1}, s_{i+1}) - c$ if player 1 chooses A and by $n(A, s_{i-1}, s_{i+1})$ if she chooses action E. Here we have specialized the term $n(a, s_{i-1}, s_{i+1})$ to refer to the number of players who choose action a among the neighbours of player i. I shall now also make the model dynamic, and suppose that time is discrete and that in each period every player gets a chance to revise her strategy. Let the configuration at time t be denoted by s^t. The decision rules (along with an initial configuration, s^1) define a Markovian dynamic process where the states of the process

[9] For a related model of cooperative behaviour with local interaction in games with conflict, see Tieman, Houba and Van der Laan (2000). In their model, players are located on a network and play a generalized (many action) version of the prisoner's dilemma. They find that with local interaction and a TIT-for-TAT type decision rule, superior payoff actions which are dominated can survive in the population, in the long run.

are the strategy profiles s. The probability of transition from s to s' is either 0 or 1. I am interested in exploring the long run behavior of the dynamic process. Recall that a state (or a set of states) is said to be absorbing if the process cannot escape from the state once it reaches it. I also note that every absorbing state (or set of states) has associated with it a corresponding stationary distribution of the Markov process.

I begin the analysis of this model by clarifying the role of local interaction. To rule out uninteresting cases, I shall suppose that $c < 1/2$.[10] Suppose that there is a string of 3 players choosing action A, and they are surrounded on both sides by a population of players choosing E. Given the decision rule, any change in actions can only occur at the boundaries. What are the payoffs observed by the player choosing action A on the boundary? Well, she observes one player choosing E with a payoff of 1, while she observes one player choosing action A with payoff $2-c$. Moreover, she observes her own payoff of $1-c$, as well. Given that $c < 1/2$, it follows that she prefers action A. On the other hand, the player on the boundary choosing action E, observes one player choosing action E, with payoff 0, one player choosing action A with payoff $1-c$ and herself with a payoff 1. Given that $c < 1/2$, she prefers to switch to action A. This suggests that the region of altruists will expand. Note however that if everyone except one player is choosing action A, then the player choosing E will get a payoff of 2 and since this is the maximum possible payoff, this will induce her neighbours to switch to action E. However, as they expand, this group of egoists will find their payoffs fall (as the interior of the interval can no longer free ride on the altruists). These considerations suggest that a long string of players choosing action A can be sustained, while a long string of players choosing E will be difficult to sustain. These arguments are summarized in the following result, due to Eshel, Samuelson and Shaked (1998).

Theorem 4.4 *Absorbing sets are of two types: one, they contain a singleton state in which all players choose either action A or action E, and two, they contain states in which there are strings of A players of length 3 or more which are separated by strings of E players of length 2 or less (on average). In the latter case, at least* 60% *of the players choose action A (on average).*

[10] The restriction that $c < 1/2$ is made to ensure that it is attractive for players to switch from E to A under some configurations.

It is worth commenting on the relative proportions of A players in mixed configurations. First note that a string of E players cannot be of size 3 or longer (in an absorbing state). If it is then the boundary E players will each have two players choosing A on one side and a player choosing E who is surrounded by E players. It is easy to show that these boundary players will switch to A. Likewise, there have to be at least 3 players in each string of A players; otherwise, the boundary players will switch to E. These considerations yield the proportions mentioned in the result above.[11]

Given the above arguments, it is easily seen that a string of 5 players choosing A cannot shrink over time. If players strategies are randomly chosen initially then it follows that the probability of such a string of A players can be made arbitrarily close to 1, by suitably increasing the number of players. This idea is summarized in the following result, due to Eshel, Samuelson and Shaked (1998).

Theorem 4.5 *Suppose that players' initial strategy choices are determined by independent, identically distributed variables where the probability of each strategy is positive. Then the probability of convergence to an absorbing set containing states with at least* 60% *of A players goes to 1, as n gets large.*

This result suggests that the set of states from which the system moves to a majority altruistic society is relatively large. Eshel, Samuelson and Shaked (1998) also study stochastic stability of different absorbing sets. They show that the states identified in the above proposition are also the stochastically stable ones. In other words, we should expect a society to spend most of the time in a state where a large share of the players are choosing the altruistic action.

The findings reported above should be seen as part of an extensive literature on spatial evolution of social norms. This literature spans the fields of biology, computer science, philosophy, and political science, in addition to economics. *First*, I note that the idea of local emergence of cooperative norms and their gradual spread in a social space has been discussed by different authors (see e.g., Ullman-Margalitt (1977), Axelrod (1997).) The model and the arguments developed above should be seen as providing a formal account of this line of reasoning.

[11] The term within brackets, 'on average' refers to the possibility of a cycle between two states, one in which there is a string of 3 players choosing E and the other in which there is a 1-player string choosing E. In this case, there are on average 2 players choosing E.

Second, I would like to explore the scope of the above argument in explaining altruism in richer interaction structures. In the model above, the persistence and spread of altruism appears to be related to the presence of A players who are protected from E players and therefore secure sufficiently high payoffs so that other A players on the boundary persist with action A as well. In larger dimensional interaction (e.g., k-dimensional lattices) or asymmetric interaction (as in a star) this protective wall may be harder to generate and this may make altruism more vulnerable. For example, in the case of star, mixed configurations are not sustainable, and it seems very easy to transit from a purely altruistic society to a purely egoist society (via the switch by the central player alone). The reverse transition require switches by at least 3 players. Thus if interaction is in a star, then we should expect to see a society of egoists only, in the long run. These arguments suggest that the robustness of altruistic behavior needs to be explored further. Existing work on this subject seems to be mostly based on simulations (Nowak and May, 1992). This work suggests that in the absence of mutations, altruism can survive in a variety of interaction settings. There is also an extensive literature in evolutionary biology on the emergence and persistence of altruistic traits in different species. In this work the spread of altruistic traits is attributed to greater reproductive success. This success leads to the larger set of altruists spilling into neighbouring areas and this in turn leads to a growth of the trait over time (see e.g., Wynne-Edwards, 1986; Eshel and Cavalli-Sforza, 1982).

Third, I would like to discuss the scope of cooperation in repeated games when these games are played on networks. While the literature on repeated games is very large, it seems that there is relatively little on repeated games with local interaction. In a recent paper, Haag and Lagunoff (2000) examine a setting in which players play the repeated prisoners dilemma games with their immediate neighbours. They examine the network architectures which support high cooperation, when discount factors vary across players. Their main result shows that under some restrictions on strategies allowed to players a desirable interaction structure has the following properties: there is a clique of patient players (who are fully linked among themselves) each of whom is linked to a limited set of impatient players. In equilibrium, the clique of patient players will play cooperatively, while the impatient players at the periphery will defect. However, given their high patience level, and in a desire to sustain cooperation with their patient partners, the core group of players will tolerate this defection by peripheral players.

5 Evolving networks

So far I have been discussing the nature of learning in the context of a given network. Our discussion suggests that in both strategic as well as non-strategic contexts, the architecture of the network has an important impact on individual learning and social outcomes. This leads us to ask: which networks are plausible? In many contexts of interest, the links that individuals have place them at a relative advantage/disadvantage. It is therefore natural to examine the incentives of individuals to form links and the implications of such link formation for social and economic interaction. This is the principal motivation for the recent research on the theory of network formation.

The process of learning in a setting where the network itself is evolving is complicated. To get a first impression of some of the issues that arise, it seems desirable to proceed in steps. In this section I will focus on the pure network formation problem: individuals learn about the optimality of links with different individuals and revise their choices in response. This leads to an evolving network. In the next section, I will examine learning in a setting where individuals decide on links as well as interact strategically with those players with whom they form links.

I will focus on the model of connections which has been extensively studied in the literature. This model is based on the notion that social networks are formed by individual decisions that trade off the costs of forming and maintaining links against the potential rewards from doing so. A link with another individual allows access, in part and in due course, to the benefits available to the latter via her own links. Thus links generate externalities and define the economic background for the network formation process. As before, let $N = \{1, ..., n\}$, with $n \geq 3$, be the set of players and let i and j be typical members of this set. A strategy of player $i \in N$ is a (row) vector $g_i = (g_{i,1}, ..., g_{i,i-1}, g_{i,i+1}, ..., g_{i,n})$ where $g_{i,j} \in \{0, 1\}$ for each $j \in N \backslash \{i\}$. I say player i has a *link* with j if $g_{i,j} = 1$. In this framework, links are *one-sided* in the sense that they can be formed on individual initiative and (as will become clear shortly) this individual pays for the costs of forming links as well. This approach to link formation was developed in Bala and Goyal (2000a).[12] A natural interpretation of links is

[12] The static model of one-sided links was introduced in Goyal (1993), while the dynamics were introduced in Bala (1997). Bala and Goyal (2000a) subsumes these earlier individual attempts.
Jackson and Wolinsky (1996) developed a closely related model of connections in which a link requires the consent of both players involved. This model is presented later in this section.

that they are information channels. A link between player i and j can allow for either one-way (asymmetric) or two-way (symmetric) flow of information. With one-way communication, the link $g_{i,j} = 1$ enables player i to access j's information, but not vice-versa. For example, i could access j's website, or read a paper written by j. With two-way communication, a link $g_{i,j}$ enables i and j to share information. An example of this is a telephone call between two players. A second interpretation is that a link reflects a social relation, which involves the giving of gifts and reciprocal favours. The set of (pure) strategies of player i is denoted by $\mathcal{G}_i$. Since player i has the option of forming or not forming a link with each of the remaining $n-1$ players, the number of strategies of player i is clearly $|\mathcal{G}_i| = 2^{n-1}$. The set $\mathcal{G} = \mathcal{G}_1 \times .... \times \mathcal{G}_n$ is the space of pure strategies of all the players.

The link $g_{i,j} = 1$ is represented by an *edge* starting at j with the arrow-head pointing at i. Figure 1 above provides an example with $n = 3$ players. Here player 1 has formed links with players 2 and 3, player 3 has a link with player 1 while player 2 does not link up with any other player. Note that there is a one-to-one correspondence between the set of all directed networks with n vertices and the set $\mathcal{G}$.

With a slight abuse of our earlier notation I can say that $N^d(i; g) = \{k \in N \mid g_{i,k} = 1\}$ is the set of players with whom i *maintains* a link. The notation "$j \xrightarrow{g} i$" indicates that there exists a path from j to i in g. Furthermore, I define $N(i; g) = \{k \in N \mid k \xrightarrow{g} i\} \cup \{i\}$. This is the set of all players whom i accesses either through a direct link or through a sequence of links. Recall that $\mu_i^d(g) \equiv |N^d(i; g)|$ and $\mu_i(g) \equiv |N(i; g)|$ for $g \in \mathcal{G}$.

I wish to model a situation where it is advantageous to access a larger number of people, and where links are costly to maintain. I will first discuss the case of one-way flow of benefits. Denote the set of non-negative integers by Z_+. Let $\phi : \mathcal{Z}_+^2 \to \mathcal{R}$ be such that $\phi(x, y)$ is strictly increasing in x and strictly decreasing in y. Define each player's payoff function $\Pi_i : \mathcal{G} \to R$ as

$$\Pi_i(g) = \phi(\mu_i(g), \mu_i^d(g)). \tag{13}$$

Given the properties I have assumed for the function ϕ, $\mu_i(g)$ can be interpreted as providing the "benefit" that player i receives from her links, while $\mu_i^d(g)$ measures the "cost" associated with maintaining them.

A special case of (13) is when payoffs are linear. To define this, I specify two parameters $V > 0$ and $c > 0$, where V is regarded as the *value* of each player's information (to himself and to others), while c is her *cost* of link formation. Without loss of generality, V can be normalized to 1. I now define $\phi(x, y) = x - yc$, i.e.

$$\Pi_i(g) = \mu_i(g) - \mu_i^d(g)c. \tag{14}$$

In other words, player i's payoffs are the number of players she observes less the total cost of link formation. I identify three parameter ranges of importance. If $c \in (0, 1)$ then player i will be willing to form a link with player j for the sake of j's information alone. When $c \in (1, n-1)$, player i will require j to observe some additional players to induce him to form a link with j. Finally, if $c > n - 1$ then the cost of link formation exceeds the total benefit of information available from the rest of society. Here, it is a dominant strategy for i not to form a link with any player.

Given a network $g \in \mathcal{G}$, let g_{-i} denote the network obtained when all of player i's links are removed. The network g can be written as $g = g_i \oplus g_{-i}$ where the '$\oplus$' indicates that g is formed as the union of the links in g_i and g_{-i}. Recall that a strategy profile g^* is a Nash equilibrium if $\Pi_i(g_i^* \oplus g_{-i}^*) \geq \Pi_i(g_i \oplus g_{-i}^*)$, for all $g_i \in \mathcal{G}_i$, and for all $i \in N$. A *strict* Nash equilibrium is a Nash equilibrium in which each player gets strictly higher payoffs with her current strategy than he would with any other strategy.

I first study the static network formation problem. Recall that a wheel is a network in which each player forms exactly one link, represented by an arrow pointing to the player. The following result, due to Bala and Goyal (2000a), provides a complete characterization of equilibrium networks in the above model.

Theorem 5.1 *Let the payoffs be given by (13). Then a strict Nash network is either the wheel or the empty network.*

In particular, in the context of the linear model, the wheel is the unique equilibrium network if $c < 1$, the wheel and the empty network are the only equilibria for $1 < c < n - 1$, while the empty network is the unique network for $c > n - 1$. I now provide a sketch of the main arguments underlying the above theorem. The *first* step in the proof is to show that a Nash network is either connected or empty. Consider a non-empty Nash network, and suppose that player i is the player who observes the largest number of players in this network. Suppose i does not observe everyone. Then there is some player j who is not observed by i and who does not observe i (for otherwise j would observe more players than i). Since i gets value from her links, and payoffs are symmetric, j must also have some links. Let j deviate from her Nash strategy by forming a link with i alone. By doing so, j will observe strictly more players than i does, since she has the additional benefit of observing i. Since j was observing no more players than i in her original strategy, j increases her payoffs by her deviation. This contradiction implies that i must observe every player in the society. It then follows that every other player will have an incentive to either link with i or to observe him through a sequence of links, so that the network is connected. If the network is not minimally connected, then some player could delete a link and still observe all players, which would contradict Nash. The *second* step exploits the refinement of strictness and is based on the following observation: if two players i and j have a link with the same player k, then one of them (say) i will be indifferent between forming a link with k or instead forming a link with j. We know that Nash networks are either connected or empty. This means that in the one-way flow model a (non-empty) strict Nash network has exactly n links. Since the wheel is the unique such network, the result follows.

Theorem 5.1 shows that individual incentives restrict the range of possible network architectures quite dramatically. This characterization of equilibrium networks, however, raises the issue: how do individuals choose links if they start with some different network and is there some pressure moving them to the equilibrium networks identified above. In other words, will individuals learn to coordinate their links and arrive at a wheel network?

Bala and Goyal (2000) introduced the study of dynamics in the formation of networks and I will follow their approach here. They used a variant of the myopic best response dynamic to study the above question. Two features of the process are important: one, there is some probability that an individual exhibits *inertia*, i.e., chooses the same strategy as in the previous period. This ensures that players do not perpetually miscoordinate. Two, if more than one strategy is optimal for some individual then she *randomizes* across the optimal strategies. This requirement implies, in particular, that a non-strict Nash network can never be a steady state of the dynamics. The rules on individual behavior define a Markov chain on the state space of all networks; moreover, the set of absorbing states of the Markov chain coincides with the set of strict Nash networks of the one-shot game. The following result, due to Bala and Goyal (2000a), shows that the learning process converges to the equilibrium networks identified above.

Theorem 5.2 *Let the payoff functions be given by equation (13) and let g be some initial network. Then the dynamic process converges to the wheel or the empty network in finite time, with probability 1.*

In the context of the linear model, the process converges to the wheel for $0 < c < 1$, the wheel or the empty network for $1 < c < n-1$ and the empty network for $c > n-1$. The proof of the above theorem exploits the idea that well connected individuals generate positive externalities. Fix a network g and suppose that there is a player i who accesses all people in g, directly or indirectly. Consider a player j furthest away from i in the network g, in other words $j \in \text{argmax}_{k \in N} d_{i,k}(g)$. This also means that player j is not critical for player i in the network g, i.e. player i is able to access everyone even if player j deletes all her links. (It is easy to see that there will always exist such a player j.) Allow player j to move; she can form a single link with player i and access all the different individuals accessed by player i. Thus if forming links is at all profitable for player j then one best-response strategy is to form a single link with player i. This strategy in turn makes player j well-connected. We now consider some person k who is not critical for j and apply the same idea. Repeated application of this argument leads to a network in which everyone accesses everyone else via a single link, i.e. a wheel network.

I now consider the case where the flow of benefits is two-way. It this case, benefits flow between two players so long as one of the two has formed with the link with the other. To capture this two-way flow I define $\hat{g}_{i,j} = \max\{g_{i,j}, g_{j,i}\}$. The link $g_{i,j} = 1$ is represented by an *edge* between i and j: a filled circle lying on the edge near player i indicates that it is this player who has initiated the link. Figure 4 below depicts the example of Figure 1 for the two-way model. As before, player 1 has formed links with player 2 and 3, player 3 has formed a link with agent 1 while agent 2 does not link up with any other agent.[13] Given the strategy profile g it is now straightforward to define a network $\hat{g}$, using the above operation. Every strategy profile g has a unique representation in the manner shown in the figure.

[13]Since players choose strategies independently of each other, two players may simultaneously initiate a two-way link, as seen in the figure.

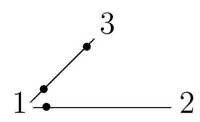

Figure 4

I can extend the notion of path as follows: there is a *tw-path* (for two-way) in g between i and j if either $\hat{g}_{i,j} = 1$ or there exist agents $j_1, \ldots, j_m$ distinct from each other and i and j such that $\hat{g}_{i,j_1} = \ldots.. = \hat{g}_{j_m,j} = 1$. I write $i \xrightarrow{\hat{g}} j$ to indicate a tw-path between i and j in g. Let $N^d(i;g)$ and $\mu_i^d(g)$ be defined as in the earlier model above. The set $N(i;\hat{g}) = \{k \mid i \xrightarrow{\hat{g}} k\} \cup \{i\}$ consists of agents that i observes in $\hat{g}$ under two-way communication, while $\mu_i(\hat{g}) \equiv |N(i;\hat{g})|$ is its cardinality. In the two-way flow model, the pay-off to player i with a strategy g_i, faced with a profile g_{-i} is given by

$$\hat{\Pi}_i(g) = \phi(\mu_i(\hat{g}), \mu_i^d(g)). \tag{15}$$

where $\phi(.,.)$ is as in the one-way flow model. The case of linear payoffs is $\phi(x,y) = x - yc$ as before. I obtain, analogously to (14):

$$\hat{\Pi}_i(g) = \mu_i(\hat{g}) - \mu_i^d(g)c. \tag{16}$$

The parameter ranges $c \in (0,1)$, $c \in (1, n-1)$ and $c > n-1$ have the same interpretation as above. A centre-sponsored star is a star in which the centre forms and hence pays for all the links. The following result, due to Bala and Goyal (2000a), provides a complete characterization of the architecture of strict Nash networks in the two-way flow case.

Theorem 5.3 *Let the payoffs be given by (15). A strict Nash network is either a centre-sponsored star or the empty network.*

In particular, in the linear model given above, the centre-sponsored star is the unique (strict) equilibrium network for $0 < c < 1$, while the empty network is the unique equilibrium for $c > 1$. I now sketch the main arguments underlying the proof. The first step in the proof is to show that a Nash network is either empty or connected. The second step in the proof exploits the following observation: if player n has a link with j then no other player can have a link with j. The idea behind this is that if two agents i and j have a link with the same agent k, then one of them (say) i will be indifferent between forming a link with k or instead forming a link with j. As a Nash network is connected, this implies that n must be the center of a star. A further implication of the above observation is that every link in this star must be formed or "sponsored" by the center. The dynamics of the two-way flow model are well behaved. The following result, due to Bala and Goyal (2000a), provides a characterization.

Theorem 5.4 *Let the payoff functions be given by equation (15) and let g be some initial network. If $\phi(x+1, y+1) > (<) \phi(x, y)$ for all $y \in \{0, 1, ..., n-2\}$ and $x \in \{y+1, y+2, ..., n-1\}$ then the dynamic process converges to the center-sponsored star (empty network) in finite time, with probability 1.*

The above model of connections supposes that links can be formed by single individuals. In many economic applications – such as two firms collaborating on a project or two countries signing a bilateral trade agreement – it is natural to suppose that the formation of a link requires the acquiscence of both the players who are directly involved. This leads us to a model of two-sided link formation. Jackson and Wolinsky (1996) develop a general model of two-sided link formation. They also develop a solution concept for such games: *pair-wise stability.* A network g is said to be pair-wise stable if no individual has an incentive to delete any link that exists in the network and no pair of players has an incentive to form a link that does not exist in the network. I now present a variant of the above model of connections to illustrate the role of different link formation assumptions. I will then turn to the issue of learning and dynamics in two-sided networks.

I follow Jackson and Wolinsky (1996) in the following exposition. A link is two-sided and the flow of information or value is also two-way. A link is denoted by $\bar{g}_{i,j} \in \{0, 1\}$, where it is assumed that $\bar{g}_{i,j} = \bar{g}_{j,i}$. A network $\bar{g} = (\{\bar{g}_{i,j}\})_{i,j \in N}$; the space of all networks $\bar{\mathcal{G}}$, is equivalent to the set of all undirected networks on n vertices. I can extend the notions of path, component, distance and connectedness to this set of networks in a natural manner. I suppose that the value of every player is 1, and I shall also assume that there is decay in information or value in indirect links which is related to the distance between players. Let $d_{i,j}(\bar{g})$ refer to the geodesic distance between two players i and j, in network $\bar{g}$ and let $\mu_i^d(\bar{g})$ refer to the number of players with whom player i forms a link in network g. The decay rate is summarized in a number $\delta \in [0, 1)$. Given this notation, I can now state the payoffs to player i in a network $\bar{g}$:

$$\Pi_i(\bar{g}) = 1 + \sum_{j \in N(i;\bar{g})} \delta^{d_{i,j}(\bar{g})} - \mu_i^d(\bar{g})c. \tag{17}$$

The following result, due to Jackson and Wolinsky (1996), offers a partial characterization of pair-wise stable networks. In the following result I use the shortened expression *pw-stable* network to refer to pair-wise stable network.

Theorem 5.5 *Suppose payoffs are given by (17). A pw-stable network has at most one non-singleton component. For $c < \delta - \delta^2$, the unique pw-stable network is the complete network g^c. For $\delta - \delta^2 < c < \delta$, a star is a pw-stable network (it is not the unique stable one). For $\delta < c$, the empty network is pw-stable and any pw-stable network which is non-empty is such that each player has least two links.*

The argument underlying the first part relies on the symmetry of the game.[14] Suppose that $\bar{g}$ is a pair-wise stable network and that there are two non-singleton components with (say) players i and j having a link in one component while players k and l have a link in the other component. From the definition of stability it follows that each of the players are better off with the link. Given the payoff function (17) it follows then that player i and player l would have a strict incentive to form a link with each other, thus contradicting the hypothesis of stability. The rest of the proposition is straightforward to verify.

As in the one-sided link model, I now explore the dynamics of network formation in the two-sided link formation case. I suppose that in each period one pair of players is randomly picked and has a choice of forming (or not forming) a link. If a link is already present then the players can decide whether to sever it. If no link exists, they can decide to form a link and at the same time each of the players can also delete any subset of the links that they currently maintain (so long as both players agree to this). In line with most of the work reported in this survey, I assume that they make decisions on the basis of myopic payoff

[14] In case of $\delta = 1$ the following result holds: for $0 < c < 1$, a pair-wise stable network is minimally connected; the star and the line network are two possible stable networks. If $c > 1$ then the empty network is uniquely stable.

maximization. The following result, due to Watts (2001), presents a partial characterization of the behaviour of the above dynamic process.

Theorem 5.6 *If $0 < c < \delta - \delta^2$ then every link will form and the process converges to the complete network in finite time, with probability 1. If $\delta - \delta^2 < c < \delta$, then there is a positive probability that the star network will emerge. However, this probability is decreasing in the number of players and goes to 0 as the number of players gets large. If $c > \delta$ then no link will be formed and the network remains empty, with probability 1.*

The arguments underlying the first and third statements follow directly from the assumption on parameters and the myopic decison rule. The argument for the formation of the star is as follows: fix a player 1, and get each of the other players, starting from 2, 3, and so on until n, to have an opportunity to form links with player 1. It follows from the facts that $c < \delta$, the initial network is empty and that players are myopic that every pair of players who have an opportunity to form a link will do so. The star emerges at the end of period $n-1$ and given the rules about matching, this pairing sequence occurs with positive probability. The result on the decreasing probability of a star with respect to number of players exploits the following observation: for the star to form, every player $j \neq 1$ must meet player 1 before she meets any other player. This is true because if (say) players 1 and 2 meet in period 1 and players 3 and 4 meet in period 2, then at the end of period 2, there will be two linked pairs. Now, suppose player 1 meets player 3 in period 3. Clearly, they will form a link, under the assumption that $c < \delta$ and players are myopic. To get a star with player 1 at the center, we must have players 1 and 4 meet and form a link. However, if they meet, player 1 will not agree to form a link with player 4, since the net payoff $(\delta - c) - \delta^2 < 0$. The proof follows by noting that the probability that every player meets player 1 before meeting any other player goes to 0 as n gets large.

The theory of network formation is a very active field of research currently; in addition to the papers mentioned above, recent work includes Aumann and Myerson (1989), Boorman (1975), Calvo (2000), Corominas (1998), Dutta, van den Nouweland and Tijs (1995), Goyal and Joshi (2002), Goyal and Moraga (2001), Kranton and Minehart (2001). Recent book length treatments of this work are contained in Dutta and Jackson (2001) and Slikker and van den Nouweland (2001).

I briefly discuss the issues related to learning that are being explored. *One* concern has been the rate of convergence of the dynamics. This is an important issue since the state space of the dynamic process is large. There are $2^{n(n-1)}$ networks with n agents; so, for instance, in a game with 8 players, there are $2^{56} = 7 \times 10^{16}$ possible directed networks, which suggests that a slow rate of convergence is a real possibility. The dynamics are Markovian and so convergence to the limiting distribution is exponential. As noted above in the section on coordination games, however, this allows for a wide range of speeds of convergence. The rate of convergence is studied using simulations in Bala and Goyal (2000a). They find that both in the one-way flow model as well as the two-way flow model discussed above, the dynamics converge rapidly to the limit network. For instance, in a game with 8 players, the average time to convergence is less than 250 periods in the one-way flow model, while the average time to convergence is 28 periods in the two-way model. (These numbers are for a high probability of strategy revision in every period; the rates of convergence remain fast as this probability is varied.)

A *second* concern is the relation between individual incentives and social efficiency. In some models (the one-sided connections model is an example) equilibrium network architectures coincide with social efficient ones, while in others (the two-sided connections model is an example) there is a conflict between individual incentives and social efficiency. This has motivated a general exploration of the set of circumstances under which there is a conflict. Most of this research has been done in the two-sided link formation framework; Jackson (2001) provides a survey of this body of work.

A *third* concern has been the value of indirect links and the nature of decay as distance between players increases in a network. The discussion of the Jackson and Wolinsky (1996) model reflects this idea as the value of connections decline exponentially with respect to distance. A similar formulation of decay has been studied in the one-sided link framework by Bala and Goyal (2000a) as well. More generally, the issue of decay is related closely to idea of spillovers and externalities. (Notice that $\delta = 0$ corresponds to the case where indirect links are worthless and hence there are no externalities in network formation.) The model of connections discussed above reflects positive spillovers and is not appropriate for studying settings where congestion and competition for partners are central. The study of such models remains an open problem.

Some other issues that have been explored are heterogeneity in valuations and costs of forming links (Galeotti and Goyal, 2002), imperfect reliability of links (Bala and Goyal 2000b; Haller and Sarangi, 2001), multiple equilibrium (Jackson and Watts, 2001), link formation by far-sighted players (Aumann and Myerson, 1989; Currarini and Morelli, 2000; Dutta, Ghosal and Ray 2002), and imperfect information on structure of network (McBride, 2002).

6 Optimal links and actions in games

I now turn to an exploration of models in which individuals choose links (and thereby shape the network) and also choose actions in strategic games they play with those they have formed connections. There are a variety of economic examples that fit naturally in this framework. Firms collaborate on research with each other as well as compete in the product market subsequently, while individuals invest in relationships and then choose actions in the social interaction (the action could be whether to smoke or not, to indulge in criminal activity or not or pertain to the choice of learning different softwares or languages). As before, I will consider both coordination games and games of conflict.

6.1 Coordination games

I present a simple model in which players choose their partners and choose an action in the coordination games they play with their different partners. This framework allows us to endogenize the nature of interaction and study the effect of partner choice on the way players coordinate their actions in the coordination game.

The issue of endogeneous structures of interaction and the impact of endogeneity on coordination choices has been explored in early papers by Bhaskar and Vega-Redondo (2002), Ely (1996), Mailath, Samuelson and Shaked (1994), and Oechssler (1997). They use a framework in which players are located on islands. Moving from an island to another implies severing all ties with the former island and instead playing the game with all the players in the new island. Thus neighborhoods are endogenized via the choice of islands. In my exposition I will follow some of the more recent papers in this subject, since they fit in more naturally with the framework that I am using in this the survey. I will briefly discuss the relationship between the different approaches later in this section.

As before I shall suppose that $N = \{1, 2, \ldots, n\}$ is the set of players, where $n \geq 3$. Each player has a strategy $s_i = \{g_i, a_i\} \in \mathcal{G}_i$, where g_i refers to the links that she forms while $a_i \in A_i$ refers to the choice of action in the accompanying coordination game. Any profile of link decisions $g = (g_1, g_2 \ldots g_n)$ defines a directed network. Given a network g, I say that a pair of players i and j are directly linked if at least one of them has established a linked with the other one, i.e. if $\max\{g_{ij}, g_{ji}\} = 1$. To describe the pattern of players' links, I shall take recourse to our earlier notation and define $\hat{g}_{ij} = \max\{g_{ij}, g_{ji}\}$ for every pair i and j in N. I refer to g_{ij} as an active link for player i and a passive link for player j. I will say that a network g is essential if $g_{ij}g_{ji} = 0$, for every pair of players i and j. Also, let $G^c(M) \equiv \{g : \forall i, j \in M,\ \hat{g}_{ij} = 1,\ g_{ij}g_{ji} = 0\}$ stand for the set of complete and essential networks on the set of players M. Given any profile $s \in S$, I shall say that $s = (g, a) \in S^h$ for some $h \in \{\alpha, \beta\}$ if $g \in G^c$ and $a_i = h$ for all $i \in N$. More generally, I shall write $s = (g, a) \in S^{\alpha\beta}$ if there exists a partition of the population into two subgroups, N^α and N^β (one of them possibly empty), and corresponding components of g, g^α and g^β, such that: (i) $g^a \in G^c(N^\alpha)$, $g^\beta \in G^c(N^\beta)$; and (ii) $\forall i \in N^\alpha,\ a_i = \alpha$; $\forall i \in N^\beta$, $a_i = \beta$.

Individuals located in a social network play a 2×2 symmetric game with a common action set. The set of partners of player i depends on her location in the network. In particular, I assume that player i plays a game with all other players in the set $N^d(i; \hat{g})$. The bilateral game is the same as the coordination game discussed in section 3 above. Recall that there is a common set of actions $A = \{\alpha, \beta\}$. The payoffs in the single game are as follows:

1 \ 2	α	β
α	a, a	d, e
β	e, d	b, b

Figure 5

I will assume that the restrictions stated in (7) hold. Recall that these restrictions imply that there are two pure strategy equilibria of the game: (α, α) and (β, β) and that coordinating on either of them is better than not coordinating at all. I shall also assume that $b \geq a$ in what follows.

Every player who establishes a link with some other player incurs a cost $c > 0$. Given the strategies of other players, $s_{-i} = (s_1, \ldots s_{i-1}, s_{i+1}, \ldots s_n)$, the payoffs to a player i from playing some strategy $s_i = (g_i, a_i)$ are given by:

$$\Pi_i(s_i, s_{-i}) = \sum_{j \in N^d(i;\hat{g})} \pi(a_i, a_j) - \mu^d(i; g) \cdot c \tag{18}$$

I make two remarks about this formulation. One, individual payoffs are aggregated across the games played by him. This seems appropriate in the present setting since the number of games an individual plays is endogenous, I want to explicitly account for the influence of the size of the neighborhood and thus choose the aggregate-payoff formulation. Two, the description of strategies and the paypoff formulation reflects the assumption that every player i is obliged to choose the same action in the (possibly) several bilateral games that she is engaged in. This assumption is natural in the present context: if players were allowed to choose a different action for every two-person game they are involved in, this would make the behaviour of players in any particular game insensitive to the network structure.

I start with the following result, due to Goyal and Vega-Redondo (2002), which provides a complete characterization of equilibrium outcomes.

Theorem 6.1 *Suppose payoffs satisfy (7) and $b > a$. (a) If $c < d$ then the set of equilibrium profiles $S^* = S^\alpha \cup S^\beta$, (b) if $d < c < a$, then $S^\alpha \cup S^\beta \subset S^* \subset S^{\alpha\beta}$, the first inclusion being strict for large enough n, (c) if $a < c < b$, then $S^* = S^\beta \cup \{(g^e, (\alpha, \alpha, ..., \alpha))\}$, (d) if $c > b$, then $S^* = \{g^e\} \times A^n$.*

I provide a sketch of the main arguments underlying this result. The first step derives restrictions on equilibrium network architectures implied by individual incentives. If costs of link formation are low (say $c < e$) then a player has an incentive to link up with other players irrespective of the actions the other players are choosing. On the other hand, when costs are quite high (specifically, $a < c < b$) then everyone who is linked must be choosing the efficient action. This, however, implies that it is attractive to form a link with every other player and I get the complete network again. In contrast, if costs are at an intermediate level ($d < c < a$), a richer set of configurations is possible. On the one hand, since $c > d(> e)$, the link formation is only worthwhile if other players are choosing the same action. On the other hand, since $c < a(< b)$, coordinating at either of the two equilibria (in the underlying coordination game) is better than not playing the game at all. This allows for networks with two disconnected components in equilibria. In view of these considerations, parts (a) and (c) follow quite directly. I now elaborate on the coexistence equilibria identified in part (b). In these equilibria, there are two unconnected groups, with each group adopting a common action (different in each group). The strategic stability of this configuration rests on the appeal of 'passive' links. A link such as $g_{ij} = 1$ is paid for by player i, but both players i and j derive payoffs from it. In a mixed equilibrium configuration, the links in each group must be, roughly, evenly distributed. This means that all players enjoy some benefits from passive links. In contrast, if a player were to switch actions, then to derive the full benefits of this switch, she would have to form (active) links with everyone in the new group. This lowers the incentives to switch, a consideration which becomes decisive if the number of passive links is large enough (hence the requirement of large n).

The above result says that individual incentives restrict the range of network architectures quite sharply in this setting and this has a bearing on the extent of heterogeneity that is permissible. However, there is a residual multiplicity: in parts (a) and (c), this multiplicity permits alternative states where either of the two actions is homogeneously chosen by the whole population, while in part (b), the multiplicity allows for a wide range of possible states where neither action homogeneity nor full connectedness necessarily prevails. Are these outcomes all equally stable? This question leads me to examine the stochastic stability of different equilibria.

Time is discrete $t = 1, 2, 3, \ldots$. At each t, the state of the system is given by the strategy profile $s(t) \equiv [(g_i(t), a_i(t))]_{i=1}^{n}$ specifying the action played, and links established, by each player $i \in N$. I will suppose that the decision rules are the same as in the one-sided link formation model described above in section 5. In every period, there is a positive independent probability that any given individual gets a chance to revise her strategy. If she receives this opportunity, I assume that she chooses a best-response to what other players chose in the preceding period. If there are several strategies that fulfill then any one of them is taken to be selected with, say, equal probability. This strategy revision process defines a simple Markov chain on $S \equiv S_1 \times ... \times S_n$. I then define the perturbed dynamics as in section 4, with a small "mutation" probability $\epsilon > 0$. For any $\epsilon > 0$, the process defines a Markov chain that is aperiodic and irreducible and, therefore, has a unique invariant probability distribution. Let us denote this distribution by μ_ϵ. I analyze the form of μ_ϵ as the probability of mistakes becomes very small, i.e. formally, as ϵ converges to zero. Define $\lim_{\epsilon \to 0} \mu_\epsilon = \hat{\mu}$. When a state $s = (s_1, s_2, \ldots, s_n)$ has $\hat{\mu}(s) > 0$, i.e. it is in the support of $\hat{\mu}$, I say that it is *stochastically stable.*

The following result, due to Goyal and Vega-Redondo (2002), provides a complete characterization of stochastically stable social networks and actions in the coordination game.

Theorem 6.2 *Suppose (7) and* $b > a$. *There exists some* $\bar{c} \in (e, a)$ *such that if* $c < \bar{c}$ *then* $\hat{S} = S^\beta$ *while if* $\bar{c} < c < b$ *then* $\hat{S} = S^\beta$. *Finally, if* $c > b$ *then* $\hat{S} = \{g^e\} \times A^n$.

This result illustrates that the *dynamics* of link formation play a crucial role in the model. I observe that the only architecture that is stochastically stable (within the interesting parameter range) is the complete one, although players' behavior in the coordination game is *different* depending on the costs of forming links. However, if the network were to remain fixed throughout, standard arguments indicate that the risk-dominant action must prevail in the long run (cf. Kandori, Mailath and Rob, 1993). Thus it is the link formation *process* that, by allowing for the *co*-evolution of the links and actions, shapes individual behavior in the coordination game.

I now briefly provide some intuition on the sharp relationship found between the *costs* of forming links and the corresponding behavior displayed by players in the coordination game. On the one hand, when the cost of forming links is small, players wish to be linked with everyone irrespective of the actions they choose. Hence, from an individual perspective, the relative attractiveness of different actions is quite *insensitive* to what is the network structure faced by any given player at the time of revising her choices. In essence, a player must make her fresh choices as if she were in a complete network. In this case, therefore, the risk-dominant (and possibly) inefficient convention prevails since, under complete connectivity, this convention is harder to destabilize (through mutations) than the efficient but risk-dominated one. By contrast, if costs of forming links are high, individual players choose to form links only with those who are known (or perceived) to be playing the same action. This lowers the strategic uncertainty in the interaction and thus facilitates the emergence of the efficient action.

I comment on two features of the model: the one-sided link formation and the simultaneous choice of links and actions in the coordination game. It is possible to show that the main result on stochastically stable networks and the relation between costs of link formation and coordination remains essentially unchanged if link formation is two-sided and costs are borne by both players. However, a model with two-sided links and sequential choice of links and actions yields quite different outcomes. This formulation has been explored by Jackson and Watts (2001). They consider a two-sided link formation model, in which pairs of players are given an opportunity to form links. They make a decision on forming links under the assumption that their actions (in the corresponding coordination game) will remain the same as in the previous period. Once the network is in place, a player is chosen at random to choose an action in the cordination game. This player chooses an action that maximizes current payoff, given the current network and the actions of players in the previous period. The choices of players can be perturbed/randomly altered with small probability and I look for stochastically stable outcomes. The following result, due to Jackson and Watts (2002), provides a complete characterization of stochastically stable outcomes in this setting.

Theorem 6.3 *Suppose that (7) hold and $a < b$. Suppose that link formation is two-sided and choice of links and actions in the coordination game is sequential. (a) If $0 < c < e$ then a complete network with everyone choosing α is the unique stochastically stable outcome, (b) if $e < c < a$ then a complete network with everyone choosing either α or β can be the stochastically stable outcome, and (c) if $a < c < b$ then a complete network with everyone choosing β or the empty network with everyone choosing α are the stochastically stable outcomes.*

The ideas behind parts (a) and (c) are close to the ones mentioned in the case of one-sided links. The impact of the sequentiality of links and actions is reflected clearly in part (b) and I discuss this result now. Suppose we are a complete network with everyone choosing action α. Now let there be 2 trembles which lead to two players switching actions to β. Then players choosing action α would like to delete their links with these two players, and a component of 2 players choosing β arises. This component can grow with single mutations to β until it takes over the population. Jackson and Watts show that the process of transition from an all β state to an all α state is symmetric and so the complete network with everyone choosing action α or β are both stochastically stable outcomes. In this process, notice that players who have trembled from (say) α to β cannot switch actions and simultaneously offer to form links with the other players. If this simultaneous change was possible then the players will want to switch actions and form links with the other $n-2$ players and these other players will accept these links (under the parameter conditions given and the process will revert back to an all α state). Thus the assumption of sequentiality of link formation and actions choice is central to the above result.

In a closely related paper, Droste, Gilles and Johnson (2001) explore a model of networks and coordination in which there is an ex-ante spatial structure with all players located around a circle and the costs of link formation being higher for players who are further away. In this case, long run outcomes have richer spatial interaction structures but the risk-dominant action prevails in the interesting parameter ranges.

I now discuss the connections between the above results and the earlier work of Ely (1996), Mailath, Samuelson and Shaked (1994), Oechssler (1997), and Bhaskar and Vega-Redondo (2002). The basic insight flowing from the earlier work is that, if individuals can easily separate/insulate themselves from those who are playing an inefficient action (e.g., the risk-dominant action), then efficient "enclaves" will be readily formed and eventually attract the "migration" of others who will adopt the efficient action eventually. One may be tempted to identify *easy* mobility with *low* costs of forming links. However, the considerations involved in the two approaches turn out to be very different. This is evident from the differences in the results: recall that in the network formation approach, the risk-dominant outcome prevails if the costs of forming links are small. There are two main reasons for this contrast. First, in the network formation approach, players do not *indirectly* choose their pattern of interaction with others by moving across a *pre-specified* network of locations (as in the case of player mobility). Rather, they construct *directly* their interaction network (with no exogenous restrictions) by choosing those agents with whom they want to play the game. Second, the cost of link formation is paid per link formed and thus becomes truly effective only if it is high enough. Thus it is precisely the restricted "mobility" that high costs induce which helps insulate (and thus protect) the individuals who are choosing the efficient action.

If the costs of link formation are low, then the extensive interaction this facilitates may have the unfortunate consequence of rendering risk-dominance considerations decisive.

I now briefly discuss some other issues that have been studied in the literature. In the above framework, individuals choose links and actions to maximize current payoffs. An alternative formulation consisting of reinforcement in link formation along with the imitation of the highest payoff action has been studied in a recent paper by Skyrms and Pemantle (2000). They study the stag game (in which the payoffs to a player from the risk-dominant action are independent of the choice of the partner). The dynamics of links are driven by reinforcement: a link becomes likely in the future if the current experience is positive. They find that if the speed of adjustment of actions is relatively slow as compared to speed at which reinforcement works on links, then all players converge to the efficient (but risk-dominated) action. On the other hand, the likelihood of choosing the efficient action declines quite sharply as the speed of switching actions increases. These are interesting findings and seem to be in line with the earlier work on mobility reported above, where efficiency in coordination games becomes more likely as mobility – in the sense of switching partners – becomes easier.

I conclude this section by briefly discussing some work on learning of optimal links and actions in games of congestion. Recall these are two player two action games in which there are two pure strategy equilibria $\{\alpha, \beta\}$ and $\{\beta, \alpha\}$. Bramoulle, Lopez, Goyal and Vega-Redondo (2002) study these games. They show that the density of the network varies inversely with respect to the costs of forming links: for low costs the equilibrium network is complete, for moderate costs the equilibrium network is a bipartite graph, while for high costs, the equilibrium network is empty. Moreover, the relative proportions of the individuals choosing the different actions depends crucially on the cost of forming links. For low costs, only the complete and essential graph with a unique proportion of individuals doing each action arises in equilibrium. However, for moderate costs there is a wide variety of proportions that can arise in equilibrium. These proportions have very different welfare properties and typically equilibrium networks are not efficient. Finally, they find that in contrast to the coordination game analysis above, all equilibria of the game are stochastically stable.

6.2 Games of conflict

There appears to be very little formal work on network formation and games of conflict.[15] In this section, I discuss a recent paper by Vega-Redondo (2002); this paper raises a variety of issues and also employs techniques slightly different than those I have used in the survey so far, so my discussion of this paper will be brief. Suppose players can form links as well as choose strategies in an infinitely repeated prisoner's dilemma. There are two additional features of the model: payoffs change over time and players are allowed the freedom to choose different actions across the bilateral games they engage in. Thus the network of interaction has two functions. One, it defines the pairs of players who play the game and two, it shapes the flow of information about individual behavior across a collection of players. The main results pertain to the effect of the fluctuation in payoffs on a variety of issues such as the shape of the network (degree, average distance and the size of largest component), and the level of aggregate payoffs obtained. I report the findings in the simulations. It should be noted that the dynamic process of link formation and strategy choice is ergodic and so the average behavior in simulations is a good indicator of the general properties of the model. The first finding is that greater volatility in payoffs leads to lower connectivity of the network. The second finding is that an increase in volatility leads to a fall in average distance, in other words an increase in the cohesiveness of the society. The third finding pertains to the relative size of the two largest components: most of the players belong to the largest component. The final finding pertains to the average payoff per link: this average payoff declines as the volatility increases. Taken along with the first finding on average degree of network, this result illustrates the role of network effects in supporting cooperative behaviour. The paper also obtains some analytical results concerning average connectivity and payoff volatility.

7 Concluding remarks

Traditionally, economists have studied social and economic phenomena using models with centralized and anonymous interaction among individual entities. Moreover, prices have been the main coordinating device for the interaction among individuals. In recent years, we have developed a richer repertoire of models which allow for decentralized (and local) interaction among individuals, and coordination is attained via a variety of non-price mechanisms. The

[15] There is however some work on cooperation induced through the threat of exclusion/ostracism which is related; see e.g., Hirshliefer and Rasmusen (1989), Annen (2003).

research on learning in networks surveyed in this paper should be seen as part of this general research programme. In this section I summarize what we have learnt so far and also propose some open questions.

I started the paper with results on learning for a given network of interaction. Existing results show clearly that, both in strategic as well as non-strategic settings, the network structure influences the actions individuals make and this in turn has serious implications for the level of welfare they can hope to attain in the short run as well as the long run. In the context of non-strategic learning, we know that individuals with similar ranking of actions who are located in a connected society will eventually all earn the same payoff. However, the level of this payoff – whether it is the maximum attainable – will depend on architecture of the network of communication. Existing work has explored the possibilities of complete learning in societies with an infinite number of players. The main finding here is that a desirable communication network should allow for 'local independence' – people having distinct sources of information which facilitates experimentation and gathering of new information – and at the same time it should have channels of communication across these localities to spread successful actions. The optimality of this network specially with regard to finite societies remains an open question.

In the context of coordination games played on networks the main finding is that in random matching models or in complete networks, equilibrium outcomes correspond to the pure strategy of the underlying game. Moreover, the stochastically stable outcome is typically the risk-dominant equilibrium. However, with local interaction, a variety of outcomes are possible and this is also true when we restrict attention to stochastically stable outcomes. The stochastic stability results have been obtained for myopic best response decision rules and equiprobable deviations. These results can be greatly strengthened and refined if we suppose that deviations from best response follow a log linear function (which make deviations from optimal actions sensitive to payoff differences). In particular, it can be shown that dynamics always select the risk-dominant equilibrium and that convergence to this outcome is rapid and essentially independent of the size of the population for a wide class of local interaction networks. The class of interaction networks identified here also emphasizes close-knit local communities, which is similar in spirit to the local independence idea developed in the non-strategic learning setting. The differences in results between best response dynamics and log-linear decision based dynamics lead me to wonder if there are close relations between interaction structure and decision rules. For instance, is imitation easier and more effective in local interaction settings. My discussion of games of conflict played on networks focused on prisoners dilemma games and also highlights this connection between decision rules and interaction structure. The main result obtained here is that a combination of imitation based decision making and local interaction ensures high levels of cooperation in the society. This result has been obtained for interaction on a circle only, and an interesting open question here is the robustness of altruism and the possibility of cooperation in more general networks of interaction.

I then moved to a discussion of learning in a setting where the network of structure itself is evolving as a result of individual decisions on links with others. This is a relatively new strand of the research and there are a number of fascinating questions that are open. My presentation focused on the connections model of network. There are two general findings here: the first finding is that strategic link formation implies sharp predictions on equilibrium network architectures and the second finding is that the dynamics of link formation based on individual learning have strong self-organizing properties. The process converges, the limit networks can be explicitly characterized (and have simple and classical architectures such as stars and wheels) and the rate of convergence to the limit networks is fast. The model of connections if a natural model for networks with non-rival goods as it captures one type of positive externality of linking activity. But there are a variety of other externalities – positive and negative as well as a mixture of the two – that arise in networks. A systematic analysis of these spillovers and their bearing on network architectures is an open problem.

In the final part of the survey, I studied games where the strategies which involve links as well as other actions. The study of these models is still at an early stage. To the best of my knowledge there is no formal model of link formation and learning optimal actions in a non-strategic setting (as explored in section 3). This seems to be a very promising area for further work, and seems to be closely related to both the traditional learning models as well as the search theory literature. In the context of coordination games, the main finding of the existing work is that the costs of forming links decisively shape individual behavior and determine long run equilibrium. This finding should be seen as a first step in the exploration of a variety of issues relating to evolving social structures and economic performance. For instance, one may wonder if it is desirable to allow persons to freely form links or if some restrictions on link activity may be socially desirable. The existing models are very stylized and richer models can be useful to study phenomenon such as ghettos and the formation of exclusive clubs. Lastly, I discussed work on games of conflict with endogenous networks. This work is just beginning and I believe that further research in this area will deepen our understanding of the origins of social capital and trust.

References

[1] Albert, R. and A-L. Barabasi (2002), Statistical mechanics of complex networks, *Review of Modern Physics*, 74.

[2] Allen, B. (1982), Some stochastic processes of interdependent demand and technological diffusion of an innovation exhibiting externalities among adopters, *International Economic Review*, 23, 595-607.

[3] Anderlini, L. and A. Ianni (1996), Path dependence and learning from neighbours, *Games and Economic Behavior*, 13, 141-177.

[4] Annen, K. (2003), Social Capital, Inclusive Networks, and Economic Performance, *Journal of Economic Behavior and Organization*, 50, 449-463.

[5] Aumann, R. and R. Myerson (1989), Endogenous formation of links between players and coalitions: an application of the Shapley Value, in A.Roth, (ed), *The Shapley Value*, Cambridge University Press.

[6] Axelrod, R. (1997), *The complexity of cooperation: agent-based models of competition and collaboration*. Princeton University Press. Princeton, New Jersey.

[7] Bala, V. (1997), Notes on network formation, mimeo, McGill University.

[8] Bala, V. and S. Goyal (2001), Conformism and diversity under social learning, *Economic Theory*, 17, 101-120.

[9] Bala, V. and S. Goyal (2000a), A Non-Cooperative Model of Network Formation, *Econometrica*, 68, 5, 1181-1231.

[10] Bala, V. and S. Goyal (2000b), A strategic analysis of network reliability, *Review of Economic Design*, 5, 3, 205-229.

[11] Bala V. and S. Goyal (1998), Learning from Neighbours, *Review of Economic Studies* 65, 595-621.

[12] Banerjee, A. (1993), A Simple Model of Herd Behavior, *Quarterly Journal of Economics*, 107, 797-817.

[13] Banerjee, A. and D. Fudenberg (1999), Word-of-mouth learning, *mimeo*, MIT.

[14] Banks, D. and K. Carley (1996), Models of social network evolution, *Journal of Mathematical Sociology*, 21, 173-196.

[15] Baron, R., J. Durieu, H. Haller and P. Solal (2002), Control costs and potential functions for spatial games, *mimeo*, VPI.

[16] Bergin, J. and B. Lipman (1996), Evolution of state dependent mutations, *Econometrica*, 64, 943-56.

[17] Bikhchandani, S., D. Hirshliefer and I. Welch (1992), A theory of fads, fashion, custom, and cultural change as informational cascades, *Journal of Political Economy*, 100, 992-1023.

[18] Billingsley, P. (1985), *Probability and Measure*. Wiley. New York.

[19] Bhaskar, V. and F. Vega-Redondo (2002), Migration and the evolution of conventions, *Journal of Economic Behavior and Organization*, forthcoming.

[20] Blume, L. (1993), The statistical mechanics of strategic interaction, *Games and Economic Behavior*, 4, 387-424.

[21] Blume, L. and D. Easley (1995), What has the rational learning literature taught us?, in A. Kirman and M. Salmon (eds), *Learning and Rationality in Economics*, Oxford University Press. Oxford.

[22] Bramoulle, Y. (2001), Congestion and social networks: an evolutionary analysis, *mimeo*, University of Maryland.

[23] Bramoulle, Y., D. Lopez, S. Goyal and F. Vega-Redondo (2002), Social interaction in anti-coordination games, *mimeo*, Alicante University.

[24] Bramoulle, Y., and R. Kranton (2003), A network model of public goods: experimentation and social learning. *mimeo*, University of Toulouse and University of Maryland.

[25] Brock, W. and S. Durlauf (2001), Interactions-based models, *Handbook of Econometrics: Volume V*, Heckman, J and Leamer, E (eds). North Holland. Amsterdam.

[26] Boorman, S. (1975), A combinatorial optimization model for transmission of job information through contact networks, *Bell Journal of Economics*, 6, 1, 216-49.

[27] Burke, M. and P. Young (2001), Competition and cutsom in economic contracts: a case study of Illinois agriculture, *American Economic Review*, 91, 559-573.

[28] Burt, R. (1994), *The social structure of competition*, Harvard University Press. Cambridge, MA.

[29] Calvo, A. (2000), Bargaining in networks, *Unpublished Ph.D thesis*, University Pompeu Fabra, Barcelona.

[30] Chwe, M (2000), Communication and Coordination in Social Networks, *Review of Economic Studies*, 67, 1, 1-17.

[31] Coleman, J. (1990), *The Foundations of Social Theory*, Harvard University Press. Cambridge, MA.

[32] Coleman, J. (1966), *Medical Innovation: A Diffusion Study*, Second Edition, Bobbs-Merrill, New York.

[33] Corominas, M. (1998), On two-sided network markets, *Unpublished Ph.D thesis*, University Pompeu Fabra, Barcelona.

[34] Currarini, S. and M. Morelli (2000), Network formation with seuqential demands, *Review of Economic Design*, 5, 229-251.

[35] van Damme, E. and J. Weibull (2002), Evolution and refinement with endogenous mistake probabilities, *Journal of Economic Theory*, forthcoming.

[36] Doreian, P. and F. Stokman (2001), Evolution of Social Networks, Special Volume: *Journal of Mathematical Sociology*, 25.

[37] Dorogovtsev, S.N and J.F.F Mendes (2001), Evolution of networks, *Advances in Physics*, 51, 1079-1187.

[38] Droste, E, R. Gilles, and K. Johnson (1999), Endogenous interaction and the evolution of conventions, *mimeo*, Tilburg University.

[39] Dutta, B., S. Ghosal and D. Ray (2002), Far sighted network formation, *mimeo*, Warwick University.

[40] Dutta, B. and M. Jackson (2001), (eds) *Models of the Strategic Formation of Networks and Groups.* Springer-Verlag, Berlin.

[41] Dutta, B., A. van den Nouweland, and S. Tijs (1995), Link formation in cooperative situations, *International Journal of Game Theory*, 27, 245-256.

[42] Easley, D. and N. Kiefer (1988), Controlling a stochastic process with unknown parameters, *Econometrica*, 56, 1045-1064.

[43] Elias, N. (1978), *The history of manners*, Part 1 of *The Civilizing Process.* Pantheon Books, New York.

[44] Ellison, G. (1993), Learning, Local Interaction, and Coordination. *Econometrica*, 61, 1047–1071.

[45] Ellison, G. and D. Fudenberg (1993), Rules of Thumb for Social Learning, *Journal of Political Economy*, 101, 612-644.

[46] Ellison, G. and D. Fudenberg (1995), Word-of-mouth communication and social learning, *Quarterly Journal of Economics,* 109, 93-125.

[47] Ely, J. (1996), Local Conventions, *mimeo*, Northwestern University.

[48] Eshel, I. and L. Cavalli-Sforza, (1982), Assortment of encounters and evolution of cooperativeness, *Proceedings of the National Academy of Sciences*, 79, 4, 1331-35.

[49] Eshel, I., Samuelson, L., and Shaked, A., (1998), Altruists, egoists, and hooligans in a local interaction model, *American Economic Review*, 157-179.

[50] Freidlin, M. I. and A. D. Wentzell (1984), *Random Perturbations of Dynamical Systems.* Translation of Freĭdlin and Ventt͡sel′ (1979) by J. Szücs. Springer-Verlag, New York.

[51] Fudenberg, D. and D. Levine (1998), *The theory of learning in games*, MIT Press, Cambridge.

[52] Glaeser, E., B. Sacredote and J. Scheinkman (1996), Crime and Social Interactions, *Quarterly Journal of Economics*, 111, 507-548.

[53] Glasear, E. and J. Scheinkman (2001), Measuring social interactions, *mimeo*, Princeton University.

[54] Goyal, S. (1993), Sustainable communication networks, *Tinbergen Institute Discussion Paper*, Rotterdam, TI 93-250.

[55] Goyal, S. (1996), Interaction Structure and Social Change, *Journal of Institutional and Theoretical Economics*, 152, 3, 472-495.

[56] Goyal, S. (1999), *Networks, learning and equilibrium,* Inaugural Lecture, Erasmus University, Eburon Press, Delft.

[57] Goyal, S. and S. Joshi (2003), Networks of collaboration in oligopoly, *Games and Economic Behavior*, 43, 57-85.

[58] Goyal, S. and J.L. Moraga (2001), R& D Networks, *Rand Journal of Economics,* 32, 4, 686-707.

[59] Goyal, S. and F. Vega-Redondo (2002), Network formation and social coordination, mimeo, Queen Mary, University of London and University of Alicante.

[60] Granovetter, M. (1974), *Getting a Job: A Study of Contacts and Careers*, Harvard University Press, Cambridge MA.

[61] Granovetter, M (1985), Economic Action and Social Structure: The Problem of Embeddedness, *American Journal of Sociology,* 3, 481-510.

[62] Griliches, Z. (1957), Hybrid Corn: An exploration in the economics of technological change, *Econometrica*, 25, 501-522.

[63] Haag, M. and R. Lagunoff (1999), Social Norms, Local Interaction, and Neighborhood Planning, *mimeo*, Georgetown University.

[64] Hagerstrand, T. (1969), *Innovation diffusion as a spatial process.* University of Chicago Press. Chicago.

[65] Haller, H. and S. Sarangi (2001), Nash networks with heterogeneous agents, *mimeo*, VPI Working Paper.

[66] Hirshleifer, D., and E. Rasmusen (1989), Cooperation in a Repeated Prisoner's Dilemma with Ostracism, *Journal of Economic Behaviour and Organization*, 12, 1, 87- 106.

[67] Ianni, A., (2001), Correlated equilibria in population games, *Mathematical Social Sciences*, 42, 3, 271-294.

[68] Jackson, M. (2003), A survey of models of network formation: stability and efficiency, to appear in *Group Formation in Economics: Networks, Clubs and Coalitions*, edited by G. Demange and M. Wooders. Cambridge University Press.

[69] Jackson, M. and A. Watts (2002), The evolution of social and economic networks, *Journal of Economic Theory*, 106, 2, 265-295.

[70] Jackson, M. and A. Watts (2002), On the formation of interaction networks in social coordination games, *Games and Economic Behavior*, 41, 2, 265-291.

[71] Jackson, M. and A. Wolinsky (1996), A Strategic Model of Economic and Social Networks, *Journal of Economic Theory*, 71, 1, 44-74.

[72] Judd, K. and L. Tesfatsion (2005), *Handbook of Computation Economics II: Agent based computational economics.* North Holland, Amsterdam.

[73] Kandori, M. (1997), Evolutionary game theory in economics, in D. Kreps and K. Wallis (eds), *Advances in Economics and Econometrics: Theory and applications.* Cambridge University Press.

[74] Kandori, M., G. Mailath and Rob (1993), Learning, mutation and long run equilibria in games, *Econometrica*, 61, 29-56.

[75] Kirman A.P., C. Oddou and S. Weber (1986), Stochastic communication and coalition formation, Econometrica, 54, 1, 129-138.

[76] Kirman, A. (1997), The Economy as an Evolving Network, *Journal of Evolutionary Economics*, 7, 339-353.

[77] Kirman, A., and Zimmermann, J.-B. (2001), *Economics with Heterogeneous Interacting Agents*, Series: Lecture Notes in Economics and Mathematical Systems. Springer Verlag.

[78] Kranton, R. and D. Minehart (2001), A Theory of Buyer-seller networks, *American Economic Review,* 91, 485-508.

[79] Lee, I. H. and A. Valentinyi (2000), Noisy Contagion without Mutation, *Review of Economic Studies*, 67, 1, 17-47.

[80] Lee, I., A. Szeidl and A. Valentayini (2002), Contagion and State dependent Mutations, *mimeo*, Southampton University.

[81] Liggitt, T (1985), *Interacting Particle Systems*, New York, Springer-Verlag.

[82] Mailath, G. Samuelson, L. and Shaked, A., (1997), Correlated equilibria and local interaction, *Economic Theory*, 9, 551-568.

[83] Manski, C. (2000), Economic analysis of social interactions, *mimeo*, Morthwestern University.

[84] Marimon, R. (1997), Learning from learning in economics, in D. Kreps and K. Wallis (eds), *Advances in Economics and Econometrics: Theory and applications.* Cambridge University Press.

[85] McBride, M. (2002), Position-specific information in social networks, *mimeo*, University of California, Irvine.

[86] McLennan, A. (1984), Price dispersion and incomplete learning in the long run, *Journal of Economic Dynamics and Control*, 7, 331-347.

[87] Morris, S. (2000), Contagion, *Review of Economic Studies*, 67, 1, 57-79.

[88] Munshi, K. (2003), Networks in the Modern Economy: Mexican Migrants in the U. S. Labor Market, *Quarterly Journal of Economics*, 118, 2, 549-597.

[89] Nowak, M. and R. May, (1992) Evolutionary games and spatial chaos, *Nature*, 359, 826-29.

[90] Oechssler, J (1997), Decentralization and the coordination problem, *Journal of Economic Behavior and Organization* 32, 119-135.

[91] Raub, W. and J. Weesie (1990), Reputation and efficiency in social interactions, *American Journal of Sociology,* 96, 626-655.

[92] Robson, A. and F. Vega-Redondo (1996), Efficient equilibrium selection in evolutionary games with random matching, *Journal of Economic Theory*, 70, 65-92.

[93] Rothschild, M. (1974), A two-arm bandit theory of market pricing, *Journal of Economic Theory*, 9, 185-202.

[94] Ryan, B. and N. Gross (1943), The diffusion of hybrid seed corn in two Iowa communities, *Rural Sociology*, 8, 15-24.

[95] Slikker, M. and A. van den Nouweland (2001), *Social and Economic Networks in Cooperative Game Theory.* Kluwer Academic Publishers. Boston.

[96] Skyrms, B. and R. Pemantle (2000), A dynamic model of social network formation, *Proceedings of the National Academy of Sciences,* 97, 16, 9340-9346.

[97] Samuelson, L. (1997), *Evolutionary Games and Equilibrium Selection.* MIT Press, Cambridge, MA.

[98] Schelling, T. (1975), *Micromotives and macrobehaviour*, Norton. New York.

[99] Smallwood, D. and J. Conlisk (1979), Product quality in markets where consumers are imperfectly informed, *Quarterly Journal of Economics*, 93, 1-23.

[100] Smelser, N. and R. Swedberg (1994), *The Handbook of Economic Sociology*, Princeton University Press, Princeton, New Jersey.

[101] Snijders, T. (2001), The statistical evaluation of social network dynamics, *Sociological Methodology*, 31, 361-395.

[102] Taylor, R. (1979), *Medicine out of control: The anatomy of a malignant technology.* Sun Books. Melbourne.

[103] Tieman, A., H. Houba and G. van der Laan (2000), On the level of cooperative behavior in a local interaction model, *Journal of Economics*, 71, 1-30.

[104] Ullman-Margalit, E. (1977), *The emergence of norms*, Clarendon Press. Oxford.

[105] van den Nouweland, A. (2003), Models of network formation in cooperative games, to appear in *Group Formation in Economics: Networks, Clubs and Coalitions*, edited by G. Demange and M. Wooders. Cambridge University Press.

[106] Vega-Redondo, F. (2002), Building social capital in a changing world, *mimeo*, Alicante.

[107] Watkins, S. (1991), *Provinces into Nations: Demographic Integration in Western Europe, 1870-1960*, Princeton University Press, Princeton. New Jersey.

[108] Wynne-Edwards, V. (1986), *Evolution through group selection*, Oxford. Blackwell Publishers.

[109] Young, P. (1998), *Individual Strategy and Social Structure*, Princeton University Press. Princeton.

[110] Young, P (1998), Diffusion in social networks, *mimeo*, Johns Hopkins University.