Maastricht University

Networks as emergent structures from bilateral collaboration

Citation for published version (APA):

Cowan, R., Jonard, N., & Zimmermann, J-B. (2004). Networks as emergent structures from bilateral collaboration. (MERIT Research Memoranda; No. 017). Maastricht: MERIT, Maastricht Economic Research Institute on Innovation and Technology.

Document status and date:
Published: 01/01/2004

Document Version:
Publisher's PDF, also known as Version of record

Please check the document version of this publication:

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

Link to publication

Networks as Emergent Structures from Bilateral Collaboration

Robin Cowan, Nicolas Jonard & Jean-Benoît Zimmermann

2004-017

Networks as Emergent Structures from Bilateral Collaboration*

R. Cowan[†], N. Jonard[‡], and J-B. Zimmermann[♮]

[†]MERIT, University of Maastricht, P.O. Box 616
6200 MD Maastricht, The Netherlands
[‡]CNRS, CREA Ecole Polytechnique, 1 Rue Descartes
75005 Paris, France
[♮]CNRS, GREQAM, EHESS, 2 Rue de la Charité
13002 Marseille, France

August 31, 2004

Abstract

In this paper we model the formation of innovation networks as they emerge from bilateral actions. The effectiveness of a bilateral collaboration is determined by cognitive, relational and structural embeddedness. Innovation results from the recombination of knowledge held by the partners to the collaboration, and the extent to which agents' knowledge complement each others is an issue of cognitive embeddedness. Previous collaborations (relational embeddedness) increase the probability of a successful collaboration; as does information gained from common third parties (structural embeddedness). As a result of repeated alliance formation, a network emerges whose properties are studied, together with those of the process of knowledge creation. Two features are central to the innovation process: how agents pool their knowledge resources; and how agents derive information about potential partners. We focus on the interplay between these two dimensions, and find that they both matter. The networks that emerge are not random, but in certain parts of the parameter space have properties of small worlds.
JEL Classification: L14, Z13, O3
Keywords: Networks, Innovation, Network Formation, Knowledge

1 Introduction

One of the effects of the recent rapid technical advance is a change in the technological structure of many firms. As new technologies have emerged and been integrated into existing products and technology spaces, the successful firm now must have access to expertise covering a much larger scope of technologies than in the past. We have observed a rapid increase in the prevalence of "multi-technology" firms. Firms have had to incorporate many new types of expertise both in their production activities, and in innovation (see for example Powell et al., 1996, Grandstand and Sjolander, 1990; Grandstand, 1996; Teece and Pisano, 1989). This raises the difficulty that knowledge and technology necessary for innovation may lie outside a firm's traditional core competence. A now common strategy for addressing this problem, adopted by more and more firms, is to form alliances, both with competitors and with non-competing firms and institutions. 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. Consequently, these co-operative agreements for R&D have grown dramatically in number since the 1970s.[1]

The increase in alliances, which have been created for purposes of gathering or exchanging knowledge and information, has led to a new view of industry structure. In the past we have observed hierarchies and markets as dominant structures. Implicit is a well-defined notion of the boundary of the firm. As knowledge (and indeed many other things) are passing between firms in what cannot be described as purely market transactions, analysts have begun to discuss the network-based organization (see Powell, 1990 for example). Networks differ from markets and hierarchies in a variety of ways, but can be seen as depending on particular types of interactions between pairs of agents within the economy. Within a network structure, firm boundaries are relatively porous, and a firm survives by having good contacts with other firms who hold complementary assets.

Clearly, the strategic alliance is a central part of the process creating a network structure within an industry. For any firm seeking to expand its innovation capabilities through alliance formation, though, there is the question of choosing a partner. Simply partnering with every other firm is not feasible for many reasons (there are time, resource and management constraints; competitive considerations would militate against it and so on), so selecting a partner, acknowledging that many other firms are doing the same, becomes a strategic issue that firms must deal with. A basic way of analysing this issue is to observe that alliances do not exist in a vacuum, but are embedded in a variety of other issues. Analysts have focussed on three: relational embeddedness; cognitive embeddedness; and structural embeddedness. Each of these features of a possible relationship lends it value, and they play a central role in determining the desirability of different potential partners.

In this paper we examine the creation and evolution of a network of firms through bilateral alliance formation. The micro-economics of our model deal with the firms' immediate goal, which is to find an alliance for purposes of innovation, and in the model each firm's choice of partner is driven explicitly by concerns of cognitive, structural and relational embeddedness. These micro-economic activities, however, create an emergent

---

[1]See Hagedoorn (2001) for a review and discussion of this trend.

network structure in which firms are connected to each other through bilateral links.[2]

In the context of finding an alliance partner, cognitive embeddedness refers to two agents' abilities to integrate effectively their respective knowledge. Empirical analyses of alliance formation conclude that firms look for partners with the best technological fit in the sense of providing "missing resources".[3] If we consider that a firm has a knowledge profile, we can think of it as being located in some knowledge or competence space. When two firms innovate jointly, the action will be most effective if their locations complement each other. In the literature a Euclidean space has been used, and there is a consensus that the effectiveness of cooperation has an inverted U-shape in cognitive distance. If firms are too close together, their knowledge overlaps too much and there is little point in sharing; if they are too far apart they have difficulty understanding each other, and so sharing is too difficult. The arguments are very appealing intuitively (see for example Grant, 1996 or Nooteboom, 1999), and Mowery et al. (1998) find this effect empirically. In the model below we take a more explicit notion of complementarity, using a richer notion of complementarity than a simple Euclidean distance.

One consequence of a knowledge partnership is that partners will develop closer cognitive ties. That is, their knowledge profiles will, in general become more similar. Mowery et al. (1998), for example find that "technological overlap between joint venture partners after alliance formation is greater than their pre-alliance overlap" (p. 517). (See also Dyer and Nobeoka, 2000.) This has the feature of increasing embeddedness, but after a time may make firms less attractive to each other, since as they become similar, there is less to share. This intuition may be misleading in some cases, and as we discuss below, it depends heavily on the nature of the innovation process, and how firms integrate their competencies in that process.

Cooperation between firms is risky, and is marked by uncertainty regarding the skills of the partner, your joint ability to work together, the potential partner's reliability, his goals and so on. (See Powell 1990, p. 318 for a discussion of these risks.) This can be cast as an issue of incomplete information, and the most obvious way to reduce the uncertainty is to improve the information used in choosing a partner. There are two sources of this type of information: experience and other agents. The first relates to relational embeddedness, the second to structural embeddedness (see Uzzi 1996, 1997 for discussion).

Past experiences with an agent will both improve abilities to cooperate, and yield information about that agent. Cooperation implies mutual knowledge and sharing of routines, representations, ways of thinking, the ability to share tacit knowledge and so on: in general the creation of some common ground on which the cooperation can be built. Galison (1999), for example, shows that in experimental physics, cooperation between theorists,

---

[2]To model strategic technological alliances in their entirety is far beyond the scope of this paper. Technological alliances can be very rich and varied, and firms have many motivations for entering into them. See Oliver (1990) for a discussion of the motivations of firms to form outsider relationships. But for our purposes we focus on a single effect, namely the production of shared knowledge, and how firms' behaviour in this regard leads to the emergence of networks.

[3]This idea has a long history in the management literature, going back to Penrose (1959). See Chung et al. (2000); Doz (1988); Hamel et al. (1989); Narula (1999); Richardson (1972); Teece (1986) among others. In this regard, one difficult issue has to do with how resources of partners are combined, which in turn determines what "complementary" means. This is an issue we address directly below.

experimentalists and instrument makers is made possible by the emergence of some sort of "creole", an intermediary level knowledge specific to a given pair of actors, that has been built through their repeated interactions. This common ground can be built through the experience of cooperation (Garcia-Pont and Nohria, 2002). In addition, repeated interaction creates trust in the broadest sense (both in terms of motives and in terms of competencies) and the importance of predictability implies that trust construction and learning are strongly related (Sako, 1991; Dodgson, 1996). All of these considerations will clearly create an inertia in partnership formation, and a stability in the emergent network structures: agents will, all else equal, prefer partners they have worked with in the past. Powell et al. (1996) show that firms that have engaged in partnerships in the past are more likely to engage in them in the future. This is a general result. Chung et al. (2000), Gulati (1995) and Roijakkers (2003) all find that if two particular firms have allied with each other in the past, these two firms are more likely to have an alliance together in the future.

The second source of information about potential partners is other agents (see Kogut et al. 1992, for example). Those who have worked with a firm will have experience that they can, in principle, share with others who might be considering working with that firm. This is captured by the idea that many alliances exhibit structural embeddedness: there is some tendency for firms to find partners that are close to them in network space. In the model we develop below, this source of information is included explicitly: a firms perceived value to me as a partner increases if my previous partners have had good experiences with that firm. My network of immediate contacts is a source of information about possible future contacts.

The entanglement of these different effects — learning about partners and learning from partners — makes analysis of network formation tricky. Both Chung et al. (2000) and Gulati (1995), examining very different industries, find that both a history of direct interactions and indirect ties affect the probability that two firms will form a partnership in the future. As the number of past ties between two firms increases the probability that they will form a partnership together in the future increases and then decreases. There seems to be a concave relationship. While it seems reasonable that the effect of a shared history should taper off, it is not immediately clear why it should turn negative at any point. Similarly, the number of indirect links between two firms has an effect on the probability that they will form a direct link or partnership in the future: the effect is initially positive, but again tapers off, Chung et al. suggesting that their data display an inverted-U.[4] That a shared history or shared neighbours could have a negative effect seems odd. The explanation may be in the difficulty of measuring their other explanatory variable, namely strategic complementarity. Both papers acknowledge that this is a difficult variable to measure, and both provide a *static* measure of it. This means that in their implicit model, when two firms interact there is no effect on the strength of their strategic complementarity. We have argued that this is likely not to be the case, and in fact that strategic complementarity should weaken after a point, as two firms continue to interact

---

[4]This claim should be treated cautiously, since the peak occurs at 185 alliances according to their calculations.

and learn from each other, their knowledge bases will eventually have too big an overlap (and indeed this convergence is what Mowery et al., 1998, and Uzzi, 1997, find). Such an effect would explain a decrease in strength, and even a negative effect, of past interactions on the probability that two firms interact in the future. The model that we develop below has this built in, and under certain conditions about the nature of the complementarity, this effect is clear.

In this paper we design an agent-based model of network formation in which agents are repeatedly forming bilateral pairs for the purposes of creating new knowledge.[5] We abstract from pure "network-oriented" strategic motives of firms, such as filling structural holes (Burt, 1992), or increasing their positions of centrality (Podolny, 1993). We focus instead on the effects of firms' motivations driven by short-term innovation concerns. Partnerships are embedded cognitively, structurally and relationally, and this embeddedness directly influences the choice of partner by every firm.[6] Over time this bilateral link formation process results in an emergent structure which is characterised as an industrial network.[7] We analyse the effects of the nature of the innovation process, and the way firms can effectively pool their knowledge on the structure that emerges and on the nature of the knowledge that is held by the network.

2 The model

We present here a schematic description of the model before turning to the formal description. Each period every firm forms an alliance with one other firm, based on the output that they expect from the collaboration. If preferable a firm can also choose isolation. An alliance having formed, the firms pool their knowledge and use the joint knowledge stock as inputs into new knowledge production. If the alliance is successful, and new knowledge is created, this is added to each firm's existing stock of knowledge, and then the partnership is dissolved. In the next period firms form new alliances, possibly with previous partners,

---

[5]There is now a growing literature in economics on network formation, but the majority of it treats the problem in game-theoretic terms, looking for stable structures that emerge from agents's one-time decisions about whether to form links. The concern in this literature tends to be whether the stable networks are efficient. (For a recent survey see Jackson and Dutta, 2003.) This work tends not to address the evolutionary nature of network formation and operation, and tends to have a thin model of firm behaviour. For an example of adaptive networks in the context of customer loyalty, see Kirman and Vriend (2001).

[6]These three types of embeddedness span the two types of explanations for partner choice in the literature, namely resource complementarity and social structural context. In this regard the model we develop here is in the tradition of empirical work such as Gulati (1995) and Chung et al. (2000) which find support for both explanations.

[7]Recent empirical analysis of economic networks has focused to a large extent on their structural properties, in particular whether various networks are small worlds. In general the answer is yes. Co-authorship in a variety of academic disciplines (Newman, 2001); patent citation in US biotech (Johnson and Mareva, 2002); interlocking corporate directorships in the US (for example, Davis et al., 2003); technology alliances (Duysters and Verspagen, forthcoming); the BRITE/EURAM network and the 5th Framework TSER network (Cowan and Jonard, 2004) all have small world properties, and there is a consensus that small worlds are pervasive. A second structure that has received attention recently is the scale free network (Barabasi and Albert, 2000). Riccaboni and Pammolli (2000) find that networks in the life sciences and ICT industries have scale free properties.

4

possibly with new partners. This repeated bilateral partnership formation and dissolution generates the emergence and evolution of an economy-wide network structure, fosters knowledge growth, and changes the knowledge endowments of the firms in the economy.[8]

Formally, consider a finite population of agents denoted $S = \{1, \ldots, n\}$. Each agent $i \in S$ is characterized by a knowledge endowment of several types of knowledge. This is represented as a vector $v(i) = (v^1(i), \ldots, v^{\mathbf{m}}(i))$, where each element $v^m(i) \geq 0$ represents the amount of knowledge of type $m = 1, \ldots, \mathbf{m}$ held by agent i. This representation permits us to treat agents as located in an $\mathbf{m}$-dimensional knowledge space. Knowledge is thus treated as a form of human capital, of which distinct types exist. This notation describes the agents, and we use it throughout. We treat individually the two parts of the model: knowledge creation, and partnership formation.

2.1 Knowledge production

Both knowledge and innovation are very difficult to characterize satisfactorily, and any characterization has its weaknesses. We adopt the simple vector characterization of knowledge described above, but focus more attention on the properties of joint innovation. A representation of the innovation process should satisfy several minimal requirements (for supporting evidence see Mowery et al., 1998; Uzzi, 1997). Consider two individuals i and j who innovate together. After innovation has taken place, one would expect the following to be true: the knowledge amounts held by i and j have increased; the knowledge "profiles" of i and j have changed; and the similarity of the knowledge profiles of i and j (that is the relative distance between them in the knowledge space) has fallen.

Operationally, each pair of agents (i, j) creates an amount of new knowledge determined by a production function, and this amount is simply added to both partners' existing knowledge endowments. This process is very simple in principle and satisfies the three requirements just described. Recall that each period agents form pairs (in a process discussed below), and the knowledge that they bring to the pair is pooled, then serving as input to the innovation process. The pooling is done through

$$v^m(i, j) = (1 - \theta) \min\{v^m(i), v^m(j)\} + \theta \max\{v^m(i), v^m(j)\}, \ m = 1, \ldots, \mathbf{m}. \quad (1)$$

The pooling of course remains virtual, as each of the partners remains the owner of his own skills. Nevertheless, it permits a useful formalization, permitting a concise representation of the inputs available for the joint innovation production process.

In Equation (1), θ reflects the nature of the knowledge pooling which the knowledge creation task demands. If it is possible to separate the sub-tasks in the innovation process, the agents will specialize, each agent doing some sub-tasks, and bringing the results together at the end to create the complete innovation. Here, since specialization is possible, the better

---

[8]We have assumed here that agents are pursuing knowledge for its own sake. This is unrealistic in general for firms, who pursue knowledge more generally for the sake of profits. To incorporate that explicitly in the model adds significant complication, demanding a fully blown goods market with production and consumption. We avoid that by this simplifying assumption, which, in an industry involved in rapid technical change, will be behaviourally quite adequate.

econometrician will do the econometrics; the better growth theorist will do the growth theory, and so on, and thus the pooled knowledge vector is the element-wise maximum of the individual vectors. θ is close to one. By contrast, if the tasks are not separable, and both partners must be involved in every sub-task, then the weaker partner will be a bottleneck: the pooled knowledge vector will be the element-wise minimum, and θ is close to zero. The advantage of this formalization is that we can avoid specifying a particular distance function, and an optimal distance (the peak of the inverted-U). More specifically, it permits us to have a richer model of joint innovation than is usually considered (for example by Mowery et al., 1998; Nooteboom, 2000; Alstyne and Brynjolfsson, 1996, 1997; Peretto and Smulders, 2002) in which a simple Euclidean distance in knowledge space determines the effectiveness of a partnership. Putting knowledge together involves considering deeper issues of complementarities, captured by the production function approach.

Knowledge pooling is where strategic complementarity is defined. What counts as a complement depends on the nature of the innovation process, and thus on the nature of the way knowledge can be pooled. On this interpretation θ turns out to be a measure of the taste for dissimilar partners. If θ is close to 0 then for any element m, in which $v^m(j) < v^m(i)$, agent j reduces the effectiveness of i. The converse is true as well. In this case, agents will be driven to find partners similar to themselves, so they create as little drag on each other as possible, and possibly agents end up alone. By contrast, if θ is close to 1, agents look for partners whose endowments tend to complement their own, since they can benefit from each others' strengths. Implicitly, they search for partners who are different from themselves in the sense of being good where they are bad. Naturally every agent would still prefer to be with someone better than him in every discipline, but then the dominant agent would refuse the partnership.

The pooled knowledge vector serves as the vector of inputs to the innovation production function. To formalize this, we use a standard constant elasticity of substitution (CES) production function $\phi : \mathbf{R}_+^{\mathbf{m}} \to \mathbf{R}_+$, with

$$\phi\left(v\left(i,j\right)\right) = \left\{ \sum_m \left(v^m\left(i,j\right)\right)^\beta \right\}^{1/\beta}. \tag{2}$$

The parameter $\beta \leq 1$, $\beta \neq 0$, is an inverse measure of the elasticity of substitution across knowledge types, which is written as $1/(1-\beta)$. To see how it affects the type of partnerships wanted, consider an agent with asymmetric profile (i.e. marked strengths and weaknesses) and remark that ϕ is symmetric in its arguments and homogeneous. In general this agent would like to find a partnership such that the joint profile is more evenly distributed and so a higher isoquant can be reached. How desirable this is varies with the degree of substitutability across knowledge types. When β is small the agent is very eager to find a partnership and many different partners will complement him in such a way that his innovation output increases. By contrast, when β is close to 1, being evenly skilful becomes less important since substitution between different knowledge types gets easier. Consequently, less networking should be observed.

6

2.2 Innovative success and experience

An innovation project may fail, and the projected new knowledge may not be created. Here history is important since project success is driven in part by familiarity of the partners, and the nature of the embeddedness of the partnership, as discussed above.

Suppose i and j have had partnerships in the past. The probability that i and j, will have a successful collaboration is determined by the information they have about each other at the time of the collaboration. As discussed in the introduction, information has two sources: their shared history, and mutual former partners. Our focus will be on how the balance between my own history (direct information) and my partners' histories (indirect information) affects my decisions.

Conditional to what has happened up to period $t-1$, define the a priori probability that a collaboration between i and j in the next period, t, is successful as $\pi_t(i,j) = \pi_t(j,i)$. Suppose further that their previous collaboration took place $s(i,j)$ periods ago (so in period $t - s(i,j)$). (Henceforth we drop the t subscript for readability.) Now define the *direct credit* of the pair (i,j) in period t to be

$$\gamma(i,j) = \rho^{s(i,j)}\chi(i,j), \tag{3}$$

where $0 < \rho < 1$ is a discount factor, and $\chi(i,j) = 1$ if i and j had a successful interaction when they last tried $s(i,j)$ periods ago, while $\chi(i,j) = 0$ otherwise. Equation (3) captures the simple assumption of learning to collaborate by successfully collaborating. The more distant our previous success is in time, the less credit we have in each other's eyes (and the less likely it is that our next attempt is a success). In case of a failure zero credit obtains. Focussing solely on Equation (3) would give rise to a situation in which partners are a source of (possibly) complementary knowledge, but no information about other agents, and my own interactions are the only basis on which to judge the value of other agents as potential partners.

In a similar way, define now the *indirect credit* of the pair (i,j) in period t to be

$$\eta(i,j) = \sum_{k \neq i \neq j} \rho^{s(i,k)}\chi(i,k) \cdot \rho^{s(k,j)}\chi(k,j), \tag{4}$$

where $\chi(i,k)$ and $\chi(k,j)$ are defined as earlier. The product $\chi(i,k)\chi(k,j)$ is non zero only if the last attempt of i and k was a success, and the last attempt of k and j was a success. Through the decay term $\rho^{s(i,k)}\rho^{s(k,j)}$ indirect credit also weakens as time passes without new successes. A sketch of the way this works would be as follows. Suppose I am interested in agent j. I look through my list of past partners, and when I find one with whom I was successful in our most recent interaction, I ask, "In your latest interaction with j, were you successful [which determines $\chi(k,j)$], and if so, when was that interaction [which determines $\rho^{s(k,j)}$]?" The sum of these over all my previous partners determines the indirect credit between me (i) and j. Focussing solely on Equation (4) would imply that partners, in addition to being a source of (possibly) complementary knowledge, are the only source of information.

In general both direct and indirect information or credit will be considered, so to inter-

polate between the two extremes the credit of the pair (i, j) will be a weighted sum of the two, according to

$$c\left(i,j\right) = \frac{\alpha\gamma(i,j)}{\alpha + (1-\alpha)\sum_{k\neq i}\chi_{i,k}\chi_{k,j}} + \frac{(1-\alpha)\,\eta\left(i,j\right)}{\alpha + (1-\alpha)\sum_{k\neq i}\chi_{i,k}\chi_{k,j}}, \tag{5}$$

with $0 \leq \alpha \leq 1$ the weight of the agent's own opinion. This way $c\left(i,j\right)$ accounts for both i's direct knowledge of j (as captured by $\gamma(i,j)$) and the knowledge accessed by i via neighbours who are common to j, as measured by $\eta\left(i,j\right)$. Then the probability that the next collaborative attempt is a success is simply assumed to be

$$\pi\left(i,j\right) = \pi_L + (\pi_H - \pi_L)c\left(i,j\right), \tag{6}$$

and the expected amount of knowledge produced by a cooperation between i and j can be expressed as

$$F(i,j) = \pi\left(i,j\right) \cdot \phi(v\left(i,j\right)). \tag{7}$$

This is the amount of knowledge produced in case of success multiplied by the probability that the cooperation succeeds. (Thus agents are assumed to be risk-neutral.) Firms that innovate in isolation have one source of risk removed, namely that associated with having to work with a partner. This does not make autarchic innovation a sure thing though: we assume that a firm innovating alone is successful with probability π_H. Given this, π_L plays a central role in a firm's decision to operate in autarchy rather than collaborating. A larger baseline probability of failure associated with the partner's contribution (lower π_L) will obviously imply a stronger tendency for firms to conduct innovative projects in isolation. Note that $F\left(i,j\right) = F\left(j,i\right)$, i.e. the score function is symmetric.

If the innovation project is successful, and new knowledge is created, it is added to each of the partners' knowledge vectors. The general intuition is that as an agent uses knowledge or is exposed to it, he will assimilate at least part of it, and will thereby change the precise area of his expertise. As the argument of the production function is the joint knowledge profile, it seems natural to let this joint profile also determine the type of knowledge produced.[9] It is assumed that the probability, conditional to the collaboration being successful, of the new knowledge being of type m is

$$\frac{v^m\left(i,j\right)}{\sum_m v^m\left(i,j\right)}. \tag{8}$$

If the collaboration fails, both agents get 0.

---

[9]We have explored other variants in which a share of the new knowledge is allocated according to the joint profile to a category common to both participants, while the remaining part is allocated according to each partner's profile. As long as the share of new knowledge allocated according to the joint profile is not negligible, the results remain qualitatively unchanged.

2.3 Pair formation and equilibrium

We draw on the literature on matching problems for our basic model of pair formation. Because we consider a single population of firms rather than two populations (of jobs and workers for instance) matching here is a roommate problem, rather than a standard marriage problem. A one-sided, roommate, matching problem is defined as follows (see Gale and Shapley, 1962). Each individual $i \in S$ has a strict preference ordering $\succ_i$ over all the individuals in $S - \{i\}$. All preferences are complete and transitive. Let $\succ = \{\succ_i, i \in S\}$ denote the profile of the preference orderings of all the individuals in S. We generalize the standard roommate problem to include the possibility of self-matching. This is done in a straightforward way: the preference ordering is over the entire set S. The pair $(S, \succ)$ is a generalized roommate matching problem, and a matching is a partition of S into q singletons and $(n-q)/2$ pairs of roommates, that is to say a bijection $\mu : S \to S$ such that $\mu(\mu(i)) = i$ for all $i \in S$.

Having defined a matching, the equilibrium concept can be presented. A matching μ is said to be stable in $(S, \succ)$ if there is no $(i, j) \notin \mu$ such that both $j \succ_i \mu(i)$ and $i \succ_j \mu(j)$. Put another way, stability is characterized by the non-existence of blocking pairs. In the particular problem examined here, the preference profile $\succ$ is generated by the expected output of a pairing $F : S^2 \to \mathbf{R}_+$, which associates to any pair of individuals (i, j) a value that represents the expected innovative output of this pair. In the event that $i = j$ the pooled vector is simply the vector of i, and production remains defined as it was above. The profile of preference orderings of $i \in S$ is then defined for all (j, k) in S^2 by

$$j \succ_i k \text{ iff } F(i, j) > F(i, k), \tag{9}$$

i.e. preferences are directly derived from expected output.

Before turning to the emergence of network structure and the associated knowledge dynamics, we discuss in further detail the market clearing mechanism present in this model. Because agents in any pair assign the same cardinal value to their match, a unique stable matching always exists, i.e. the market for alliances always possesses a unique equilibrium. We prove this by construction.

Proposition 1 *The roommate matching problem $(S, \succ)$ with preferences given by Equation (9) has a unique stable matching μ.*

The formal proof is given in the appendix. It is quite simple, as is the intuition. Of all possible pairs of agents in the economy one pair produces the biggest innovation. The two agents in that pair will block any matching in which they are not together, because each prefer each other to anyone else. That pair of agents must be in the stable match. Consider the sub-population which excludes that pair. Within that sub-population there is a most innovation potential pair. They too must be joined in the stable match. This recursive argument generates a unique stable matching. As we allow for self-matching the argument naturally extends to the case of isolated agents maximizing innovative output.

It is worth noting that unlike the standard matching problems, here agents are permitted to remain as individuals, innovating autarchically. What drives this is the trade-off between

an agent's ability to substitute between different types of his own knowledge, and what he gets by joining another. In a partnership, an agent will have access to superior knowledge of type m, but then will also necessarily be forced to accept a lower value of type ℓ, because of the partner's inferior knowledge of that type. This trade-off is evaluated differently depending on the elasticity of substitution in production (β), and on the divisibility of the innovation task (θ) and thus on how heavily is the maximum weighted in knowledge pooling.

To see how this works, consider agent i as a point in knowledge space: $\mathbf{R}_+^{\mathbf{m}}$. No collaboration will take place with any j if i dominates j in every knowledge type, since combining knowledge with j will necessarily put i on a lower isoquant. Similarly, j will reject a partnership with i if j dominates. So i must search for a partner in the region $\mathbf{A}_i \subset \mathbf{R}_+^{\mathbf{m}}$ such that $v^m(i) > v^m(j)$ for some m and $v^m(i) < v^m(j)$ for others. In words this means that a necessary condition for collaboration potentially to form is that there is a complementarity between the potential partners: I must be stronger than you in some way, and you must be stronger than me in some other way. Now consider $\theta = 0$ (pooling is on the minimum). Even for potential partners in $\mathbf{A}_i$ agent i will refuse all collaborations, as he loses in the categories where $v^m(i) > v^m(j)$, and gains nowhere. Symmetrically if $\theta = 1$, i will accept any collaborator in $\mathbf{A}_i$ and will be accepted by any $j \in \mathbf{A}_i$, since both of them gain in at least one knowledge type. This suggests that more collaborations should be possible when θ is larger.

Proposition 2 *As θ increases, the possibilities of collaboration as captured by the number of acceptable partners increase for any agent $i \in S$.*

The formal proof is given in the appendix; again it is simple and intuitive. There is an indifference frontier such that i will consider collaboration with $j \in \mathbf{A}_i$ if j's endowments are above the frontier, and not if they are below the frontier. As θ increases from 0 to 1 the frontier moves lower and so more people become acceptable to i. As this effect of θ is true for anyone, i also becomes acceptable to more people.[10]

These two results will be useful in guiding our interpretation of the results of the dynamic system. First we know that in any period there is a unique set of partnerships that will emerge. Second we know that potentially, more partnerships are available to each individual as θ increases, which implies that there is a static effect whereby for larger values of θ there will be fewer isolated individuals each period. In addition there is a dynamic effect of θ. As discussed above, after two agents have jointly innovated, they add the same amount and type of knowledge to their respective stocks. That is, they move the same distance, parallel to one axis in the $\mathbf{R}_+^{\mathbf{m}}$ knowledge space. Both the isoquant map from the innovation production function, and the map of the indifference frontiers determining the area of knowledge space in which fruitful partners can be found, are families of convex curves, becoming "more parallel" to the axes as quantities increase. This implies that by moving in lock-step parallel to the axes, eventually one agent moves out of the acceptable

---

[10]Implicit in the way this is worded would be an assumption that agents are distributed uniformly over $\mathbf{R}_+^{\mathbf{m}}$. This is purely to facilitate exposition. The intuition generated through this linguistic shortcut is not misleading.

area of the other. When this happens, they search for new partners. If θ is large, new partners are relatively easy to find, and we would expect the result is that agents have relatively many distinct partners over the history of the economy.

3   Numerical experiment

The model just developed represents a complex dynamic process. As such it is impossible to track analytically, so we use numerical experiments to examine how the performance of the system responds to different parameters. We are concerned with two parameters here: the nature of the innovation task (whether divisible or not) and consequently the nature of knowledge pooling as measured by θ; and the relative importance of direct versus indirect information about potential partners as measured by α. These parameters are varied in the numerical experiments below, in order to understand their effects on network structure and knowledge growth and distribution.

We study a population of $n = 100$ firms. At the outset individual knowledge endowments are randomly drawn from a uniform distribution over [0,1], independently for every element $v^m(i)$ in each agent's knowledge vector. Each period, the market for collaborative agreements is activated and firms form pairs (or stay alone) in order to innovate. The pairing results in a stable matching where stability is defined as above (everybody is as satisfied as possible, given everyone else's preferences), and where the value of a pair is equal to the expected amount of knowledge produced by that pair. After innovation, the new knowledge is added to the firms' knowledge stocks; the firms' knowledge types change, as described previously; and so does accumulated experience. At the end of the period all pairs disband, and the process begins again in the following period. We iterate this process for 1,000 rounds, recording data for the entire history of the industry. In the numerical experiments reported below we chose the initial probability of a success to be $\pi_L = .9$, and learning gradually increases it to values close to $\pi_H = .99$. We fix the elasticity of substitution in the innovation production function at $\beta = 1/4$. Discounting is performed with a discount factor $\rho = .95$. The parameters we examine are α and θ. Regarding pooling (θ) we consider 100 randomly generated values uniformly distributed over $[0,1]$. Similarly α takes 100 uniformly distributed values in $[0,1]$. We thus have a large data sample on which we apply a simple non-parametric estimation technique — Kernel regression (Yatchew, 1998 for a comprehensive presentation) — which basically amounts to local averaging.

Regarding the properties of knowledge accumulation, we are interested in the allocation of knowledge both across individuals, and "within" individuals. Knowledge is used for innovation, so we define as our knowledge measure an agent's "innovative potential", that is, how much innovation an agent could produce on his own. In this way, rather than simply summing an agent's different knowledge types, we aggregate into a scalar which is easier to interpret. Letting $\phi = \sum_i \phi(v(i))$ designate total innovative potential, equity in allocation across individuals can be assessed by considering the coefficient of variation

$$v = \frac{\sqrt{\sum_i \phi(v(i))^2 / n - (\phi/n)^2}}{\phi/n}. \qquad (10)$$

Large values of v indicate the coexistence of rich and poor knowledge agents, while lower values of v indicate a more even distribution.

At the individual level symmetrically the specialization index s_i of individual i can be defined via the coefficient of variation in his endowments $s(i) = \sigma(i)/\bar{v}(i)$, where $\bar{v}(i)$ is the average knowledge level of i and $\sigma(i)$ the standard deviation. The larger this index is the more of a specialist and less of a generalist i is. Summing across the population produces a normalized specialization index $s = \sum_i s(i)$, for which large values indicate a population of experts, while low values indicate a population of generalists.

Regarding the network, in any period t the static network consists of q isolated agents and $(n-q)/2$ disconnected pairs, as given by the stable matching μ_t. To study the properties of the dynamic network, we record the list of connections active over time. This generates a weighted graph, in which the weight of an edge indicates how frequently the two firms have interacted in the history.

Denote (S, V_t) the graph associated with the stable matching achieved at time t, with $V_t(i,j) = 1$ if $(i,j) \in \mu_t$, and $V_t(i,j) = 0$ otherwise. The weighted graph recording past interactions is denoted (S, W_t), where $W_t(i,j)$ is the frequency of activations of the connection between i and j, obtained as $W_t(i,j) = \sum_{1 \leq s \leq t} V_s(i,j)/t$. For this graph we study the properties of the distribution of collaborative links, specifically the average path length and the clustering coefficient (cliquishness in Watts and Strogatz, 1998). To move from W_t to a 0/1 graph, distances must be computed first. The distance $d(i,j)$ between two nodes i and j is the number of edges in the highest frequency path linking them. Indeed any path $i_0, i_1, \ldots, i_z$ with $i_0 = i$ and $i_z = j$ has an associated frequency $\prod_{l=1,\ldots,z} W_t(i_{l-1}, i_l)$ and a length $z \geq 1$. Thus between two agents i and j a path with maximum frequency exists, and its length is denoted $d(i,j)$. The average path length is then

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

and simply measures how distant vertices are on average (its inverse is sometimes referred to as closeness centrality in social networks analysis). Denoting $\Gamma(i) = \{j \neq i : d(i,j) = 1\}$ the neighbourhood of vertex i and $n_i = \#\Gamma(i)$ the size of i's neighbourhood, the average degree of the graph is

$$\Delta = \frac{1}{n} \sum_i n_i.$$

The average degree Δ measures the density of the interaction structure.

Another interesting feature of network structure is the extent to which there is asymmetry in the connectivity of agents, that is, the extent to which stars exist (agents having much larger neighbourhoods than others). The measure we use is network centralization, i.e. the variance in the distribution of links over agents normalized by the variance that would obtain in a perfect star network of the same size.

The clustering (or cliquishness) coefficient is the share of active links between any given

vertex's neighbours, averaged over the system. It is written

$$c = \frac{1}{n} \sum_i \sum_{j,l \in \Gamma(i)} \frac{X(j,l)}{n_i(n_i - 1)},$$

where $X(j,l) = 1$ if $d(j,l) = 1$ and 0 otherwise.

4 Results

In the sections that follow, the results are displayed as shaded contour plots, which should be read as maps in an atlas: darker greys imply higher values on the z axis. This provides a compact display of the relationship between the relative importance of my information (α), the pooling parameter (θ), and the performance measures we are concerned with.

4.1 Network

In this section we examine how several measures of network structure respond to changes in the two parameters in the model.

4.1.1 Degree

Figure 1 displays the number of connections held by the average agent, that is, how many distinct partners an agent has, on average, over the history of the economy. There is a region below $\theta = .25$ in which autarchic innovation prevails, so the degree of the network is consistently 0. This follows from the intuitions of Proposition 2. Once this critical value of θ is crossed, degree increases both when the decomposition of tasks becomes easier (θ increases) and when the relative importance of direct (versus indirect) credit (α) decreases.

The first effect stems largely from the properties of the innovative process. In terms of Proposition 2 this can be seen by considering the conditions under which agents prefer collaboration to autarchy. Statically, as θ increases, the range of other knowledge vectors that would generate an improvement over autarchy is larger, so there are more prospective partners. This implies that in each period, there are likely to be fewer agents innovating in isolation. In addition, increasing θ increases the propensity of the stable matching to create pairs of unlike agents. But repeated interaction makes agents more alike in knowledge profiles, so with higher θ they are likely to switch partners more frequently. This dynamic effect combined with the static effect explains the effect of θ on average degree in Figure 1. The other effect is the effect of α, the importance of the private signal versus the social, constructed one. When own experience is weighted heavily ($\alpha \geq .75$), an agent's networking is focussed on a very small number of people (never more than 5). On the contrary when I take into account the information provided by my partners about their partners, indirect credit accumulates and increases the attractiveness of new people. In essence, previous partners are valuable in introducing an agent to new potential partners. This is what we observe in Figure 1 when moving right to left for high enough θ. When I value my partners' information more heavily than my own, ($\alpha \leq .5$), we observe a steady increase in
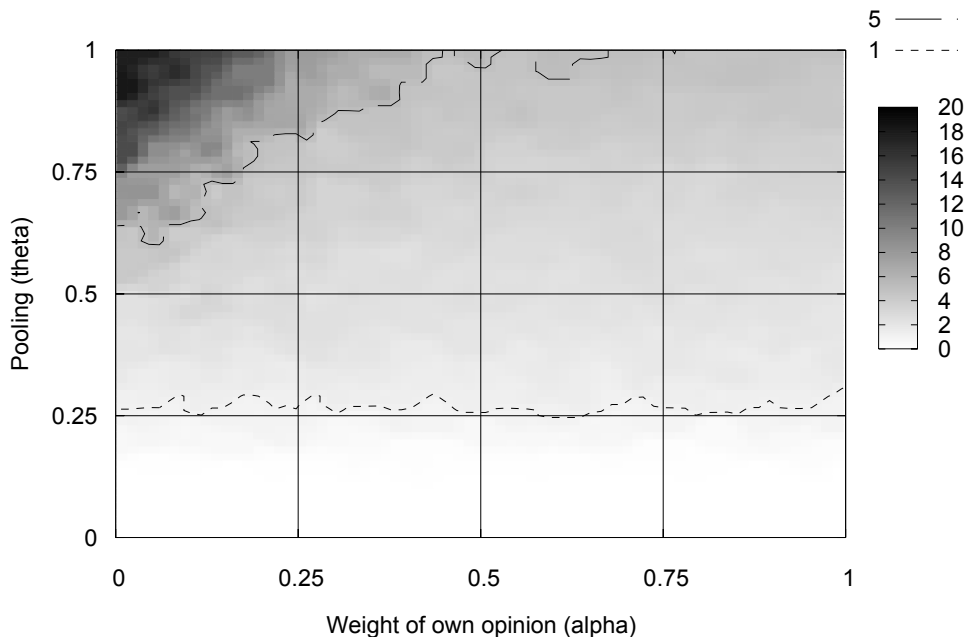
Figure 1: Degree of the emerging network in the (α, θ)-space.

the average number of partners as secondary information increases in relative value. In the extreme, when an agent disregards his own information and uses only that acquired from his previous partners, agents form alliances with fully 20 percent of the population.

4.1.2 Connectedness

Figure 2 is a binary indicator of the connectedness of the network (0 for disconnected; 1 for connected; intermediate grey levels are artefacts of the Kernel smoothing procedure). The pattern is driven by θ: when θ exceeds a critical level of about .75, a single connected component emerges. This effect works through the nature of optimal partnerships and the dynamics of knowledge as discussed in the previous paragraph. It is a consequence of rising degree, roughly independent of α.

Looking at Figures 1 and 2 together we can see that for $.25 \leq \theta \leq .75$ we have a disconnected graph with low degree (between 1 and 5) over the entire α-range. Within this zone the network consists of connected sub-components; that is, as the network evolves over time, small isolated groups form. We will see below that these groups are relatively densely connected.

4.1.3 Path Length

Figure 3 depicts average shortest path lengths (distances) between agents in the network. Because a disconnected graph has some agents who are infinitely distant from each other, the averages shown in Figure 3 are computed only for the networks that are connected, i.e.
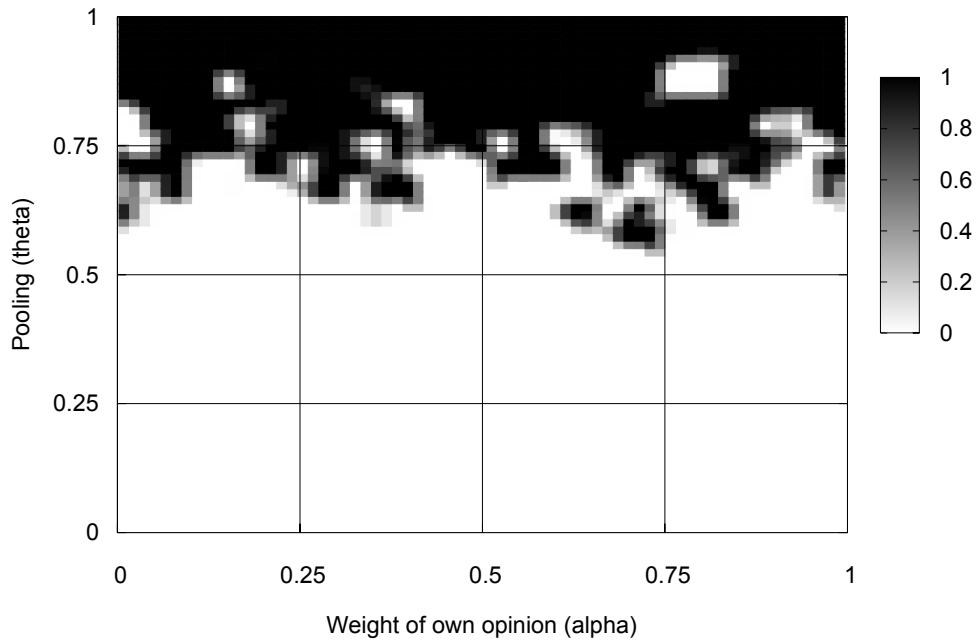
Figure 2: Connectedness of the emerging network in the (α, θ)-space.

the black parts of Figure 2. Figure 3 is the mirror image of Figure 1. Path lengths fall as density increases, as one would expect, and this is the dominating effect.
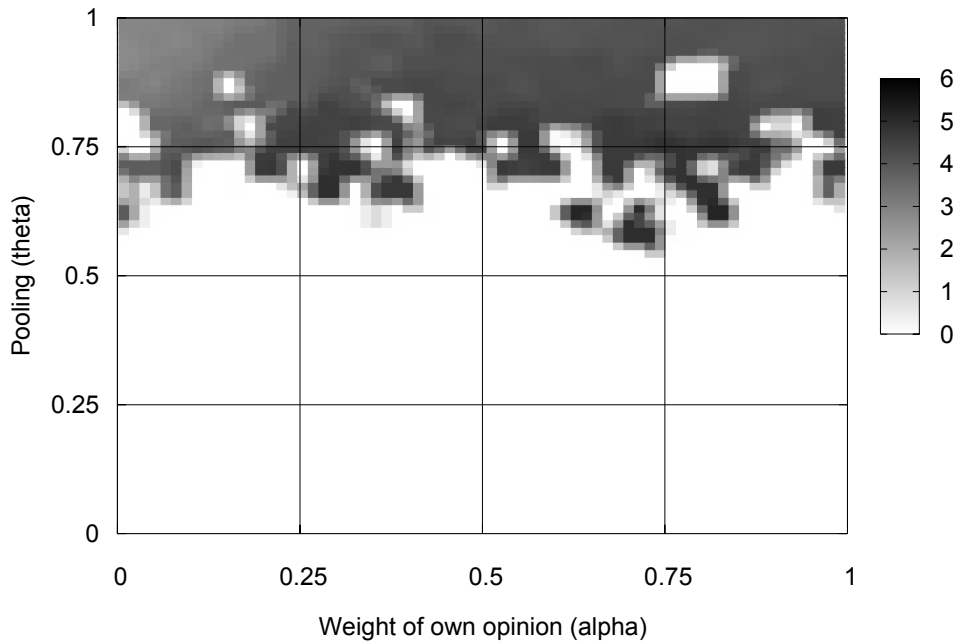

Figure 3: Average path length of the emerging network in the (α, θ)-space.

For the sake of comparison, we can consider random graphs of equivalent degree. From Figure 1, the degree of our networks fall from roughly 20 to 2 when $\theta = 1$ and from 5 to 2 when $\theta = .75$. A random graph of uniform degree Δ has average path length of $\ln n / \ln \Delta$, which implies that if our networks were random, path lengths would rise from about 1.5 to 6.6 as α increases from 0 to 1 for $\theta = 1$, but from 2.6 to 6.6 when $\theta = .75$. The path lengths in our networks are slightly longer than those of equivalent random graphs, which suggests that there is more structure than a random connection model would predict. The nature of that structure can further be characterised by looking at cliquishness.

4.1.4 Cliquishness

Figure 4 displays the average share of an agent's neighbours who are also neighbours of each other (the number of triangles divided by the number of possible triangles). Thus this coefficient measures the degree of local transitivity emerging in the network. In the white region at the bottom of the graph all firms always innovate as individuals rather than as part of a pair. There clustering is defined to be zero. Above the collaborative threshold and in the region $\alpha \leq .5$ we observe relatively high levels of cliquishness, getting higher as we move left. This time the pattern is weakly increasing with θ and markedly decreasing with α.
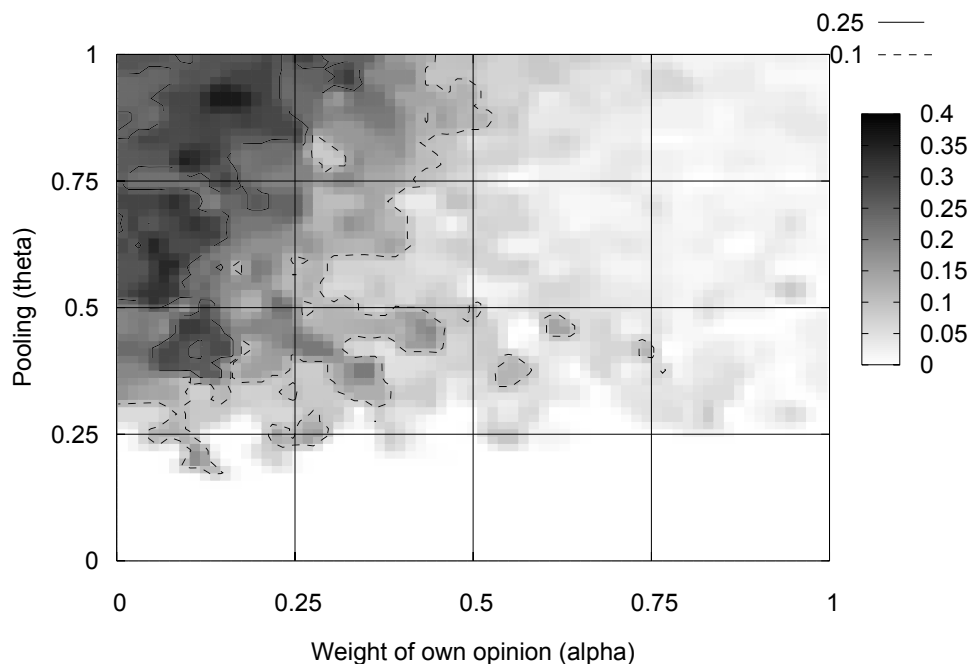

Figure 4: Cliquishness of the emerging network in the (α, θ)-space.

This measure of cliquishness can be misleading, though, since it is strongly correlated to the degree of the graph. As agents acquire more links, even if they are acquired at random, the network will become more dense locally. To get a better measure of the structure of the graph, it is necessary to compare cliquishness with that of a network of known structural

properties as we did for distance. Excess cliquishness, as shown in Figure 5, re-scales the measure from Figure 4 to make it comparable to a random graph with the same average degree (which has cliquishness of approximately Δ/n.) Figure 5 shows the ratio of observed over "predicted" cliquishness. Values significantly larger than 1 would indicate a structure richer than a random graph. There is a very clear effect of α, and also a visible effect of θ.
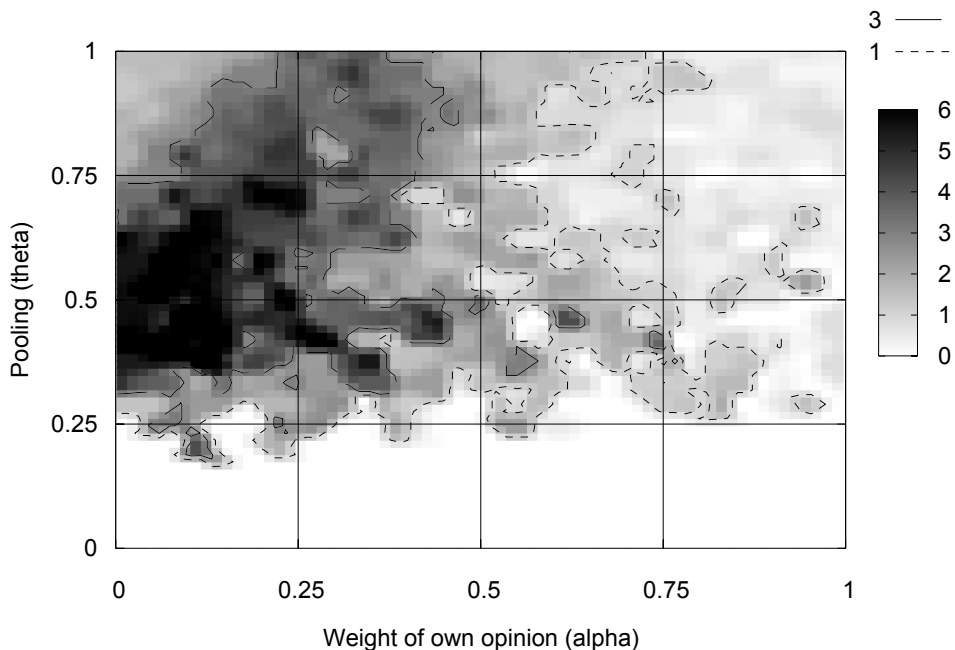

Figure 5: Excess cliquishness of the emerging network in the (α, θ)-space.

When the source of information about possible partners is mostly my past partners (say $\alpha < .5$), excess cliquishness is much bigger than 1: it is often larger than 3 and in the zone $(\alpha, \theta) \in [0, .25] \times [.25, .75]$ it is even larger than 6. This says that transitive groups are forming in zones with relatively low degree (between 1 and 10), while the graph as a whole is still not connected. Above $\theta = .75$ the graph connects and even there, as long as α is less than .5, normalized cliquishness significantly exceeds 1 and often exceeds 3. Recalling the discussion of Figure 3, in which we observed path lengths slightly larger than those of a random graph, the networks that emerge here, ($\theta > .75$, $\alpha < .5$), have the characteristic properties of small worlds: they are sparse, cliquish and have short paths. The general pattern is that as "second hand" information becomes increasingly important in evaluating potential partners, the network becomes more cliquish. The apparent exception, in the upper left corner of the space, is an artefact of the high density in that region. When agents have many partners, even a random network is highly cliquish (in the extreme case of a complete graph, cliquishness is at it's maximum of 1). Thus it will be difficult for any network to have values of cliquishness that are much higher than the equivalently dense random network (yet values larger than 1 are observed).

This suggests three regions, and three types of structures. When the innovation task is not separable, (very low θ) agents act in isolation. As separability becomes possible,

but is not extreme, the network resembles a caveman graph: we see small groups of agents who are densely connected within themselves, but isolated from the rest of the economy. When innovation consists of largely seprable tasks, small world structures begin to emerge, provided indirect sources of information remain relatively important.

What this result suggests is that the process generating small worlds depends heavily both on the cognitive properties of innovation, and also on structural embeddedness. "Second hand" information about prospective partners generates structural inertia in network formation, as an agent is more likely to pair with partners that have been recommended by previous partners. Thus our formal model generates results that are consistent with other discussions of alliance formation (Kogut et al., 1992 for instance, or Uzzi, 1996, for an empirical study that finds this effect.).

4.1.5 Centrality

Finally, we turn to network centralization: to what extent do the emergent networks contain stars? The answer is essentially that they do not. There is no pattern to network centralization over the (α, θ) parameter space, and the measures are uniformly low: in the range of 4 percent. The distribution of links over agents is relatively, though not completely, uniform. No agent is able to capture a dominant position in these networks.

4.1.6 Graphical representation of the networks

In order to illustrate more visually the emerging patterns of networking, Figure 6 shows typical networks obtained in 4 different regions of the (α, θ)-space. We consider two values of each of the parameters: $\alpha \in \{.2, .8\}$ and $\theta \in \{.4, .8\}$. This illustrates regions of the parameter space in which the emergent networks differ in terms of degree and cliquishness. Panels in the figure are arrayed correspondingly to the (α, θ)-values. Table 1 summarizes the essential features of the networks that have formed, giving the size of the giant connected component; average degree; cliquishness; excess cliquishness; and network centralization.

$\theta\backslash^\alpha$	.2	.8
	Giant comp.: 98	Giant comp.: 95
	$\Delta = 3.12$	$\Delta = 2.88$
.8	Cliq.: .094	Cliq.: .009
	Excess Cliq.: 1.54	Excess Cliq.: .19
	Net. Centr.:4.0%	Net. Centr.: 3.22%
	Giant comp.: 55	Giant comp.: 49
	$\Delta = 1.64$	$\Delta = 1.66$
.4	Cliq.: .108	Cliq.: .077
	Excess Cliq.: 1.18	Excess Cliq.: .83
	Net. Centr.: 4.49%	Net. Centr.: 3.44%

Table 1: Summary statistics.

None of the graphs shown here is connected, but we can see that the giant component is large for large θ and small for low θ. Note also that autarky is very common for low

values of θ. Many agents actually pursue their innovative activities on their own, without ever forming alliances with other firms. Excess cliquishness is higher for higher θ values, which is where we observe more structure than would be present in a random network.[11]
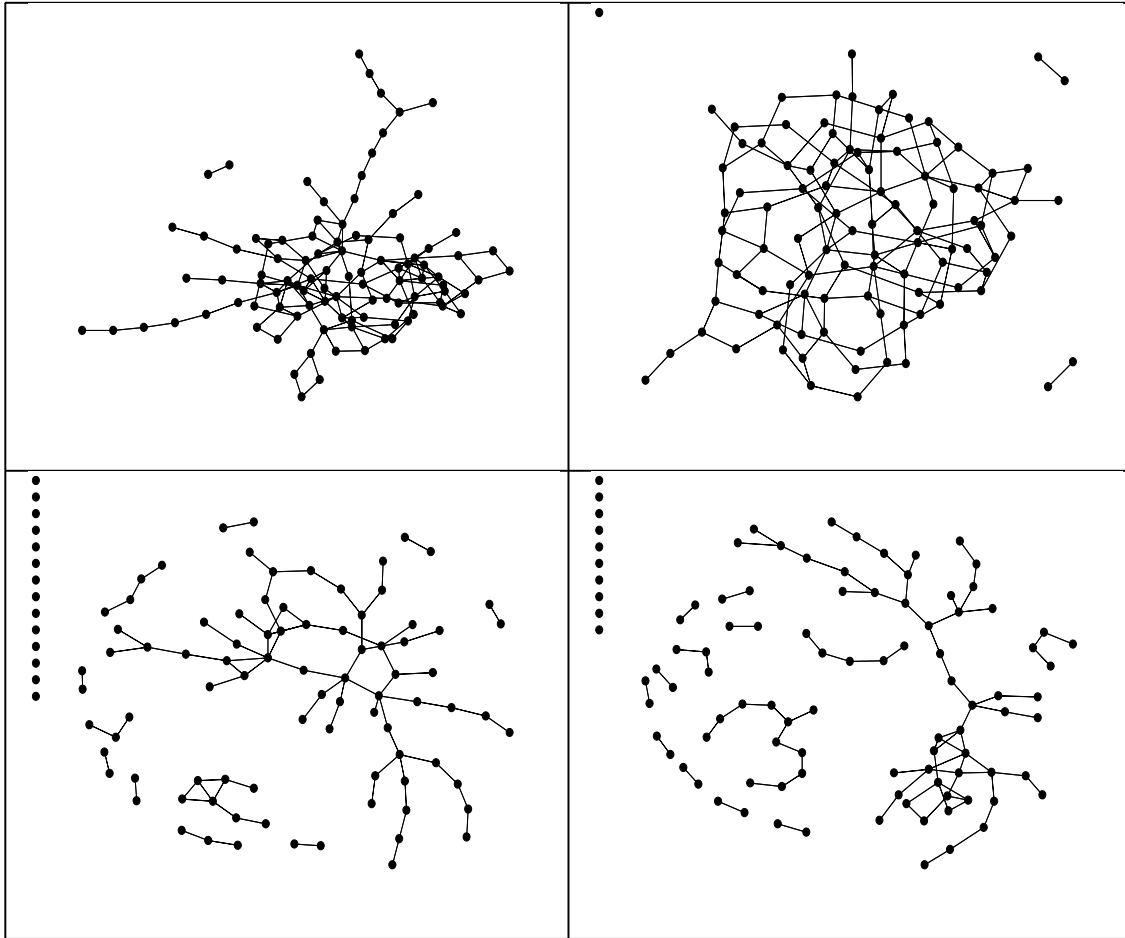

Figure 6: Four typical networks.

4.2 Knowledge

In this section we examine two measures of knowledge performance: the distribution of knowledge over individuals, and the extent to which individuals specialize in one type of knowledge.[12]

---

[11]Visually, the networks shown here exhibit structures similar to networks from different eras in the pharmaceutical biotechnology industry, as shown in Roijakkers, (2003), chapter 6.

[12]We cannot sensibly examine knowledge levels, as there is no way of normalizing the output to compare different parts of the parameter space. The pattern of knowledge levels is dominated by the effect of θ: if pooling is done using the maximum, inputs to innovation are in general larger, innovations are larger, and knowledge levels grow faster.

4.2.1 Knowledge distribution

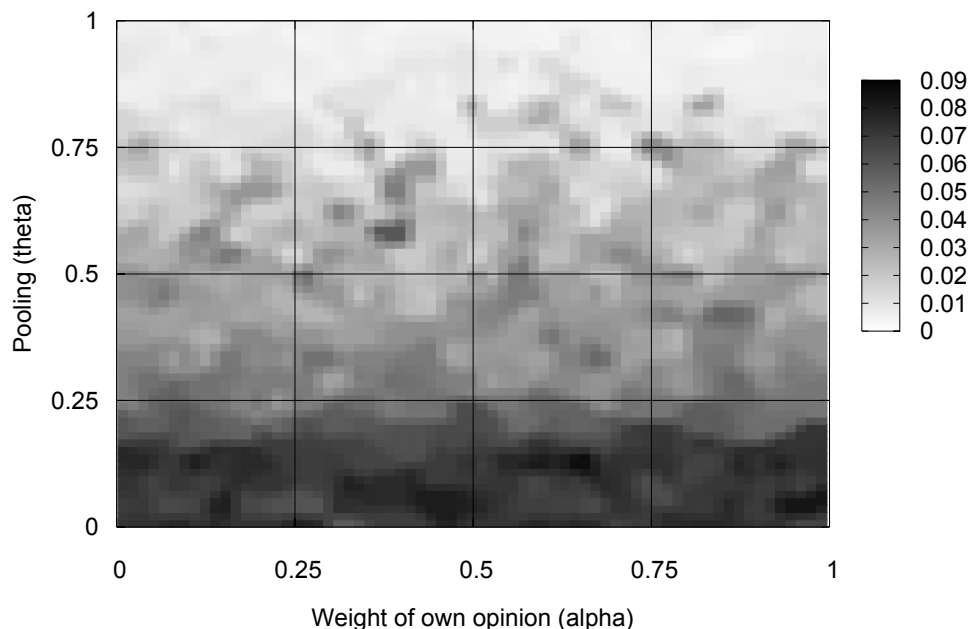Figure 7 depicts the relationship between the long run coefficient of variation of individual innovative potential and the parameters α and θ.

Figure 7: Knowledge coefficient of variation in the (α, θ)-space.

The extent to which indirect knowledge is used to evaluate potential partners (α) has little effect on this measure, but there is a clear negative relationship between the weight of the maximum in pooled knowledge (θ) and the equality in distribution of knowledge over agents. In the region of isolated innovators ($\theta < .25$), inequality in knowledge distribution is driven by initial conditions: agents who received a large knowledge endowment in the initial random assignment are able to make large innovations. Their knowledge stocks increase rapidly. Those with small initial endowments make small innovations and theirs increase slowly, thereby magnifying initial differences. The parameter α has no effect because there is no networking. As networking becomes more intense, in the sense that more agents regularly have partners, the distribution of innovative potential becomes more equal. Joint innovation implies that partners move towards each other in knowledge space, both in terms of where their expertise lies and (in relative terms) how much knowledge they possess. The more intense is the networking activity, the stronger is this effect. It is compounded by the fact that when θ is large, the stable matching mechanism favours pairs whose partners have highly complementary knowledge profiles. When this happens, the knowledge accumulation process we have modelled has the effect of decreasing the relative gaps between agents relatively rapidly, which in turn implies a rapid decrease in the gap between their innovative potentials. This convergence is exactly what the coefficient of variation captures, as the average relative distance to the average individual.

4.2.2 Specialization

As agents innovate, their knowledge profiles change, and they can become generalists or specialists. In Figure 8 we display the average individual specialization as measured by the individual coefficient of variation on knowledge endowments.
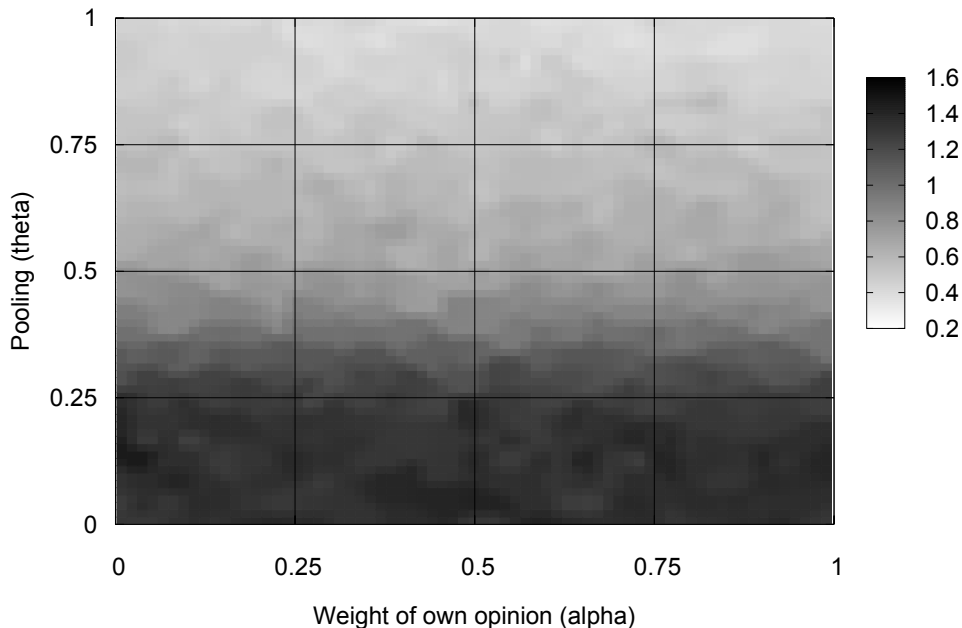

Figure 8: Expertise (specialization) in the (α, θ)-space.

The idea here is that an agent's knowledge is spread over the 5 categories, and if he has significantly more in one category than in the others, he can be called a specialist, whereas if the agent has roughly equal amounts in each of the categories, he would be a generalist. The Herfindahl index provides a natural measure for this specialization. The individual index has an analytical lower bound of 0 when all components of the endowment are identical (the agent is a generalist), and an upper bound of 2 when all the knowledge held by an agent is in one category (the agent is a pure expert). The effect of θ is clear, with the degree of specialization falling as pooling moves towards the maximum of the partners' endowments.[13] This figure shows essentially a mirror image of the pattern in Figure 1, the density of the network. Agents become highly specialized when they innovate largely as individuals. What drives this is that the type of knowledge produced is probabilistically the same as the knowledge input. Thus an agent is likely to innovate where he has most

---

[13]Large θ also implies that the innovation task is separable, so here specialization could be efficient. Extreme specialization would not be efficient in the context of the model, however, since there are 5 types of input to the production function, and partnerships are restricted to two members. Generalizing the model to permit larger coalitions would make specialization efficient in this region. It is difficult to sustain, though, because of the nature of knowledge accumulation — all members of a partnership learn the same thing, which immediately dilutes the specialization for most of them.

knowledge. In expected value, this process will lead an agent to innovate always in the same knowledge type, and so drive extreme specialization. When alliances form in a more systematic way (larger θ) there is a rapid emergence of generalist profiles. Sometimes an agent will innovate in his speciality, sometimes in his partner's, and as θ increases, an agent will have many different partners. This will both smooth the agent's profile immediately, and possibly even shift his area of expertise. This sort of mixing produces much flatter profiles, and the more partners an agent has, the more this mixing will take place.

5   Conclusion

In this paper we have developed a stylized but rich model of the innovation process in which the underlying principle of innovation is the recombination of existing knowledge. Heterogeneous knowledge actors form partnerships which are embedded cognitively, structurally and relationally, and this embeddedness directly influences the choice of partner by every firm. Over time this bilateral link formation process results in an emergent network structure. We have analyzed the properties of this network and the corresponding patterns of knowledge accumulation, and studied how they respond to changes in the cognitive division of labour performed by the knowledge agents (the pooling parameter θ), and changes in the relative importance of the sources of information accessed by firms in the network (α).

First we find that there is a critical value of the knowledge pooling parameter θ where the process tips from a world of firms operating in isolation when the innovation process is not separable, to a collaborative world with intensive networking as innovation can be more effectively sub-divided into independent sub-tasks. Connectedness emerges for larger values of the pooling parameter. As we move along the pooling dimension, the tendency of the network is thus to evolve from a sea of isolated agents to an archipelago of pairs and small subgroups, and finally to one big continent. This tendency is stronger when the weight of "second hand" information is larger, i.e. α is small, and agents rely on past partners for information regarding the value of potential future partners. This reliance generates structural inertia in network formation. The emergent networks retain a relatively low degree regardless of the parameters values: the average number of partnership reaches 20% of the population in a small part of the parameter space, but otherwise the networks stay fairly sparse. When we look for structure by benchmarking the networks against the "equivalent" random ones, we find that average distances between agents are only slightly larger than those of a random graph, but that the measure of transitivity (the extent to which "my friends are friends of each other") significantly exceeds that of a random graph in the zone in which my partners' opinions are taken seriously ($\alpha \leq .5$) and where a cognitive division of labour is feasible ($\theta \geq .4$). This suggests very clearly a network with more structure than a random one, and more specifically a small world type of structure.

Regarding the properties of knowledge accumulation, the most marked result is the decline of both inter-individual aggregate differences, and individual specialization as knowledge pooling becomes easier (θ increases). It may seem at odds with the general intuition that as labour division gets easier specialization should increase, but is easily understood

when considering the way new knowledge is allocated to individuals: the field in which novelty is created is common to both partners, and this precisely prevents specialization.[14]

Discussions of networks that arise from alliance formation, and empirical investigations of real networks emphasize the importance of embedded interactions. The formal model presented here shows a process by which the embeddedness of interactions translates into different network structures. Small worlds are commonly found in alliance networks, but they are not the only structure present in empirical results. The model we have developed can generate different network structures, depending on paramaters, and in particular shows when small worlds can be expected, depending not only on the nature of the embedded relations, but also on the nature of the innovation task itself. Our results underline the observation that understanding embedded relations is central in understanding the network stsructures that we observe, but they also show the importance of understanding how firms actually combine their knowledge to create innovations. Our model of this process provides a way to address the issue of knowledge combination as an important aspect of innovation.

6 Appendix

In this appendix we provide the proofs for the analytical results on a single round of matching.

Proposition 1 *The roommate matching problem* $(S, \succ)$ *with preferences given by Equation (9) has a unique stable matching* μ.
Proof: The algorithm to construct the stable matching is as follows. Let $S_0 = S$ and $\mu_0 = \{\varnothing\}$. There exists a pair (a_1, b_1) such that $F(a_1 b_1) = \max_{(i,j) \in S^2} F(i, j)$. Then (a_1, b_1) must belong to any stable assignment, as b_1 is preferred by a_1 to any other partner, and this preference is reciprocal. No matching which does not involve this pair could be stable. Hence the pair (a_1, b_1) is necessarily part of a stable matching.Let then $\mu_1 = \mu_0 + \{(a_1, b_1)\}$ and $S_1 = S_0 - \{a_1, b_1\}$. It is possible that $a_1 = b_1$, in which case $\mu_1 = \mu_0 + \{(a_1, a_1)\}$ and $S_1 = S_0 - \{a_1\}$. Again there is a single pair within S_1 maximizing the innovation function. Proceed recursively this way. Denote p the smallest integer such that $S_p = \{\varnothing\}$. Then $\mu_p \equiv \mu$ is a stable matching and $q = 2p - n$ is the number of agents that have preferred self-matching to cooperation. □

In case of a tie (that is to say when individual i can achieve the same innovative output with two or more different partners), we apply an arbitrary rule to guarantee that the score function is still strict: if $F(i, j) = F(i, k)$ then $j \succ_i k$ if and only if $j > k$.

Proposition 2 *As* θ *increases, the possibilities of collaboration as captured by the number of acceptable partners increase for any agent* $i \in S$.
Proof: We provide the derivation in 2 dimensions. It extends to any $\mathbf{m} > 2$ but the calculations are more tedious. Take $x = (x_1, x_2)$ as a point in $\mathbf{R}_+^2$. First we can restrict

---

[14]In Cowan et al. (2003) a situation in which people were innovating most often in their domain of expertise was explored, and there of course as division got easier specialization increased. The contrast between these two results indicates the importance of understanding exactly how new knowledge is absorbed by economic actors.

the set of x's potential collaborations to $\mathbf{A}_x = \{y \in \mathbf{R}_+^2 : \#\{m : x_m < y_m\} = 1\}$. Consider $y = (y_1, y_2) \in \mathbf{A}_x$, with $x_1 > y_1$ and $x_2 < y_2$. Agent x is then indifferent between autarky and collaboration with y iff

$$(\theta x_1 + (1 - \theta) y_1)^\beta + (\theta y_2 + (1 - \theta) x_2)^\beta = x_1^\beta + x_2^\beta.$$

This equality implictly defines a curve in $\mathbf{R}_+^2$, the indifference frontier. The frontier goes through x, since any agent is indifferent between autarchy and partnering with an indetical agent. Any y lying above this curve is preferred by x to isolation. Using the implicit function theorem it is straightforwardly seen that

$$\frac{\partial y_1}{\partial \theta} = -\frac{(\theta x_1 + (1 - \theta) y_1)^\beta (\theta y_2 + (1 - \theta) x_2) (x_1 - y_1)}{(\theta x_1 + (1 - \theta) y_1)^\beta (\theta y_2 + (1 - \theta) x_2) (1 - \theta)} \\ -\frac{(\theta y_2 + (1 - \theta) x_2)^\beta (\theta x_1 + (1 - \theta) y_1) (y_2 - x_2)}{(\theta x_1 + (1 - \theta) y_1)^\beta (\theta y_2 + (1 - \theta) x_2) (1 - \theta)},$$

which is negative. The same calculation shows that also $\partial y_2 / \partial \theta < 0$. Thus as θ increases from 0 to 1 the frontier gets lower everywhere and thus more people become acceptable to i. $\square$

References

[1] Alstyne, M. V. and E. Brynjolfsson (1996), "Widening Access and Narrowing Focus: Could the Internet Balkanize Science?", *Science*, 271, 1479-1480.

[2] Alstyne, M. V. and E. Brynjolfsson (1997), "Electronic Communities: Global Village or Cyberbalkans?" Working Paper MIT Sloan School.

[3] Barabasi, A. and R. Albert (2000), "Statistical Mechanics of Complex Networks", *Review of Modern Physics*, 74, 47-97.

[4] Burt, R. S. (1992), *Structural Holes: The social structure of competition,* Harvard University Press, Cambridge, Mass.

[5] Chung, S., S. Singh and K. Lee (2000), "Complementarity, status similarity and social capital as drivers of alliance formation", *Strategic Management Journal*, 21, 1-22.

[6] Cowan R. and Jonard N. (2004), "Network Structure and the Diffusion of Knowledge", *Journal of Economic Dynamics and Control*, 28, 1557-1575.

[7] Cowan, R., N. Jonard and J-B. Zimmermann (2003), "La Dynamique Conjointe de la Connaissance et des Réseaux", *Revue d'Economie Industrielle*, 103, 253-274.

[8] Davis, G. F., M. Yoo, M. and W. E. Baker (2003), "The Small World of the American Corporate Elite, 1982-2001", *Strategic Organization*, 1, 301-326.

[9] Dodgson M. (1996), "Learning, trust and inter-firm technological linkages: some theoretical associations", in Coombs R., Richards A., Saviotti P.P. and Walsh V. (Eds.), *Technological Collaboration, the Dynamics of Cooperation in Industrial Innovation*, Edward Elgar

[10] Doz, Y.L. (1988), "Technology partnerships between larger and smaller firms: Some critical issues", in F.J. Contractor and P. Lorange (Eds.), *Cooperative Strategies in International Business*, Lexington Books, Lexington, Mass.

[11] Duysters, G. and B. Verspagen (forthcoming), "The Small worlds of Strategic Technology Alliances", *Technovation.*

[12] Dyer, J.H. and K. Nobeoka (2000), "Creating and Managing a High Performance Knowledge Sharing Network: The Toyota Case", *Strategic Management Journal*, 21, 345-367.

[13] Gale D. and Shapley L. (1962), "College admissions and the stability of marriage", *American Mathematical Monthly*, 69, 9-15.

[14] Galison, P. (1999), "Trading zone: Coordinating action and belief", in M. Biagioli (Ed.), *The science studies reader*, 137-160, New York: Routledge.

[15] Garcia-Pont, C. and N. Nohria (2002), "Local versus Global Mimetism: The dynamics of alliance formation in the automobile industry", *Strategic Management Journal*, 23, 307-21.

[16] Grandstand, O. (1996), "Towards a Theory of the Technology-Based Firm", *Research Policy*, 27, 465-489.

[17] Grandstand, O. and S. Sjolander (1990), "Managing Innovation in Multi-Technology Corporations", *Research Policy*, 19, 35-60.

[18] Grant, R.M. (1996), "Towards a Knowledge-Based Theory of the Firm", *Strategic Management Journal*, 17, 109-122.

[19] Gulati, R. (1995), "Social structure and alliance formation patterns: A longitudinal analysis", *Administrative Science Quarterly*, 40, 619-652.

[20] Hagedoorn J. (2001), "Inter-firm R&D partnerships. An overview of major trends and patterns since 1960", in J.E. Jankowski, A.N. Link and N.S. Vonortas (Eds.), *Strategic research partnerships: Proceedings from a NSF workshop*, Washington DC, National Science Foundation.

[21] Hamel, G. Y. Doz, and C.K. Prahalad (1989), "Collaborate with your competitors and win", *Harvard Business Review*, 67, 137-139.

[22] Jackson, M.O. and B. Dutta (2003), "On the Formation of Networks and Groups", in Dutta, B. and M. Jackson (Eds.), *Networks and Groups: Models of Strategic Formation*, Springer Verlag, Heidelberg.

[23] Johnson, D.K.N. and M. Mareva (2002), "It's a Small(er) World: The role of geography and networks in biotechnology innovation", Wellesley College Working Paper 2002-01.

[24] Kirman, A. and N. Vriend, (2001), "Evolving Market Structure: An ACE model of price dispersion and loyalty", *Journal of Economic Dynamics and Control, 25*, 459-502.

[25] Kogut, B. W. Shan and G. Walker (1992), "Competitive cooperation in biotechnology: Learning thought networks?" in N. Nohria and R. Eccles (Eds.) *Networks and Organizations: Structure, Form and Action*, Harvard Business School Press, Boston.

[26] Mowery, D.C., Oxley, J.E. and B.S. Silverman (1998), "Technological Overlap and Interfirm Cooperation: Implications for the Resource-based View of the Firm", *Research Policy*, 27, 507-523.

[27] Narula, R. (1999), "Explaining Strategic R&D alliances by European firms", *Journal of Common Market Studies*, 37, 711-723.

[28] Newman, M. E. J. (2001), "The structure of scientific collaboration networks", *Proc. Natl. Acad. Sci. USA,* 98, 404-409.

[29] Nooteboom B. (1999) *Inter-firm Alliances. Analysis and Design*, Routledge, London.

[30] Nooteboom B. (2000), *Learning and Innovation in Organizations and Economics*, Oxford University Press.

[31] Oliver, C. (1990), "Determinants of inter-organizational relationships: integration and future direction", *Academy of Management Review*, 15, 241-265.

[32] Penrose, E. (1959), *The Theory of the Growth of the Firm*, Basil Blackwell, Cambridge, Mass.

[33] Peretto, P. and S. Smulders (2002), "Technological Distance, Growth and Scale Effects", *Economic Journal*, 112, 603-624.

[34] Podolny, J.M. (1993), "A status-based model of market competition", *American Journal of Sociology*, 98, 829-872.

[35] Powell, W.W. (1990), "Neither markets nor hierarchies: Network forms of organization" in B.M.Shaw and L.L. Cummings (Eds.) *Research in Organizational Behavior,* 12, 395-336.

[36] Powell, W., Koput, K., W., Smith-Doerr, L. (1996), "Inter-organizational collaboration and the locus of innovation: Networks of learning in biotechnology", *Administrative Science Quarterly*, 41, 116-145.

[37] Riccaboni, M. and F. Pammolli (2002), "On Firm Growth in Networks", *Research Policy*, 31, 1405-1416.

[38] Richardson, G. B. (1972), "The organization of industry", *Economic Journal*, 82, 883-897.

[39] Roijakkers, N. (2003), *Inter-firm cooperation in high-tech industries: a study of R&D partnerships in pharmaceutical biotechnology*, PhD thesis Maastricht University, Maastricht, Datawyse/Universitaire Pers Maastricht.

[40] Sako M. (1991), "The role of trust in Japanese buyer-supplier relationships", *Ricerche Economiche*, 45, 449-474.

[41] Teece, D. (1986), "Profiting from technological innovation: Implications for integration, collaboration, licensing and public policy", *Research Policy*, 15, 285-305.

[42] Teece, D. and G. Pisano (1989), "Collaborative Arrangements and Global Technology Strategy: Some Evidence from the Telecommunications Equipment Industry" in R.A. Burgelman and R.S. Rosenbloom (Eds.), *Research on Technological Innovation, Management and Policy*, 4, 227-256, Greenwich, CT: JAI Press.

[43] Uzzi, B. (1996), "The sources and consequences of embeddedness for the economic performance of organizations: The network effect", *American Sociological Review,* 96, 674-698.

[44] Uzzi, B. (1997), "Social structure and competition in interfirm networks: The paradox of embeddedness", *Administrative Science Quarterly*, 42, 35-67.

[45] Yatchew, A. (1998), "Nonparametric Regression Techniques in Economics", *Journal of Economic Literature*, 36, 669-721.

[46] Watts D. and Strogatz S. (1998), "Collective Dynamics of Small-World Networks", *Letters to Nature*, 393, 440-442.