

On the Creation of Networks and Knowledge*

Robin Cowan[†], Nicolas Jonard[‡] and Jean-Benoît Zimmermann[§]

May 19, 2002

Abstract

This paper examines the evolution of networks when innovation takes place as a result of agents bringing together their knowledge endowments. Agents freely form pairs creating a globally stable matching. Paired agents combine their existing knowledge to create new knowledge. We study the properties of the dynamic network formed by these interactions, and the resultant knowledge dynamics. Each agent carries an amount of knowledge of a certain type, and the innovative output of a pair is a function of the partners' endowments and types. We find evidence that the pattern of substitution between quantity and type of knowledge in the innovation function is vital in determining the growth of knowledge, the emergence of expertise and the stability of a number of network structures. Network structure itself exhibits a phase change when the relative importance of diversity compared to quantity increases beyond a threshold value.

DRAFT: Not to be cited.

Keywords: Knowledge, Networks, Innovation, Stable Matching, Small Worlds.

*Part of this research was done while Jonard was a visiting scholar at the International Institute on Infonomics, and we thank that institute for its support. We also thank the French Embassy in Canada for having supported this research through a grant from the France-Canada research fund.

[†]MERIT, Maastricht University.

[‡]CNRS, BETA, Université Louis Pasteur.

[§]CNRS, GREQAM, Université Aix-Marseille II.

1 Introduction

In the new economy, knowledge takes a great importance as an economic good. Not only is knowledge tied to innovation and technological change, it also becomes, perhaps as information, an economic good in its own right. In both these roles, the production and distribution of knowledge and information has become a central part of the analysis and discussion of the economy. The central concern of this paper is precisely this production and distribution.

Economists have typically shied away from explicit discussions of knowledge creation, as it often seems to demand non-trivial inputs of creativity, and this is something about which economists have nothing to say. Thus knowledge creation is often discussed simply in terms of cost-reducing, quality-improving, or productivity-enhancing technical innovations. As such, it is often modelled as a random event — a firm makes a draw from a probability distribution to determine next period’s costs (or productivity or quality). If firms are thought to be able to affect their innovative ability through expenditure, this can be modelled as changing the mean or variance of the probability distribution.

In this paper we present a more explicit and less “black-boxed” model of knowledge creation. We take seriously the idea of knowledge creation as knowledge re-combination. There have been many empirical (largely case) studies of innovation within firms showing that the lion’s share of economic value from innovation arises from a process that is largely the re-combination of existing knowledge.¹ Creating knowledge most often involves putting together knowledge that already exists: knowledge is produced out of knowledge. While there is now significant empirical evidence supporting this view, there has been virtually no theoretical or abstract work on this issue. The present paper begins to address that lacuna.

Early models of diffusion, of which two early examples are Griliches (1957) or Mansfield (1961), use the epidemic paradigm. In these models populations are seen as “mixing” in the sense that any agent is equally likely to come in contact with any other. While this has many advantages as a modelling strategy (not least being that it fits macro-structures relatively well) it is somewhat unrealistic. More recently, economists have begun to pay

¹See for example Kodama (1986,1992); Gibbons et al. (1994); Sutton and Hargadon (1996) and Hargadon and Sutton (1997).

attention to the networks over which information diffuses. In other words, economists recognize that much knowledge transmission takes place in bilateral (or if multi-lateral, then only multi-lateral among a small number of agents) interactions. When this is the case, the network over which these communications take place can be an important locus of investigation.² In many network models, though, the network structure is given, and even if agents have many interactions, any agent interacts with a fixed, and typically small, subset of the population. But this seems unreasonable in many instances — an agent will typically change the set of agents with whom he interacts as he gains experience and learns about where the highest rents lie.

There are also models of evolving networks, but this area of inquiry has remained relatively unexplored until the last few years. One approach in the literature is the formation of networks (or coalitions) based on pure cost considerations. Costs are described by a topological structure based on the characteristics of the set of individuals: the closer two individuals are, the less expensive it is for them to establish a relationship (Debreu, 1969; Haller, 1994). Another approach insists on the benefits from coalescing, for which cooperative game theory is a natural framework. Later developments allow us to distinguish three broad categories: purely cooperative approaches, purely non-cooperative approaches, and mixed approaches, in particular the notion of pairwise stability. Qin (1996) explores the cooperative situation, while the non-cooperative approach is used in Bala and Goyal (1998, 2000). There it is shown that when players can create one-sided links by simply making the investment (i.e., an individual cannot prevent another one from connecting him), a variety of Nash networks can emerge. Jackson and Wolinsky (1996) consider the equilibrium concept of pairwise stability (individuals can both sever and form connections), and show that only the star and the complete network are equilibrium outcomes.³ Finally, in a very different framework in which interactions take place among adaptive individuals, a fully dynamic account of social network evolution can be found in Plouraboué et al. (1998), and a specific application to buyer-seller networks is in Weisbuch et al. (1997).

The present paper contributes to the literature on network formation and evolution in the context of knowledge creation and diffusion. We construct

²See for example David and Foray (1992); Valente (1995); Steyer and Zimmermann (1998); Cowan and Jonard (1999 and 2000)

³For a good summary of some of the important existing work on game theoretical analysis of network formation, see Jackson and Dutta (2001).

a model in which firms form alliances in order to create new knowledge and in so doing transmit knowledge between them. We use this structure to inquire about the relationship between aspects of knowledge creation and the types of network structures that emerge. The dynamic model we develop lies between the pure search for Nash equilibria of evolutionary game theory and the melioration models of social network dynamics.

2 Creating knowledge through interactions

Strong evidence exists regarding the importance of recombination as a form of innovation and knowledge creation. Economists have pointed out another recent change in the knowledge environment of firms. This is referred to as “the expanding knowledge base” of firms and industries.⁴ The general idea is that in most industries today the technologies both being used and being produced involve technological expertise that covers a much broader range of “disciplines” than has hitherto been the case (see Smith, 2000). What this implies is that types of knowledge necessary to innovate and compete successfully can lie outside a firm’s main area of expertise. A now common way of coping with this problem is to form an alliance with a firm that has the missing expertise. These cooperative agreements for R&D have grown dramatically in number since 1976.⁵ Antonelli (1999) argues that inter-firm cooperation can be extremely effective in increasing the circulation of tacit knowledge, and in creating possibilities for a firm to acquire knowledge outside its boundaries. Further though, cooperation can enhance both “the opportunity for accelerated recombination of the bits of codified knowledge generated by each cooperating firm; [and] the scope for capitalizing on potential complementarities between the variety of firms and between the different R&D activities performed by each firm.” (Antonelli, 1999, p. 10).

To model strategic technological alliances in their entirety is far beyond the scope of this paper, and in particular we make no attempt to include any issues of intellectual property rights in our discussion. We are interested, rather, in the nature of the networks of innovators that emerge from a process in which firms engage in bilateral cooperation to produce knowledge. Technological alliances can be very rich and varied, not only in terms of outputs but also in inputs and objectives of the participants, but for our

⁴See Zimmermann (1995); Cowan and van de Paal (2000).

⁵See Hagedoorn (2001) for a review and discussion of this trend.

purposes we can focus on a single effect, namely the production of shared knowledge; and consider only one type of input, namely the pre-existing knowledge held by the participants.

In this paper we develop a model in which each period every firm in the population seeks a partner for knowledge creation. We restrict attention to bilateral partnerships. Firms come together and pool their existing knowledge to create new knowledge which they then both absorb. Partnerships last for one period exactly, at the end of which each firm seeks a new partnership, possibly with the previous partner. That firms search for (possibly) new partners each period implies that over time a network of knowledge flows can form, as firms create links with several different firms. The central question in this paper has to do with the structural properties of these networks and whether the properties change under different knowledge production technologies.

When innovation is the outcome of agents' pairwise interactions, rational agents will seek partners to maximize the expected outcome of the partnership. Put another way, at each period of time, an agent will rank each of his potential partners on the basis of the expected benefits from that interaction. The issue then is to divide the population into $n/2$ pairs such that each agent's partner is the best he can have, given everyone else's preferences.

Problems of this sort have been addressed in the literature as matching problems, as defined by Gale and Shapley (1962). Most of the literature on matching has focused on desirable (stable) associations between agents of two distinct populations (men and women, firms and workers, ...), where each individual tries to find a match with one or several individuals of the other population. In the case at hand, however, individuals belong to a single population of even cardinality n and search for a match within this population. This problem, known as the roommate, or one-sided matching problem (as opposed to the standard two-sided marriage problem), was first mentioned in the original paper by Gale and Shapley (1962) (see also Knuth, 1976).

A one-sided stable matching is a partition of the population into $n/2$ pairs such that no two individuals can be found who would prefer each other to their actual partners in the matching. Such a matching does not necessarily exist (contrasting with the standard marriage problem, where a stable matching always exists).⁶

⁶According to Bartholdy and Trick (1986), a sufficient condition for the existence and

3 The model

The model is two-part: innovation takes place as a result of the joint effort of pairs of agents, while in parallel a process of network formation and evolution is at work. Regarding innovation, individuals' knowledge is combined to produce new knowledge which is added to their existing knowledge stocks. A production function encompassing both agents' amounts and types of knowledge determines how much new knowledge is generated. How pairs form is the second important issue in the model. The ranking agents make of each other is a function of their joint knowledge profile (as explained below). Hence rankings change over time because innovation changes individual knowledge profiles. A stable matching always exists (see Proposition 1) so each period this stable match determines which pairs of agents combine their knowledge. A number of issues are of interest in this model, falling within two broad categories: the process of network formation and evolution; and the properties of knowledge growth. Before examining them, we discuss both matching and knowledge production in detail.

3.1 One-sided matching

Consider a finite set $S = \{1, \dots, n\}$ of individuals engaged in repeated interactions. Each individual $i \in S$ is characterized by a real-valued knowledge endowment, v_i , in the form of a quantity of knowledge $\rho_i > 0$ and a composite index representing a type of knowledge $\theta \in (0, \pi/2)$: $v_i = (\rho_i, \theta_i)$. This allows a very simple polar representation of individuals, as well as the existence of a continuum of knowledge types. Individual knowledge level and type are randomly drawn from a uniform distribution over $(0, 1) \times (0, \pi/2)$ at the beginning of the process. At each period, agents seek partners with whom to innovate and improve, as much as possible, their current knowledge profiles. Once knowledge has been created and absorbed, agents start again a round of matching, based on their new knowledge profiles.

Innovation then, is a process that combines the knowledge of two agents and produces new knowledge. In this process, all else equal, more knowledge

uniqueness of a stable matching in the roommates problem is that individuals can be ranked such that each agent prefers a roommate who is closer to him to anyone placed further away in the ranking. Irving (1985) has proposed an efficient algorithm for finding a stable matching (provided one exists). By taking advantage of certain features of our structure we are able to improve significantly on Irving's algorithm.

will be better. But diversity of knowledge inputs may also matter. Two agents whose expertise lies in exactly the same field are likely not to find many synergies. Thus in general diversity is beneficial to the innovative process, and enters positively into the production function. To fix ideas we model this as a constant returns Cobb-Douglas function, the arguments being the lengths of the knowledge vectors of the two agents (how knowledgeable they are), and the extent to which their expertise differs, that is, the difference between their knowledge “angles”. The extent to which diversity matters is a parameter we vary as the exponent on the “difference in angles” argument in the Cobb-Douglas function.

A one-sided (or roommate) matching problem is defined as follows. Each individual $i \in S$ has a strict preference ordering \succ_i over all the individuals in $S - \{i\}$ (a list of his most preferred first). All preferences are total and transitive, and \succ denotes the profile of the preference orderings of the individuals. The pair (S, \succ) is called a roommates matching problem. A matching is a partition of S into $n/2$ disjoint pairs of roommates, i.e. a bijection $\mu : S \rightarrow S$ such that $\mu(\mu(i)) = i$ for all $i \in S$, and $\mu(i) \in S - \{i\}$, for all $i \in S$. (Sometimes $\mu(i) = j$ (or equivalently $\mu(j) = i$) is written $(i, j) \in \mu$.) A matching μ is stable in (S, \succ) if there is no $(i, j) \notin \mu$ such that

$$j \succ_i \mu(i) \text{ and } i \succ_j \mu(j).$$

When this fails, that is, when $j \succ_i \mu(i)$ and $i \succ_j \mu(j)$, both i and j would leave their partners to match with each other.

In the particular problem examined here, generating the preference profile \succ is done via a symmetric function $r : S^2 \rightarrow \mathbb{R}_+$ which associates to any pair of (distinct) individuals a value that represents the innovative output of this pair.⁷ The profile of preference orderings \succ is then defined by

$$j \succ_i k \Leftrightarrow r(i, j) > r(i, k), \forall i, j, k \in S, i \neq j \neq k. \quad (1)$$

Disregarding cases of indifference (when there are i, j and k such that $r(i, j) = r(i, k)$), it will be shown that a unique stable matching always exists.

⁷As we have modelled it, r is a constant returns Cobb-Douglas function in amounts and diversity of knowledge.

3.2 Knowledge and innovation

There are many ways to characterize knowledge, none of them without its pitfalls. For our purposes two aspects of knowledge are important—quantity and type. Recent views of innovation as recombination make it imperative that any formal characterization of knowledge permit that equally knowledgeable agents may know different things. That is, the knowledge of an agent comprises many different types of knowledge. One approach then would be to formalize knowledge as a long vector, each element representing a different type of knowledge. Agents differ then along many dimensions. This proves cumbersome in implementation however, and is more detailed than necessary for the issues we are exploring here. A simplification of this approach is to represent an agent’s knowledge as a pair: one element signifying a quantity of knowledge; the other signifying a type. This representation has the drawback that it takes what is clearly a complex thing, an agent’s knowledge stock, and represents it as a pair. But it has the great benefit of simplicity; it is easily generalized; and does capture some important elements of the nature of knowledge.

A representation of the innovation process must satisfy several minimal requirements. Consider two individuals i and j with knowledge stocks (ρ_i, θ_i) and (ρ_j, θ_j) respectively. Suppose they are paired in a stable matching. As innovation is jointly conducted, after innovation has taken place, the following should be true:

- the knowledge amounts held by i and j have increased;
- the knowledge types of i and j have changed;
- the distance between the knowledge types of i and j has fallen.

In a geometric sense, both vectors are longer after the innovative episode, and the angle between them has decreased, but this still does not give us a precise idea of the direction of change. A fourth assumption is necessary for the description of innovation to be complete:

- most often, innovation produces knowledge of a type which is “between” the types of the two contributors; however, innovation sometimes produces very different knowledge.

in an individual's knowledge stock.

The standard deviation of $\tilde{\theta}$ is a parameter we vary. If a technological system is deeply entrenched in a trajectory, we would expect the standard deviation to be small. On the other hand, in a case in which there is wide scope for technological discovery in many areas, the standard deviation of this distribution would be large. The smaller the standard deviation the more likely the new knowledge will lie "between" the knowledge of the partners. Varying this parameter permits us to explore the dynamics of these different situations.

Formally, the innovation function is $r : S^2 \rightarrow \mathbb{R}_+$, with

$$r(i, j) = |\theta_i - \theta_j|^\beta (\rho_i \rho_j)^{1-\beta}, \beta \in (0, 1).$$

Parameter β measures the importance of diversity in knowledge types in the innovative process. Individual i 's new knowledge endowment is simply $\rho_i(t+1) = \rho_i(t) + r(i, j)$, and given the innovation is in type $\tilde{\theta}$ -knowledge, i 's new knowledge type is

$$\theta_i(t+1) = \frac{\rho_i(t)}{\rho_i(t+1)}\theta_i(t) + \frac{r(i, j)}{\rho_i(t+1)}\tilde{\theta},$$

the weighted mean of his previous knowledge type and the new knowledge type, where the weights are the relative amounts of knowledge involved.

3.3 The knowledge dynamics

Each period pairs are formed, innovation takes place and endowments change. At the end of the period the partnership is dissolved. This process is repeated, and we investigate the long run behaviour of the system. In principle knowledge vectors grow in length indefinitely, but types of knowledge are bounded between 0 and $\pi/2$. Eventually, in the creation of knowledge, only amounts matter as that argument will dominate in the production function. This is clearly unreasonable. To solve this problem we re-normalize every period so that the population average knowledge vector has length one. What this implies is that data on aggregate knowledge levels have no meaning. However, we can still reasonably measure growth rates, and variance in knowledge levels across agents. In examining knowledge, two issues are of concern: its growth and its distribution. and allocation of knowledge in this economy. The economy's knowledge growth rate is simply the growth rate of the average knowledge level per period. It will be evaluated along the simulation

horizon. Since the mean level of knowledge is by definition 1, the variance of knowledge levels $\sigma^2 = \frac{1}{n} \sum_i \rho_i^2 - 1$ provides a good measure of the equality of the distribution.

3.4 Network dynamics

A relational network exists in this economy, but in a dynamic rather than a static sense. To study its properties, the list of connections active after the last ℓ periods of the simulation horizon are recorded (in practice, we will drop the first 100 realizations and ℓ will be equal to 2,000). This we do because the one-period graph only consists of isolated pairs, whereas over the last ℓ periods a weighted graph can be constructed by counting the number of activations for any connection. Let $G(S, V_t)$ be the graph associated with the stable matching achieved at time $t = 0, 1, 2, \dots$, that is $V_t(i, j) = 1$ if $(i, j) \in \Gamma_t$ and $V_t(i, j) = 0$ otherwise. The weighted graph recording past interactions is denoted $G(S, W)$, where $W(i, j) = W(j, i)$ is the frequency of activation of the connection between i and j over the final ℓ periods. For this graph several quantities are of interest. We study the frequency distribution of collaborations and, following Watts and Strogatz (1998), two structural parameters: the average path length and the average cliquishness.⁸

To examine network structure directly, we use path length and cliquishness. Weighted-graph equivalents of path length and cliquishness are constructed first.⁹ A measure of local connectivity bearing a strong resemblance to cliquishness can be computed as

$$C_W = \sum_i \sum_{j \neq i} \sum_{l \neq j} W(i, j) W(i, l) W(j, l). \quad (2)$$

The product involved in this triple sum is only non-zero when i , j and l have all been in contact with each other. So C_W is a measure of the extent to which an individual's partners are partners of each other. Unfortunately C_W does not take its values in a well defined range, so its interpretation is slightly more difficult than the standard cliquishness index. Similarly, a proximity index can be defined as the largest value along any path of the

⁸If 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 the friends of one agent are also friends of each other.

⁹In what follows the time index is dropped for the sake of simplicity.

product of the number of meetings. The general proximity index of i and j is equal to

$$L_W(i, j) = \max_{(i_1, \dots, i_k)} W(i, i_1) W(i_1, i_2) \dots W(i_k, j), \quad (3)$$

where the maximum is taken over all the possible paths, and we are interested in $L_W = \sum_i \sum_{j \neq i} L_W(i, j)$. Proximity is roughly the inverse of path length, but an interpretation problem analogous to the one mentioned for C_W arises with this statistic.

To address these interpretation problems we consider the original measures produced by Watts and Strogatz (1998) in a slightly modified context. Instead of working directly on W , redefine $W(i, j)$ to be equal to 1 if $W(i, j) \geq W^*$ and 0 otherwise, where W^* is defined as follows. As a null hypothesis assume that the distribution of partnerships is defined as a uniform multinomial with $n - 1$ categories. This defines a mean and standard deviation for each element $W(i, j)$. Then choose W^* to be equal to the mean minus 2 standard deviations.¹⁰ We can now use the standard definitions of cliquishness and path length.

Formally, defining $d(i, j)$ as the distance or length of the shortest path between i and j , the average path length is

$$L = \frac{1}{n} \sum_i \sum_{j \neq i} \frac{d(i, j)}{n - 1} \quad (4)$$

and simply measures how distant vertices are on average, which is a global property of the graph. Average cliquishness C is a measure of local connectivity capturing the share of active links between any given vertex's neighbours. It is written

$$C = \frac{1}{n} \sum_i \sum_{j, l \in W(i)} \frac{W(j, l)}{n_i (n_i - 1) / 2}, \quad (5)$$

where $n_i = \#\{j \mid d(i, j) = 1\}$ is the size of i 's neighbourhood. These four statistics together give a reasonably complete description of the structural properties of the underlying network. We add one simple measure, namely the average degree of the graph $D = \frac{1}{n} \sum_i n_i$, as a measure of the density of the interaction structure.

¹⁰The value of this algorithm is that it provides a systematic way of deciding whether $W(i, j)$ is close enough to zero that it can be considered unimportant from the point of view of network structure. Using the mean minus a single standard deviation as the cut-off produces qualitatively identical results.

4 Existence: Stable matching

Before turning to the emergence of network structure and the associated knowledge dynamics, we discuss the matching mechanism present in this model. Because any pair of agents assigns the same cardinal value to their match, a unique stable match always exists. We can prove this proposition by construction.

Proposition 1 *For any innovation function $r : S^2 \rightarrow \mathbb{R}_+$ generating a strict preference ordering \succ , the matching problem (S, \succ) has a unique stable matching μ .*

Proof The algorithm to construct the stable matching is as follows. Let $S_0 = S$ and $\mu_0 = \{\emptyset\}$. Consider the roommate matching problem (S_0, \succ) , where the profile of preference orderings \succ is defined by the innovation function r , as stated in equation (1). As preferences are strict, there is a single pair (a_1, b_1) such that $r(a_1, b_1) = \max_{(i,j) \in S_0^2, i \neq j} r(i, j)$. No matching which does not involve this pair could be stable, as a_1 and b_1 both prefer each other to any other person they might be matched with. Hence the pair (a_1, b_1) is necessarily part of a stable matching. Let $\mu_1 = \mu_0 + (a_1, b_1)$ and $S_1 = S_0 - \{a_1, b_1\}$. Consider then the new matching problem (S_1, \succ) . It only involves $n - 2$ individuals with exactly the same preferences (modulo those concerning a_1 and b_1) as before this operation. Again there is a single pair (a_2, b_2) such that $r(a_2, b_2) = \max_{(i,j) \in S_1^2, i \neq j} r(i, j)$. Let then $\mu_2 = \mu_1 + (a_2, b_2)$ and $S_2 = S_1 - \{a_2, b_2\}$. Repeat until everyone belongs to a pair: the set $\mu \equiv \mu_{n/2}$ is a stable matching as no pair can blockade it, and it is unique as the sequence $(a_1, b_1), (a_2, b_2), \dots, (a_{n/2}, b_{n/2})$ is uniquely defined. \square

In case of a tie (that is to say when individual i can achieve the same innovative output with two or more different partners), existence is still guaranteed but uniqueness is lost. The elimination algorithm to find a stable matching is unchanged, except that when two (or more than two) pairs achieve the same output, one of them is picked up randomly and the procedure just iterates. It is worth mentioning that the algorithm for getting a stable matching in this particular matching problem is much simpler and faster than the multi-pass algorithm proposed by Irving (1985). However computational constraints remain heavy: the algorithm for finding a stable matching is still polynomial $O(n^2)$ and it is repeated every time period; and

a non-polynomial algorithm is necessary for computing both the path length L and the proximity index L_W .

5 Numerical analysis

There are in the model two parameters that could affect the evolution of the network structure: s , the standard deviation of the distribution from which the angle of innovation is drawn; and β , the value of diversity in knowledge production. In the experiments that follow we simply examine two representative values for s , running each experiment twice, once for high and once for low s . The effects of β we examine in more detail. Random values of β are drawn from a uniform distribution over $(0, 1)$. For each β value $n = 100$ individuals repeatedly interact over a horizon of length 3,000 periods. We build a data sample of 100 different β -values and two s -values and the associated statistics. We extract the underlying structure from these data using non-parametric estimation techniques.

5.1 Network dynamics

The emergence of network structure can be described using three indicators: local order or cliquishness; path lengths; and density or degree of the graph. We examine the effects of two parameters, β , a measure of the importance of diversity in innovation, and s , which measures the extent to which new knowledge lies between the expertise of the innovators. In the discussion that follows, we restrict ourselves to the 0/1 graph. Analyzing the weighted graph, using the measures discussed above creates a problem of interpretation, since the measures are not bounded. Graphically, however, the patterns are identical to those produced by analysing the 0/1 graph.

5.1.1 Density

Figure 2 plots the relationship between density, as measured by the average degree, and β for two values of s , $s = 0.1$ (focussed trajectory) and $s = 3$ (diversified trajectory).

What we observe is that there is a phase change in the structure of the graph for $\beta^* \approx 0.4$. When $\beta \leq \beta^*$ the degree is very low. This implies that agents in general have had very few partners over history, and have tended to




Figure 2: Average number of connections per individual

interact repeatedly with the same partners. For $\beta > \beta^*$, the network is very different. Agents interact with many different partners over time, no longer sticking to the same ones. It is worth mentioning that for small values of β (for both values of s) and for β close to 0.25 in the case $s = 0.1$, the degree can be significantly larger than 1 although not systematically. This suggests a small region (roughly speaking for $0.2 < \beta < 0.4$) where two types of regimes can exist: a world of pairs (degree equal to one) and a more complex world of degree above 1 and below 15, depending on the path taken by the process. After β^* by contrast there is only one regime with degree around 50. Also note that when β is further increased after 0.75, the degree tends to decline when innovation is very localized (small s), while it remains constant when innovation is less focussed (large s).

In Figure 3 we depict the number of individuals involved only in pairs that are stable over time, i.e., of individuals having a unique life-long innovation partner.




Figure 3: Number of individuals involved in pairs

The phase change that exists for the degree is also present when the number of paired individuals is considered, with a sharp decline from a large number (between 96 and 100) below β^* to 0 for $\beta > \beta^*$. There is also some noise in the region $[0.2, 0.4]$, with a number of paired individuals that doesn't follow a clear pattern, a phenomenon similar to what we observed for the degree on Figure 2. Above β^* persistent pairs do not form. Interestingly, for very small β and a diversified innovative trajectory (large s), no pairs form and, as we will see when the path length is considered, we even get a connected graph.

5.1.2 Path Length

We see in Figure 4 the same phase change, at the same value of β^* in average path length as in the density graph.

For $\beta < \beta^*$, the graph is disconnected. That is, there are isolated groups of agents who interact within groups but not across groups. In principle, this

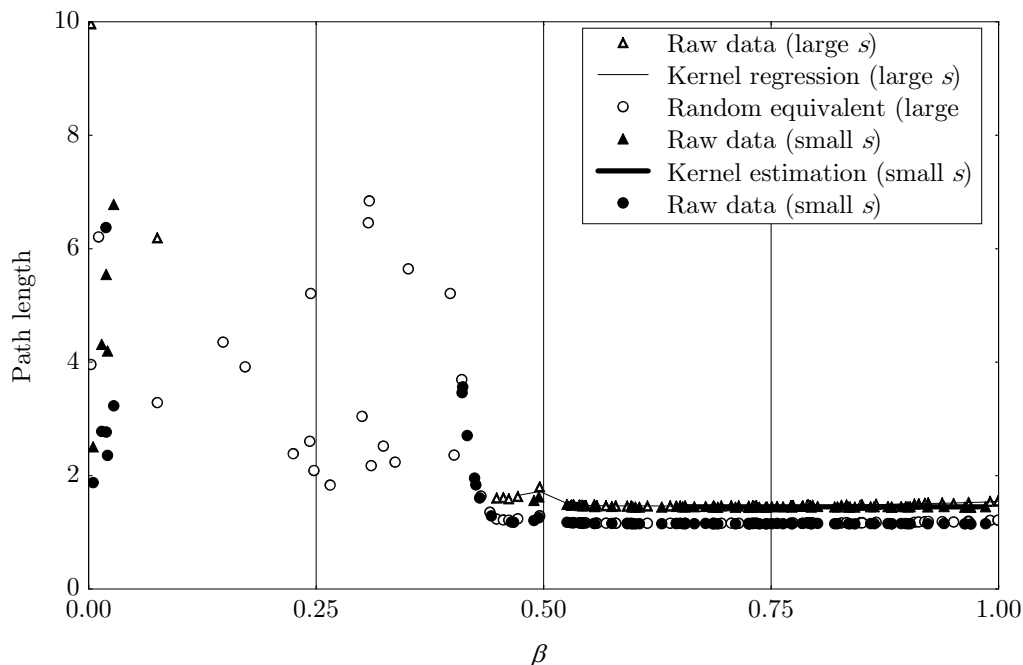


Figure 4: Average path length between any pair of individuals

makes the average path length infinite.¹¹ To keep the figure informative, we only consider a small range over the y -interval so between 0 and β^* most of the points are invisible. When $\beta > \beta^*$, path length is very small. The lowest possible path length is 1, and exists only in a complete graph. The networks we obtain are thus not complete (their degree is below 55, and a complete graph must have degree 99 in our case) but do have relatively short paths. What this implies is that knowledge flows relatively rapidly around the graph, and indirect exploitation of distant knowledge, through successive rounds of innovation, can be an active feature of the economy. It is known that for a fixed density, random graphs have very short average path lengths. The formula (which is valid for infinite size graphs) for the average path length

¹¹Recall that average path length refers to the average path length between pairs of nodes in the graph. If there is no path between two agents, this creates an infinitely long path between them, making the average path infinite. In our calculation of path length, two disconnected agents have a path length of 10,000 rather than infinity. This makes the calculation possible.

over a random graph with n vertices and average degree d is $\ln(n)/\ln(d)$. In Figure 4 we represent the path length of the equivalent random graph (that is the random graph with 100 vertices and a degree a given in Figure 2). The path length in our case is consistently above the path length of the equivalent random graph. The networks that emerge from our dynamic process are not random, but do have path lengths approaching those of a random graph of the same degree. Interestingly, and related to the previous paragraph, it is worth mentioning that in a diversified trajectory (and to a smaller extent in a focussed one) there are situations with low β ($\beta \approx 0$) in which path length is finite (between 2 and 7 in the figure).

5.1.3 Local Order and Network Structures

To describe local order we use the standard measure of cliquishness: the share of my partners who are also partners of each other.

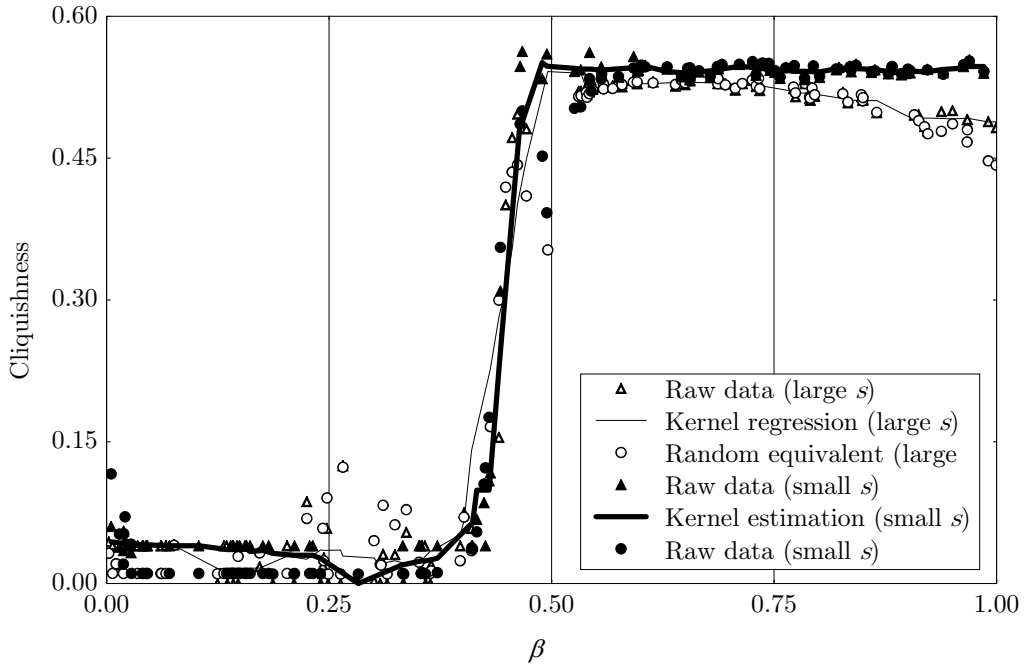


Figure 5: Average cliquishness

Like the previous structural measures, cliquishness shows a phase change

at the same value of β^* . Cliquishness is not defined for a graph of disjointed pairs, so our convention is to assign a value of zero for cliquishness of a pair. At the other extreme, a complete graph has a value of 1, and so does a collection of complete disconnected components. In figure 5 we see for small values of β many networks with C close to 0, but most often strictly positive. Since we know that the graph is disconnected here (and thus not complete) this means that we must be observing the emergence of a “caveman graph”, that is, a graph of disconnected, complete subgraphs. From figure 2 it is clear that the caves are very small, since the average number of connections per agent is between 1 and 3. When the value of the average cliquishness gets very small, is that there are many isolated pairs in the network. In this part of the parameter space “only length matters” so the two agents with the most knowledge will partner, and create the largest innovation, thus maintaining their positions as the most knowledgeable agents. Their partnerships is stable, as is the partnership of the next most knowledgeable pair, and so on. There is a small range, roughly $0.3 < \beta < 0.45$ in which the graph consists of several, but a falling number of, subgraphs which are themselves highly cliquish. We no longer have a caveman graph, but rather several connected, but not complete subgraphs.

The cliquishness of the equivalent random graph almost perfectly tracks the observed cliquishness, whatever the value of s .

5.2 Knowledge dynamics

Figure 6 depicts the relationship between the long-run economy-wide knowledge growth rate and β .

For both values of s there seem to be, again, two regimes of knowledge growth — one associated with a disconnected graph, and one with a connected graph. In both regions knowledge growth rates as diversity becomes more important. The contrast between large and small s is explained by the fact that with a small value of s , there is a natural convergence over time in the knowledge type of the population. Since innovations tend to occur “between” the knowledge of the members of the innovating pair, knowledge types converge over time. When β is large, diversity is important, and this convergence effect reduces growth. When s is large, however, diversity is preserved over time, since there is a large variance in the type of knowledge created, so the natural convergence is mitigated by the randomness here.

Again, though, there is a small ambiguous region, for $0.25 < \beta < 0.4$,




Figure 6: Growth rate of the average knowledge level

where there is some intermediate behaviour between that of the connected and disconnected regimes. As discussed above, in this region we see the emergence of something resembling a caveman graph — small connected (sometimes complete) sub-graphs within the disconnected total graph. Within each cave, we have the behaviour of a connected graph, which, holding β constant, is good for knowledge production.

The issue of equity is addressed by considering $\gamma = \sigma/\bar{\rho}$, the coefficient of variation of individual knowledge levels. Coefficient γ measures how dispersed around the average individual the population is, relative to this average individual. As we re-normalize at each time step, the coefficient of variation reduces to the standard deviation. Figure 7 depicts the way γ varies with β .

For $\beta < \beta^*$, γ is approximately equal to 7, which corresponds to a situation in which knowledge is concentrated in the hands of exactly 2 individu-




Figure 7: Coefficient of variation of individual knowledge levels

als.¹² There is a first fall in dispersion a bit before the critical β , and then a very marked one down to a point where under both small and large s things stabilize. For small s the coefficient of variation stays around 1 whereas it stabilizes in the vicinity of 0.5 for large s . Hence dispersion decreases with β in a non-smooth way, and the decrease is more marked when the trajectory is more dispersed, undergoing a phase change which again is very comparable to what could be observed concerning network structure. In the beginning of the parameter space “only length matters” so the two agents with the most knowledge will partner, and create the largest innovation, thus maintaining their positions as the most knowledgeable agents. Their partnerships is stable, as is the partnership of the next most knowledgeable pair, and so on. As renormalization carries on, we end up with all the “renormalized” knowledge

¹²Given that we renormalize knowledge levels each period, the standard deviation is equal to 7 if and only if 98 individuals have 0 (or almost 0) and 2 have 50, for an average 1.

being held by the two most knowledgeable individuals. It takes a significant weight on variety to combat this tendency and to allow for pairs to fluctuate, entailing that homogeneity in terms of levels decreases. It is worth mentioning that at the other end of the parameter space, “only angle difference matters”, which means that the two most distant (in angle) individuals will be matched, so will the next more different individuals, and so on. However the size of their innovation does not matter, only angles do. So with time, because innovation tends to take place “between” the knowledges of the innovating pair, there is a trend of decreasing diversity in the population, as all agents lie between the two most extreme agents of the previous period. This effect is mitigated, of course, as s increases.

6 Conclusion

In this paper we have focussed on issues of knowledge dynamics and emerging network structures when agents create knowledge through partnership agreements. We have emphasized that knowledge creation through cooperation can be a complex process that involves pooling the competencies of the partners in different ways, depending on the nature of the innovation process. Similarly, given a pooling structure, different types of knowledge can be either substitutes or complements for each other in the innovation process itself. The former consideration is represented by the θ parameter in our model, the latter by the elasticity of substitution β in a constant elasticity of substitution production function. Both parameters affect the results, both in terms of the rate of knowledge production, and in terms of the network structure. Knowledge creation is fastest when the knowledge creation production function exhibits complementarity in its inputs. The nature of knowledge pooling has no noticeable effect however. By contrast, heterogeneity among agents regarding their knowledge levels is, generally, decreasing in θ . As it becomes more possible for one partner to dominate in the pooling process (or perform more of the innovation activity) heterogeneity across agents decreases.

The nature of knowledge pooling turns out to be crucial in determining the emergent structure of the network. When θ is very low, stable matchings create pairs of agents whose expertise lies in the same fields — agents can be seen as substituting for each other. Innovation tends to occur in that single field of expertise, and the two agents remain good partners for each

other in the next round. This implies the emergence of stable pairs of agents; the network becomes a set of $n/2$ disconnected pairs. By contrast, when θ approaches 1, stable matchings create pairs of agents whose respective expertise lies in different fields. Innovation in this case does not reinforce existing expertise of the two agents and thus make them more similar, but rather reducing the differences between them. Thus a pair matched in this round, because they are more similar to each other after innovation are less likely to be a good match in the next round. What this implies is that pairs of agents will form and disintegrate rapidly, agents will constantly search for new partners, and the emergent network will approach a random network. With θ values between these extremes, the network structure too is between the extremes. Here knowledge pooling creates a mixture of the two agents' knowledge with each knowledge type. As θ grows from its minimum value the diameter of the stable groups or cliques increases from two. The emergent cliques tend to be relatively complete, but globally the graph remains disconnected. Cliquishness increases, and is maximal at $\theta = 0.5$. For intermediate values, $0.5 < \theta < 1$, cliquishness falls, as the islands of communicating agents make external connections, opening themselves to each other by establishment of "shortcut" links. Global path length falls in this region and it may be that small worlds, as defined by Watts and Strogatz (1998), emerge.

References

- [1] C. Antonelli. Communication and innovation : The evidence within technological districts. Presented at the International Conference "Knowledge Spillovers and the Geography of Innovation: A Comparison of National Systems of Innovation", Université Jean Monnet, Saint Etienne, France, 1999.
- [2] V. Bala and S. Goyal. Learning from neighbours. *Review of Economic Studies*, 65:595–621, 1998.
- [3] V. Bala and S. Goyal. A non-cooperative model of network formation. *Econometrica*, 68:1181–1229, 2000.
- [4] R. Cowan and N. Jonard. Network structure and the diffusion of knowledge. Technical Report 99028, MERIT, Maastricht University, 1999.

- [5] R. Cowan and N. Jonard. The dynamics of collective invention. Technical Report 20-018, MERIT, Maastricht University, 2000.
- [6] R. Cowan and G. Van de Paal. Innovation policy in a knowledge-based economy. European Commission, Enterprise Directorate-General, 2000.
- [7] P. A. David and D. Foray. Markov random fields, percolation structures and the economics of document standards diffusion. Technical Report 326, CEPR, Stanford University, 1992.
- [8] G. Debreu. Neighbouring economic agents. *La Decision*, 171:85–90, 1969.
- [9] D. Gale and L. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 69:9–15, 1962.
- [10] M. Gibbons, C. Limoges, H. Nowotny, S. Schartzman, and M. Scott, P. AndTroiw. *The New Production of Knowledge : The Dynamics of Science and Research in Contemporary Societies*. SAGE, London, 1994.
- [11] Z. Griliches. Hybrid corn: An exploration in the economics of technological change. *Econometrica*, 25:501–522, 1957.
- [12] J. Hagedoorn. Inter-firm R&D partnerships: An overview of major trends and patterns since 1960. *Research Policy*, 2001. forthcoming.
- [13] H. Haller. Topologies as infrastructures. In R. Gilles and P. Ruys, editors, *Imperfections and Behavior in Economic Organizations*. Kluwer, 1994.
- [14] A. Hargadon and R. Sutton. Technology brokering and innovation in a product development firm. *Administrative Science Quarterly*, 42:716–749, 1997.
- [15] R. Irving. An efficient algorithm for the “stable roommates” problem. *Journal of Algorithms*, 6:577–595, 1985.
- [16] M. Jackson and B. Dutta. On the Formation of Networks and Groups. In M. Jackson and B. Dutta, editors, *Models of the Strategic Formation of Networks and Groups*, Springer, 2001.

- [17] M. O. Jackson and A. Wolinsky. A strategic model of social and economic networks. *Journal of Economic Theory*, 71:44–74, 1996.
- [18] D. E. Knuth. *Mariages Stables*. Presses Universitaires de Montréal, 1976.
- [19] F. Kodama. Japanese innovation in mechatronics technology. *Science and Public Policy*, 13:44–51, 1986.
- [20] F. Kodama. Technology fusion and the new research and development. *Harvard Business Review*, 22:70–78, 1992.
- [21] E. Mansfield. Intrafirm rates of diffusion of an innovation. *Review of Economics and Statistics*, 45:345–359, 1961.
- [22] F. Plouraboue, A. Steyer, and J.-B. Zimmermann. Learning induced criticality in consumers adoption pattern: A neural network approach. *Economics of Innovation and New Technology*, 6:73–90, 1998.
- [23] C. Qin. Endogenous formation of cooperation structures. *Journal of Economic Theory*, 71:218–226, 1996.
- [24] K. Smith. What is the ‘knowledge economy’? Knowledge-intensive industries and distributed knowledge bases. Presented to DRUID Summer Conference on “The Learning Economy — Firms, Regions and Nation Specific Institutions”, June 15-17, 2000
- [25] A. Steyer and J.-B. Zimmermann. On the frontier: Structural effects in a diffusion model based on influence matrixes. In P. Cohendet, P. Llerena, H. Stahn, and G. Umbhauer, editors, *The Economics of Networks: Interaction and Behaviours*, pages 118–135. Springer, 1998.
- [26] R. Sutton and A. Hargadon. Brainstorming groups in context: Effectiveness in a product design firm. *Administrative Science Quarterly*, 41:685–718, 1996.
- [27] T. W. Valente. *Network Models of the Diffusion of Innovations*. Cresskill, NJ: Hampton Press., 1995.
- [28] D. Watts and S. Strogatz. Collective Dynamics of Small-World Networks., *Letters to Nature*, 393, May, 1998.

- [29] G. Weisbuch, A. Kirman, and D. Herreiner. Market organization. In R. Conte, R. Hegselmann, and P. Terna, editors, *Simulating Social Phenomena*, pages 221–240. Springer, 1997. Lecture Notes in Economics and Mathematics, Volume 456.
- [30] J.-B. Zimmermann. Le concept de grappes technologiques: Un cadre formel. *Revue Economique*, 1995.