

Network models of innovation and knowledge diffusion

Citation for published version (APA):

Cowan, R. (2004). *Network models of innovation and knowledge diffusion*. MERIT, Maastricht Economic Research Institute on Innovation and Technology. MERIT-Infonomics Research Memorandum Series No. 016 <https://doi.org/10.26481/umamer.2004016>

Document status and date:

Published: 01/01/2004

DOI:

[10.26481/umamer.2004016](https://doi.org/10.26481/umamer.2004016)

Document Version:

Publisher's PDF, also known as Version of record

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

MERIT-Infonomics Research Memorandum series

Network models of innovation and knowledge diffusion

Robin Cowan

2004-016



*MERIT – Maastricht Economic Research
Institute on Innovation and Technology*

PO Box 616
6200 MD Maastricht
The Netherlands
T: +31 43 3883875
F: +31 43 3884905

<http://www.merit.unimaas.nl>
e-mail: secr-merit@merit.unimaas.nl



International Institute of Infonomics

c/o Maastricht University
PO Box 616
6200 MD Maastricht
The Netherlands
T: +31 43 388 3875
F: +31 45 388 4905

<http://www.infonomics.nl>
e-mail: secr@infonomics.nl

Network models of innovation and knowledge diffusion*

Robin Cowan
MERIT
University of Maastricht
P.B. 616
6200MD Maastricht, Netherlands
r.cowan@merit.unimaas.nl

July, 2004

* Much of the material in this paper arises from the collaborative work done with Nicolas Jonard, and owes much to him. I also acknowledge the very helpful comments of Muge Ozman and encouraging words of Franco Malerba.

1 INTRODUCTION

Much of modern micro-economics is built from the starting point of the perfectly competitive market. In this model there are an infinite number of agents — buyers and sellers, none of whom has the power to influence the price by his actions. The good is well-defined, indeed it is perfectly standardized. And any interactions agents have is mediated by the market. That is, all transactions are anonymous, in the sense that the identities of buyer and seller are unimportant. Effectively, the seller sells “to the market” and the buyer buys “from the market”. This follows from the standardization of the good, and the fact that the market imposes a very strong discipline on prices. Implicit here is one (or both) of two assumptions. Either all agents are identical in every relevant respect, apart, possibly, from the prices they ask or offer; or every agent knows every relevant detail about every other agent. If the former, then obviously my only concern as a buyer is the prices asked by the population of sellers since in every other way they are identical. If the latter, then each seller has a unique good, and again what I am concerned with is the price of it. In either case, we see that prices capture all relevant information and are enough for every agent to make all the decisions he needs to make.

It is commonplace to observe that we seldom see a perfectly competitive market. We never see infinite numbers of agents, and obviously someone is able to influence prices, else how would they be set, and how would they change? (See the debate between Kirzner, 1962 and 1963; and Becker, 1962, 1963, on this issue.) But that alone does not say that the model is a bad approximation. More important is that while it may be possible in principle for any agent to interact with any other, in fact this is far from the case in practice. Typically any agent only interacts with a very small proportion of the population. The “global market” does not really exist. Interactions tend to be quite localized in many senses. It is also the case, though, that we don’t have perfect information about other agents, neither is it the case that they are identical. Moreover, this non-identity may matter — it may be economically relevant. Agents are in different physical locations, one from another. In some ways this is easily handled by adding to the definition of a good its location, and then including the price of transportation in the information set. But more complex changes have to be made for example when agents are competing: address models take account of the fact that when agents are close together they may have to compete harder with each other than when they are far apart.¹ Agents may use different standards

or technologies. If a good exhibits network externalities or increasing returns to adoption, two agents will create externalities for each other if they use the same standard or technologies, but not if they do not. But often this is not an anonymous effect: it is true that there is a global network externality to telephone use, but there is a much stronger local externality: the value of a telephone is much higher if my friends own them than if they do not, regardless of aggregate behaviour. The identity of who is using which technology makes a difference to its economic value to any particular agent. Similarly, the history that a pair of agents shares can have a significant effect on the nature of their interactions. If we have interacted many times and the experience has been good, trust is created and we can have certain kinds of interactions (based on a hand-shake, for example). If we have no history, or even more so if we have had bad experiences in the past, handshakes will not be enough, and there may be some types of interactions that we simply cannot arrange. In each case, more than just the price of the good matters.

In each of these cases, with the possible exception of history, the phenomenon can be represented as an issue of space and distance. Geographical space is the most natural. Issues of transportation costs are easily understood in this context, and the degree of competition among sellers is relatively well-understood as it is affected by distance (See Anderson et al. 1992.) However, technologies or standards exist in a space of technical attributes; preferences exist in a Lancasterian characteristic space and so on. The point then can be seen that the distance between two agents affects their interactions. If, for example, agents are close together in technology space, they may be able to share hardware, software or information. These interactions may be impossible if the agents are far from each other in technology space — the software is not compatible across hardware problems; the technologies have different operating features, so information about how to make one of them efficient will not apply to the other, and so on.

As pointed out above, address models can be used to capture some of these phenomena. But when coupled with other issues that have been studied by economists in recent decades, there may be a more natural way to approach some of these problems.

1. See Hotelling (1929) or Salop (1979) for example. But see also Anderson et al. (1992), especially chapter 4, to illustrate how quickly address models become extremely difficult in terms of tractability.

2 KNOWLEDGE DISTRIBUTION

Distance and transportation costs are clearly relevant when talking about physical goods. Surprisingly, they are also relevant when discussing something as intangible as knowledge or information. In many economic models, knowledge is assumed to be a public good: once discovered and publicised, it is freely available to all in the economy. The presence of the internet and the world wide web should make the “globally available” aspect of knowledge even stronger. Consider patents. Patents are highly codified data and are by definition stored in a public database, the most important of which (USPTO and EPO) are now searchable online. Thus anyone with an internet connection can read any patent granted in the US or Europe (and many other jurisdictions as well). Diffusion of the knowledge or information contained in patents should be a truly global phenomenon. But the work of Jaffe et al. (1993) shows quite clearly that this is not the case. The diffusion of information is a process embedded in location. Distance is important in knowledge diffusion — the closer we are to the location of the originator of knowledge, the sooner we learn it.

This result has been revisited by Breschi and Lissoni (2003). Jaffe et al. looked at diffusion over geographical distance; Breschi and Lissoni add social distance. Knowledge diffusion is not only a geographically spatial phenomenon, it is also a “socially spatial” phenomenon. That is, the more closely connected socially I am to the originator of a piece of knowledge, the more quickly I will learn about it. What these results strongly suggest is that knowledge diffusion is taking place to a very large extent by personal contact. Knowledge is passed informally through face-to-face contacts, and direct interactions. One explanation for this is that much relevant knowledge is tacit, or at least not codified, which means that it cannot be broadcast, and can only be transmitted locally, through direct interactions. So space matters for knowledge diffusion, and social space may matter as much as or more than geographic space. “Social space” can only be sensibly understood in the context of social networks.

3 BEHIND THE CURRENT INTEREST IN NETWORKS

In economics, interest in networks has been fuelled by three different developments all of which occurred in the last decades of the twentieth century.

3.1 Network Technologies

Since the early 1980s, economists have been interested in technologies that function as part of a network. The prototypical example is the telephone. A single telephone has no value, but when connected to a network of other telephones, it does. We see many examples of these sorts of technologies, particularly associated with information and communication: telephone; telefax; email and so on. Two of the main issues in the economics literature on this subject have been the presence of network externalities and standardization. For an artefact to join a network, it must be able to communicate with the other artefacts that make up that network. Some form of inter-operation standards must be in place. In the early days of email, the first large commercial network, CompuServe, was unable to connect to the BITNET, which was the main email form used over the internet. Thus a potential email user had to decide which network to join. More generally, who you are connected to, and who you want to be connected to is a question of which standard to adopt. It is worth observing, since this issue will arise again below, that many of these network technologies are concerned with the distribution of information. Information is diffused using these technologies, and physical networks form as agents decide which standard to adopt, and thus which other agents it will be possible to communicate with. Networks form as agents adopt technologies. Some of these networks reside in physical infrastructure, others are more virtual, for example the network of people sharing information about a computer operating system or programming language. Network issues are thus at the heart of many of the new technologies.

3.2 Boundaries of the firm

In the last couple of decades there has been a resurgence of interest in the boundaries of the firm. (See for example Holmstrom and Roberts, 1998.) There seems to have been a trend whereby the monolithic, vertically integrated firm has slowly dissolved as more and more of its activities are out-sourced. This began simply with parts, as a firm faced the “make-or-buy” decision for the parts of its assembly. The “raw materials” that came into the firm became less “raw”. This is a relatively minor change, but does involve creating relationships with suppliers based not only on price but also on trust, since, a low quality shipment of parts (whether deliberate or not, on the part of the supplier) could be disastrous. Firm boundaries have since become even more porous: firms now out-source not

only parts, but also design. Parts manufacturers are likely to know more than parts users about how to design parts efficiently and for efficient production, so it is sensible that they have a strong hand in this design process. Again, though, even more than with parts supply, with design supply firms and their suppliers need a particular kind of relationship. Long-standing relationships, partial cooperation, joint ventures all are forms that emerged as the former vertically integrated firms attempted to take advantage on the human capital and competencies of their suppliers. These are market relationships only to a small degree. They demand direct interactions, and a significant amount of stability, to ensure that the transactions can take place effectively.

3.3 The expanding knowledge base

One explanation for the deeper out-sourcing is the expanding knowledge base of industries. Granstrand (1998) summarizes several studies that reach the following conclusions: In recent decades, large, technology-based firms in particular have expanded their technology bases into a wider range of technologies. This technological diversification has been one of the fundamental causes of firm growth. In addition, because an expanding technology base demands an expanding knowledge base, this technological diversification has been a leading cause not only of an increase in R&D expenditure, but also of an increase in external technology sourcing. (See also Grandstand and Sjolander, 1990.) Generally, we can see that as the range of a firm's technologies expands, the firm must learn about them, how to use the new technologies, how to develop them, how to integrate them into existing products and how to create new products from them. Thus not only does a firm need more knowledge, it needs more types of knowledge. Initially in any case, it can be difficult to develop this knowledge in-house. Because of its nature, knowledge, and especially new knowledge can be difficult to acquire in the market, so firms seek some form of collaboration with other firms that already have the knowledge. Firms sought arrangements to outsource not only production and design, but now also R&D. The rise of strategic R&D alliances since 1980 has been very striking. (See Hagedoorn, 2002.) Having a portfolio of alliances acts as insurance for a firm wishing not to be taken by surprise by new technology developments. It also gives the firm rapid access to disparate and varied information should it need it. (For a discussion of the motivation behind the rise in strategic alliances see for example Narula and Hagedoorn, 1999.) Firms act to create links

with other firms with the goal to transmitting knowledge across those links. In a more extreme version, we see the rise of the networked organization. This is sometimes viewed as an intermediate organizational form, between markets on the one hand, and hierarchies on the other. The central idea is that the “organization” takes advantage of a large amount of autonomy between nodes, but at the same time uses links between those nodes to transmit value between them. Clearly, this organizational form precludes the idea of global interaction in the sense that any firm can interact with any other. Indeed, part of the structure demands that firms interact only with a relatively circumscribed subset of the entire population.

What we observe, taking these issues together, is occasions in which an agent interacts with only a small part of the population. Further, these interactions are not (entirely) market mediated, but take place rather over some more direct connection, and long-lived connections can be more effective than temporary ones. Finally, the identity of the agents can make a difference to the nature and success of the interaction. This sounds much like the description of a network: agents are nodes, they are linked to particular other nodes, and interactions take place only over these direct links.

4 THE ECONOMICS OF NETWORKS

4.1 Empirically

There is now a sizeable empirical literature that uses this approach for industry studies: Silicon Valley, and the computer industry more generally (Saxenian, 1991; Langlois and Robertson, 1992); Japanese business groups (Imai, 1989; Lincoln and Gerlach, 1998); the biotech industry (Orsenigo et al., 2001; Powell et al., 1996); Italian industrial districts (Lazerson & Lorenzoni, 1999; Pyke et al.) the automobile industry (Dyer, 1996); the fashion industry (Uzzi, 1997). All of them observe that the “peculiar” properties of these industries can be understood by analysing the agents’ behaviour as part of a network structure, and that an important part of the way these industries function is determined by precisely who interacts with whom. Social capital, developed by direct interactions between agents, permits the agents to circumvent many of the problems implicit in anonymous markets such as opportunistic behaviour, imperfect information and incomplete contracts,

the transmission of tacit knowledge, knowledge spillovers and so on. The social capital developed through network interaction therefore reduces transaction costs, and can serve to make the industry more efficient in many ways.

4.2 Theoretically

In addition to the relatively recent empirical work on networks, there has been a slightly longer tradition of abstract network analysis. Roughly speaking, analytical models have been of two types: Ising models, and random graphs.

4.2.1 Ising Models

In an Ising model agents are located at fixed points in a regular integer space, for example at the integers on a line or at integer coordinates in Euclidian two-dimensions. Agents are directly connected to their n nearest neighbours in the physical space. These structures are often called nearest-neighbour networks or graphs. The networks are clearly very regular, and have high levels of local coordination, and typically, are locally very dense. They do, however, have relatively long average paths between pairs of agents. The general model was developed in 1925 (Ising, 1925) to study ferro-magnets, and therefore its analytic properties are relatively well-understood. In economics it has been used to study a wide variety of situations: macro dynamics (Durlauf, 1993); technology diffusion (Allen, 1982, or An and Kiefer, 1993); criminal behaviour (Glaeser et al. 1996); the effectiveness of prices (Föllmer, 1974). This work all takes advantage of the fact that in the Ising model local interactions are very neatly defined, and only distance, between locations, rather than absolute location matters. Nonetheless, complex patterns of aggregate behaviour can emerge from local interactions.

4.2.2 Random Graphs

At the other extreme of network structures, economists have used random graph models to study, for example, coalition formation (Kirman 1983, or Ionnides, 1990); technology diffusion (Steyer and Zimmermann, 1998), learning (Ellison, 1993); strategy revision (Blume, 1995). In these graphs, any agent is connected with some probability to any other agent in the population, regardless of location. Consequently, the networks that emerge have no spatial patterns in the space in which agents are physically located. An alternative interpretation is that agents are not located in physical space, but rather only in network

space, so any notion of distance is only distance in the network. In addition, there is essentially no local coherence, and generally graphs are not locally dense, even in network space. On the other hand, random graphs are known to have low average path lengths: the path in network space between two agents is, on average, low.

From the point of view of knowledge creation and diffusion, the focus on these extreme forms of network structure, draws attention to the contrast in the properties of local density and path length. It is often asserted that innovation is facilitated by agglomeration of human capital. A location in which many agents are working on related issues creates a critical mass of knowledge workers so that they become an epistemic community, in which a common language emerges, problem definitions become standardized, and problem solving heuristics emerge and are developed. In this type of situation, Marshall's idea that knowledge is in the air applies strongly. Agglomerations having these features foster innovation. (See the empirical studies referred to above.) Applying the network model, this situation is present if the network of agents exhibits high local density, or equivalently high "cliquishness". That is, if there are (relatively small) groups of agents who are closely and heavily interconnected. As we have seen, this is true in a nearest-neighbour graph, but not in a random graph.

By contrast, knowledge diffusion is most rapid when path lengths are short. Diffusion is about spreading a piece of knowledge to all agents in the economy, and this will happen fastest (and with least degradation) when it takes few steps between the originator of the knowledge and its recipients. Short path lengths, associated with the random graph structure make this happen most quickly.

Thus there appears to be a tension between knowledge creation and knowledge diffusion. The tension is resolved by the network structure known as the small world. This is a network structure which is at the same time highly cliquish and with short paths between agents. It lies between the nearest-neighbour network and the random network.

4.2.3 Small worlds

Economists focused on Ising and random graph models partly because they had been extensively explored in other disciplines, but more because there was no obvious way to characterise systematically structures that lay between these two extremes. Watts and Strogatz (1998) develop a one-parameter class of random graph models in which the parameter could be used to scale between the regular graph of the Ising model and the random graph

of random graph theory. The algorithm for generating a graph was as follows: create a regular nearest-neighbour graph. Fix the parameter p . Examine each link in the network, and with probability p rewire one end of that link, connecting it to an agent chosen randomly from the population (in the process checking to ensure that agents are not connecting to themselves, and that any pair of agents is connected at most once.) This algorithm scales between the regular, nearest-neighbour graph (if $p = 0$) and the random graph (if $p = 1$), but also has the feature that it does so without changing the density, or total number of connections in the graph. This is illustrated in Figure 1.




Figure 1: Transition from a locally ordered structure to a random graph, dependent on p .

This algorithm provides a way of finding the small world. Small worlds are characterized by two structural parameters: cliquishness and path length. Cliquishness captures the idea that “my friends are friends of each other”, that is, of local density. Formally, average cliquishness is defined as

$$C = \frac{1}{N} \sum_i \sum_{(j,l \in \Gamma_i)} \frac{X(j,l)}{\|\Gamma_i\|(\|\Gamma_i\| - 1)/2}$$

where $X(j,k)=1$ if $j \in \Gamma_l$ and $X(j,k)=0$ otherwise, and Γ_i is the set of neighbours of i (those to whom he is directly connected), and $\|\Gamma_i\|$ is the size of that neighbourhood.

Average path length is formally defined as

$$L(p) = \frac{1}{N} \sum_i \sum_{j \neq i} \frac{d(i,j)}{N-1}$$

where $d(i,j)$ is the length of the shortest path between i and j .

Not surprisingly, as p increases, since the local coherence necessarily falls with the introduction of random links, cliquishness decreases. Similarly, since a random link is a potential short-cut, average path length also falls. Crucially, though, the latter falls much faster than the former. This is illustrated in Figure 2, where we observe that for values of p between roughly 0.01 and 0.1, cliquishness is close to that of a nearest neighbour graph, while path length is close to that of a random graph. In this region, conditions are favourable both for knowledge creation and for knowledge diffusion.




Figure 2: Cliquishness and path length as determined by network structure

5 SMALL WORLDS AND INNOVATION

Diffusion involves passing information or knowledge from one agent to another. There are many possible micro-economic models of this process. Empirical studies have revealed two striking ones. Allen, (1983) observes what he calls “collective invention” in the steel industry in Cleveland (UK) in the 19th century. (There are many other significant examples of the same phenomenon, see Cowan and Jonard, 2004 for a brief discussion.) In a situation of collective invention, competitors in an industry share technical information relatively freely with each other. Technical advance by one firm is essentially broadcast to other firms in the region. The broadcast takes place sometimes through local publication (Nuvolari, forthcoming), or through presentations at meetings (Allen, 1983). The most popular modern example is open source software development. Again, technical advances

made by one developer are broadcast to a group of fellow-developers, and so a common undertaking to improve the technology exists.

The second micro-economic model of knowledge transfer is identified by von Hippel (1998). He observes a form of knowledge barter in a variety of industries. Technical managers of plants meet informally and discuss the technical problems they are faced with. Information passes among them outside any sort of market context. There is a caveat, however, which is that for information passed there is a *quid pro quo* — it seems to be accepted that if a person receives information but does not pass any out, he is eventually excluded from the interactions. Dyer and Nobeoka (2000) find similar strictures against free riding in the Toyota network—firms that do not freely share their knowledge are sanctioned. In the Toyota network sanctions are formal whereas in von Hippel’s work they seem to be socially governed. In both cases, though, knowledge transfers can be seen as a form of barter, in which agents effectively trade knowledge and information.

Recent papers by Cowan and Jonard (2003, 2004) examine these knowledge transfer mechanisms using an underlying network structure to govern who interacts with whom.

5.1 Knowledge transfer as barter

Consider a population of N agents. Each agent has direct connections to, on average, n other agents. An agent is characterized by the set of agents to whom he is directly connected, and a knowledge endowment. Knowledge is of several types, so agents’ endowments are represented by a vector. When agents meet, they trade knowledge in a barter arrangement, exchanging knowledge of different types. The “exchange” however, only increases knowledge levels due to the non-rivalrous property of knowledge as a good. It is also the case that knowledge transfer is not complete. Absorptive capacity is not perfect. (See Cohen and Levinthal, 1989, for a discussion of this issue.)

Specifically, consider two agents i and j . Their respective knowledge endowments are $v_i=(v_{i,1},v_{i,2} \dots v_{i,K})$ and $v_j=(v_{j,1},v_{j,2} \dots v_{j,K})$. If i and j meet they make all possible trades, and we can describe how their knowledge vectors evolve in the following way. The K knowledge types are partitioned into pairs (m,n) such that $v_{i,m} > v_{j,m} + c$ and $v_{j,n} > v_{i,n} + c$. Then

$$\begin{aligned}
v_{i,m}(t+1) &= v_{i,m}(t); \\
v_{i,n}(t+1) &= v_{i,n}(t) + a(v_{j,n}(t) - v_{i,n}(t)); \\
v_{j,m}(t+1) &= v_{j,m}(t) + a(v_{i,m}(t) - v_{j,m}(t)); \\
v_{j,n}(t+1) &= v_{j,n}(t).
\end{aligned}$$

where a is the absorptive capacity, and c is the per trade transaction cost. Absorptive capacity is less than one, which implies that the post-trade knowledge levels do not converge completely. Agent j is able to learn only part of what i has to teach. A corollary of this is that as knowledge travels along a multi-agent chain, from i to j to k and so on, the knowledge degrades. Thus transmitting knowledge over a long chains is costly not only in terms of time, it is costly in terms of the diminution of the quantity of knowledge.

The dynamics of the model are very simple: Each period one agent is chosen randomly. At random he selects one of his direct connections, and with that agent makes all possible trades. Over time, agents' knowledge endowments change as they trade. The question of interest is how aggregate knowledge levels grow with these underlying micro-economics, and how aggregate growth is affected by network architecture. To answer that question, the model is simulated, using the Watts-Strogatz algorithm to create the networks over which this process takes place.

At the beginning of each run of the simulation a network of agents is created using the Watts-Strogatz mechanism for a given value of p . Agents are endowed with a knowledge vector $v_i=(v_{i,1},v_{i,2} \dots v_{i,K})$, where $v_{i,k}$ is drawn from a uniform $[0,1]$ distribution. Each period in simulation time a randomly drawn agent attempts trade with one of the agents to whom he is directly connected. The simulation is run until all possible trades have been exhausted, and then re-started with a new p value.

In a model like this, where no new knowledge is being created, it is clear that short path lengths should facilitate aggregate knowledge growth. Short path lengths imply that the knowledge that does exist is spread rapidly and widely through the population, and this would constitute a source of rapid growth in aggregate knowledge. This implies that a random network should be highly efficient. The results are different however. It is the small world network that is most effective in long run aggregate knowledge growth.

The industrial districts literature argues that cliquishness is good for knowledge creation, as described above. But none of those effects is explicitly included in this model. The value of cliquishness arises from a different source. One of the disadvantages of a bar-

ter economy is the requirement that potential traders satisfy a double coincidence of wants. When this fails, potentially improving transfers do not take place. Cliquish network structures mitigate this problem, because within a clique there are typically several indirect but still short paths between any two agents. Thus if i and j cannot trade due to a failed double coincidence of wants, the desired knowledge of j can still be transmitted to i if they share a common neighbour. This will be the case if the network is cliquish. Thus the local redundancy of cliquish graphs is an effective way of circumventing problems arising from the double coincidence of wants constraint.

The time series of knowledge levels for different network structures shows interesting patterns. In a random network, there is no discernable geographic structure, but path lengths are short. Here, diffusion is very rapid early in the history, but the process is exhausted at relatively low knowledge levels as the double coincidence of wants constraint starts to be violated. In a network that is very regular, such as the nearest-neighbour network, the graph is locally dense but path lengths are long. This implies that diffusion is slow. But because of the local density, the double-coincidence of wants issue does not bite, so aggregate knowledge growth can continue for a long time. In the small world region, the network is locally dense, but has relatively short paths. The latter should imply rapid initial growth; the former should imply continued growth. Indeed we observe these features of the small world structure. Growth in the small world is relatively rapid early on (though not quite as rapid as the random world) and it continues for a very long time.

The small world results are affected by absorptive capacity. The small world has value because it combines short path lengths with local redundancy. Redoing the experiment with different values for absorptive capacity shows that the optimal network structure, as measured by p in the re-wiring algorithm, becomes more and more random, that is, the optimal value of p increases with absorptive capacity. This occurs because the relative importance of path length and cliquishness changes as absorptive capacity changes. This can be seen by considering the extreme case of $a = 1$. In this case if i trades with j , *the* two become identical. A third agent, k , therefore, is indifferent between trading with i and trading with j . Thus if k has a link to j , there is no value to k to have a link to i . This says that there is no value to cliquishness — k would be better to have a link to someone that j is not connected to.

5.2 Knowledge transfer as broadcast

The model just described saw knowledge transfer as taking place through a bi-lateral barter exchange. A contrasting situation exists where agents engage in local broadcasting of their new knowledge. This can be modelled using a similar underlying structure, but changing the knowledge dynamics (Cowan and Jonard, 2003).

Consider a population of N agents each having direct connections to, on average n other agents. An agent is characterized by the set of agents with whom he has direct connections, referred to as his neighbours, and by a scalar knowledge stock. Each period, one agent innovates and broadcasts his knowledge to his neighbours. This knowledge is received by his neighbours, and partially absorbed. The population of agents is heterogeneous in two respects: ability to innovate and ability to absorb. Thus the characterization of an agent also includes his idiosyncratic innovation ability, b_i and his idiosyncratic absorptive capacity, a_i . The knowledge of an agent can increase for two reasons: the agent can innovate, a process which takes place exogenously; or the agent can receive new knowledge from a recent innovator. The knowledge dynamics for an innovator are simply

$$v_i(t+1) = (1+b_i)v_i(t).$$

If j is a neighbour of i , and i innovates, i then broadcasts to j . Agent j 's knowledge increases as:

$$v_j(t+1) = \max(v_j(t), v_j(t) + a_j(v_i(t) - v_j(t))),$$

and the knowledge of i is unchanged.

As in the previous section, in each run of the simulation a network is created for a particular value of p ; agents are assigned random absorptive capacities, random innovation abilities, and random initial knowledge levels. The simulation is run for 50,000 periods, and long run knowledge levels are recorded. This is performed for 1000 p values, and 4 sets of innovation and absorption levels.

The model can be used for two purposes: to compare the performance characteristics of different network structures, and to ask how innovation and diffusion over fixed network compares with random diffusion such as an epidemic model. A slight modification of the model lets us approximate an epidemic model of diffusion. In the model as described any time a particular agent innovates, he broadcasts his knowledge to the same set of n agents. These agents are his (unchanging) neighbours. Suppose instead that when an agent broadcasts, he broadcasts to n agents chosen at random from the entire population.

Here agents do not have a fixed set of neighbours, so the distribution of knowledge in the first stage is completely random, as it is in the standard epidemic model. Comparing this model to the one described above will give an indication of the value of a fixed network as a diffusion vehicle.

In this model, similarly to the barter model discussed above, the magnitude of absorptive capacity has an important effect on the results. When absorptive capacity is low ($a < 0.5$) small worlds dominate in terms of long run knowledge levels. The region of $0.01 < p < 0.1$ produces the highest long run aggregate knowledge levels. However, when absorptive capacity is high ($a > 0.5$) the random world ($p = 1$) dominates, and the extent to which it dominates increases with a .

Absorptive capacity also affects the comparison between networked and random, epidemic diffusion. If absorptive capacity is low, a fixed network performs better than random diffusion, no matter the value of p . By contrast, when absorptive capacity is high, random diffusion is better than almost all network structures. The value of the network comes from the repetition of interaction, which is most valuable when to convey knowledge effectively takes more than a single interaction.

Small worlds perform well as structures over which to diffuse knowledge and support innovation. But there are situations in which the small world is not optimal. If the knowledge and the community of agents surrounding it is such that transmitted knowledge is easy to absorb, then a random world, which has short paths between agents, is most efficient. However, a more general result is that networks per se are a good thing — they perform better than a world in which knowledge is diffused to random agents each time. It is important to note here that this is all in the context of networks with constant density. So random diffusion does not mean broadcasting to the entire economy, but rather to a randomly chosen subset of it.

Both of the models just described were models of diffusion. In a variation of the barter model Cowan et al. (2004), introduce innovation by the knowledge recipients. Here agents are not merely passive receivers of knowledge but they actively attempt to use their new knowledge to make innovations. In this model, if i innovates and then broadcasts to j , then if $v_i(t) > v_j(t)$ the knowledge level of j changes as:

$$v_j(t) = v_j(t) \left[1 + \left(\frac{v_j(t)}{v_i(t)} \right)^\gamma \left[1 - \left(\frac{v_j(t)}{v_i(t)} \right)^\gamma \right] \right]$$

If the recipient is well below the broadcaster in knowledge level, it may be difficult for the recipient to understand the message, and so absorption is low. At the other extreme, if the recipient is close to the broadcaster, there is little new in the message. At intermediate positions, it is possible, (particularly if γ is large), for the recipient to receive the information, innovate and finally achieve a knowledge level higher than the broadcaster. We understand this situation to represent an industry in which innovative potential is large.

In this model the small world effect is not present. The value of cliquishness tends to be outweighed by the effects of short paths. However, the value of short paths depends heavily on the parameter γ . When γ is small, so the industry has little innovative potential, knowledge advance is largely through imitating those who exogenously innovate. Short paths are valuable, and the random graph is most efficient. However, if γ is large, the reverse is true. In this case leap-frogging becomes an important phenomenon and agents who do so make big advances. They can rapidly leave the rest of the population behind. This means that their broadcasts are too esoteric for others to understand, and thus contribute little to aggregate population growth. This feature is highly attenuated, though, if the network is cliquish. Within a clique agents tend to be similar to each other, so relatively few agents get left behind, and cliques advance rapidly together. Cliquish structures produce the highest growth rates.

The same argument explains the comparison of a network structure with random broadcasts. Again the value of γ is important. When innovation potential is high, networks perform better than a homogeneous mixing. When γ is low, the reverse is true. Again it has to do with agents getting too far ahead of the rest of the population, and so contributing little to aggregate growth.

This can be interpreted as relating to the industry life cycle. In a new industry, innovative potential is high, as there are many avenues unexplored, and many discoveries yet to be made. This corresponds to a high γ , and so suggests that the optimal network structure for communication will be cliquish. By contrast, in a mature industry, dominant designs exist, innovation is more exploitation than exploration, and so innovative potential is much lower. This can be represented by a low γ , and so the optimal network structure will be random. These results fit well with the discussion of absorptive capacity. In a new industry, typically knowledge is less-codified, and so absorptive capacity is low. One of the features of the emergence of a dominant design is that the cost-benefit calculation for

codification changes, and codification tends to be a more common activity. This will increase absorptive capacity, and so change the optimal network structure to be more random.

In spite of the wide play the notion of small worlds has had in general, in economics the efficiency properties of small worlds have been little explored. One exception is White (2001), who develops a trade model in which agents must trade only over a network structure, but in a context of imperfect information. He compares the performance of the economy, in terms of whether or not it reaches a Pareto optimal equilibrium, and importantly, how long it takes to get there, accumulating transactions costs along the path. He finds that when agents trade over a small world network structure the economy performs better on these criteria than it does using different trading infrastructures: agents in a small world rapidly reach the Pareto optimal equilibrium (as does a complete network, which mimics perfect competition) but with very small search costs.

Recent empirical analysis of economic networks has focused to a very large extent on the network structure. In particular, economists have asked whether various networks have the small world structure. In general the answer is “yes”.² Many examples of small world networks in economics have been observed: patent citations among Dutch firms (BRITE/EURAM network, the 5th Framework TSER network, Dutch patent citations (Cowan and Jonard, 1999); Strategic alliances Duysters and Verspagen 2002; LANL co-authorship in a variety of academic disciplines (Newman, 2001; Barabasi et al. 2001); patent citation in US biotech (Johnson and Minerva, 2002), interlocking corporate directorships in the US (for example, Davis et al., 2003). There has been considerable work now describing economic networks, and it is clear that the small world structure is pervasive, yet there has been little empirical work addressing the efficiency of different structures. Research along these lines could make an important contribution to our understanding of different economic structures.

6 NETWORK FORMATION

In much of the literature that uses networks as an analytical tool the network structure is taken as given. Agents have links to particular other agents and these links do not change

2. One suspects that there may be an observation bias here though — it may be difficult to publish a paper saying that a particular network of economic agents is not a small world.

over time. This is clearly a very special case, as most networks evolve continuously and nodes create and destroy links to other nodes.

Roughly two strands of literature have emerged in economics to study this phenomenon. A literature based on a game-theoretic approach defines a notion of stability and asks which network structures are stable under different conditions.³ Here the aim is to characterize agents' incentives to form or break links, and then to look for configurations in which no agent wants to form new links or break existing one. A central question here is whether stable networks are efficient. A second literature has a much more evolutionary approach, in which in principle every agent is connected to every other, but links are activated probabilistically. The probability of i activating a link to j depends on whether this link has been valuable in the past, so as experience is gained with different links, the probability of choosing them changes. This model is very similar in spirit to the Hebb learning model in psychology (Hebb, 1955).

6.1 Pairwise Stability

In the game-theory tradition authors look for stable network structures. Stability is determined at the agent level. If no agent wishes to destroy any of his existing links, and no pair of unconnected agents wish to form a link between them, then the structure is stable. An early example of such a model is the connection model of Jackson and Wolinsky (1996). In this model an agent pays a cost c to participate in a link. If agents i and j are connected they receive benefits from that connection of $\delta^{d(i,j)}$, where $\delta < 1$ and $d(i, j)$ is the number of links in the shortest path between i and j . In this simple model, which network structures are stable is clearly determined by the relative sizes of c and δ . If $\delta + \delta^2(N - 2)/2 < c$ costs are too big and no connections are formed—the empty network is stable. If $c < \delta - \delta^2$, connections are cheap, and the complete network is stable. For intermediate values of c , there are many possible stable networks, one of which is the star network.

Many models build on this simple connection model, and the results are relatively robust. The continued emergence of empty, complete and star networks is striking (and

3. See for example Jackson and Wolinsky (1996); Bala and Goyal (2000); Goyal and Vega-Redondo (2000) Jackson and Watts (2001); Watts (2001); Jackson and Dutta (2000). Jackson (2004) is a nice survey.

in some conditions a circle or a wheel emerges as well). Very roughly speaking, the density of the stable network increases as costs of forming links decreases. But because of the nature of network structures, there tend to be critical values across which we observe dramatic changes in the emergent network architecture. Generally, stable networks are efficient, but there is a region of the parameter space (for intermediate values of the costs of forming links) in which there could be a tension between stability and efficiency: some stable networks may be efficient, but others are not.

Pairwise stability is both a strong and a weak equilibrium concept. It is strong in the sense that it characterises a stationary state. It describes an unchanging network. But it could be weak in that it can describe fragile networks. Carayol and Roux (2004) argue that introducing a small amount of noise into the link formation algorithm is more realistic and implies that the stochastically stable state is a more reasonable solution concept. Further, they introduce the assumption that costs of link formation will be lower when partners are located near to each other (equivalently, the probability of meeting a potential partner increases as the distance between them decreases).⁴ Introducing this type of preferential meeting into a simulation model of network formation implies first that many pairwise stable states (and in particular the star) are not stable, and second that the stochastically stable networks often have small world properties.

A paper that deserves particular attention here is Goyal and Joshi (2002). This is one of the first papers to introduce solid micro-economics into a network formation model (another early paper is Goyal and Moraga, 2001). Here firms form bilateral relations to reduce costs, and then compete on a product market. As is observed in most industries, firms can have more than one partner, forming additional partnerships without the permission of current partners. Both the costs of forming links and the nature of product competition determine the nature of the emerging network. Using the definition of stability introduced by Jackson and Wolinsky (1996), Goyal and Joshi show that when the costs of link formation are low, the complete network is both uniquely stable and uniquely efficient under quantity competition. By contrast, if the firms are competing on prices, the empty network is uniquely stable, whereas a linked double star is uniquely efficient. If competition is described more generally by the conditions under which firms can make positive profits, then it is the strength of competition that matters. When all firms make

4. See also Johnson and Gilles (2000), who also discuss spatially dependent costs.

positive profits but lower cost firms make higher profits, then under strong but not unreasonable conditions the complete network is uniquely stable. When only the lowest cost producer(s) make positive profits, then the network takes on a dominant group architecture: there is a (small) complete sub-network of the low-cost producers, accompanied by the remainder of firms all of which are isolated. This network structure also emerges when costs of link formation are high, and competition is Cournot. What this paper shows is that even with minimal structure relating the network structure and firms' costs (a firm's costs fall linearly with the number of its partners), the nature of competition is crucial in determining emergent network structures. Thus we should expect very different structures to emerge in different industries.

6.2 Network formation by learning

Several models of network formation have been developed in which agents make connections to other agents probabilistically, increasing the probability of connecting to a particular agent if experiences with him are positive. (See for example Weisbuch et al. 2000; Kirman and Vriend, 2001, or Plouraboué et al. 1998, for example.) This approach has been used to study very different phenomena, and has produced a wide array of results. And while the models within this tradition are too varied to have a "standard, baseline model", they tend to have the feature that even though in principle networks evolve forever, they do tend to settle down and exhibit aggregate regularities. For example, Kirman and Vriend (2001) in a market with buyers and sellers, find that, as in real markets, the network structure stabilizes to a state in which there is significant buyer-seller loyalty, which most buyers tending to use only one seller. At the same time, though, the network retains some fluidity, as there is a group of non-loyal buyers, who even in the long run continue to patronize a large number of sellers. This structure of a stable, relatively unchanging core of activity, accompanied by enough fluidity to prevent the network from ossifying is a typical result in these models

There is a long line of empirical work in the strategic management literature focussed on the inertia in alliance formation. It is now generally accepted that the current state of the network is an important variable for explaining future link formation. The general conclusion is that this is the result of accumulated joint experience between part-

ners, which generates both trust and a shared cognitive space. (See for example Garcia-Pont and Nohria, 2002; Gulati, 1999; or Powell et al. 1996.)

6.3 Stable matching and learning

There have been attempts to strike a balance between these two extremes, in which agents are relatively rational, yet the market can, in principle, evolve forever. One example (Cowan et al. 2000) looks at network formation in the context of innovation and diffusion. In that model agents have a knowledge endowment characterised by an amount and a type. Agents form pairs and as a pair innovate, adding the new knowledge to their respective stocks. Over time, pairs form and dissolve, and a network emerges. The pair formation is treated as a “matching problem”, and in each period the set of pairs is stable.⁵ A stable match is one in which no two agents prefer each other to their assigned partners (if they did they would leave their assigned partners and form a pair). Put another way, a matching assignment is stable if there is no blocking coalition of size 2 (any larger blocking coalition would include a blocking coalition of size 2). The structure of the network depends heavily on the nature of innovation. If innovation is driven largely by knowledge quantities the network tends to evolve to isolated pairs of agents — partnerships never dissolve. On the other hand, if diversity of knowledge is important in innovation, agents have many different partners over time, and the resulting network is very dense. What drives this result is the assumption that if two agents work together to innovate, they learn from each other, and so become more similar to each other than they were at the beginning of the episode. This implies that if diversity in knowledge is important, if two agents repeatedly innovate together, eventually they will be too similar to be valuable to each other. This change between regimes is very abrupt, almost a phase change. There is a small region of the parameter space, though, where intermediate structures seem to emerge. In this intermediate region, where both diversity and quantity are important, small cliques of (more than 2) agents form which permit enough diversity to be preserved. This is the region in small worlds might appear. These results fit well with the static results described above — empty networks (which in this model are approximated by isolated pairs) and complete networks (which in this model are approximated by networks in which agents are connected to roughly 60 percent of other agents) emerge. We observe no stars here,

5. See Gale and Shapley (1963) on matching problems.

however, and the region in which structures other than these two extremes are possible is very small.

An interesting empirical question about network formation links to the discussion of small worlds. The motivation of firms to create links that generate a cliquish structure is well-known from the literature on agglomeration and industrial districts.⁶ But a more interesting question is the mechanism or microeconomics motivation behind the formation of inter-clique links. One paper that examines this issue is Baum et al. (2003). Using data on syndicates formed in the merchant banking industry in Canada between 1952 and 1990, Baum et al. ask about the motivations of banks to form clique-spanning ties. It is possible to characterise links in a small world as one of two types — intra-clique links, and inter-clique links. Clique formation is based on the belief in the value of multi-lateral (as opposed to purely bilateral) arrangements. Resource pooling, risk sharing, the formation of critical masses and so on, provide incentives to create a group of interlinked agents rather than a simple pair. To form a link with a particular agent, or to include him or her in a group, it is necessary to know about his competencies, needs, intentions and resources. With this information one can judge both whether that agent would bring value to a connection and whether he would be interested in forming the connection. There are two sources for this information. First, past ties with that agent generate a relational embeddedness, which contributes to stability in the network structure. Second, common partners yield information, and this is seen as a source of structural embeddedness. Motives underlying intra-clique ties seem relatively straightforward, and have been explored in other literatures, for example work on agglomeration economies. More problematic are clique-spanning ties.

Baum et al. suggest three possibilities: chance, “insurgent partnering” and “control partnering”. “Insurgent partnering” refers to the action of a peripheral firm which is attempting to find a more central, and thus more important place in the industry. In social network analysis terms, it observes a structural hole in the network of merchant banks, and attempts to fill it by connecting with banks on either side of it. “Control partnering” by contrast, occurs when an already central firm attempts to strengthen its position as a leader by finding a new source of information. That is, a firm that is at the centre of a sub-net-

6. See for example, Uzzi(1997) or Baker (1990), on strategic issues of link formation

work connects with another sub-network in order to access the information there, and to control its flow throughout the larger network.

To establish which of these strategies accounts for clique-spanning ties, Baum et al. regress a bank's rate of clique-spanning tie formation on measures of a bank's centrality, the whether a bank occupies structural holes and several bank-specific attributes. They find that all three strategies are present, but that chance and insurgent partnering dominate. That is, a network changes from a cliquish but not globally well-connected network to a small world because peripheral firms attempt to find a more central place in the structure. They see opportunities to gain power in the industry by bringing together two more or less existing cliques.

Goyal et al. (2004) look at co-authorship of academic economics papers. What they observe is a striking change in the network over the last 30 years. They find that as the prevalence of co-authored papers increases the size of the giant connected component likewise increases. It remains the case, though, that the average distance between economists in the connected component remains small. The frequency distribution of links is very skewed, and the structure of the graph is one of interlinked stars. It appears that a small world is emerging in economics itself.

7 CONCLUSIONS

The use of network analysis in economics is growing rapidly. Network models offer ways to overcome several of the shortcomings of more traditional market models. When those models do not apply, network modelling may provide advantages. This is particularly the case in the context of innovation, when the transmission of knowledge plays such an important role. Codified knowledge can be diffused globally, without reference to underlying connection structures, but tacit knowledge has very different properties. Its diffusion depends heavily on the network of agents through which it will spread. To understand the dynamics of innovation and diffusion network models provide a promising avenue.

One recent development in network analysis more generally has been the introduction of the concept of a scale-free network (see for example Barabasi and Albert, 2000). Scale-free networks are characterised through the frequency distribution of links, which follows a power law. This is a highly skewed distribution, with a high variance: most nodes have very few links, a few nodes have many links. This type of network, which has been

observed in many natural phenomena, has the property of being very robust to the failure of random nodes. However, it is not robust to directed attacks. That is, failures of particular nodes can be catastrophic to the performance of a scale free network. This development is only just beginning to come into economics, but it fits well with the recent work on small worlds: many networks appear to be both scale free and small world (as one example, see Riccaboni Pammolli, 2002, who find that networks in the life sciences and ICT industries have scale free properties). Thus one interesting, and unaddressed question is how these two concepts fit together in the context of economics, and whether the issue of robustness has any relevance for how an industry functions. As yet, we have little idea about the efficiency properties of scale-free networks.

The next major challenge for economists lies in the empirical analysis of network structures, and particularly in understanding the efficiency properties of existing structures. Beyond case studies, the data requirements for this type of exercise are very heavy. Patent data-bases are rich enough, but provide evidence of only one type of information flow, and so construction of large, detailed data-bases is a task ahead.

Research on networks is being done in a wide variety of fields: psychology, sociology, strategic management, physics and so on. Economics is one branch with particular interests, and so it is facile to say that economics must look outside for inspiration. Nonetheless, the interconnection between economic activity, networks, and other aspects of network-mediated behaviour creates a rich research environment in which this approach will continue to yield valuable insights into the way modern economies function and perform.

8 REFERENCES

- Allen, R., 1983. "Collective Invention" *Journal of Economic Behavior and Organization*, 4:1-24.
- Allen, B., 1982. "Some Stochastic Processes of Interdependent Demand and Technological Diffusion of an Innovation Exhibiting Externalities among Adopters", *International Economic Review*, 23, 595-608.
- An, M. and N. Kiefer, 1993. "Local Externalities and Societal Adoption of Technologies", *Journal of Evolutionary Economics*, 5, 103-117.
- Anderson, S.P., A. de Palma, and J.F. Thisse, 1992. *Discrete Choice Theory of Product Differentiation*. MIT Press, Cambridge, Massachusetts.
- Arora, A. and A. Gambardella, 1990. "Complementary and External Linkages: The Strategies of the Large Firms in Biotechnology", *Journal of Industrial Economics*, 38(4), 361-379.
- Baker, W.E., 1990. "Market Networks and Corporate Behaviour" *American Journal of Sociology*, 96:589-625.
- Bala, V. and S. Goyal, 1998. "Learning from Neighbours". *Review of Economic Studies*, 65:595-621.
- Bala, V. and S. Goyal, 2000. "A Non-cooperative Model of Network Formation", *Econometrica*, 68: 1181-1229.
- Barabasi, A. and R. Albert, 2000. "Statistical Mechanics of Complex Networks", *Review of Modern Physics*, 74: 47-97.
- Baum, J.A.C., A. V. Shipilov, and T. J. Rowley, 2003. "Where Do Small Worlds Come From?", *Industrial and Corporate Change*, 12 (4).
- Becker, Gary S. 1962. "Irrational Behavior and Economic Theory" *Journal of Political Economy*, 70(1) pp. 1-13.
- Becker, Gary S. 1963. "Rational Action and Economic Theory: A Reply to I. Kirzner," *Journal of Political Economy*, 71(1) pp. 82-83.
- Blume, L. 1995. "The Statistical Mechanics of Best-Response Strategy Revision", *Games and Economic Behavior*, 11(2),111-145.
- Breschi, Stefano, and Francesco Lissoni, 2003. "Mobility and Social Networks: Localised Knowledge Spillovers Revisited," CESPRI Working Papers 142, CESPRI, Centre for Research on Innovation and Internationalisation Processes, Universita' Bocconi, Milano, Italy.
- Carayol, N. and P. Roux 2003. "Self-Organizing Innovation Networks: When do Small Worlds Emerge?," Working Papers of GRES - Cahiers du GRES 2003-8, Groupement de Recherches Economiques et Sociales.
- Cohen, W. and D. Levinthal, 1989. "Innovation and Learning: The Two Faces of R&D", *The Economic Journal*, 99, 569-596.
- Cowan, R. and N. Jonard 2003. "The Dynamics of Collective Invention", *Journal of Economic Behavior and Organization*, 52(4): 513-532.
- Cowan, R. and N. Jonard 2004. "Network Structure and the Diffusion of Knowledge" *Journal of economic Dynamics and Control*, vol 28: 1557-1575
- Cowan, R., N. Jonard and M. Ozman, 2004. "Knowledge Dynamics in a Network Industry", *Technological Forecasting and Social Change*.
- Cowan, R., N. Jonard and J.-B. Zimmermann, 2003. "On the Creation of Networks and Knowledge" in Gallegati, M., A.P. Kirman, and M. Marsili (eds.), *The complex Dynamics of Economic Interaction: Essays in Economics and Econophysics*, Berlin: Springer.
- Davis, G. F., M. Yoo, M. and W. E. Baker, 2003. "The Small World of the American Corporate Elite, 1982-2001", *Strategic Organization*, 1(3): 301-326.
- Durlauf, S., 1993. "Non-ergodic Economic Growth", *Review of Economic Studies*, 60, 349-366.
- Dyer, J. H. 1996. "Does Governance Matter? Keiretsu Alliances and Asset Specificity As Sources of Japanese Competitive Advantage", *Organization Science*, 7(6) 649-666.
- Dyer, J. H. and K. Nobeoka, 2000. "Creating and Managing a High Performance Knowledge-Sharing Network: The Toyota Case", *Strategic Management Journal*, 21, 345-367.

- Ellison, G. 1993. "Learning, Local Interaction and Coordination", *Econometrica* 61(5), 1047-1071.
- Föllmer, H. 1974. "Random Economies with many Interacting", *Journal of Mathematical Economics*, 1, 51-62.
- Gale, G. and L. Shapley, 1962. "College Admissions and the Stability of Marriage", *American Mathematical Monthly*, 69: 9-15.
- Garcia-Pont, C. and N. Nohria, 2002. "Local versus Global Mimetism: The dynamics of alliance formation in the automobile industry", *Strategic Management Journal*, 23:307-21.
- Glaeser, E., Sacerdote, B. and J. Scheinkman, 1996. "Crime and Social Interaction", *Quarterly Journal of Economics*, 111, 507-548.
- Goyal, S. and F. Vega-Redondo, 2004. "Network Formation and Social Coordination", *Games and Economic Behavior*, forthcoming.
- Goyal, S., Marco van de Leij and José Luis Moraga-González, 2004. "Economics: An emerging small world?" mimeo, University of Essex.
- Goyal, S. and S. Joshi, 2002. "Networks of Collaboration in Oligopoly", *Games and Economic Behavior*, 43, 57-85.
- Goyal, S. and J-L. Moraga, 2001. "R&D Networks" *Rand Journal of Economics*, 32(4): 686-707.
- Granstand, O., 1998 "Towards a Theory of the Technology-based Firm" *Research Policy* 27(5): 465-489.
- Granstand, O. and Sören Sjölander, 1990. "Managing Innovation in Multi-technology Corporations", *Research Policy*, 19(1): 35-60.
- Gulati, R., 1999. "Network Formation and Learning: The influence of network resources and firm capabilities on alliance formation" *Strategic Management Journal*, 20:397-420.
- Hagedoorn, John, 2002. "Inter-firm R&D Partnerships: An overview of major trends and patterns since 1960", *Research Policy* 31(4), 477-492.
- Hebb, D.O. 1949. *The Organization of Behavior: A Neuropsychological Theory*, New York: John Wiley Sons.
- Holmstrom, Bengt and John Roberts, 1998. "Boundaries of the Firm Revisited", *Journal of Economics Perspectives*, (4)12: 73-94.
- Hotelling, Harold, 1929. "Stability in Competition", *Economic Journal* 39, 41-57.
- Imai, K. 1989. "Evolution of Japan's Corporate and Industrial Networks", in B. Carlsson (ed.), *Industrial Dynamics, Technological Organizational and Structural Changes in Industries and Firms*, Kluwer, Boston, 123-156.
- Ioannides, Y., 1990. "Trading Uncertainty and Market Form", *International Economic Review*, 31, 619-633.
- Ising, E. 1925. "Beitrag zur Theorie des Ferromagnetismus." *Zeitschrift fur Physik* 31, 253-258.
- Jackson, M.O. 2003. "A Survey of Models of Network Formation: Stability and Efficiency," mimeo, California Institute of Technology.
- Jackson, M.O. and Watts, A. 2001. "The Evolution of Social and Economic Networks" (revised version of California Institute Working Papers No: 1044 - 1998).
- Jackson, M.O. and A. Wolinsky, 1996. "A Strategic Model of Social and Economic Networks. *Journal of Economic Theory*, 71:44-74.
- Jackson, M.O. and B. Dutta, 2001. "On the Formation of Networks and Groups" in M. Jackson and B. Dutta (eds.), *Models of the Strategic Formation of Networks and Groups*, Springer.
- Jaffe, A, Trajtenberg, M. and R. Henderson, 1993. "Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations", *Quarterly Journal of Economics*, 63, 577-598.
- Johnson, C. and R.P. Gilles (2000) "Spatial Social Networks" *Review of Economic Design*, 5: 273-299.
- Johnson, D.K.N. and M. Mareva, 2002. "It's a Small(er) World: The role of geography and networks in biotechnology innovation", Wellesley College Working Paper 2002-01.

- Kirman, A. P., 1983. "Communication in Markets: a Suggested Approach", *Economics Letters*, 12, 101-108.
- Kirman, A. and N. Vriend, 2001. "Evolving Market Structure: An ACE model of price dispersion and loyalty" *Journal of Economic Dynamics and Control* 25 (2001): 459-502.
- Kirzner, Israel M. 1962. "Rational Action and Economic Theory," *Journal of Political Economy*, 70(4) pp. 380-385. (Reply to Becker, Gary S., "Irrational Behavior and Economic Theory,") *Journal of Political Economy*, 70(1) (Feb.) pp. 1-13).
- Kirzner, Israel M. 1963. "Rational Action and Economic Theory: Rejoinder" *Journal of Political Economy*, 71(1) pp. 84-85.
- Langlois, R. and P.L. Robertson, 1992. "Networks of Innovation in a Modular System: Lessons from the microcomputer and stereo components industries", *Research Policy*, 21(4):297-313.
- Lazerson, M.H. and G. Lorenzoni, 1999. "The Firms that Feed Industrial Districts: A return to the Italian source", *Industrial and Corporate Change*, 8(2): 235-266.
- Lincoln, J.A., and M. Gerlach, 1998. "The Structural Analysis of Japanese Economic Organization: A Conceptual Framework" in W. Mark Fruin (ed.): *Networks and Markets; Pacific Rim Strategies* New York: Oxford University Press.
- Narula, R., and J. Hagedoorn, 1999. "Innovating through strategic alliances: moving towards international partnerships and contractual agreements" *Technovation* 19: 283-294.
- Nuvolari, A. forthcoming. "Collective Invention during the British Industrial Revolution: The Case of the Cornish Pumping Engine" *Cambridge Journal of Economics*
- Orsenigo, L., F. Pammoli, and M. Riccaboni, 2001. "Technological Change and Network Dynamics: Lessons from the pharmaceutical industry", *Research Policy*, 30: 485-508.
- Plouraboué F., Steyer, A. Zimmermann, J.B. 1998. "Learning Induced Criticality in Consumer's Adoption Pattern: a neural network approach", *Economics of Innovations and New Technology*, vol. 6 p. 73-90.
- Powell, W.W., K.W. Koput and L. Smith-Doerr, 1996. "Interorganizational Collaboration and the Locus of Invention: Networks of Learning in Biotechnology", *Administrative Science Quarterly*, 41:116-145.
- Pyke, F., G. Becattini, and W. Sengenberger, (eds.), 1992. *Industrial Districts and Inter-firm Cooperation in Italy*. (Geneva: International Institute for Labour Studies).
- Riccaboni, M. and F. Pammolli, 2002. "On Firm Growth in Networks", *Research Policy*, 31(8-9):1405-1416.
- Salop, S. C. 1979. "Monopolistic competition with outside goods," *Bell Journal of Economics*, 10, 141-156.
- Saxenian, A. (1991) "The Origins and Dynamics of Production Networks in Silicon Valley", *Research Policy*, 20: 423-37.
- Steyer, A. and J.-B. Zimmermann, 1998. "On the Frontier: Structural effects in a diffusion model based on influence matrixes" in P. Cohendet, P. Llerena, H. Stahn, and G. Umbhauer (eds.), *The Economics of Networks: Interaction and Behaviours*, 118-135, Springer.
- Uzzi, B. 1997. "Social Structure and Competition in Interfirm Networks: The paradox of embeddedness", *Academy of Management Journal*, 42(1):35-67.
- Von Hippel, E, 1998. *The Sources of Innovation*. Oxford University Press.
- von Hippel, E. 1987. "Cooperation Between Rivals: Informal Know-How Trading", *Research Policy*, 16: 291-302.
- Watts, D. and S. Strogatz, 1998. "Collective Dynamics of Small-world Networks" *Letters to Nature*, 393.
- Watts, A. 2001. "A Dynamic Model of Network Formation", *Games and Economic Behaviour*, 34: 331-341.
- Weisbuch, G., A. Kirman, and D. Herreiner, 2000. "Market Organisation and Trading Relationships" *Economic Journal* 110:411-436.
- Wilhite, A. 2001. "Bilateral Trade and 'Small-World' Networks" *Computational Economics* 18: 49-64.