

Network structure and the diffusion of knowledge

Citation for published version (APA):

Cowan, R., & Jonard, N. (1999). *Network structure and the diffusion of knowledge*. MERIT, Maastricht Economic Research Institute on Innovation and Technology. MERIT Research Memoranda No. 026 <https://doi.org/10.26481/umamer.1999026>

Document status and date:

Published: 01/01/1999

DOI:

[10.26481/umamer.1999026](https://doi.org/10.26481/umamer.1999026)

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.

Network Structure and the Diffusion of Knowledge*

Robin Cowan and Nicolas Jonard

MERIT

PO Box 616, 6200 MD Maastricht, The Netherlands

July 29, 1999

Abstract

This paper models knowledge diffusion as a barter process in which agents trade different types of knowledge. It captures the observed practice of informal knowledge exchange among agents with related, though different knowledge. Agents are located on a lattice and are directly connected with a small number of other agents. Agents repeatedly meet those with whom they have direct connections and trade if mutually profitable trades exist. In this way knowledge diffuses throughout the economy. We examine the relationship between lattice structure and diffusion performance. A one dimensional, periodic lattice has as one extreme a completely regular structure—every agent is connected to n nearest neighbours. At the other extreme every agent is connected to, on average, n agents located at random on the lattice. We use a re-wiring procedure to examine the space of lattice structures that fall between these extremes. We find that the performance of the system exhibits clear ‘small world’ properties, in that the steady state level of average knowledge is maximal when the lattice structure is that of a small world (that is, when most connections are local, but roughly 5 percent of the connections are long distance). A surprising result is that the variance of knowledge levels among agents is also maximal in the small world region. We explain both these results as reflecting the dynamics of knowledge transmission as affected by the nature of connections among agents.

*We thank participants of seminars and conferences in Marseille (GREQAM), Maastricht, Munich and Genova, as well as T. Ziesemer, J.-B. Zimmermann, S. Luchini and A. Kirman for helpful comments. Authors E-mail addresses are racowan@mud.cgl.uwaterloo.ca and n.jonard@merit.unimaas.nl.

1 Introduction

One of the remarkable features of a market is its ability to process information. Even with all of the myriad factors that can influence economic phenomena, in principle, in a well-functioning market economy, agents need only pay attention to one piece of information for each good, namely its price. An implication of this feature of markets is that, to the extent that agents interact with each other, interactions are anonymous. Any interaction between two agents is mediated by the market, and I need know nothing about you except the price you are charging or asking, which, in equilibrium, will be the same as that charged or asked by every other agent.¹ Recently however, some economists have been examining phenomena in which interactions are not anonymous. Agents differ from each other in significant ways. In particular, any agent only interacts with a small subset of the set of all agents. This set may have different general characteristics from the population average—different locations, endowments, technologies or tastes. The particular characteristics of an agent’s ‘neighbours’ will clearly influence the behaviour and welfare of the agent. Non-anonymity becomes important.

While economists understand that, and in an abstract sense how, a market processes information, there is little explicit examination of it. In particular, in this tradition there is no mention of the network or communication structures under which agents operate and transmit or exchange knowledge and information. By contrast, in the literature on interacting agents, communication structures play a central role. The details of who is connected to whom affects what type of information is passed, how much, and how efficiently. All of these can have an effect on the aggregate performance of the system being modelled.

Observing both the empirical importance of communication network structures, and the theoretical result that key to the behavior of the models have been the information transmission structures imposed on them, two questions immediately arise: first, if the network structure is exogenous, precisely how do the structural properties of the network affect aggregate outcomes ? and second, if network formation is endogenous, what structures are likely to emerge ? The current paper makes some progress in answering the first of these questions, namely how the communication network structure affects the aggregate performance of the system.²

This paper addresses explicitly the nature of the social communication network and its effect on the aggregate performance of the system. We

¹This is of course an extreme characterization of standard economic models, but nonetheless represents a significant part of the economics literature.

²See Bala and Goyal (1998) or Plouraboue et al. (1998) e.g. on the second question.

take as the motivating application a system of knowledge diffusion. In our model knowledge diffuses through barter exchange among pairs of agents. Aggregate performance is measured as the mean knowledge level over all agents. The parameter we vary is the degree of spatial randomness in the links between agents through which knowledge can pass. At one extreme we have an entirely local network—every agent is connected to his n nearest neighbours. At the other extreme agents are connected randomly to, on average, n other agents, located anywhere. We examine the space of network structures between these extremes. In terms of the aggregate performance of the system, one region of the space stands out. This region corresponds to the “small world” network structure as defined formally by Watts and Strogatz (1998), and informally by Granovetter (1973). We discuss this below.

While the body of literature on interacting agents uses different interaction structures to model the network through which agents communicate, attention to their effects has been limited. Two extreme structures are common. At one extreme, in generalized Ising models all interactions are local and the structure of the interactions completely regular. Every agent is directly connected to the same small number of his nearest neighbours. Neighborhoods tend to be small but the intersection of the neighborhood of two nearby agents is relatively large. Put in other words, the structure is very ‘cliquish’ in the sense that most of my neighbours are neighbours of each other. On the other hand, though, the average path length between two agents is relatively large. This structure has been used to examine macroeconomic dynamics (Durlauf, 1991 and 1993); technology diffusion (Allen, 1982; An and Kiefer, 1993; David et al. 1998); the effectiveness of prices in stabilizing an economy (Föllmer, 1974); criminal behaviour (Glaeser et al., 1996). Clustering in behaviour is a common result: in steady states we observe neighbourhoods in which agents’ behaviour is similar to that of other agents in the neighbourhood, but different from that of agents in other neighbourhoods.³

At the other extreme of interaction or network structures are models based on random graph theory. Here, in principle every agent could be connected to every other, and there is no spatial structure imposed on the connections. Thus interactions are global in one of two senses. In some models a mean field approximation is used, in which agents are assumed to be affected by the average behaviour of the population. Alternatively, in some models connections are random, exhibiting no spatial order, and links are thus global in the sense that any agent could in principle be connected to any other, regardless of location. Here, the structure is not at all cliquish, on

³A recent survey on these and related topics can be found in Kirman (1998).

average my ‘neighbours’—those agents to whom I am directly connected—have nothing to do with each other but average path lengths are short because any long-distance link can serve as a ‘short-cut’ between agents. This general structure of global or non-ordered connections has been used to examine coalition formation (Kirman, 1983; Ioannides, 1990); innovation diffusion (Plouraboue et al., 1998).

A common way to address structural influence is to vary the network by changing the number of connections per agent. Results in percolation theory can be interpreted this way: increasing either the link probability or the node probability effectively increases the number of links, on average, that are activated. It is now well-known that aggregate performance (for example extent of diffusion of an innovation) undergoes a phase change as the structure changes along this dimension. (See for example David et al., 1998). In virtually all of the work that explicitly addresses the relationship between network structure and aggregate performance, though, the emphasis has been on the density of the connections, rather than on their spatial structure.⁴ The focus in this paper, by contrast, is specifically on the spatial structure of the connections, and we deliberately hold the number of connections constant throughout our analysis.

2 Knowledge

The production and diffusion of knowledge has long been viewed as a vital component of economic growth. Empirical growth accounting studies, work on national systems of innovation, and endogenous growth models all make this point in different ways. In this line, taking knowledge accumulation as being associated with technological change, there has been a significant amount of empirical, econometric research on the diffusion of specific technological innovations. But we know that this is not the only form that knowledge accumulation takes. Learning by doing (Arrow, 1962) and learning by using (Rosenberg, 1982) are perhaps the best examples of knowledge accumulation that is not associated with new technologies. The existence of learning curves, which describe progress within a technology, indicates some sort of knowledge accumulation that does not fit the “new technology” description. The kind of learning associated with learning curves is often not

⁴One exception is Midgeley et al. (1992) who have a model with spatial characteristics—agents are located in cliques, within an industry, or outside the industry. Again they find that increasing the number of connections in the model increases diffusion speeds, though different types of connections (intra-clique; inter-clique; outside the industry) have different effects.

codified (Cowan and Foray, 1997), and thus is only transmitted in face to face interactions. If knowledge generated by learning by using and doing is diffused, therefore, models of its diffusion must take explicit account of the structure of connections between agents. The position that some knowledge is diffused only through face-to-face interactions is consistent with the work of Jaffe et al. (1993) on patent citation, in which patent citation is shown to be a geographically localized phenomenon.⁵ There are several possible explanations for this, but one is that there is some knowledge that is not being globally diffused.

While not denying the importance of other kinds of knowledge and modes of diffusion, in this paper we focus on this particular kind of knowledge—knowledge that is exchanged in face-to-face communications. In-depth empirical studies indicate that there is a well-established informal network through which knowledge is traded (von Hippel, 1987; Hicks, 1995; Schrader, 1991). Even among competitors knowledge is exchanged, but in a barter arrangement. Conferences, publications and conversations in the bar are all situations in which agents give up information they have generated. It is this kind of knowledge trading and the network structure in which it takes place that we model in the next section.

3 The model

The model we design explores how the topology of agent interactions influences aggregate regularities. It is one in which agents are characterized by knowledge endowments that evolve over time through a simple process of barter exchange. This model is particularly well-suited for representing knowledge ‘trading’ and diffusion among heterogenous agents. Note that since agents trade knowledge, models in which knowledge is treated as a scalar cannot capture an important aspect of knowledge diffusion. In contrast to most other models here knowledge is represented as a vector (we can think of this as accounting for several distinct types of knowledge).

Agents exist on a graph, connected with a relatively small number of other agents. Whenever two agents (i and j) meet, that is if a connection between them is activated, they trade knowledge. Trades are only possible, though, if i has superior knowledge of one type while j has superior knowledge of another type. In the model links are activated sequentially and randomly, so that agents are acting in a changing world. Agents i and j meet and make all possible trades, but before they meet again, j may have met k , traded, and

⁵For similar evidence see for example Prevezer and Swann, 1996; Feldman, 1994.

in doing so have created more possible trades with i . It should be clear that we are not concerned with radical innovations, and particularly not product innovations; the model captures effects of incremental innovations.

3.1 The interaction structure and the rewiring of connections

We consider a population of agents endowed with a network structure. The undirected graph associated with this network is written $\mathcal{G}(I, \Gamma)$, where $I = \{1, \dots, N\}$ is the set of agents and the correspondence $\Gamma = \{\Gamma(i), i \in I\}$ gives the list of vertices to which each vertex is connected. Formally $\Gamma(i) = \{j \in I \setminus \{i\} \mid d(i, j) = 1\}$, where $d(i, j)$ is the length of the shortest path from vertex i to vertex j . Only agents separated by one edge can interact.

3.1.1 Randomly ‘rewired’ graphs

A random graph is usually a graph in which there is either (i) a given probability (usually uniform) associated to the existence of an edge between any pair of vertices (cf. Kirman, 1983) or (ii) a given number of connections to be allocated according to a particular procedure among randomly chosen pairs of vertices (Ioannides, 1990). In either case, the correspondence Γ is randomly constructed, the tuning parameter being (i) the probability of observing an edge between two randomly chosen vertices or (ii) the total number of edges. The family of graphs we consider here belongs to this second category, but exhibits richer features. It contains a constant number $n \cdot N/2$ of edges, but in general these edges are not uniformly allocated among random pairs of vertices. We start with vertices arrayed on a circle with each vertex being only connected to its n nearest neighbors (n is even), and progressively disrupt this local order by randomly rewiring the edges. The tuning parameter is the probability $p \in [0, 1]$ of any link being rewired. Graphs that belong to this family are labelled $\mathcal{G}(I, n, p)$. Figure 1 depicts three example graphs for $N = 16$ and $n = 4$.

A re-wiring algorithm allows us to examine networks that are intermediate between regular, locally ordered, and totally random.⁶ We begin with a completely regular graph: a circular lattice with N vertices, each vertex having edges only to its n nearest neighbours (n is even). We operate on each edge of the vertex sequentially, but in a particular order. Begin with vertex one, and the edge connecting it to its nearest neighbour clockwise. With

⁶This is the procedure used in Watts and Strogatz (1998), but our results would hold for a broader class of rewiring procedures.




Figure 1: The transition from a locally ordered structure to a random one and the small world phenomenon for $\mathcal{G}(I, n, p)$ graphs.

probability p break the connection to the neighbour and re-connect that edge to a vertex chosen uniform randomly over the entire lattice. With probability $1 - p$, leave the edge unchanged. Progress around the lattice clockwise, considering one edge per vertex, and dis-allowing duplicate edges. After one complete round, repeat the procedure, considering the second nearest clockwise neighbour. Simply repeat the process, considering progressively more distant neighbours, until every edge has been considered once. By tuning the parameter p , we can vary the graph structure from completely regular ($p = 0$), through intermediate states, ($0 < p < 1$), to completely random ($p = 1$). This procedure creates variation in the number of edges per agent, but maintains an average of n connections per agent.

Figure 1 depicts three illustrative configurations with increasing disorder as p is increased, for $N = 16$ and $n = 4$. For small but non-zero values of p , the graph is highly clustered like a regular graph but has the low average path length of an almost-random graph. Hence, despite the fact that the number of edges is kept constant throughout the experiment (there are exactly $n \cdot N/2$ edges whatever p), tuning p exerts a crucial influence on the relational density of the interaction structure as well as the length of the path between any pair of vertices.

3.1.2 Structural properties

Watts and Strogatz (1998) point out that the structural properties of $\mathcal{G}(I, n, p)$ graphs are intuitively captured by the concepts of average path length and average cliquishness. If one thinks of social networks representing friendship, both have intuitive interpretations. The path length is the number of friendships in the shortest chain connecting two agents. Cliquishness reflects the extent to which friends of one agent are also friends of each other. Formally, defining $d(i, j)$ as the length of the shortest path between i and j , the average path length $\mathcal{L}(p)$ is

$$\mathcal{L}(p) = \frac{1}{N} \sum_{i \in I} \sum_{j \neq i} \frac{d(i, j)}{N-1} \quad (1)$$

and simply measures how distant vertices are on average, which is a global property of the graph.⁷ By contrast, average cliquishness $\mathcal{C}(p)$ is a measure of local connectivity in that it captures the share of active links between a given vertex's neighbors.⁸ It is written

$$\mathcal{C}(p) = \frac{1}{N} \sum_{i \in I} \sum_{j, l \in \Gamma(i)} \frac{X(j, l)}{|\Gamma(i)| (|\Gamma(i)| - 1) / 2}, \quad (2)$$

where $X(j, l) = 1$ if $j \in \Gamma(l)$ and $X(j, l) = 0$ otherwise. When $p = 0$ and N is large enough, $\mathcal{C}(0) = 3/4 \cdot (n-2)/(n-1)$ and $\mathcal{L}(0) \sim N/(2n)$. Conversely, $\mathcal{L}(1) \sim \ln N / \ln n$ and $\mathcal{C}(1) \sim n/N$ for large N . One might conjecture that large cliquishness is always associated with large path length and low cliquishness with low path length. Actually, as emphasized by Watts and Strogatz (1998), there is a non-negligible interval for p over which $\mathcal{L}(p) \simeq \mathcal{L}(1)$ yet $\mathcal{C}(p) \gg \mathcal{C}(1)$. This interval constitutes the ‘small world’ region (see figures 1 and 2). It follows from the fact that introducing a long range edge provides a shortcut not only between the two vertices that this link connects, but also for their immediate neighbours, the neighbours of those neighbours and so on. Thus when the number of long distance links is small their marginal effect on average path length is large. The existence of the small world region is a generic property of large sparse connected networks, Watts and Strogatz claim. The evolution of path length and clique size with p is depicted on figure 2, for a graph of $N = 500$ vertices, each vertex having on average $n = 10$ nearest neighbours. For the sake of clarity, we plot the normalized values $\mathcal{L}(p)/\mathcal{L}(0)$ and $\mathcal{C}(p)/\mathcal{C}(0)$.

⁷Note that an alternative measure would have been the diameter of the graph, i.e. the maximal path length over the graph.

⁸In graph theory, a clique is a set of vertices, any two of which are adjacent.

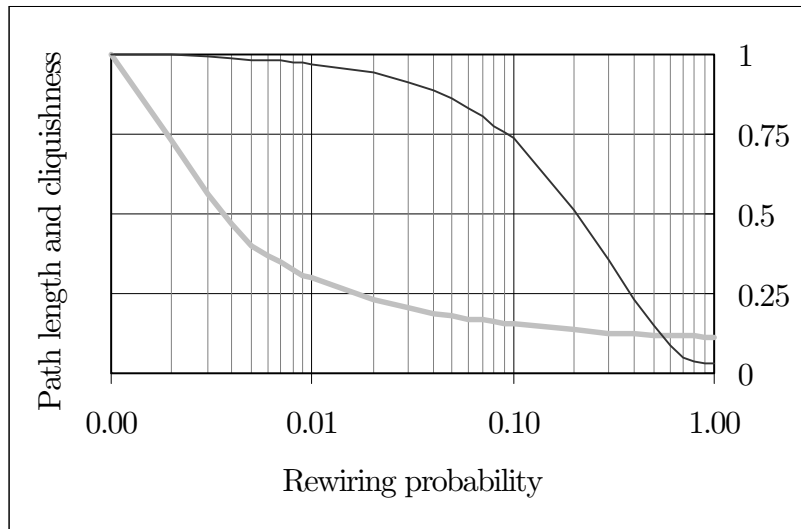


Figure 2: Average cliquishness and average path length as functions of p

The upper curve (thin black) on figure 2 is the normalized average cliquishness index $\mathcal{C}(p)/\mathcal{C}(0)$ for $p \in [0, 1]$. It remains almost constant when p is reasonably small and falls slowly for large values of p . By contrast, average path length (thick grey) as measured by $\mathcal{L}(p)/\mathcal{L}(0)$ falls quickly for very small p values and flattens out near 0.01. Hence we see that for $p \in [0.005, 0.1]$, clique size and path length diverge, creating a small world region in the space of network structures.

3.2 Knowledge interaction

Each agent is characterized by a knowledge vector that evolves over time through a barter exchange. Formally, let $V_i(t) = (V_{i,k}(t); k = 1, \dots, K)$ denote agent i 's knowledge endowment at time t . Agents i and j might interact provided $j \in \Gamma(i)$ (equivalently $i \in \Gamma(j)$ since the graph is non-directed) and if trading is mutually advantageous. Let $n(i, j) = |\{k \mid V_{i,k}(t) > V_{j,k}(t)\}|$ be the number of knowledge categories in which i strictly dominates j , for any $(i, j) \in I^2$. Then clearly trade can take place if and only if

$$n(i, j) \cdot n(j, i) \neq 0. \quad (3)$$

Condition 3 states that there should be some double coincidence of wants. An agent whose knowledge endowment totally dominates his counterpart's endowment has little to learn from the latter, hence little to gain from the

interaction. In that case barter simply does not take place. By contrast, as soon as equation 3 holds, trade is engaged and all possible trades are exhausted. Each agent gives and receives some knowledge in a number of knowledge categories equal to $\min\{n(i, j), n(j, i)\}$. We have assumed that knowledge is only partly assimilable,⁹ hence trading, when it occurs, results in a gain equal to a share α of the knowledge differential. Formally, category k knowledge spills over from the more wealthy agent according to $V_{i,k}(t+1) = V_{i,k}(t) + \alpha \cdot (V_{j,k}(t) - V_{i,k}(t))$. Note that ‘trade’ here is a barter—one agent transfers part of his knowledge to another one and is paid back with knowledge of a different category. As long as there is one category in which agent i dominates agent j and one in which j dominates i , both agents perceive trade as mutually advantageous, even if observation from an absolute point of view would suggest that one has less to gain than the other. A rule is necessary to solve cases in which equation (3) holds but $n(i, j) \neq n(j, i)$. In that particular case, it is simply assumed that the categories involved in trading are randomly chosen, with uniform probability.¹⁰

We add one more feature to the system we have just described, namely the presence of ‘experts’. There are a small number of agents who, initially, have much better information or knowledge than the population average. For each category of knowledge, there is a small number of experts, and the diffusion and absorption of that expertise is subject of the experiments in the next section.

4 Simulation results

We are interested in evaluating the efficiency, also sometimes referred to as ‘distribution power’, of the system (David and Foray, 1995) as it responds to the structural properties of the graph. We look, therefore, at knowledge level. Agent i ’s average knowledge level is $\bar{\mu}_i(t) = \sum_k V_{i,k}(t) / K$. The average level

⁹The strictly partial assimilation arises from the fact that tacit knowledge is needed to assimilate and use fully any piece of information (cf. Cowan and Foray, 1997). Further, the value of a piece of information lies in great part in its integration with other information, so typically it is not possible simply to ‘add’ a piece of information to an existing information structure. Put another way, absorptive capacity is never perfect (Cohen and Levinthal, 1990). Note also that with $\alpha < 1$ the model has the property that knowledge ‘degrades’ as it is transmitted. Thus, the longer the path a piece of information travels, the less value it is to the recipient.

¹⁰As a particular example, if $V_i(t) = (5, 6, 3, 5, 2)$ and $V_j(t) = (8, 2, 2, 9, 2)$, four types of knowledge are exchanged ($n(i, j) = n(j, i) = 2$) and assuming $\alpha = 0.5$ yields $V_i(t+1) = (6.5, 6, 3, 7, 2)$ and $V_j(t+1) = (8, 4, 2.5, 9, 2)$.

of knowledge in the economy at time t is

$$\bar{\mu}(t) = \frac{1}{N} \sum_{i \in I} \bar{\mu}_i(t)$$

and the standard deviation in knowledge allocation is

$$\sigma(t) = \frac{1}{N} \sum_{i \in I} \bar{\mu}_i^2(t) - \bar{\mu}^2(t).$$

In order to collect data on distribution power, we consider here a knowledge exchange economy with a population of $N = 500$ agents and $n = 10$ links per agents (hence a total number of 2500 edges). Each agent is endowed with a 5-category knowledge vector that is randomly initialized by giving each category a uniformly drawn value between 0 and 10. Parameter α , which accounts for the ease with which new knowledge is assimilated and integrated to the existing knowledge stock of an agent (and could well represent the extent to which knowledge is codified in the particular economy we consider) is set to 0.5. Hence, an agent is only able successfully to integrate fifty percent of the knowledge differential that exists between him and his counterpart each time this pair meets.¹¹ A large number of meetings is considered ($\mathbf{T} = 50000$) so as to ensure that each link (each edge of the graph) is activated frequently enough (this gives an average of 20 activations per link). The share of actual experts in the population is kept low (25 experts hence 5% of the whole population) and these are endowed, in one knowledge category for each, with a component equal to 30.

Each meeting consists in picking at random an agent in I and (randomly) activating one of his links (choosing an edge emanating from him), letting then all possible trades take place.¹²

To isolate the small world phenomenon, the rewiring probability p is varied from 0.001 to 1 by adding increments that allow us to represent results on a logarithmic axis. For each p value, forty replications are run and averaged to provide statistically significant evidence on the aggregate behavior of the system.

4.1 Small worlds and knowledge distribution

The expression ‘strength of weak ties’ coined by Granovetter to mark the crucial importance of indirect relationships in collective phenomena led to the

¹¹Increasing α changes the speed of the process but leaves the qualitative results unchanged (provided it remains strictly less than one).

¹²Not every meeting will result in trades—one member of the meeting pair may dominate the other in all knowledge categories.

acknowledgment of the fact that ‘the world is actually a small world’. This property, that Watts and Strogatz (1998) formally identify as the conjunction of a short average path length and a high degree of cliquishness, is clearly a property of the structure and not of agents’ particular modes of interaction. Thus it may not resist certain modes of interaction or, put another way, there may be no ‘isomorphism’ between the properties of the structure and that of the interactions it conveys. One might reasonably conjecture that decreasing path length is invariably associated with faster and wider diffusion. This is certainly true, but it should be kept in mind that the particular structure we employ here is one in which the number of links to be allocated among agents is kept constant. Hence a vertex that is ‘missing’ between two adjacent nodes, or equivalently a communication link that is broken between adjacent agents, harms the local efficiency of diffusion, thereby reducing the probability of seeing geographical clusters of ‘high-knowledge’ agents emerge. This suggests that path length and clique size act in opposite directions, cliquishness being a source of high redundancy and fast diffusion at the local level, while path length indicates how easily novelty can be obtained from distant parts of the economy. While this may be less important for the short run properties of the diffusion process, it makes a significant difference in the long run.




Figure 3: The economy-wide average knowledge level as a function of the rewiring probability p

Figure 3 depicts how varying p affects the long-run performance of the

economy in terms of knowledge diffusion. The long-run average level of knowledge $\bar{\mu}(\mathbf{T})$ is clearly a non-monotonic function of p , with a sharp efficiency peak in the small world region. As seen in figure 2, the small world region (wherein clique size is relatively large but path length is relatively short) lies roughly between $p = 0.005$ and $p = 0.1$. As measured by average knowledge levels, this system performs best at $p = 0.06$. Hence the small world structural property has a dynamic counterpart in the knowledge exchange economy that we consider. A clear implication of this is that the average path length or the diameter of a network of relationships is not an unequivocal measure of the ‘distribution power’ of this structure. Diminishing the distance between members of an organization by reallocating the links is far from being a sufficient measure, since how these links are allocated also matters.

4.2 Small worlds and heterogeneity

Examining inter-agent heterogeneity in the steady state provides further insights into the workings of our economy.




Figure 4: Agent heterogeneity

Figure 4 shows the evolution of the long-run variance $\sigma(\mathbf{T})$ of knowledge within the system as a function of the rewiring probability p . Surprisingly heterogeneity and efficiency respond in a similar manner to changes in p . Again a single-peaked curve obtains, the small world region corresponding

to the highest level of disparities in terms of knowledge allocation. Hence there is a region in which the structural properties of the network of agents' relationships is at the same time producing the best overall performance in terms of how much knowledge diffuses through the system, and the worse overall performance when homogeneity of the allocation is considered. The conclusions one draws from this will depend, of course, on whether efficiency or equity of knowledge diffusion is considered more important.

A tentative explanation to this apparent paradox might be advanced here. With zero or very small p , the network is extremely dense at the local level, with strong redundancy implying fast access to locally available knowledge, but slow access to more distant one. There will be two types of regions on the graph: those without local experts, where knowledge diffuses homogeneously but remains at a relatively low level; those with local experts, where expertise creates a knowledge cluster. There will be a relatively small number of the latter but they will achieve, within the cluster, higher levels of knowledge. In a small world, though, agents with direct connections to experts relatively rapidly become experts themselves, thus forming new (possibly distant) high-knowledge clusters. The high cliquishness of the small world ensures that agents near experts reinforce their knowledge levels. This keeps them distinct from regions that are not closely connected to experts. The small world conditions are such that there is at the same time an (almost) intact local redundancy but an increased number of areas in the economy where expert knowledge becomes available. This effect continues to the point when further increasing the share of long distance links implies that accessing expert knowledge has almost no local effects since no local order exists anymore. The peak in the variance of knowledge endowments therefore results from the multiplication of access points to expert knowledge, while at the same time keeping local interrelated clusters with homogenous levels of knowledge.

4.3 Small worlds and diffusion dynamics

To this point, our analysis has focussed on long-run properties. We now turn attention to the transitory properties of the model. The speed at which knowledge diffuses is also a major policy concern and we can use the model to examine how this is affected by network architecture. Figure 5 depicts the time series of the average knowledge level for three characteristic values of the rewiring probability p .

The lowest p -value (0.007) is located in the regular (locally ordered) network zone, in which diffusion is locally efficient though there are only a small number of access points to expert knowledge. Obviously the economy-wide




Figure 5: Diffusion time path

diffusion of information is sluggish and only dwindling pieces of knowledge are exchanged as time passes. The highest p -value (0.7), by contrast, belongs to an area where many long-distance vertices coexist with a small number of short-range links. In this region, sizable informational ‘jumps’ are accomplished during a large number of meetings before trading possibilities become exhausted. Finally, the small world region shows its distributional superior strength but figure 5 also reveals that its dominance takes time to emerge. If short-run distribution properties are considered important, structures in which the average path length is as short as possible definitely dominate.

4.4 On Clique Size

The results presented indicate that the behaviour of our knowledge diffusion system change as the degree of randomness of the network architecture changes. The architectural changes of the network have been encapsulated in two structural parameters, namely average path length and cliquishness. The effect of average path lengths is relatively obvious, intuitively—a short path can only increase the diffusion power of the system; knowledge will move to different parts of the graph more quickly. But what effect does clique size have? A cliquish graph has the property of local redundancy.¹³

¹³Young (1998, chapter 6), in a model of games played on graphs, explores the importance of network structure in the diffusion of particular strategies. He uses the concepts

What this implies is that if agents i and j are in the same clique, there are many paths by which knowledge can pass between them. This turns out to be important in this model due to the imperative of the double coincidence of wants. If i has knowledge that would benefit j but there is a failure of the double coincidence of wants between i and j , there are many other possible paths along which the knowledge can be transmitted. This implies a certain role for the clique structure. In a relatively cliquish graph, an agent should have knowledge that is very similar to that of the agents with whom he is connected. In a non-cliquish graph, an agent will, in general, have knowledge that is relatively dis-similar to that held by those to whom he is directly connected. In the latter case, there will be knowledge held by one agent which would be of value to an agent directly connected to him which cannot be transmitted. Speaking loosely, “potentially beneficial transfers” are not made. Since transfers only increase average knowledge, the effect of reduced cliquishness is to reduce average knowledge levels. This is illustrated in Figure 6 which shows average knowledge dis-similarity as a function of the rewiring probability p .

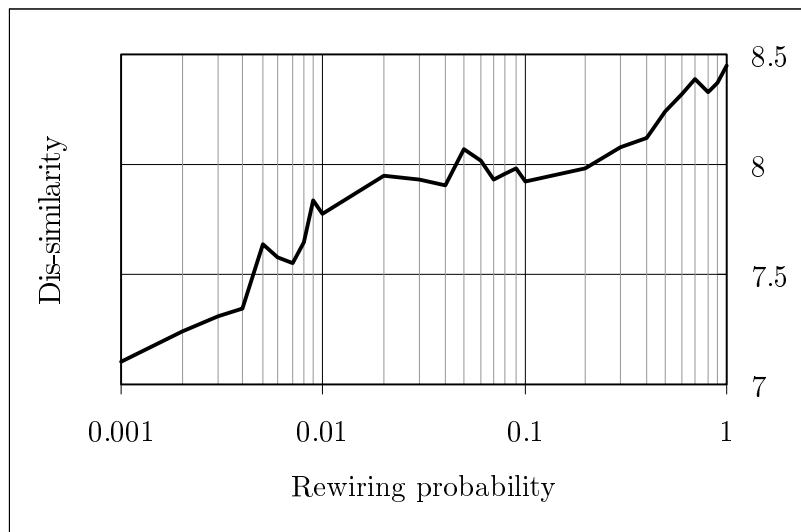


Figure 6: Agents dis-similarity

Here, knowledge dis-similarity refers to a global measure of how distant any agent and his neighbours are. It is defined as the average over neighbouring agents of pairwise distances in the knowledge space. In particular,

of ‘cohesiveness’ and ‘close-knittedness’ to characterize the geometry of the network.

i 's dis-similarity index is

$$\mathcal{D}(i) = \frac{1}{|\Gamma(i)|} \sum_{j \in \Gamma(i)} \frac{1}{K} \sum_k |V_{i,k} - V_{j,k}|,$$

where $V_{i,k}$ is the level of knowledge of type k held by agent i . Index $\mathcal{D}(i)$ equals 0 whenever i and all his neighbours are identical and increases as they become increasingly dissembling. We plot the average over agents.¹⁴

5 Discussion and Conclusions

In this paper we have modelled one of the processes by which knowledge is diffused. As more and more policy makers come to be concerned with the so-called knowledge economy, this diffusion becomes a central policy issue. We have shown that the extent of diffusion is clearly affected by the structure of the network over which the diffusion takes place, and that there is an identifiable region of the space of structures in which the diffusion is much more complete than elsewhere in the space. This “small world” region exists where the proportion of long distance, as opposed to short distance links is between 5 and 10 percent of all direct links between agent pairs.

Technology and knowledge policy, especially in Europe, is currently going through a phase (one might be tempted to say craze) in which clustering and localization is seen as an extremely important phenomenon. Policy makers are very keen to find policies which will encourage clustering, and create new “Silicon Valleys” in new places and different industries. This craze is based on the belief that the knowledge transmission mechanisms modelled here are extremely important in any innovation system. Our results indicate, though, that this policy goal must be treated with some circumspection; it is possible to have too much clustering. It is very important to maintain or even build strong links outside the cluster.

Our results also raise a problem for any policy maker involved in regional technology or knowledge policy. This is the very old efficiency/equity trade-off. In the region of the space of network structures in which diffusion results in a high average knowledge level, it also results in a high variance among agents. That is, the distribution of knowledge levels is relatively unequal. To the extent that distribution remains a policy concern, if knowledge is

¹⁴This result is robust to the definition of dis-similarity. For example, defining dis-similarity as $\frac{1}{K} \sum_k \frac{1}{|\Gamma(i)|} \sum_{j \in \Gamma(i)} \Phi(V_{i,k}/V_{j,k})$, where $\Phi(x) \equiv \min\{x, 1/x\}$, we find again a strong positive relationship between dissimilarity and p .

considered a key input to wealth or income, policies aimed at inducing efficient knowledge diffusion will have to address the consequent distribution of income.

The model could be extended in several obvious ways. We have taken the network structure as given, and have examined its effect on the knowledge diffusion process. But clearly, networks evolve in response to agents' experiences. Thus an agent who has had a very successful exchange with one particular agent is likely to try to return to that agent in the future. By this sort of mechanism the strengths of links between agents will change as they gain experience with the network. Policy would be very interested in the circumstances in which a network evolves into a small world network, with the associated 'nice' properties.

In the model in this paper, there is in effect no innovation, only diffusion, which we can interpret as the consequence of a single innovative episode. In general, though, innovations happen continually. Including innovation introduces a potentially complex feedback. Part of the rationale for technology policy focussing on spatial clusters is that it is believed that within a cluster, ideas meet and interact, producing an environment conducive to rapid creation of new ideas. Thus introducing ongoing innovation introduces another trade-off—in a very cliquish graph, this sort of local interaction conducive to idea creation will be significant. But our results show that highly cliquish networks have poor diffusion properties. There is, therefore a new form of the trade-off between production of innovations, which in general is a good thing for wealth creation, and the diffusion of those ideas, which is a similarly good thing.

The model as constructed is very general in its form and can be extended and applied in a variety of ways. It shows again the value of being well-connected, but also the value of having shortcuts available. Applying the message to the scientific enterprise, we should specialize, but not be closed to ideas that seem to come from very far afield.

References

- 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 Kiefer, N., 1993. "Local Externalities and Societal Adoption of Technologies", *Journal of Evolutionary Economics*, 5, 103-117.

- Arrow, K., 1962. "The Economic Implications of Learning-by-doing", *Review of Economic Studies*, 29, 155-173.
- Bala, V. and S. Goyal, 1998. "Learning from Neighbours", *Review of Economic Studies*, 65, 595-621.
- Cohen, W. and D. Levinthal, 1989. "Innovation and Learning: the Two Faces of R&D", *The Economic Journal*, 99, 569-596.
- Cohen, W. and D. Levinthal, 1990. "Absorptive Capacities: a New Perspective on Learning and Innovation", *Administrative Science Quarterly*, 35, 128-152.
- Cowan, R. and D. Foray, 1997. "Quandaries in the Economics of Dual Technologies and Spillovers from Military to Civilian Research and Development", *Research Policy*, 24, 851-868.
- David, P. and D. Foray, 1995. "Accessing and Expanding the Science and Technology Knowledge Base", *STI Review*, 16, 13-68.
- David, P., Foray, D. and Dalle J.-M., 1998. "Marshallian Externalities and the Emergence and Spatial Stability of Technological Enclaves", *Economics of Innovation and New Technology*, 6, 147-183.
- Durlauf, S., 1991. "Multiple Equilibria and Persistence in Aggregate Fluctuations", *American Economic Review P&P*, 81, 70-74.
- Durlauf, S., 1993. "Non-ergodic Economic Growth", *Review of Economic Studies*, 60, 349-366.
- Föllmer, H. 1974. "Random Economies with many Interacting Agents", *Journal of Mathematical Economics*, 1, 51-62.
- Feldman, M., 1994. *The Geography of Innovation*, Kluwer Academic Publishers.
- Glaeser, E., Sacerdote, B. and Scheinkman, J., 1996. "Crime and Social Interaction", *Quarterly Journal of Economics*, 111, 507-548.
- Granovetter, M., 1973. "The Strength of Weak Ties", *American Journal of Sociology*, 78, 1360-1380.
- Hicks, Diana, 1995. "Published Papers, Tacit Competencies and Corporate Management of the Public/Private Character of Knowledge", *Industrial and Corporate Change*, 4, 401-424.

- Ioannides, Y., 1990. "Trading Uncertainty and Market Form", *International Economic Review*, 31, 619-633.
- Jaffe, A, Tratjenberg, M. and Henderson, R., 1993. "Geographic localization of knowledge spillovers as evidenced by patent citations", *Quarterly Journal of Economics*, 63, 577-598.
- Kirman, A. P., 1983. "Communication in Markets: a Suggested Approach", *Economics Letters*, 12, 101-108.
- Kirman, A. P., 1998. "The Economy as an Interactive System", 491-531, in *The Economy as an Evolving Complex System II*, Arthur, W. B., Durlauf, S. and Lane, D. (eds.), Addison Wesley.
- Midgley, D, D. Morrison, and J. Roberts, 1992. "The Effect of Network Structure in Industrial Diffusion Processes", *Research Policy*, 21, 533-552.
- Plouraboue, F., A. Steyer and J-B Zimmermann, 1998. "Learning Induced Criticality in Consumers Adoption Pattern: A Neural Network Approach", *Economics of Innovation and New Technology*, 6, 73-90.
- Prevezer, M. and P. Swann (1996). "A Comparison of the Dynamics of Industrial Clustering in Computing and Biotechnology", *Research Policy* vol.(25)7 pp. 1139-1157.
- Rosenberg, N., 1982. *Inside the Black Box: Technology and Economics*, Cambridge University Press.
- Schrader, S. 1991. "Informal Technology Transfer Between Firms: Cooperation through Information Trading", *Research Policy*, 20, 153-170.
- von Hippel, E., 1987. "Cooperation Between Rivals: Informal Know-how Trading", *Research Policy*, 16, 291-302.
- Young, P. H., 1998. *Individual Strategy and Social Structure*. Princeton University Press.
- Watts, D. and S. Strogatz, 1998. "Collective Dynamics of Small-World Networks", *Letters to Nature*, 393, May.