

Structural holes, innovation and the distribution of ideas

Robin Cowan · Nicolas Jonard

Published online: 25 July 2007
© Springer-Verlag 2007

Abstract We model knowledge diffusion in a population of agents situated on a network, interacting only over direct ties. Some agents are by nature traders, others are by nature “givers”: traders demand a quid pro quo for information transfer; givers do not. We are interested in efficiency of diffusion and explore the interplay between the structure of the population (proportion of traders), the network structure (clustering, path length and degree distribution), and the scarcity of knowledge. We find that at the global level, trading (as opposed to giving) reduces efficiency. At the individual level, highly connected agents do well when knowledge is scarce, agents in clustered neighbourhoods do well when it is abundant. The latter finding is connected to the debate on structural holes and social capital.

1 Introduction

This paper examines the relationship between the architecture of an industrial R&D network and efficiency in knowledge distribution, both from the point of view of individual firm performance, and at the level of the system.

Recent technological changes have had the effect of creating multi-product firms: the knowledge base on which both production and innovation are founded has, in general, become much broader, covering more, and different types of knowledge

R. Cowan
BETA, Université Louis Pasteur, 61 Avenue de la Forêt Noire, 67000 Strasbourg, France

R. Cowan (✉)
UNU-MERIT, Maastricht University, P. O. Box 616, Maastricht, 6200 MD, The Netherlands
e-mail: r.cowan@merit.unimaas.nl

N. Jonard
Université du Luxembourg, 162a Avenue de la Faiencerie, 1511 Luxembourg, Luxembourg

(Granstrand and Sjolander 1990). As a consequence, firms increasingly discover that their in-house knowledge is not sufficient for efficient production or innovation. This had driven them to seek the knowledge they need outside, in other firms. However, and precisely due to the nature of knowledge, this task is difficult to achieve through pure market interaction. Thus, firms are now forming relatively long term alliances, formal and informal, with other, often competing, firms. This has led to the networked organization, a hybrid organizational form lying between the market and a pure hierarchy (Powell 1990), which takes advantage of both market and non-market interactions. Networks supply firms with rapid, flexible access to resources outside their core competencies. Strong, stable contacts with other firms can provide a firm with the knowledge it needs for its immediate production or innovation without navigating the difficulties of market transactions for knowledge. In addition, contacts of this type can also provide a form of insurance—giving a firm rapid access to information about developments taking place in other firms or related industries. These observations are particularly relevant to knowledge-intensive, and science-based industries.

There is now a large literature, much located in the management field, examining the structural properties of innovation or R&D networks. Three properties recur. Networks tend to be sparse. That is, of the total possible connections between agents, the actual connections constitute a small proportion. Networks tend to be locally dense. Local clusters of closely interconnected agents are common. In addition though, the local clusters tend to be only sparsely connected to each other. Finally, the distribution of links over agents tends to be highly skewed. It is probably too strong to say there is common evidence of power law distributions, but relatively heavy tails do exist. The discussion in the literature concerns how these properties arise, and how firms' performance is affected by them.

As to the relationship between network position and performance, there are roughly speaking two competing views. On the one hand, following Coleman (1988), it is possible to argue that dense sub-groups are a source of social capital.¹ A group of highly inter-connected agents generates trust, common languages and problem-solving heuristics, social disapprobation for opportunistic behaviour and so on. If firms i and j are linked, they can share information about a common partner k . This reduces significantly the incentives for k to behave opportunistically against, j , even if he will never see j again, since information about his behaviour will travel rapidly to i and to all the other members of the clique. More positively, if i is working on a problem, using information gathered from j and k , or in discussion with j and k , an ability of j and k to discuss the problem with each other, or exchange information about it, can only have a positive effect on i 's ability to solve his problem. These considerations imply that structurally embedded partnerships will be important sources of value for a firm, and redundant links are privately, and probably also socially valuable.² Thus for a firm, a useful link formation strategy is to close open triangles, and create strong cliques.

¹ See also Walker et al. 1997 on the same subject.

² The third link that closes a triangle can be seen as redundant since its effect is simply to create a path of length one between two agents where a path of length 2 already existed. Notice here that "redundant" is only strongly applicable if this reduction in path length serves no purpose, that is, if in general path lengths are not (privately) important.

The value of this strategy is observed in empirical studies by [Dyer and Nobeoka \(2000\)](#), on the automobile industry; [Gulati and Gargiulo \(1999\)](#) in a study of alliance formation in several industries; [Powell et al. \(1996\)](#) who study the impact of network position on innovation performance in the biotechnology sector; and [Rowley et al. \(2000\)](#) in a study of strong and weak ties in innovation networks in the steel and semiconductor industries.

On the other hand, [Burt \(1992\)](#) argues that dense local links are redundant in a strong sense, that the existence of structural holes in a firm's ego network is efficient, and that locally dense networks can be a source of rigidity. A structural hole exists if two of my neighbours are not linked to each other. Through these two neighbours I am connected to different parts of the larger network, and thus have access to different sources of dispersed information. Thus if a firm is to form a new link, closing a structural hole is less valuable than finding a partner to whom none of my current partners is currently connected. This is closely related to the argument of [Podolny \(1993\)](#), that firms attempt to increase their betweenness centrality. If many shortest paths between firms go through firm *i*, then *i* can exert considerable control over knowledge flows. Particularly in the knowledge economy, control of knowledge flows can be translated into rents. This is an argument that clique spanning ties are valuable. The value of structural holes has been examined empirically by, for example, [Ahuja \(2000\)](#) in the context of the international chemical industry (structural holes have a negative impact on industry performance, whereas indirect and direct ties have a positive impact on firm innovative performance); [Gargiulo and Bennassi \(2000\)](#) who find in a study of an Italian IT firm that dense local networks do not respond well to change (they find a trade-off associated with the safety conferred by cohesive ties (social capital) and the flexibility conferred by ties that connect different parts of a network); [Baum et al. \(2003\)](#) study the sources of inter-clique link formation in the Canadian merchant banking industry.

This debate between social capital and structural holes is formalised through the notion of clustering. An agent's ego network is clustered if many of its partners are partners of each other. The structural holes argument claims that a highly clustered ego network is bad for performance; the social capital position argues the opposite. At an aggregate level, individual clustering levels can be averaged to describe an industry (or sector or economy) network. By extension, the structural holes argument implies that unclustered networks will perform well, whereas the social capital position argues that locally dense networks, which by definition are highly clustered, will perform well. One way in which these positions are sometimes reconciled (see for example [Rowley et al. 2000](#)) is that they apply to different moments in an industry life cycle. When an industry is young, technologies are being explored, and many different avenues of advance are potentially fruitful. Here, it is important to have rapid access to "distant" (both in geographical and technological space) information. Thus redundant ties are less valuable than ties that connect to different parts of the network. Structural holes are desirable. However, in a more mature industry, there are fewer technological surprises, so exploitation is more common. Here, a dense core of agents addressing similar issues creates the critical mass that is necessary to make further progress along the chosen path. Social capital becomes more valuable.

While clustering has received much attention in the literature, there is a second aspect of structure that is now considered important. This is the distribution of edges over nodes, and in particular the extent to which this distribution is skewed, with a few nodes having many links and the majority having few. Though the skewness of the degree distribution is relatively well established (see Powell et al. 2005 for example), its implications for R&D networks are not well understood. The genesis of a skewed distribution lies in some form of “preferential attachment”: firms with many existing links are likely to be valuable partners, and so attract more links. On the face of it, this is intuitively appealing, since a large number of partnerships indicates that a firm both has useful knowledge and knows how to collaborate. But a skewed distribution implies the presence of stars in the network—agents through whom many (short) paths run. Consequently, while this structure is robust to random failures (since the failure of a randomly selected node affects most likely affects only a few other nodes), it is very fragile to specific failures (if a star fails, many other nodes, and paths between many pairs of nodes, are affected). Stars in a network can serve as important centres of knowledge distribution, and so a highly skewed network may be conducive to very rapid diffusion of knowledge. On the other hand, though, if a star ceases to participate in the system for some reason, this can cause a serious disruption in the distribution system, and if the situation is such that agents withdraw from the system from time to time (for whatever reason) a flatter link distribution may be preferable.

While network structure will clearly have an effect on the efficiency of diffusion, structure may interact with the micro-specifics of exchange. The transmission of knowledge, particularly among competing firms, is a challenge for economists, especially if there is no market for knowledge. Two patterns of transmission have been observed empirically. Allen (1983) describes “collective invention” in which knowledge is given away as a (local) gift. In the steel industry in Cleveland UK in the mid nineteenth century, for example, steel producers met regularly under the auspices of societies like the Cleveland Institute of Engineers, the South Wales Institution of Engineers or the national Iron and Steel Institute and disclosed their own recent technological developments. As a producer made an advance in furnace height or temperature, for example, that producer would document the change—how it was accomplished, the technical effects and so on—and present this to other local firms. Knowledge was essentially given away to competitors within the local cluster, and as a consequence, the technology developed rapidly.³ von Hippel (1987) on the other hand documents a barter exchange. Technical managers of steel mini-mills in the US exchange technical information and explicitly help each other solve problems.⁴ But here the transfer is not a gift: there is a quid pro quo. While the interaction is not market-based, there are social sanctions if an agent routinely receives but does not give knowledge. In essence,

³ McGaw (1987) finds a similar pattern in paper manufacturing in New England in the early 1880s, and Lamoureaux (1999) cites other examples from the eighteenth and nineteenth centuries in the US.

⁴ von Hippel also found the same phenomenon in aerospace and waferboard industries. Powell et al. (1996) document a similar phenomenon in biotech.

knowledge is bartered. In both of these cases, knowledge transmission is local, taking place in face-to-face interactions. If this represents the nature of knowledge diffusion, then the structure of local interactions will play a central role in the process through which a “piece of knowledge” moves from one geographic location to another, or more generally how it diffuses throughout an economy.⁵

These are the issues we take up in this paper. We are interested in how the architecture of the communication network affects its performance in terms of knowledge distribution. This relationship may change, however, depending on the details of the transmission mechanism. The model we construct below permits us to examine both architecture and transmission as variables controlling knowledge diffusion.

2 The model

In general, for any agent, more knowledge is better from the point of view of producing goods, and the communication network is the infrastructure over which agents acquire the knowledge they need in order to produce.

To capture this, we model a world in which there is a fixed, finite population of agents, and a fixed, finite number of ideas relevant to production. An agent is characterised by two properties: the set of ideas he has, and his production goal. Production is controlled by a Leontieff production function for which only a small number of ideas is necessary, different agents having different production functions. Production is done in isolation, but demands that an agent possesses all the ideas that are relevant to his productive activity, i.e. the ideas for which his production function has non zero coefficients. If one or more ideas are missing, they can be acquired via an agent’s acquaintances. Thus the set of ideas held by any agent evolves over time, and we shall assume this takes place through a simple process of one-to-one exchange or gift.

Ideas have the feature that agents do not lose by giving. Thus it might be argued that agents could well give without asking for reciprocity. If agents have to compete in a second step, however, even though an agent’s knowledge level does not decrease by giving away ideas, his competitive position might. Different industrial contexts will display different “terms of trade”. In the present paper, we will consider a world of knowledge traders, a world of knowledge givers, and a mixed situation in which both co-exist.

2.1 Network structure

Let $G(V, N)$ be the undirected graph representing the industry network, with $V = \{1, \dots, n\}$ the set of agents and $N = \{N_i, i \in S\}$ the correspondence specifying, for each $i \in V$, the neighbourhood N_i of i . The degree of firm i is the number of direct

⁵ For models examining the relationship between network structure and knowledge diffusion see Cowan and Jonard (2003) on the gift economy, and Cowan and Jonard (2004) on a barter economy.

ties of that firm

$$n_i = \#N_i.$$

Average degree in the network is then $\bar{n} = \sum_{i \in V} n_i/n$. Any $j \in N_i$ is at distance 1 from i . Indirect ties connect i to individuals at a distance strictly more than one. Define d_{ij} the distance between i and j , as the number of edges in the shortest path, or geodesic, connecting i to j . The average distance to i is then

$$d_i = \frac{1}{n-1} \sum_{j \neq i} d_{ij}$$

and average distance (characteristic path length) is $\bar{d} = \sum_{i \in V} d_i/n$.

An additional measure of the structure of local links is clustering. The extent to which i 's neighbourhood is clustered is measured by

$$c_i = \frac{2}{n_i(n_i-1)} \sum_{j,l \in N_i} \xi_{jl}$$

where $\xi_{jl} = 1$ if $j \in N_l$ and $\xi_{jl} = 0$ otherwise. This statistic measures the proportion of existing triangles among those which could involve i , given n_i . The clustering coefficient for the network as a whole is $\bar{c} = \sum_{i \in V} c_i/n$. Though more sophisticated measures of structural position can be designed, for our purpose the distribution and organization of direct and indirect ties as captured by degree, clustering and distance are sufficient.

In order to examine the relationship between structure and performance, we create an algorithm to construct a family of random graphs derived from an ordered substrate. As a first step, consider the re-wiring algorithm from [Watts and Strogatz \(1998\)](#) that has as a control parameter the probability p that a link in a periodic lattice is randomly rewired. Start from the periodic lattice with an even number m of nearest neighbours ($n_i = m$, for all $i \in V$) and sequentially consider each edge, making a decision of uniform random rewiring (with probability p) or preservation (probability $1-p$). As p increases the regular periodic lattice ($p = 0$) is left and through intermediate states ($0 < p < 1$) a random graph with uniform degree is reached ($p = 1$). This procedure creates a small amount of variation in individual degree n_i and an average degree of m is preserved as the total number of edges is kept constant (it is exactly equal to $nm/2$, half the degree sum). We know from [Watts and Strogatz \(1998\)](#) that there is an interval (the small world region) over which clustering remains high while path length has fallen close to the level of a random graph of average degree m . Indeed, when the number of random links is small the removal of a few of them has a strong effect on average path length while it has only little effect on the clustering coefficient.

Besides clustering and path length, we are interested in the effect of asymmetry in the degree distribution. There has been extensive debate (see for instance [Barabási and Albert 1999](#)) about scale-free networks (that is, networks with a power law degree distribution having exponent between 2 and 3) and the extent to which they can be

found in empirical data. Scale-free networks and power law distribution in general are not our interest here. Rather, we will simply explore the effect of having stars in the system, in addition to the possibility of having more or less clustered random structures. We do this the following way. Assume now there are two classes of agents: s stars with degree $d_i = D$ for $i \in S \subseteq V$, and $n - s$ non-stars with degree $d \leq D$. The problem is now to design a procedure analogous to the one above that permits us to tune stardom (s) and randomness (p) independently, while having a constant degree sum in the graph ($sD + (n - s)d = nm$). To that end, locate the n individuals on the circle. Take an even $d \leq m - 2$, and sequentially pick s nodes at random to form S . First the stars are taken care of: for each $i \in S$ create an even $D = (nm - (n - s)d)/s$ links. With probability $1 - p$, each of the D links is connected to one of the $D/2$ nearest nodes on each side of i on the circle. With probability p , it is connected at random. That takes care of the stars and creates also some links for the non stars. Then run across all the non stars and proceed analogously, checking that the degree constraint for non stars is also satisfied. Because only integer numbers are handled the procedure will in general not create exactly nm links. However the results will be reasonably close to that target. The 3-panel graph below summarises a few statistics for an illustration with $n = 500$, $d = 6$ and $m = 10$, i.e. a degree sum of 5,000 which the algorithm roughly preserves. The sum is displayed in the insert in the upper panel of Fig. 1. The concentration of links (upper panel) falls monotonically with s , across all p values. The two lower panels of Fig. 1 show clustering and path length versus p for $s = 21, 107$ and 500 stars. Clustering displays no significant variation with s , while displaying its usual pattern with p . Distance also behaves monotonically with degree, while the effect of s (owing to the assumption of a constant degree sum) is monotonic and quite obvious. The algorithm thus behaves well and will permit us to explore independently the effects of degree asymmetry and local disorder.

2.2 The dynamics of ideas

Over time, as ideas are exchanged, knowledge evolves. an agent's knowledge stock never shrinks, as we assume that ideas display non-rivalry (i does not lose his idea by letting j have it), thus one agent's knowledge only increases or stays constant as times passes.

2.2.1 Knowledge as sets

Agents operate in a system (an industry) where there is a finite number of existing ideas indexed by $l = 1, \dots, \ell$. Each agent is endowed with a subset of these ideas, and we denote $H_i \subseteq \{1, \dots, \ell\}$ the set of ideas held by i . Agents use ideas for the sake of production, and in that respect agents are heterogenous.

Production is done by agent i according to the Leontieff production function

$$\phi_i = \begin{cases} 1 & \text{if } P_i \subseteq H_i, \\ 0 & \text{otherwise,} \end{cases} \quad (1)$$

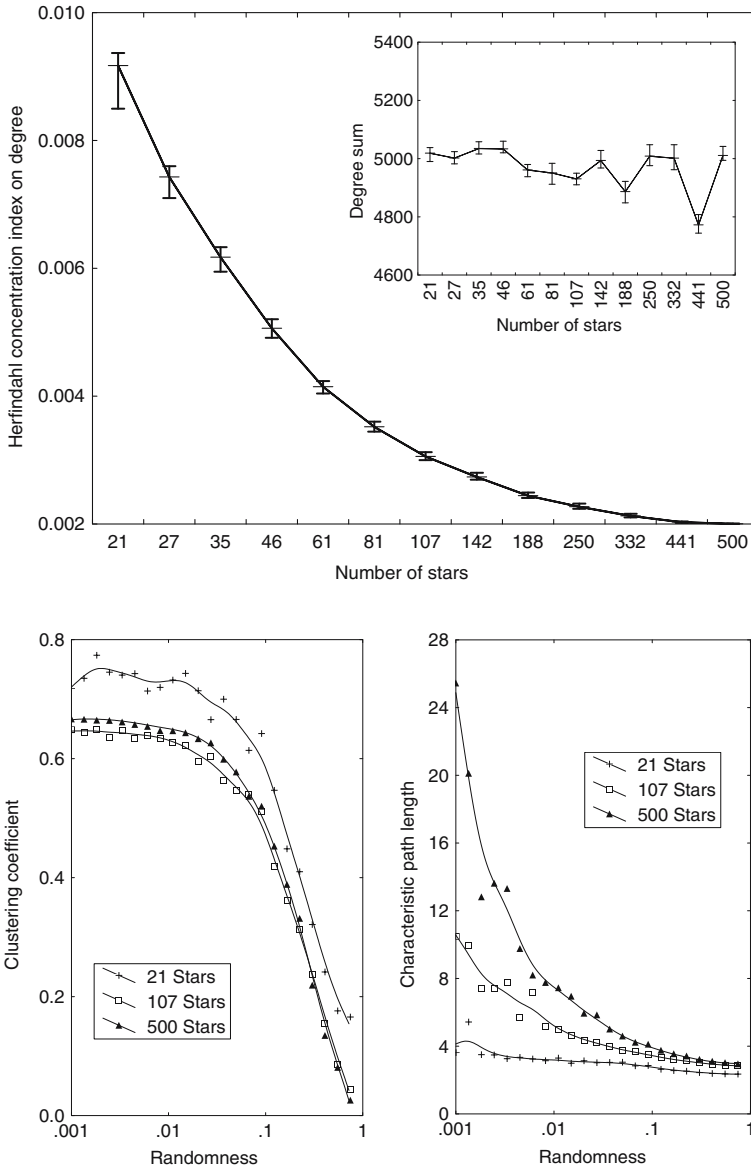


Fig. 1 Characteristics of the graph family: asymmetry in the degree distribution as measured by the Herfindahl concentration index (*upper panel*, pooled data across p , the whiskers are the minimum and maximum values), degree sum (*insert in upper panel*, pooled across p), clustering coefficient (*lower left panel*) and characteristic path length (*lower right panel*) for different numbers of stars

where $P_i \subseteq \{1, \dots, \ell\}$ is the set of ideas agent i needs to produce. Production is impossible if even one $l \in P_i \notin H_i$. As initial endowments are random, it will typically be the case that $P_i \not\subseteq H_i$, so there is room for exchange.

Consider now $j \in V_i$ (equivalently $i \in V_j$, as the graph is non-directed). Denote now $N_i = \{l \in P_i - H_i\}$ the set of ideas that i uses in production but does not have, that is, i 's needs. Then i is interested in j provided j has at least an idea that i uses but does not have (that is, if $N_i \cap H_j \neq \{\emptyset\}$).

2.2.2 Exchange

As for the “terms of trade”, a number of possibilities can be simply explored in the model presented above. Suppose the sequence of events is that each time period an agent is selected and engages in knowledge exchange with one of his neighbours j . In a *gift transaction*, upon request from i , j provides him with an element of $N_i \cap H_j$ and asks nothing in return. This captures a situation in which agents tell each other freely, creating knowledge spillovers. In a *barter transaction*, upon request from i , and provided i has an idea that j needs ($N_j \cap H_i \neq \{\emptyset\}$), an idea is exchanged for another one. This is a trading situation, in which all ideas have a common price, yielding a one-to-one exchange rate. In a mixed economy, some agents give while some trade. This is controlled by $0 \leq \pi \leq 1$ the share of knowledge givers. It is important to note that agents are not interested in all the pieces of knowledge but only in those relevant to their production activity. This implies that at some point trading can stop without every agent holding every idea.

3 Numerical experiment

A natural measure in the productive efficiency of the system is the average output

$$\phi = \sum_{i \in V} \phi_i / n,$$

the proportion of individuals who actually produce. As in this simple environment any agent either achieves production or does not, the ϕ s consist of a collection of 0s and 1s and the variance in output is equal to $\phi(1 - \phi)$, which is largest when $\phi = 1/2$ and smallest when ϕ is either 0 or 1.

The settings of the experiment are the following: a population of $n = 600$ agents and $m = 10$ links per agents (hence a total of 6,000 edges); each agent is endowed with a set of ideas that is randomly initialised by making each idea l available to i with probability $\Pr\{l \in H_i\} = q$. The production function is a 200 category one, with agents needing on average 20 ideas to produce. Which ideas an agent needs is determined randomly, independently of all other agents: $\Pr\{l \in P_i\} = \theta = 0.1$. All probabilities used to create initial knowledge endowments and production technologies are independent from each other. Each period in the simulation, one agents is selected, he activates a connection and they trade if possible. This process continues until all possible trades have been made. To examine the space of graphs, using the algorithm described in Sect. 2.1 above, we vary the rewiring probability p from 0.001 to 1. For each p value, ten different graphs are created and on each graph a single history is run, until all exchange possibilities are exhausted.

The parameters are p , the degree of disorder; s , the number of stars (with lower s values corresponding to more asymmetry in the link distribution, see Fig. 1) and π the share of agents who give rather than trade.

4 Results

In this section we present the results of our experiment. As will be seen, they connect well to some of the empirical results described in the introduction.⁶ We start by examining performance at the aggregate level before turning to firm performance. Results are presented under two initial conditions: one of scarcity, wherein initially any agent holds only 15% ($q = 0.15$) of all possible knowledge; and one of abundance, wherein an agent initially holds 85% ($q = 0.85$).

4.1 Effect of the density of traders

In the presentation of results, we suppress one parameter. The effect of the proportion π of agents who give rather than trade is monotonic throughout. As the number of givers increases, total knowledge levels, or equivalently the number of agents producing, increases for all levels of asymmetry in the degree distribution and for all values of randomness in the network. The cause is clear. The absence of a quid pro quo in exchange when givers are involved implies that when there are many of them more exchanges will take place, and so a larger proportion of firms will find the knowledge they need. In some of the other results which we discuss below, the number of traders changes the strength of the effects we observe, but in every case the direction of the effects remains unchanged. We thus set the proportion, π , of agents who give to 5%, which permits these agents to have a moderate influence on the diffusion dynamics, without swamping other effects.

In the figures below we present shaded contour plots wherein darker shades of grey indicate higher values of the variable being examined. The axes in each plot are p , the degree of randomness in network structure, and the Herfindahl concentration index, as a measure of asymmetry in the link distribution. In each figure there are two panels: the left panel shows data from the case of scarce knowledge in the initial condition; the right panel corresponds to the case of abundance.

4.2 Industry production

A measure of industry performance is the proportion ϕ of firms that have been able to acquire the knowledge they need to produce. We refer to these as producers, and to those that have not acquired the requisite knowledge as non-producers. The relationship between aggregate efficiency, the degree of randomness p and asymmetry,

⁶ We do not explicitly test the results against empirical data. One possible route to confront the model would be that described by Werker and Brenner (2004).




Fig. 2 Proportion of firms producing, with scarce knowledge (*left panel*) and abundant knowledge (*right panel*)

is shown in Fig. 2a, b. The first effect, comparing the two panels, is that not surprisingly, when knowledge is abundant more firms are able to produce. The second effect concerns the number of stars, or asymmetry of the link distribution. When knowledge is abundant, as asymmetry increases (or the number of stars falls) efficiency decreases monotonically in a very clear pattern. This is explained as follows. Increasing asymmetry means that links become more concentrated on fewer individuals. If there is a trading agent who has many links, his need for a quid pro quo can block many trades, and close paths between many pairs of agents. At the same time an agent with many links can quickly find the knowledge he needs. At this point, needing nothing, he has no reason to trade, and so withdraws, thereby closing many paths. This possibility can result in many agents not finding the knowledge they need. Clearly this possibility recedes when there are fewer agents dominating the linked distribution. When knowledge is scarce, the pattern is not monotonic: efficiency increases and then decreases as asymmetry increases, with the peak at about 400 stars. When asymmetry is at its minimum, however (when every agent is a “star”), efficiency is still much higher than when asymmetry is at its maximum (21 stars). The overall decrease in efficiency has the same explanation as in the previous case. The non-monotonicity of the relationship is explained, however, by a second effect, namely path length. The possibility that knowledge transmission is blocked as explained above, increases with the path length between sender and ultimate receiver. Thus shorter paths reduce this effect, all else equal. Increasing asymmetry thus has two effects working in opposite directions: decreasing path lengths, and concentrating links in a few agents. At an intermediate degree of asymmetry, efficiency is maximised.

It is possible to observe an effect of p in Fig. 2, though it is of smaller magnitude than the effect of s . Recall that increasing p increases the presence of structural holes, and by the same mechanisms decreases both clustering and path length. What we observe is that when knowledge is scarce, efficiency increases with p ; when knowledge is abundant, efficiency decreases with p . In the former case, at the aggregate level social capital is less valuable than are structural holes. Rapid access to distant parts of the network is highly valuable in acquiring knowledge, and at the aggregate level, short path lengths imply complete diffusion of knowledge. In the latter case,




Fig. 3 Difference in degree between producers and non-producers

with relatively high probability the needed knowledge is close in network space, so a clustered neighbourhood will imply many paths between agents that hold reciprocally desirable knowledge.

4.3 Individual performance

It is less straightforward to examine performance at the micro level as the performance of a firm must be compared to that of other firms in the same context. We are interested here in the effect of the structure of a firm's ego network on its performance.

At the end of each run in our simulation experiments, firms can be partitioned into two groups: those who have accumulated the knowledge they need to produce, and those who have not. The obvious question is whether these two groups are different from each other along interesting network dimensions. We answer this by looking at two structural parameters of each agent's ego network: degree and clustering.

In Fig. 3 we show the difference in degree between producers and non-producers, averaged over runs at each point in the parameter space. The figure shows that in terms of acquiring useful knowledge, it is valuable to have many connections—the difference in degree between producers and non-producers is always positive. When knowledge is scarce, this difference is much larger than when knowledge is abundant.⁷ In the former case, essentially only stars, having many connections, are able to produce; in the latter, many non-stars are also able to produce. This explains the magnitude of this difference. An agent with many connections rapidly acquires the knowledge he needs, at which point he withdraws from the system (completely if he is a trader, partially if he is a giver). In the worst case this disconnects the network, in the best case it makes path lengths longer. This can make it impossible for other agents to acquire the knowledge they need. This effect is severe when knowledge is scarce, much less so when knowledge is abundant since many agents will be able to produce from the initial period, and those who cannot are likely to need only a few pieces, which are relatively easy to find. The two structural parameters, asymmetry of the link distribution and randomness of the network both have visible effects when knowledge

⁷ In both cases all observed differences in means are statistically significant at the 5% level or higher.




Fig. 4 Difference in clustering between producers and non-producers

is common. The value of having many connections increases both as the randomness of the network increases and as asymmetry in the link distribution increases. The latter effect is driven largely by the fact that stars are more likely to produce than non-stars, and as asymmetry increases, stars have more neighbours. The effect of p is driven by clustering. In a barter economy, a clustered graph can alleviate the double coincidence of wants problem, because a transitive triple provides a short path over which two agents can make an indirect trade if they cannot trade directly. Thus agents with few connections can take advantage of indirect connections to get the knowledge they need. When there are few transitive triples, it is very difficult to overcome a failed double coincidence of wants, so there is a large advantage to having many potential partners, since when trade becomes impossible with one of my neighbours, I can simply turn to another.

Is having a clustered ego network valuable for an agent? Figure 4 shows the difference in ego-network clustering between those who produce and those who do not, averaged over runs at each point in the parameter space. This difference is always negative—a clustered ego-network is bad for information gathering. However, the difference is roughly an order of magnitude larger when knowledge is rare than when it is common.

This is consistent with the structural holes argument given above. Access to knowledge is vital in this economy, and in a clustered neighbourhood, links increase local density rather than connect to distant parts of the network. If an agent is part of a cluster he will, in general, have long path lengths to other agents. If there is a piece of knowledge in a distant part of the network that the agent needs, the longer the path to it, the more likely that on that path is a trader who, because he has all the knowledge he needs, has effectively stopped participating in knowledge transactions. Second, as the networks become less clustered, the negative value to an individual of being in a clustered neighbourhood also decreases, becoming statistically insignificant when there are no explicit stars. As the network itself is less clustered, the average distance from a particular agent to others in the economy falls, regardless of the extent of clustering in his ego network. This provides another route by which the effect just described is attenuated. The effects seen in these figures are largely driven by the probability that there is a failure in the path between an firm and the knowledge it needs, but when

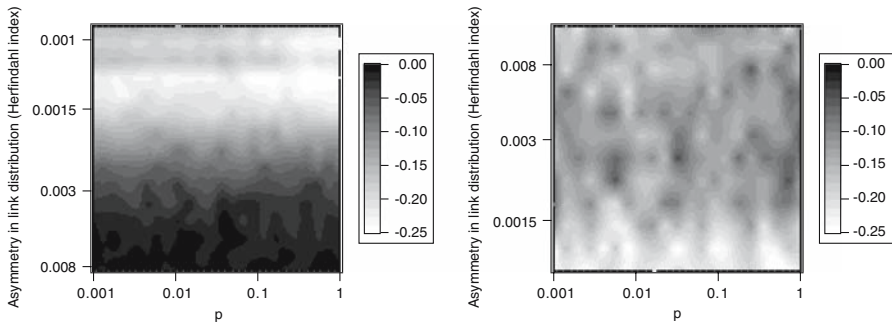


Fig. 5 Difference in performance between givers and traders

knowledge is abundant in general, these failure probabilities will be lower in general, and so the effects weaker.

4.4 Givers and traders

Figure 5 shows the difference in performance between givers and traders, measured as the difference between the proportion of traders who produce and the proportion of givers who produce. Traders always fare better, and the extent to which this is the case decreases with the asymmetry of the link distribution when knowledge is scarce, again explained by the dominance of stars among the producers. The explanation for the superior performance of traders generally lies in the nature of knowledge interactions. If a trader is involved in an interaction he always receives desirable knowledge. If a giver is involved in a transaction, he only receives information if he is the originator of the transaction. Thus the knowledge of traders grows faster than that of givers, so more of them will be able to produce in the long run.

4.5 Particular neighbours

In the structure of the model there are two types of agents that might have a significant impact on the performance of agents in their neighbourhoods: stars and givers. The effect of having a direct link to a giver is shown in Fig. 6. An agent has in his neighbourhood a proportion of givers. In these panels we display the data over the same range. We take the average difference in these proportions between producers and non-producers. In the left panel of Fig. 6 the data take values only in $[-0.01, +0.01]$, but are not statistically different from zero, whereas in the right panel all values are above 0.02, and are statistically significant. When knowledge is scarce, having a giving neighbour has no effect; when knowledge is abundant, being close to a giver significantly improves performance.

The effect of having a direct link to a star is shown in Fig. 7 using the same measure as above in Fig. 6. In both panels of Fig. 7 the data are significantly negative. Having a star as a friend is never a good thing as it is almost always a star who creates the hold-up problem by exiting from the trading process. When there are few stars, each

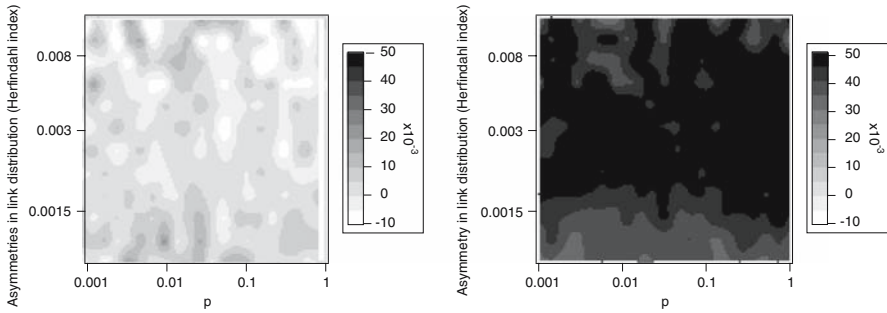


Fig. 6 Difference between the proportion of givers in the ego-network of producers and non-producers




Fig. 7 Difference between the proportion of stars in the ego-network of producers and non-producers

having high degree, the problem is strongest. The pattern in both panels essentially tracks the asymmetry of the degree distribution.

5 Conclusions

We have observed in these results that the presence of “givers” in the economy is generally a good thing. This is intuitively appealing in the first instance, as it seems natural that if agents are giving knowledge away, knowledge flows will be facilitated. Traders function differently. In a networked economy, goods and services, including knowledge flow from one agent to another through a path of intermediate agents. At each step in the path some transaction is made. If it involves a trader, then the transaction must benefit both parties. Thus if one agent on a path is no longer interested in exchange, in our case, because a trader has all the information he needs and so no longer makes trades, this eliminates any paths which flow through that agent. Thus the model illustrates that a network with a skewed link distribution, having a few stars, can be either good or bad. If the stars are givers, then their knowledge can flow rapidly out to their many partners, and from there to the rest of the economy. This will continue through the life of the economy. If the stars are traders, because they have many partners, they will rapidly acquire all the knowledge they need, and so stop trading. This blocks many paths between agents, and in the most extreme case, can disconnect the network.

In the introduction we discussed briefly the debate between the structural hole position and the social capital view. On the former, a non-clustered ego-network is good for a firm; on the latter, a densely interconnected neighbourhood provides a good knowledge environment. The model developed here permits some formalization of the resolution suggested by Rowley et al. (2000). In Sect. 4.2 we showed that when knowledge was scarce, aggregate production increased with p , whereas when it was abundant, aggregate production decreased with p . Average clustering decreases monotonically with p . Thus when knowledge is scarce a network with structural holes performs well; when knowledge is abundant a network with high social capital performs well. A situation of scarce knowledge represents a young industry, in which technologies are new, and firms are exploring the technological space to find and create the best possible variant of their products and processes. For any firm, the necessary knowledge is difficult to find, and may reside in distant parts of the economy. Here, redundant links will be less valuable than links that create short paths to other agents. A situation of abundant knowledge may represent a more mature industry. A dominant design has emerged; most firms know what it is, and have most of the knowledge needed to execute it. What is happening is that firms are exploiting their versions of the dominant design, and need details rather than new principles. Here, distant parts of the economy will have similar knowledge roughly speaking to local parts of the economy, and what is necessary is to extract the final details. Here, non-redundant links lose their advantage and we see the force of the social capital argument.

One thing that is striking in this regard is that the economy and individual firms have different responses to clustering when knowledge is abundant. At the aggregate level, to recall, when knowledge is abundant output rises as clustering increases, particularly when there are many traders. But at the individual level, being part of a cluster is almost always bad (those who produce have on average less clustered ego-networks than those who do not).⁸ The explanation may lie in the fact that when a network is highly clustered on average, a firm with low clustering can connect different parts of the network. If it is the case that firms within a cluster develop similarities, then a firm between two clusters can have access to two different types of knowledge. This provides it with an advantage in terms of finding the knowledge it requires. This appears to be so in our model even though there is no strategic knowledge acquisition or control. It has been suggested by Burt (1992), and also by Baum et al. (2003), that firms that fill structural holes in the network can control information flows. In our model there is no notion of controlling flows, but we do see the first necessary condition for such control, namely access to different knowledge pools.

Efficient knowledge diffusion is the hallmark of a healthy modern economy. What the model developed here shows is that the structures necessary to promote knowledge diffusion depend to a very great extent on the details of the industry. Industries or episodes dominated by collective invention, in which knowledge is (locally) freely given, are very different from industries in which knowledge trading is the norm. But further, how they differ depends on whether knowledge is scarce, as in a newly emerging industry, or abundant, in a mature industry. When knowledge is scarce, random

⁸ This is obviously a very rough generalization. Exceptions exist, particularly when there are many givers in the economy, or when there are no stars.

networks always perform well. When knowledge is abundant, and knowledge trading dominates, clustered networks perform best from the social point of view.

Is there a role for policy here? When knowledge is scarce, socially, random, unclustered networks perform well. This performance carries over to the individual firm level: firms with non-clustered ego networks also perform well. So we observe a coincidence between network structures that are socially and individually desirable. The coincidence disappears, though, when knowledge is abundant. Socially clustered networks are efficient, but any firm would prefer a non-clustered ego-network—one abundant in structural holes. Here, this divergence between social and private efficiency provides scope for intervention in network formation, at least in principle. What this model has shown though, is that policy-making in this area is a very delicate business—one size definitely does not fit all. The details about the state of knowledge and the social conventions regarding exchange matter a lot, and so in this area policy must be built on a very strong empirical foundation.

Acknowledgments We thank an anonymous referee for useful comments, and participants in the NEMO project (European Commission FP6, NEST-ADV 028875).

References

- Ahuja G (2000) Collaboration networks, structural holes and innovation: a longitudinal study. *Adm Sci Q* 45:425–455
- Allen R (1983) Collective invention. *J Econ Behav Org* 4:1–24
- Barabási A, Albert R (1999) Emergence of scaling in random networks. *Science* 286:509–512
- Baum JAC, Shipilov AV, Rowley TJ (2003) Where do small worlds come from? *Ind Corp Change* 12:697–725
- Burt RS (1992) Structural holes: the social structure of competition. Harvard University Press, Cambridge
- Cowan R, Jonard N (2003) The dynamics of collective invention. *J Econ Behav Org* 52:513–532
- Cowan R, Jonard N (2004) Network structure and the diffusion of knowledge. *J Econ Dyn Control* 28:1557–1575
- Dyer JH, Nobeoka K (2000) Creating and managing a high performance knowledge sharing network: the Toyota case. *Strateg Manage J* 21:345–367
- Granstrand O, Sjolander S (1990) Managing innovation in multi-technology corporations. *Res Policy* 19:35–60
- Gulati R, Gargiulo M (1999) Where do inter-organizational networks come from? *Am J Sociol* 104:1439–1493
- Lamoureaux NR (1999) Accounting for capitalism in early american history: farmers, merchants, manufacturers, and their economic worlds. Technical report, UCLA, 1999. Available at <http://www.sscnet.ucla.edu/history/activities/usccpapers/lamoreaux.html>
- McGaw JA (1987) Most wonderful machine: mechanization and social change in Berkshire paper making. 1801–1885. Princeton University Press, Princeton (1987)
- Podolny JM (1993) A status-based model of market competition. *Am J Sociol* 98:829–872
- Powell WW (1990) Neither markets nor hierarchies: network forms of organization. In: Shaw BM, Cummings LL (eds) *Research in organizational behavior*, vol 12, pp 395–336
- Powell WW, Koput KW, Smith-Doerr L (1996) Inter-organizational collaboration and the locus of innovation: networks of learning in biotechnology. *Adm Sci Q* 41:116–145
- Powell WW, White DR, Koput KW, Owen-Smith J (2005) Network dynamics and field evolution: the growth of inter-organizational collaboration in the life sciences. *Am J Sociol* 110:1132–1205
- Rowley T, Behrens D, Krachhardt D (2000) Redundant governance structures: an analysis of structural and relational embeddedness in the steel and semiconductor industries. *Strateg Manage J* 21:369–386
- von Hippel E (1987) Cooperation between rivals: informal know-how trading. *Res Policy* 16:291–302

- Walker G, Kogut B, Shan W (1997) Social capital, structural holes and the formation of an industry network. *Org Sci* 8:108–125
- Watts D, Strogatz S (1998) Collective dynamics of small-world networks. *Nature* 393:440–442
- Werker C, Brenner T (2004) Empirical calibration of simulation models. Max Planck Institute, Jena; Papers on Economics and Evolution number 2004-10