Knowledge, Innovation and Economic Geography

Richard Taylor
Centre for Policy Modelling
Manchester Metropolitan University
Manchester M1 3GH, UK
{richard@cfpm.org}

Piergiuseppe Morone
University of Napoli "L'Orientale"
Largo S. Giovanni Maggiore, 30
80134 – Napoli, Italy
{p.morone@gmail.com}

Conference paper submitted to:

Knowledge and Regional Economic Development Conference, June 9-11, 2005, Barcelona
Facultat de Ciències Econòmiques i Empresarials,
Universitat de Barcelona,
Av. Diagonal, 690
08034 Barcelona, Spain

First Draft (13-03-05)

Abstract

A central theme of the Knowledge and Regional Economic Development Conference is that of the geographical dimensions of knowledge and innovation. In this paper we are going to investigate the importance of geographical proximity for partnership formation in innovation systems, in which there is an assumed tacit aspect to knowledge flows. Once the theoretical framework has been defined, we shall present our rationale for an *agent-based modelling* approach to understand these phenomena, and then describe the model we are developing. We report on an initial series of simulation experiments we have conducted with this model, showing that when the geographical bound on the ability to use knowledge in joint innovations is relaxed, the number of partnerships increases approximately twofold. However, geographical aspects of the simulation did not change the surprising result that firms showed a large discrepancy in the number of innovations they attained.

Keywords

Agent-based, modelling, innovation networks

1. Introduction

Over the last decade, the idea that modern economies are rapidly evolving into knowledge-based or learning economies has been attributed a growing importance (OECD, 1996). In this *knowledge-based approach*, the ability to create and transfer knowledge is the crucial component in sustaining competitive advantage through innovation and other value generating activities (Pinch et al., 2003; Forsman and Solitander, 2003). A common idea is that firm's long-term competitiveness crucially depends on its ability to innovate and learn continuously (Florida, 1995; Cooke, 2001; Malmberge and Maskell, 2002). In other words, firms which handle knowledge most efficiently are also at the forefront of innovation ability and hence are *dynamically competitive* at a global scale.

The knowledge-based approach brings into view many questions about the complex interaction between individuals' capacity for learning and the necessity for firms to learn, i.e. to adapt or innovate in a changing market environment. For example: is the knowledge resource of firms embodied only in the individuals and in the relationships amongst them, or can it actually be stored in organisational routines and practices? Is the acquisition of knowledge largely controlled from the top of the management hierarchy (for example, through formal training or through prescribed programmes of research and development) or is it constructed in a more ad-hoc and problem orientated way? To what extent is knowledge transferable as compared to being context and firm-specific? Does it result largely from individuals' capacity to reflect upon their own experiences or does it depend more upon the use of interpersonal networks (think of information sharing in communities of practice)?

What most people agree upon is that individuals learn, whilst firms adapt their business practices, and that there is some kind of correspondence between the two. This leads to the concept of organisational learning (OL), where firms can be said to develop new sets of skills or competencies, as well as constructing new values, goals, and systems. In this paper we aim at developing a theoretical model able to capture some of the main features that govern organisational learning, product innovation, and the geographical distribution and clustering of firms.

We will pursue our target developing an agent-based model of interactive learning on geographical networks. Related literature has provided insight into some of the factors affecting knowledge diffusion patterns (for example Cowan and Jonard 2003 and 2004 on the effect of different network architectures, Morone and Taylor 2004a and 2004b on the importance of initial inequalities of knowledge and the geographical distribution on the network). However these models have come under criticism of being too vague in the definition of exactly what type of knowledge is being considered, and it has been questioned whether or not the current crude theoretical constructions are appropriate for considering the multifaceted phenomena of knowledge. Ultimately the problem is grounded in longstanding philosophical questions about the nature of knowledge, and debates that are unlikely to be resolved. Our contribution to this debate is to focus upon the core skills or abilities of firms as an important component of the KBV of the firm and hence, their ability to innovate. This type of knowledge, we argue, is structured in a complex way, is often firm-specific, and has geographical constraints to diffusion. We then go on to investigate how, within this framework, knowledge could be articulated and processed to generate product innovation, and how this whole process of acquiring knowledge and using it to innovate relates to geographical allocation of firms.

In section two we will first investigate the nature of knowledge, aiming at defining a possible 'taxonomy of knowledge'. Hence, we will clarify the logical connections linking different kinds of knowledge to different kinds of learning. This will pave the way to introducing the spatial dimension in the analysis in section three. Geographical clustering will then be considered as the resultant of a specific learning process directly associated to physical proximity. Once the theoretical framework has been defined, we shall present our rationale for an *agent-based modelling* approach to understand these phenomena, and then describe the model we are developing in section four. Section five reports on an initial series of simulation experiments we have conducted with this model and section six draws some preliminary conclusions.

2. Defining knowledge and different learning processes

"Knowledge can be defined as a dynamic framework or structure from which information can be sorted, processed and understood. […] Knowledge is therefore associated with a process that involves cognitive structures which can assimilate information and put it into a wired context, allowing actions to be undertaken from it" (Howells, 2002: 872). In this definition of knowledge all the relevant concepts are put foreword: knowledge is defined as a complex and dynamic structure within which information is articulated with other information through a cognitive process. This process, which we can label learning, allows actors to undertake actions which require the use of the acquired knowledge.

If the unit of our analysis is the firm, then a wide interpretation of the role of the firm is to add value to a production chain. However, modern manufacturing industry in the developed countries may be less involved in the physical processes of production than in the design and marketing phases, and more relevantly, in the *innovation of new products and production processes*. In this context, innovation can be defined as the process by which firms master and put into practice new product designs and manufacturing processes (Nelson and Rosenberg, 1993). Hence innovation is itself a process in which "new knowledge or new combination of old knowledge are embodied in products or production processes and possibly introduced into the economy. Put in a simple way, *innovation is the result of learning processes*. Learning leads to new knowledge and firms use this knowledge in an attempt to improve product and production processes" (Oerlemans, Meeus and Boekeman, 1998: 4).

As suggested by Howells, innovation involves using existing knowledge, as well as generating and acquiring new knowledge and this centrally involves learning (2002: 872). By concentrating our analysis on the concept of learning, we can define at least two broad ways in which learning can occur: *internally* and/or *externally* to the firm. "Learning to use internal resources can be accomplished in several different ways, for example through R&D activities or learning by using or doing. The external mobilisation of resources can be labelled 'learning by interacting' (Lundvall, 1988: 362), i.e. firms can use the knowledge and experience of other economic actors" (Oerlemans, Meeus and Boekeman, 1998: 3-4).

This external 'learning by interacting' has also been analysed in the social science literature where it is known as social learning. Conte and Paolucci (2001) develop a theory of social learning where they emphasize the role played by the agents' mental processes. They use cognitive explanations to distinguish between two types of social learning which they call facilitation and imitation. Facilitation is defined as "a type of social learning in which the learning agent (O) updates her knowledge base by perceiving the relationship between another agent (S) and its physical or social environment." (Conte and Paolucci: sect. 5.8) This could be an elementary type of learning in which O does not necessarily attribute to others any goals or mental states, but rather observes them, their features or behaviours.

Imitation, on the other hand, describes the situation where O intends to know what S does or looks like in order to find out norms or rules, or to adopt S's goals and/or other mental states, as long as O believes that S is an appropriate model (M) in the given context. Both kinds of social learning involve agents observing and interacting in a common environment. In our understanding, it appears that imitative learning, in the sense of Conte and Paolucci (2001), is the most likely candidate for external OL of firms.[1]

Both of the above definitions consider interaction with an external source (S) of learning and imply that firms are operating in a relational network, exchanging knowledge through direct contacts. Yet, this does not imply anything about the geographical proximity of firms.

In fact, such interactions could be mediated through information and communication technologies, hence reducing the importance of physical proximity. We could maintain that relational networks "consist of relationships connecting actors […] that are cooperating in order to acquire resources that they may not themselves possess" (Forsman and Solitander, 2003: 5). This definition takes us to the so called *economic network* approach, as defined by Håkansson (1987 and 1989) and Håkansson and Snehota (1995), which provides an interpretative framework to understand the nexus linking learning, innovation and networks (Oerlemans, Meeus and Boekeman, 1998: 3). "Håkansson's economic network model contains three main elements: actors, activities

[1] It is assumed here that firms acquisition of new skills may derive from similar processes underlying the analysis of individual learning, and that these can, presumably, also be described by means of cognitive explanations.

and resources. Actors perform activities and control or posses resources" (Oerlemans, Meeusa and Boekeman, 1998: 3). Hence, firms operating in an economic network would be the actors performing innovating activities through the acquisition and control of knowledge resources.

Along this line of research Pyka (2002) pointed out how the concept of the innovating network develops in contrast with the incentive-based neoclassical approach. Following a knowledge-based approach leads us to consider, within an evolutionary perspective, networks of firms as a central determinant in the creation of industrial novelty. From this point of view, knowledge is no longer considered as a freely available public good, but as local (technology specific), tacit (firm specific) and complex (based on a variety of scientific fields and technologies). These characteristics of knowledge are also the driving forces for the creation of innovation networks. In an earlier work Pyka (1997) showed how in regimes where technology is of major importance, the impossibility of mastering the whole required knowledge can lead to the emergence of large informal networks via self-organisation.

This approach leads to a reconceptualisation of knowledge which should no longer be considered as a homogeneous resource that can be acquired individually or interactively. This hypothesis, in fact, has been challenged by several authors who, building on the idea that knowledge is an articulated multidimensional process, have made various attempts to identify and classify different types of knowledge. When Michael Polanyi, in his seminal work (1966), first proposed a distinction between tacit and codified knowledge,[2] he observed how "we can know more than we can tell" (1966: 4), suggesting that some knowledge is intrinsically subject to some degrees of tacitness and hence can not be easily codified. Using Howells words we could hold that "explicit or codified knowledge involves know-how that is transmitted in formal systematic language and does not require direct experience of the knowledge that is being acquired and it can be transferred in such formats as blueprint or operating manual. By contrast, tacit knowledge cannot be communicated in any direct or codified way. Tacit knowledge concerns direct experience that is not codifiable via artefacts. Thus, it represents disembodied know-how that is acquired via the informal take-up of learned behaviour or procedures" (Howells, 2002: 872).

3. Dimensions of economic geography

Departing from this definition of tacit and codified knowledge we can add a geographical dimension to the analysis developed so far. Given the nature of tacit knowledge, several authors observed the importance of *trust* among the participants if genuine learning is to occur (Collins, 2001; Nonaka and Takeuchi, 1995). Conte and Paolucci argue that imitative learning also involves

[2] It is worth mentioning here that Polanyi did not stress the dichotomous nature of knowledge, pointing out, on the contrary, that the tacit/codified distinction should be seen as a continuum between wholly explicit and wholly tacit knowledge (on this point see among others: Nelson and Winter, 1982; Morgan, 2001; Howells, 2002).

the cognitive task of deciding which is a 'good' model (M) under the given circumstances, and that this decision is a reflection upon the source: "imitation implies delegation and ultimately trust: O implicitly delegates others to do (part of) the job she should do. She must trust M to some extent." (Conte and Paolucci: sect. 5.14)

Other authors emphasize that in interactions amongst firms, trust has a two-way nature that can be considered a relational asset. It is more likely to develop where the participants are engaged in several encounters, meaning that *the shadow of the future* looms larger over the present (Axelrod, 1984; Morgan, 2001). "This provides a context for reciprocity: a good deal for informal know-how trading takes place, even among rival firms, precisely because of the expectation of the information which A provides B today will be reciprocated in kind tomorrow" (Morgan, 2001: 10). This idea, first introduced by von Hippel (1987), leads to the conclusion that mutual trust and reciprocity are easier to subsist in the context of geographical proximity (Malberg, 1997).

In this perspective, tacit knowledge is considered to be 'context dependent', "being facilitated by a common language, cultural and value system. Codifiable knowledge, in contrast, can be expressed in various forms, and rapidly disseminated through various geographically dispersed user communities" (Pinch et al., 2003: 375). Such process has also been labelled *ubiquitification* of knowledge (Maskell, 1999; Maskel and Malberg, 1999).

This re-definition and re-contextualisation of tacit knowledge paves the way to a new specification of proximity. If the diffusion of tacit knowledge is facilitated by a 'shared language' (Burns and Stalker, 1961) or, as put by Dosi and Marengo, by a 'shared cognitive framework' (1994), then there is room to believe that tacit knowledge dissemination is subject to *organisational* or *relational proximity* more than to *physical* or *geographical proximity*.

However, the view which juxtaposes relational proximity on the one hand with geographical proximity on the other, has come under criticism of being "a form of *spatial fetishism* which […] lies in the assumption that there is something called *geographical proximity* which does not involve *relational proximity*, implying that the social interactions which constitute *local* actions are somehow natural, primordial or automatic, when in fact they have to be actively constructed like any other relational asset, whatever the spatial scale" (Morgan, 2001: 14). In other words, relational proximity, far from being a substitute for geographical proximity, should be considered as a complementary and reinforcing element of it.

To sum up, tacit knowledge could be defined as "person-embodied, context-dependent, spatially sticky and socially accessible only through direct physical interactions" (Morgan, 2001: 15). These are the special features of tacit knowledge which help to explain the apparent paradox of phenomena like the "economically successful industrial clusters in an age in which new telecommunication systems facilitate the transfer of ever more complex sets of knowledge at an ever-increasing rate" (Pinch et al., 2003: 375), or the spatial concentration of R&D activities in the home base of the innovating firms – defined by Kith Pavit and Pari Patel as *an important case of non-globalisation* (1991).

A final aspect related to economic geography of innovation leads us to the widely discussed rural/urban dichotomy. As argued through sections 2 and 3 firms' innovation is tightly bounded to firms' ability to acquire knowledge and, therefore, it is intrinsically a spatial-dynamic process. The literature revised so far, while providing strong argument to explain clustering, did not tackle the issue of where firms do cluster. As widely acknowledged the initial locations of new technological systems present some degrees of arbitrariness (Krugman, 1991). However, firms tend to agglomerate near innovating firms, mainly because of the need for skilled labour and information (Kangasharju and Nijkamp, 1997). In the line of Davelaar (1991), Kangasharju and Nijkamp argue that "during the incubation phase, when major product innovations are made, the swarming process of (new Schumpeterian) firms is concentrated in central (usually urban) areas, because early innovations are more dependent on the urban 'milieu' than subsequent innovations" (1997: 4). This reasoning leads us to hold that innovation clusters will often be located in urban centres, and this would be particularly true for clusters composed prevalently by small and medium enterprises (SMEs), which might depend greatly on the supportive quality of their regional environment and the innovation-relevant knowledge sources available (Koschatzky and Sternberg, 2000).

Along this line of research it was argued that *technological districts* (contrasted to traditional industrial districts) would be essentially urban. Moreover, their development, based on a wide variety of complex urban centres, would be based on the emergence of competition between cities of industrialised countries (Pecqueur and Rousier, 1991).

4. An agent-based model of learning, innovation and geographical proximity

4.1 Rationale for the modelling approach

In the field of evolutionary economics, agent-based modelling is now recognised as one of the most promising new tools of investigation. The agent-based approach allows us to capture dynamics and complexity in our models. This is exactly what is required for studying processes of innovation, firms' partnering, and formation of IDs. The objective is to understand better the relations between micro-processes (the decisions and behaviours of economic actors) and the emergence of stylised facts common across much of industry (relating to R&D and the geography of firms) in the model output. Recently there has been significant other research targeting this area, (e.g. Gilbert, Pyka, and Ahrweiler (2001); Pajares, Lopez, and Hernandez (2003)) using agent-based methodologies. Gilbert et al. (2001) suggest that it has proved difficult to analyse innovation dynamics with the traditional analytical tools and suggest as an alternative, the need for "an abstract simulation model that could constitute a dynamic theory of innovation networks" (Gilbert et al. (2001)).

In our view, this approach to the problem of innovation dynamics does offer a promising way of modelling some of these processes. The focus of this paper is upon the analysis and management of change in industrial processes, and matters of industrial policy relating to support for regional innovation networks or industrial districts. We suggest that it would be most appropriate to use agent-based modelling in conjunction with case study and microeconomic research. The findings of simulations based upon formal logic must be supported, both qualitatively and quantitatively, by empirical work. Meanwhile, we wholly agree with Pajares, et al. (2004), confirming the use of agent-based models as a 'wider laboratory' for experimental economics, and in our investigation of the model we follow the same experimental methodology. The experimental set up is discussed in section 5 of this paper. Now, for the rest of this section, we describe the agent-based model.

4.2 Firms and their social network

The unit of analysis of this model is the firm. We assume a population of N firms allocated over a *social network* which is situated upon a grid of cells. Each firm is initially assigned a random position in the grid, and interacts with its closest neighbours. Initially, the *neighbourhood* is defined as the region on the grid that includes those cells adjacent in the four cardinal directions and within the firm's visible range ν (i.e. von Neumann neighbourhood structure). Not all the cells of the grid are occupied by firms, and those occupied contain only one firm.

The social network within which firms' interaction takes place could be represented as a graph where vertices correspond to firms and edges are firms' connections. Hence, we can write: $\mathcal{G}(I, \Gamma)$, where $I = \{1, \ldots, N\}$ is the set of firms, and $\Gamma = \{\Gamma(i), i \in I\}$ gives the list of firms to which each firm is connected. This can also be written $\Gamma_x = \{y \in I \setminus \{x\} \mid d(x, y) \leq \nu\}$, where $d(x, y)$ is the length of the shortest path from firm x to firm y (i.e. the path which requires the shortest number of intermediate links to connect firm x to firm y), and ν (visibility) is the number of cells in each direction which are considered to be within the firm's spectrum. Intuitively, Γ_x defines the neighbourhood of the firm (vertex) x.

The unit of time we define in our model is called the *cycle*. In each cycle, all firms are sorted into a random order, and then each is permitted to interact with one acquaintance. Initial acquaintances are the immediate neighbours (which are those within the visible spectrum). However, a firm can learn of the existence of other firms situated outside the neighbourhood through interactions with its acquaintances (i.e. this takes place while attempting to partner to innovate - we will clarify the way in which this happens in section 4.5).

Firms allocated in this social network aim at innovating. Innovation is defined as product innovation i.e. any time an innovation occurs a new product is supplied in the market. Innovation drives the firms' search process and motivates them to partner with other organisations (generating clusters of firms). In order to accomplish a new production process, new skills are required. Hence, innovating firms will aim at developing new skills and learn how to use them in order to produce new products. In this context skills are defined as an applied and agents' embedded dimension of knowledge.

In this model each firm is initially endowed with a *Skills Profile* (SP) which is defined as a complex and interdependent structure of abilities/competencies that can mainly be acquired through *individual learning* (implemented by firms as a specialised search process which replicates the work of R&D laboratories). On the contrary, acquiring new information on skills possessed by neighbour firms and, hence, on the profitability of partnering, occurs through social interaction. In this work we make the assumption that skills are not easily transferred through interactions, i.e. that they cannot simply be diffused from one firm's profile to another. In fact, we assume that *interactive learning* (i.e. flows of skills from a firm to another) can take place only if there is a joint production (i.e. innovation partnership) among two or more firms (this will be explained in detail in section 4.6)

The model itself is structured into two separate bodies: there is a set of properties of the *system* and a set of properties of *individual firms*. For the sake of clarity we shall start describing the properties of the *system*.

4.3 Defining the Firms' Skills Universe

The system is initially endowed with a *Firms' Skills Universe* (FSU), which contains the whole knowledge of the system. In this model, the FSU is represented by a network of nodes and links: nodes in the FSU can be thought of as possible skills or technologies to be learnt by the firms, and links define the requirements of each node. The FSU structure therefore defines the way in which subsequent skills depend upon the prior acquisition of other skills. Using a similar graph notation to that used to describe the *social network* we can write: *FSU* (Σ, Ψ), where $\Sigma = \{N_0, N_1, \ldots, N_{MAX}\}$ is the set of skills, and $\Psi = \{\Psi(i), i \in \Sigma\}$ gives the list of requirements to go from one node to another.

The FSU is generated at the beginning of the execution (simulation) of the model from an initial list of x nodes N_1 to N_x (with $x > 1$) which are themselves each linked to the root node, N_0. In the initialisation sequence nodes N_1 to N_x are placed, along with N_0, in a parent list P. Then, the remaining 'child' nodes, i.e. those belonging to the set $\{N_{x+1}, \ldots, N_{MAX}\}$, are taken in turn and added to the FSU as follows:

Step 1: determine the number of parents that the child will have (must be an integer between 1 and the dimension of P)

Step 2: select those parents randomly from the set P, and make a link from each one of them to the child
Step 3: for each parent currently in P, with a small probability P_D delete that parent from P
Step 4: add the current child to the parent list P

The process is repeated until all nodes have been positioned in the FSU. Following this specification, there is a small chance that a parent node would be *sterile*, i.e. that there are no dependent child nodes and at this vertex, the FSU reaches a *dead end*.

There is one additional step which gives the FSU a more interesting structure,

Step 5: with a small probability P_S, split P into two independent parts. Subsequently, child nodes will be attached to one specific parent list (chosen randomly), from where all its parents will be drawn

Step 5 causes branching of the FSU such that different areas of it can develop independently. In other words, later nodes are positioned such that they depend on a limited number of earlier nodes that are themselves quite interdependent. This specification is intended to represent the idea that it is feasible for a firm to 'specialise', or to learn some advanced skills without first having to learn almost everything else at a more basic level.

Implementing this in the program involved using a *meta-parent list* which was a dynamic array list of variable dimension. When the original parent list split at some (randomly chosen point), then the truncated part and the second 'splitter' part would then make up the meta list. It is also possible that one of the parent lists could become empty: in this case it would simply be removed from the meta list at the end of that cycle.

In figure 1 below we reproduce a representative graph of a *Firms' Skills Universe* composed of 50 nodes (skills).

- Insert Figure 1 about here -

Figure 1: Firms' Skills Universe

For step 1 we used a very simple probability function, in which each node (with the exception of the first five nodes and the root node) will have one parent with probability $p_1=0.6$, two parents with probability $p_2=0.3$ and three parents with probability $p_3=0.1$.

4.4 Radical innovations vs. incremental innovations

The system is also endowed with an ex-ante determined *Global Innovation List* (GIL), which represents all the possible innovations that can be achieved by firms. All of the potential innovations are generated in the initialisation phase of the simulation, at the same time as the FSU is created. During each iteration where a child node N_i is placed in the FSU, there is a small probability that an innovation is generated.

In this model, the researcher must specify in advance of running the simulation how many innovations there will be in a given experiment, and the corresponding number of node-indices are then selected at random. So, for example in Figure 1 above where there are 50 nodes, the node-indices run from 1 to 50 inclusive. If we specify ten innovations in our GIL, then ten of the node-indices will be chosen, and the innovations will be generated when these corresponding nodes are placed in the FSU. This is not a trivial matter since the innovations are generated from the *current* parent lists during the set up of the industrial environment, according to the following step (Step 6).

The GIL is composed of *m* incremental innovations and *n* radical innovations. The first kind of innovation is always based upon an already existing innovation, and is created by *replacing* one existing skill with a child, or a child of a child, etc. (we should say *closest available descendent*) of that same skill. On the contrary, radical innovations are created by combining a whole new set of skills never used for previous innovations. There is one special case that is the *root innovation* – the very first innovation – which must always be a radical innovation, requiring a new combination of skills.

We define I_t as the set of incremental innovations at the initialisation time, and the single innovation $\boldsymbol{\psi}_\mathbf{t} \in I_t$ as a vector of skills. We also define R_t as the set of radical innovations at the initialisation time, and the single radical innovation $\boldsymbol{\omega}_\mathbf{t} \in R_t$ as a vector of skills.

Now we can define the GIL, at the initialisation time, using set notation as:

$$GIL_t = I_t \cup R_t \qquad (2)$$

where the set I_t is defined as $I_t = [\boldsymbol{\psi}_\mathbf{1}, \ldots, \boldsymbol{\psi}_\mathbf{m}]$; and the set R_t is defined as: $R_t = [\boldsymbol{\omega}_\mathbf{1}, \ldots, \boldsymbol{\omega}_\mathbf{n}]$. Following on from the algorithm developed in the previous section, we can now add a further step to it.

Step 6: if the node index is one of those randomly selected, generate a new innovation by either (a) INCREMENTAL method or (b) RADICAL method, as described above, and add it to the GIL

Having created the GIL and the FSU at initialisation, the model is then ready to proceed to the main simulation phase where firms' learning and innovation take place. The single objective of firms is to obtain all of the necessary skills to fulfil the requirements of an innovation. The first firm (or group of firms) to attain a particular innovation will be recorded as the *First-Mover* (FM) innovator of that product innovation. In other words, that firm (group) is recognised as the first one to develop and market the product.

In this model, as we are only concerned with the innovation process, and not the emergence of markets for those new products, when the simulation reaches the point where an innovation is accomplished and its FM identified, that innovation will be "tagged" by the innovator in the GIL. Hence, each innovation can be performed only by those firms which perform it first.

As mentioned earlier, innovations may be attained either through searching for partners with complimentary skill attributes, or through individual learning. In order to do so, in each cycle firms have an opportunity to interact as well as to carry out their own R&D. To implement this two-fold learning process will involve the use of two nested time-levels. We aim to make explicit the relation between, on the one hand, the time taken to perform R&D and individual innovation - the *individual learning cycle* – and, on the other hand, the time taken to locate potential partners, interact with them and carry out joint innovations – *the partnering cycle*. This raises the important question of how to measure the opportunity to search for partnerships – *sharing the burden* - against the opportunity to internalise the innovation process, which may be more beneficial.

4.5 Individual learning and the FSP

Firms are initially endowed with a *Skill Profile* (SP) and a list of acquaintances. The SP is composed of a fixed number of *active* skills (i.e. usable to innovate) and a variable number of *dormant* skills. The active and dormant parts are coupled in such a way that, as the simulation advances, earlier acquired skills will always be transferred from the active SP to the dormant SP (which remain mutually exclusive).

As already mentioned, the goal of each firm is to innovate. Hence, every cycle, each firm will try to find out which innovation can be performed with the possessed active skills. This is done by comparing the individual firm's (coupled) SP with the GIL. If one or more innovations can be accomplished with the possessed skills, then the firm will successfully perform them. If no innovation can be performed then the firm will try to acquire those skills which would allow innovating. This can either take place through individual learning or through interactive learning and the definition of partnership for innovating. We will now describe in more detail the individual learning process, and then in the next section we will turn to describe how partnership are established and innovations are performed.

In section 4.3 we defined the FSU as a graph composed of a list of nodes representing skills articulated in a complex and structured way. Individual learning of new skills in the FSU takes place through a search process which goes from less to more specialised skills (i.e. it is a depth-first search). The agent has list of 'child' nodes of those already in its SP, which it systematically tries to learn. Recall that child nodes can be acquired only if all of its parents have been mastered (knowledge demands knowledge in order to be acquired). If the target node cannot be immediately learnt, then the algorithm backs up one level and first tries to acquire all of the parents of that node. Following this procedure, when eventually successful, the agent then moves on to target the children of the just learnt node.

In fact, the algorithm is rather more complicated than it at first seems, for it involves distinguishing between targets (children of the just learnt node) and sub-targets (parents/ ancestors of the target). The difference arises because, unless the target turns out to be a leaf node, only the children of targets can be next-in-line targets, whereas other children of sub-targets (that are not parents of the current target) should not be added to the target list, if we wish to carry out a depth-first search of the FSU, as specified earlier.

We shall define three types of firms: small, medium and large firms. The size of the firm is proportional to the size of the active SP, which has got a fixed dimension. Since the active SP has got a fixed dimension, every time a new skill will be acquired an old skill will be reallocated to the dormant SP. It is important to mention that dormant skills are skills which have been mastered by the firm but are currently not used; hence, can be used to learn new skills but not to innovate. If the firm is able to acquire the needed skills (i.e. has got the required skills to learn the new skills), then the new skills will be mastered and then added to the active SP. Consequently, the firm will be able to innovate. The active SP aims at reproducing the dimension of R&D laboratories. Hence, as widely acknowledged in the literature, larger firms will have bigger R&D departments and, therefore, will be able to handle a wider spectrum of skills. In addition, we define an *individual learning rate* parameter which varies according to the size of the firm. A small firm can learn one new skill per cycle, medium firms two and large firms three.

The specification of a dormant SP is necessary because otherwise the specialised search process outlined above would tend to lead to firms getting trapped within a particular area of the FSU corresponding to all the active skills of that firm, and hence unable to do further individual learning. This specification clearly represents the idea that firms do not unlearn those skills they have already mastered.

4.6 Interactive learning and firms' partnerships

If the firm is not able to innovate individually, then it will select one innovation from among those not yet accomplished, for which it has one or more of the needed skills, and it will try to partner with one of the acquaintance firms and jointly innovate. The selection of the innovation is random, but with the chance of selection of each innovation weighted according to the number of skills of that innovation already possessed.

Partnering will initially happen through direct interactions among neighbour firms. In this model we will not make a clear distinction between tacit and codified knowledge, we will just assume that, due to the presence of some degrees of tacitness, interactive learning will be somehow geographically bounded: i.e. interactive learning takes place only among neighbour firms. Firms' neighbourhoods, and their opportunities to partner, are thus defined by a visibility parameter, and their relative positions on the 2D grid. The neighbour is also initially selected at random. If, combining the skills of the firm and its chosen neighbour, the partnership is able to innovate, then both firms will acquire the complementary skills (interactive learning). If individually the partners each possess some of the required skills, but in partnership they do not possess all of them, then the search process will continue with the firm contacting another of its neighbours, until either the partners can together perform the joint innovation, or until there are no more neighbours to contact. Ultimately, if this is not successful then one of the neighbour firms might work as a link between the firm initiating the interaction and one of its neighbour firms. In this way, firms might create a network of partnerships. However, we leave this intermediate linking arrangement to further exploration: in all our initial experiments we allow only the initiating firm as the contacting (and hence central) firm in the partnership.

Given such learning structure, we aim at investigating the formation of innovation networks. We shall investigate how physical proximity, through the rapid diffusion of tacit knowledge, affects the innovative capabilities of heterogeneous firms. In section 4.2 we defined Γ_x as the set of initial acquaintances of firm *x*, we can now define $\varphi_{x,t}$ as the set of firms which are not immediate neighbours but with whom we can partner through a common neighbour at time *t*, and the firm $m_t \in \varphi_{x,t}$ who is added at each *t*. Now we can define the *innovation network* for firm *x* at time *t=T* as: $\Phi_{x,T} = \Gamma_x \cup \varphi_{x,T}$, where the set $\varphi_{x,T}$ is defined as: $\varphi_{x,T} = \{m_1, \ldots, m_t, \ldots, m_T\}$.

This network structure resembles that employed in a previous work (Morone and Taylor 2004b), which possesses some nice features. In fact, such a network is dynamic by its nature as it evolves over time, reshaping itself according to the learning dynamics within it.

5. Simulation experiments with the model

In this section we make some observations on what questions will be addressed by the model we have just described. In other words, we shall clarify our understanding of what this approach might reveal about issues concerning knowledge, innovation, and economic geography. In doing so we shall first describe the experimental set-up (section 5.1) and then describe the initial results in terms of the industrial environment (section 5.2) and the performance of firms and firms' partnerships (section 5.3).

5.1 Experimental set-up

Given the above model specifications, the experimental set-up should allow us address a range of questions related to innovation systems, firms' networks, and joint production processes. In this model we see the potential to be able to describe some quite complicated decision-making micro behaviours of individual firms. Our aim is to discover whether or not these described behaviours can support the emergence of the relevant 'stylised facts' about industry performance.

The questions of interest may be categorised into those concerning the performance of firms and partnerships, and those concerning the nature of the industrial environment.

(a) Firms/Partnerships:

1) Is the innovation cycle dominated by individual firms or by partnerships?
2) Is there domination by particular individuals/groups or is there equality of innovations?
3) Do large firms show a different performance compared to small and medium-sized firms?
4) Do partnerships or clusters of firms emerge and persist – in the sense of repeat occurrences over more than one innovation?

(b) Environmental characteristics:

1) Ratio of radical to incremental innovations, in terms of the *firms/partnerships* questions (a)?
2) Does this ratio make a difference in terms of the overall rate of innovation achievement?
3) What is the effect on performance of the parameters P_S and P_D that characterise industry?
4) How do geographical aspects of the environment impact performance of firms/partnerships?
5) How does the density of agents on the geographical space impact performance?
6) Do agents' initial skill profiles and their initial social networks matter?

A central theme of the Knowledge and Regional Economic Development Conference is that of the geographical dimensions of knowledge and innovation, and that is going to be our focus in this conference paper. In particular, we shall investigate the importance of proximity for partnership formation. To do so we assume a tacit aspect to knowledge and thus a geographical bound upon joint innovation.

The simulations presented in section 5.3 therefore explore the effect of tuning the tacit dimension of knowledge required for partnering for joint innovations, which is controlled through varying the visibility parameter v. In other words, we are assuming that the requirement for geographical proximity for partnering can be satisfied only by neighbouring firms. On the contrary, distant firms are not able to overcome this requirement. By increasing or decreasing the visibility parameter we can approximate the different possible regimes corresponding to different grades of importance imputed to the tacit nature of knowledge for carrying out joint innovation.

We used the JAVA platform with the RePast (Recursive Porus Agent Simulation Toolkit, The University of Chicago's Social Science Research Computing, 2003) libraries for implementing the model and UCINET 6 software (Borgatti, 2002) for analysis of the networks data. We carried out repeated simulation experiments (batches), to identify different trajectories of model behaviour. Over the batches we then took averages for all the relevant indices, and these statistics are presented in our results section. Expected results are that (1) larger neighbourhoods (ie. with higher values of the visibility parameter) will lead to a greater number of joint innovations. Also, (2) in the case of incremental innovations dominating the dynamics we expect more persistent partnerships, and (3) in the case where we have many radical innovations, we expect the innovations to be achieved at a slower rate.

5.2 Characterising the industrial environment

Simulation experiments were carried out to explore the performance of the generating algorithm which sets up the environment, and to understand more about how some of the fixed parameters effect the way that industrial environment is defined in the model at the start of the simulation.

First of all, it should be pointed out that the earlier figure, Figure 1, was generated with the parameter $P_D = 0.1$, and *number of skills per innovation = 5*. The probability of splitting the parent list was defined as $P_S = \vert P \vert \ / \ P_{MAX}$, i.e. the current length of (number of elements in) P divided by a maximum allowed length P_{MAX}, where P_{MAX} was initially set equal to six in the simulation which produced Figure 1.

Looking at this figure, which resembles a *fused tree* structure, we can immediately see that there are lots of branches of this network, where areas of the FSU develop into independent parts. Whilst there are several ‘corner stone’ nodes that are highly connected, there are also very many ‘leaf’ nodes which appear to be end points of the structure.

Initially we were quite satisfied with this structure for the FSU as we believed it could provide a quite complicated task for the firms in terms of their individual learning. Arguably, it could also be a fairly descriptive and yet generalised abstraction of how new skills are linked to previous skills in those types of knowledge-based industries for which innovation is an important process. For our initial investigations, it should be necessary to explore the algorithm which generates this structure.

To carry out our analysis of the FSU we explored the effect of the parameters P_S and P_D. We ran one simulation with the following parameter settings: S1: $P_{MAX} = 12$, $P_D = 0.1$, S2: $P_{MAX} = 6$, $P_D = 0.2$, thus presenting a further two simulations to compare with Figure 1. The fixed parameters, i.e. those which we did not change over the set of experiments were: $N = 50$ and *number of skills per innovation = 5*.

- Insert Figure 2 about here -

Figure 2: Parameter settings (Pd, Ps) and the FSU

Figure 2 shows the result of these simulations. The upper diagram (2a) shows the case where we set $P_{MAX} = 12$ and the lower diagram (2b) where $P_D = 0.2$, otherwise we have used the same parameters as those which generated the earlier FSU of Figure 1. What is most noticeable is that the main characteristics do not change very much. However, we can see that in Figure 2a there is a greater cohesiveness of the parts of the FSU, with fewer ‘branches’ that separate out the nodes. On the other hand, Figure 2b appears to be a bit more hierarchical compared to the other figures, in the sense that more branches must be traversed in order to reach the later nodes from the root node.

5.3 Firms’ performance and partnership formation

In this set of simulation experiments we used the fixed parameters: $P_D = 0.1$, $P_{MAX} = 6$, *N=200, number of skills per innovation = 5, number of radical innovations = 10, number of incremental innovations = 40, number of cycles = 20, number of agents = 40, grid size = 20*

The visibility parameter ν was varied as integer values drawn from the set: {2,3,4,5,6}.

Rather than letting simulations run until all innovations would be eventually achieved (or fail to be achieved), it was decided to limit the number of simulation cycles. Because we think that the partnering process is the most interesting object of our study, we wanted to see what happened when individual firms could only acquire very few skills out of the possible universe, (here, 10%). This kind of arrangement, in our view, most closely resembles an industrial district environment, where many firms operate in specialised roles, and innovation is a key driver for development.

Finally, it is necessary to note that in the initial simulation experiments there are a number of specifications of the model of which we did not yet begin to investigate, namely the variable *learning rate* parameter, the specification of the *interactive learning* step, and the distinction between *active* and *dormant SP,* being different sizes for different firms. Here, it was assumed that all Firms had an individual learning rate of one skill per cycle, and that innovations could be based upon any of the Firms’ acquired skills. The results presented below describe the experiment to understand the effect of the visibility parameter as a measure of the tacit dimension of knowledge necessary to carry out joint innovation.

- Insert Table 1 about here –

Table 1: Key statistics, innovations and partnerships

- Insert Figures 3a and b about here –

Figure 3a: Average of number of individual innovations attained by an agent

Figure 3b: Average of number of joint innovations in which an agent is involved

Table 1, and Figures 3a and b show that increasing the visibility parameter from 2 to 6 results in a reduction in the average number of individual innovations from an already low number, 0.58, to 0.21. Meanwhile the average number of involvements in joint innovations per agent increases greatly from 7.31 to 16.27. This emphasises, firstly, that individual innovations are very difficult to achieve in the environment defined by the FSU and GIL: joint innovations always strongly dominate. Secondly it tells us that as the geographical bound upon joint innovation becomes less strict, we can expect fewer and fewer innovations to be attained by individual firms on their own. However, the overall number of innovations clearly increases, by well in excess of 100%. The standard deviation of number of innovations attained is very high; this could be explained by either large variation among agents, or variation between simulation runs. Finally, the average number of partnerships increases with visibility, from a value of 1.5 to 1.97. As we would expect, this reflects the larger neighbourhood pools in which to search for potential innovation partners.

- Insert Figures 4a and b about here -

Figure 4a: Average of number of innovations in which an agent is involved

Figure 4b: Standard deviation of number of innovations in which an agent is involved

Looking at the distribution of innovations within simulation runs, Figure 4 shows the average and the standard deviation across agents of the number of innovations each was involved in, either as an innovator (i.e. contacting agent) or as a partner (i.e. contacted agent). These results, which were all averaged over 100 simulations, provide more information about the distribution of innovations attained during 20 cycles of the simulation. Figure 4a shows that the average number of innovations per agent increased steadily as the visibility parameter was increased from 2 to 6, from a value of 0.46 to 1.21. This is due, as shown above, to there being an increasing number of joint innovations, involving an increasing number of partners, for a larger value of v.

Figure 4a shows a surprisingly high deviation in the number of innovations achieved among agents, considering the small average value. The standard deviation was also found to increase as we increased the visibility, going from a value of 1.71 to a value of 2.74. Further to the earlier analysis, this statistic shows that the differences across individual simulation runs is in fact due to large variation among individual agents within each simulation. This tells us that some agents happen to acquire all the right skills as well as the right partners, whereas others acquire neither.

6. Discussion and further research

To summarise what can be concluded so far about the model, it has been instructive to scrutinise the industrial environment parameters relating to the FSU characteristics because this has enabled us to get a better understanding of the generating algorithm of the model. A further way in which we can characterise the industrial environment involves the procedure by which new innovations are generated through the FSU. By setting the relative number of radical innovations and incremental innovations, we can reflect the dominant *mode of innovation* in the industry.

As regarding the initial experiments looking at firms' propensity to partner, given the geographical bound due to the essentially tacit nature of knowledge flows in innovation systems, it was found that relaxing this bound did indeed increase the number of partnerships, by a factor of more than two when we increased the visibility from 2 cells to 6. Then, perhaps the most surprising result was that firms, whilst being rather homogenous in the sense that learning rate and initial skill profiles and other starting conditions were the same, showed a large discrepancy in the number of innovations that they eventually attained.

Meanwhile individual innovation was found to be extremely difficult, and therefore was very uncommon in all simulations. Although this was not surprising given that *Firm-agents* were all designed to be specialists, by virtue of an individual learning algorithm which always tried to attain *child-skills*. We have assumed depth first individual learning, however in later experiments it could be possible to specify a different learning mechanism, where the learning target is influenced by the firm's interactions, such as the frequency of requests for the different skill components in partnership formation.

Moreover, there is substantial work to be done in order to extend the original results presented in this conference paper and make a more rigourous investigation of the model. In particular, we have not yet explored the possibility of different learning rates for small, medium-sized, and large firms, which, in our view, could represent the variability in R&D investment and absorbtive capacity of firms.

A further point to clarify is how the relationship between the time-levels upon which the individual learning and partnering cycles are based, affects outcomes. Since it seems difficult to determine what is the appropriate balance between time spent mastering new skills and searching for complementary partners for joint innovation, all the more reason to be aware of the impact of any assumptions made here.

Considering innovation planning policy, we think it is far too premature to infer anything about real-world innovation networks based upon our initial experiments with this model. In view of the rather abstract nature of the model, if we want it to become useful to stakeholders and other decision-makers then we should probably focus more on the specifics of particular industries. The departure point for this type of approach would be to now validate the model upon empirical research in order to start improving the descriptive accuracy of representations therein.

Figure 1. Firms' Skills Universe (FSU)

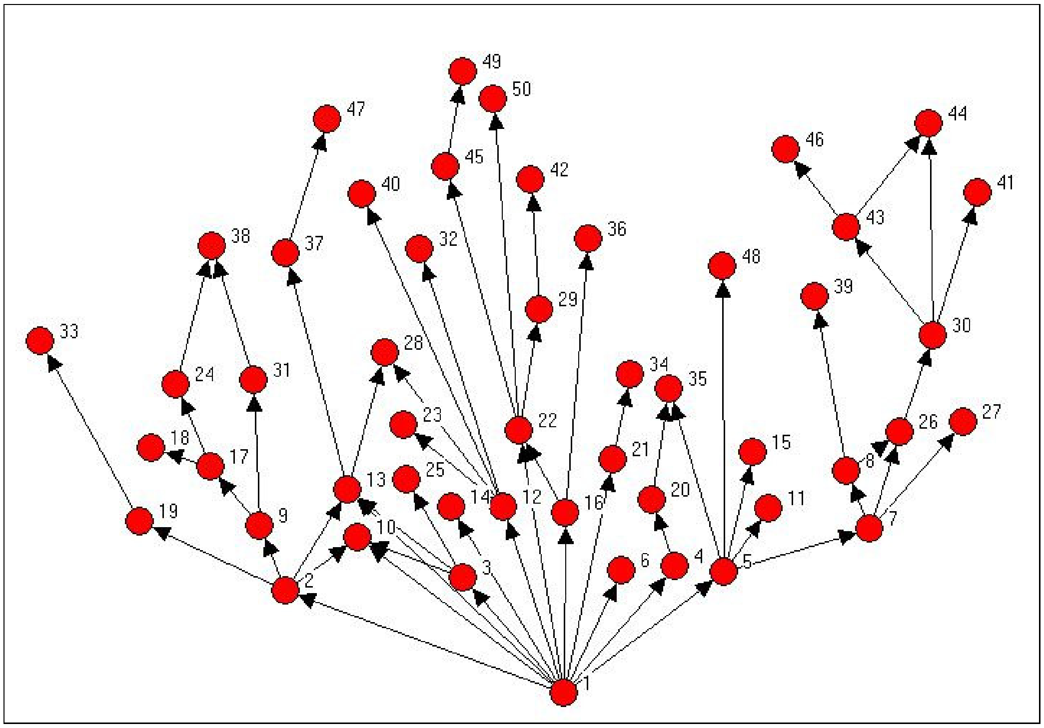

Figure 2. Parameter settings (Pd, Ps) and the FSU

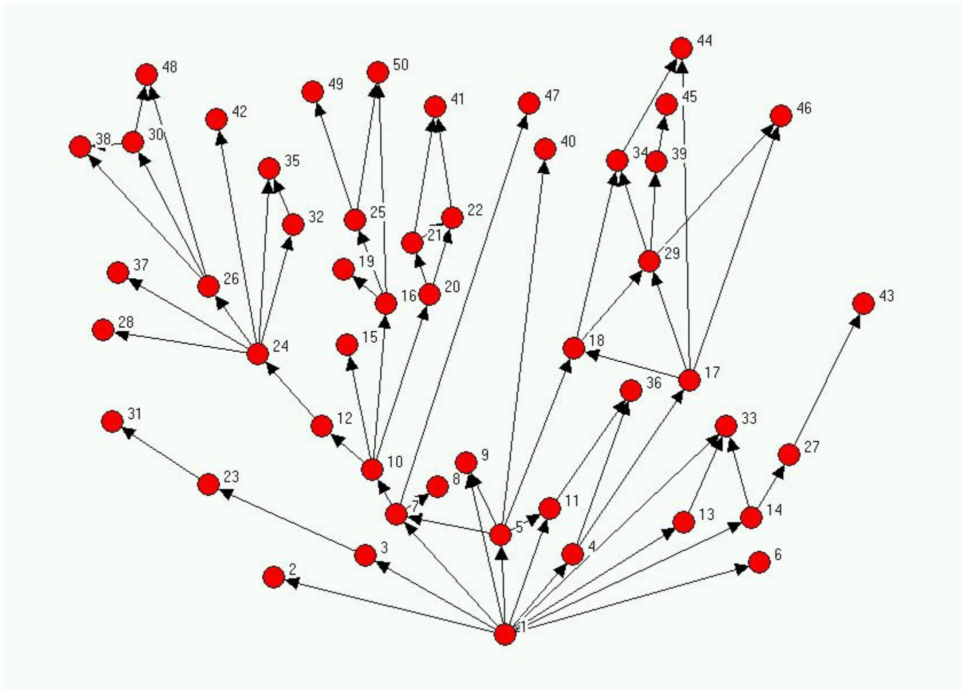

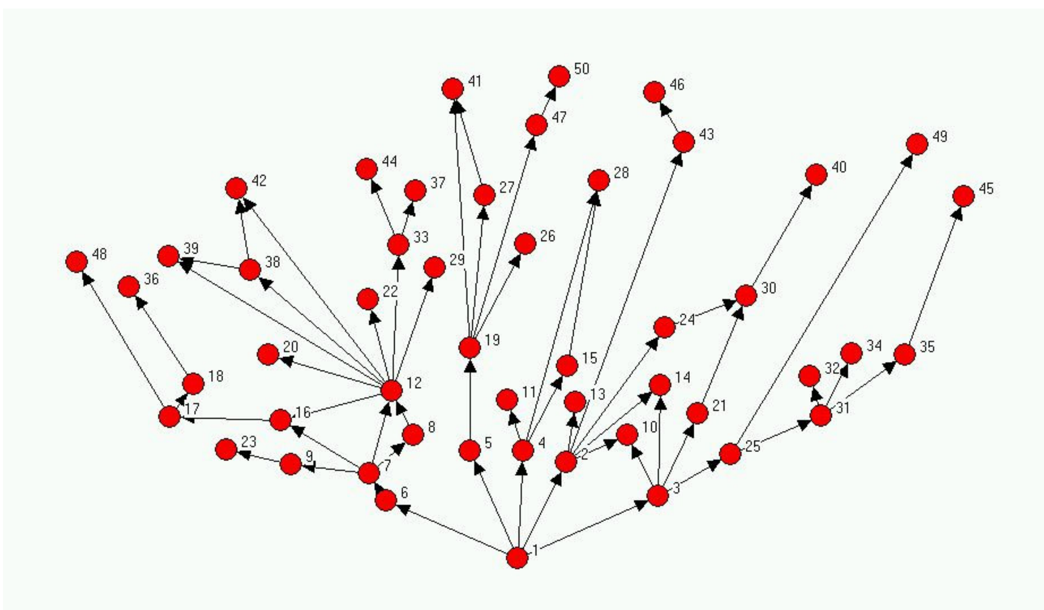

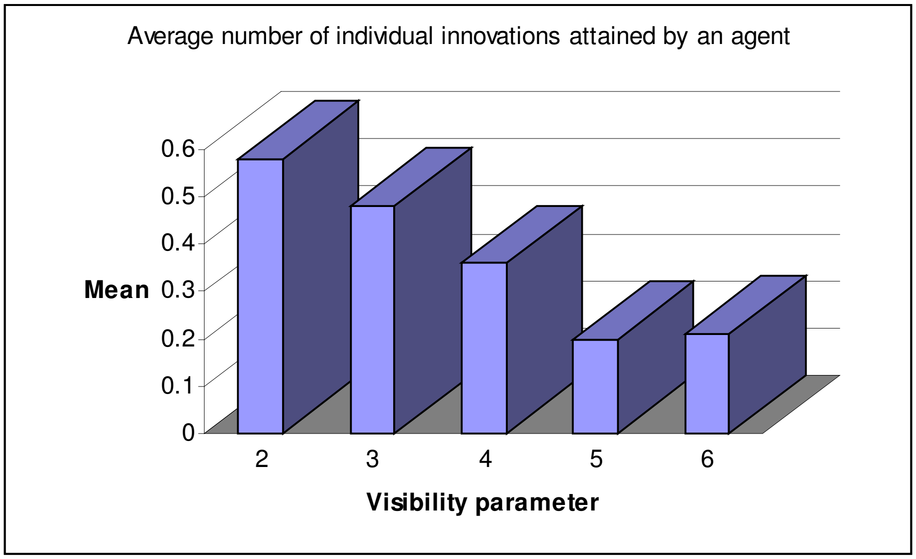

Figure 3a: Average of number of individual innovations attained by an agent

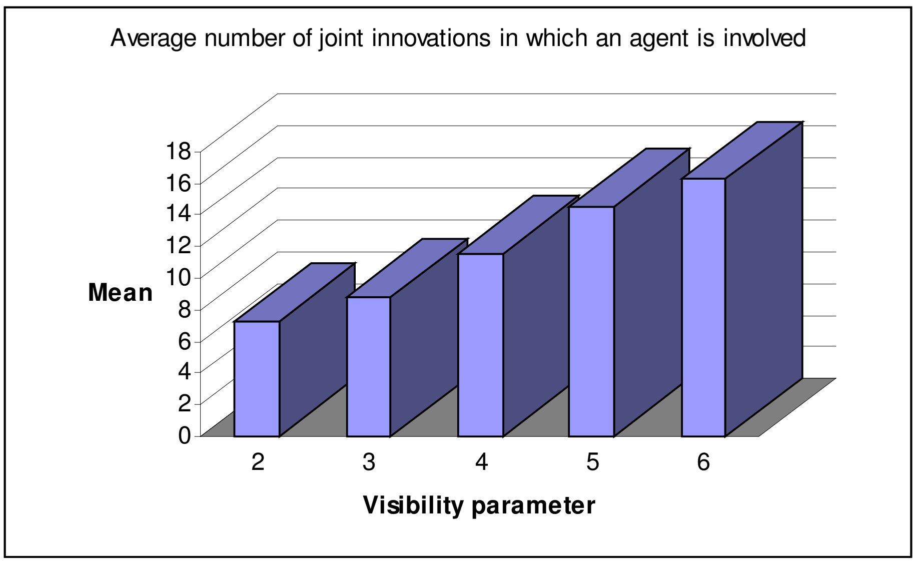

Figure 3b: Average of number of joint innovations in which an agent is involved

Visibility	Avg. indiv. innov.	Standard dev.	Avg. joint innov.	Standard dev.	Avg. number of partners
2	0.58	2.09	7.31	8.06	1.50
3	0.48	1.70	8.82	6.91	1.70
4	0.36	1.11	11.6	7.94	1.83
5	0.2	0.90	14.5	7.88	1.89
6	0.21	0.87	16.27	9.04	1.97

Table 1: Key statistics, innovations and partnerships

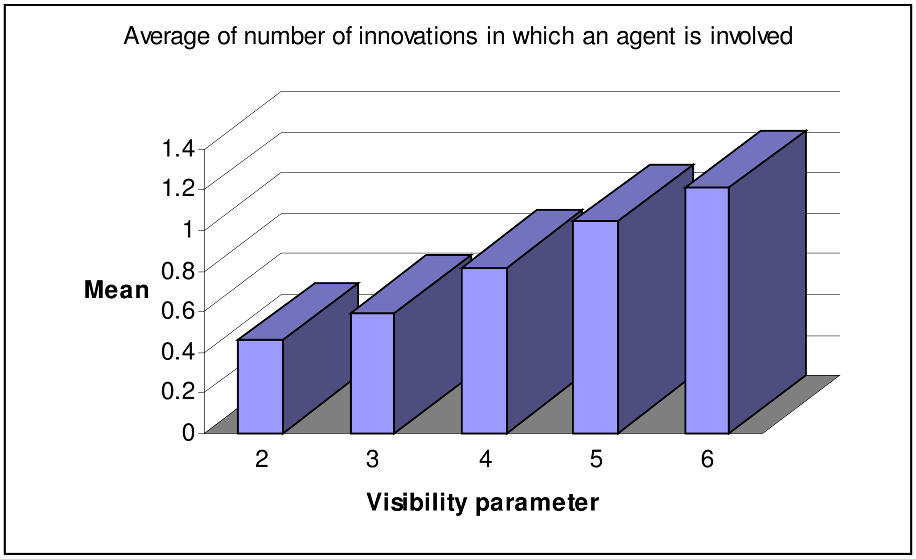

Figure 4a: Average of number of innovations in which an agent is involved (averaged over 100 simulations)

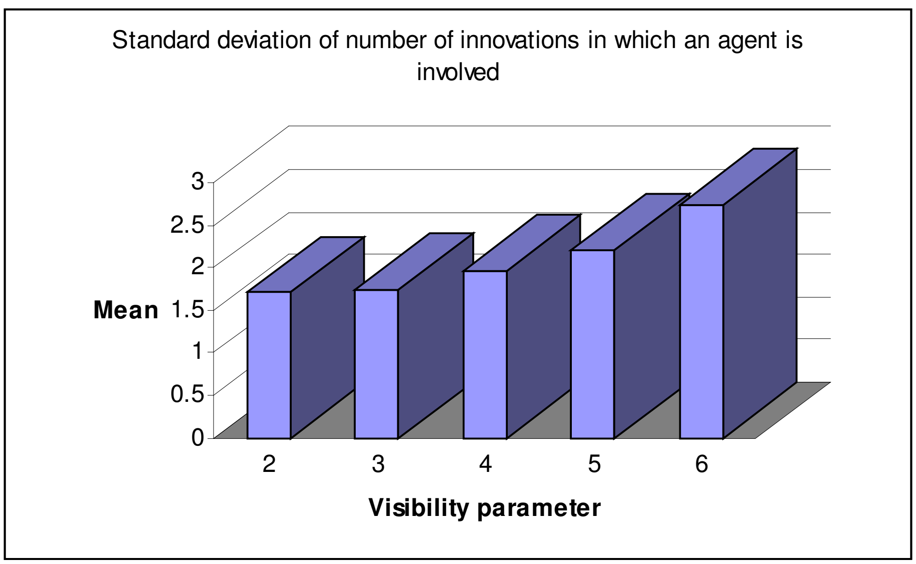

Figure 4b: Standard deviation of number of innovations in which an agent is involved (averaged over 100 simulations)

References

Axelroad, R. (1984), *The Evolution of Cooperation*, NY: Basic Books.

Borgatti, S.P., Everett, M.G. and Freeman, L.C. (2002), Ucinet for Windows: Software for Social Network Analysis. Harvard: Analytic Technologies.

Burns, T. and G. Stalker (1961), *The Management of Innovation*, Tavistock, London.

Conte, R. and M. Paolucci (2001), Intelligent Social Learning, Journal of Artificial Societies and Social Simulation, 4(1). (available online at: http://www.soc.surrey.ac.uk/JASSS/4/1/3.html)

Collins, H. (2001), "Tacit Knowledge, Trust and the Q of Sapphire", *Social Studies of Science*, 31 (1): 71-85.

Cowan, R. and Jonard, N. (2003), "The Dynamics of Collective Invention", Journal of Economic Behavior and Organization, 52 (4), 513-532.

Cowan, R. and Jonard, N. (2004) "Network Structure and the Diffusion of Knowledge", *Journal of Economic Dynamics and Control*, 28(8): 1557-1575.

Cooke, P. (2001), *Knowledge Economies : Clusters, Learning & Co-Operative Advantage*, London: Routledge.

Davelaar, E.J. (1991), *Regional Economic Analysis of Innovation and Incubation*, Avebury, Aldershot, England.

Dosi, G. and L. Marengo (1994), "Some Elements of an Evolutionary Theory of Organisational Competences", in: R .England (Ed.) *Evolutionary Concepts in Contemporary Economics*, University of Michigan Press, Ann Arbor.

Florida, R. L. (1995), "Toward the learning region", *Futures*, 27(5): 527-36.

Forsman, M. and N. Solidanter (2003), "Knowledge Transfer in Clusters and Networks", *Journal of International Business Studies – Literature Review*, n. 3, (available online at: http://copenhagen.jibs.net/LitReview/2003/2003_3_24.pdf)

Gilbert, N., A. Pyka and P. Ahrweiler (2001), "Innovation Networks – A Simulation Approach", *Journal of Artificial Societies and Social Simulation*, 4(3). (available online at: http://www.soc.surrey.ac.uk/JASSS/4/3/8.html)

Håkansson, H. (1987), *Industrial Technological Development: A Network Approach*, London: Croom Helm.

Håkansson, H. (1989), *Corporate Technological Behaviour: Co-operation and Networks*, London: Routledge.

Håkansson, H., I. Snehota (1995), *Developing Relationships in Business Networks*, London: Routledge.

Howells, J. R. L. (2002), “Tacit Knowledge, Innovation and Economic Geography”, *Urban Studies*, 39 (5-6): 871-884.

Kangasharju, A. and P. Nijkamp (1997), Innovation Dynamics in Space: Local Actors and Local Factors, Tinbergen Institute Discussion Papers, n. 97-062/3.

Koschatzky, K. and R. Sternberg (2000), R&D Cooperation in Innovation Systems - Some Lessons from the European Regional Innovation Survey (ERIS), European Planning Studies, Vol. 8, Number 4, August 2000, pp. 487-501.

Krugman, P. (1991), “History and industry location: the case of the manufacturing belt”, *The American Economic Review*, 81(2): 80-83.

Lundvall, B-Å (1988), “Innovation as a interactive process: from user-producer interaction to the national system of innovation”, in: G. Dosi, C. Freeman, R. Nelson, G. Silverberg, L. Soete (Eds.), *Technical Change and Economic Theory*, London: Pinter Publishers, pp. 349-369.

Malmberg, A. and P. Maskell (2002), “The elusive concept of localization economies - Towards a knowledge-based theory of spatial clustering”, *Environment and Planning A*, 34(3).

Malmberg, A. (1997), “Industrial Geography: Location and Learning”, *Progress in Human Geography*, 21 (4): 573-82.

Maskell, P. 1999. “Globalization and Industrial Competitiveness: the Process and Consequencesof Ubiquitification”, in: E.J. Malecki and P. Oinas (Eds.), *Making Connections: Technological Learning and Regional Economic Change*, UK: Ashgate Publishing.

Maskell, P. and Malmberg A. (1999), “The Competitiveness of Firms and Regions: ‘Ubiquitification’ and the Importance of localised learning”, *European Urban and Regional Studies*, 6: 9-25.

Morgan, K. (2001), “The Exaggerated Death of Geography: Localised Learning, Innovation and Uneven Development”, *Paper presented at The Future of Innovation Studies Conference*, Eindhoven University of Technology, 20-23 September 2001.

Morone, P. and Taylor, R. (2004a), "Small World Dynamics and the Process of Knowledge Diffusion. The Case of the Metropolitan Area of Greater Santiago De Chile", *Journal of Artificial Societies and Social Simulation,* 7:2.

Morone, P. and R. Taylor (2004b), "Knowledge Diffusion Dynamics of Face-to-Face Interactions", *Journal of Evolutionary Economics*, 14: 327-351.

Nelson, R. and N. Rosenberg (1993), “Technical innovation and national systems”, in: R. Neslon (Ed.), *National Innovation Systems: A Comparative Analysis*. Oxford: Oxford University Press.

Nonaka, I. and H. Takeuchi (1995), *The Knowledge-Creating Company*, OUP, Oxford.

OECD (1996), Employment and Growth in the Knowledge-Based Economy, Paris.

Oerlemans, L., M. Meeus and F. Boekema (1998), "Innovation: some empirical explorations of spatial embeddednes", in: J. van Dijk, F. Boekema (Eds.), *Innovation in firms and regions* (in Dutch), Assen: Van Gorcum, pp. 31-61.

Pajares, J., A. Lopez and C. Hernandez (2003), "Industry as Organisation of Agents: Innovation and R&D Management", *Journal of Artificial Societies and Social Simulation*, 6(2). (available online at: http://www.soc.surrey.ac.uk/JASSS/6/2/7.html)

Pavitt, K. and P. Patel (1991), "Large Firms in the Production of the World's Technology: An Important Case of Non-Globalisation", *Journal of International Business Studies*, 22: 1-21.

Pinch, S., N. Henry, M. Jenkins, S. Tallman (2003), "From 'industrial districts' to 'knowledge clusters': a model of knowledge dissemination and competitive advantage in industrial agglomeration", *Journal of Economic Geography*, 3: 373-388.

Pecqueur, B. and N. Rousier (1991), "Les districts technologiques, un nouveau concept pour l'etude des relations technologies/territories", in : Plan Urbain (Ed.) *Colloque de recherche sur les technopoles et les autres actions territoriales visant a favoriser les transferts de technologie*.

Polanyi, M. (1966), *The Tacit Dimension*, Routledge: London.

von Hippel, E. (1987), "Co-operation Between Rivals: Informal Know-How Trading", *Research Policy*, 16: 291-302.