Tuesday, 28 November 2000 First Draft

Preliminary and Incomplete
Comments Warmly Welcome

Knowledge Diffusion Dynamics and Network Properties of Face-to-Face Interactions*

Piergiuseppe Morone
SPRU – Science and Technology Policy Research
University of Sussex, Mantell Building
Brighton BN1 9RE, England
E-mail: P.Morone@sussex.ac.uk

Richard Taylor
CPM – Centre for Policy Modelling
Manchester Metropolitan University
Manchester M1 3GH, England
E-mail: r.i.taylor@mmu.ac.uk

ABSTRACT: This paper aims to understand some of the mechanisms which dominate the phenomenon of knowledge diffusion in the process that is called 'social learning'. We examine how knowledge spreads in a network in which agents interact by word of mouth. The social network is structured as a network graph consisting of agents (vertices) and connections (edges) and is situated on a wrapped cellular automata grid forming a torus. The target of this simulation is to test whether knowledge diffuses in a homogeneous way or whether it follows some biased path towards convergence or divergence.

JEL classification: D63, O30, R10

Keywords: Diffusion, Inequality, Knowledge, Network, Small world.

* The authors are grateful to Nick von Tunzelmann, Scott Moss, Robin Cowan and Ed Steinmueller for helpful comments in different stages of the work. The usual disclaimers apply.

"Knowledge is like light. Weightless and intangible, it can easily travel the world, enlightening the lives of people everywhere. Yet billions of people still live in the darkness of poverty - unnecessarily."

World Development Report, 1998/1999.

1. Introduction

Over the last decade, the importance of the dynamics of knowledge diffusion in the field of economics research has been increasingly recognised (World Bank, 1999; UNDP, 1999). As suggested in the 1998/1999 *World Development Report*, Ghana and South Korea had virtually the same income per capita back in the 1960s. 'By the early 1990s Korea's income per capita was six times higher than Ghana's. Some reckon that half of the difference is due to Korea's grater success in acquiring and using knowledge' (*World Development Report*, 1999: 1). It is unquestionably true that knowledge is not equally distributed among nations and, within nations, among people. This unequal distribution reflects, to some extend, the unequal distribution of wealth. From this it follows that closing the *knowledge gaps* is a priority for boosting an even process of development and growth. In this paper we will examine how people can acquire knowledge and how this is spread among a heterogeneous population which interact by face-to-face communications. Very schematically, each individual can accumulate new knowledge in two different ways: through a process of *individual learning* and/or through a process of *social learning*. The former process has been the main target of general human capital theory investigations. Classic studies of the learning process, though, suggest that *social learning* amounts to much more than formal schooling, and indicate the key importance of associating with peer groups, for example in actual working environments (e.g. Arrow, 1962; Rosenberg, 1982).

In the following sections we will present a formal model which investigates how knowledge spreads among agents situated in a network composed of several neighbourhoods. Section 2 discusses previous related work. Then, section 3 presents our model of knowledge diffusion. The last section shows the results of a simulation exercise based on the model. The main target of the model and the simulation is to observe the knowledge distribution dynamics and the network properties of the model. In other words, we examine under which conditions knowledge diffuses in a homogeneous way and under which conditions it does not. Furthermore, borrowing some concepts from the theory of graphs, we will investigate whether the network has the properties of a normal graph, of a random graph, or, as we expect, of a small world.

2. Previous work on the knowledge diffusion model

Cowan and Jonard (1999; C&J from now on) have already studied the phenomenon of knowledge diffusion in a network structure. We will start by describing the main features and results of this model to benchmark our model. C&J develop a model in the framework of graph theory. The aim of their model is to capture effects of incremental innovation and their diffusion over a network of heterogeneous agents. Agents are arranged in one-dimensional space, where each agent occupies one vertex and may interact with their *k* nearest neighbours on either side. The population of individuals is endowed with different levels of initial knowledge. Knowledge in this model is considered as a vector of values. A small number of agents are 'expert' and are endowed with a high level of knowledge in at least one value of the vector. All individuals interact among themselves, exchanging information via a simple process of barter exchange. Barter can only take place if the first individual has superior knowledge of one type and the other individual has superior knowledge of another type. Knowledge is a non-rival good and can be traded without decreasing the level of knowledge possessed by each trader. Individuals do not exchange information with everybody indiscriminately, but can only interact with their *k* nearest neighbours. With a certain probability *p*, though, any individual will "break the connection to the neighbour and re-connect that edge to a vertex chosen uniform randomly over the entire lattice. With probability *1-p*, this leaves the edge unchanged" (C&J, 1999: 5). As the value of *p* changes, the structure of the network changes. When the probability of breaking the local connection is set to zero the lattice will look perfectly regular, having each individual connected only with the closest neighbours (*k* is set equal to 4). This situation generates local clusters related with each other in a consecutive way. By choosing a fairly small value of *p* (however, strictly different from 0) a different situation is achieved, in which the lattice still looks fairly regular and highly clustered, but "because each non-local connection is a potential short-cut between two vertices the graph has the low average path length of an almost-random graph" (C&J, 1999: 6). The concept of average path length introduced here is a measure of the efficiency of the model, giving the average number of steps required to connect each pair of vertices in the lattice. To sum up, the scenario identified by low levels of *p* (the so-called small world) is one in which a short average path length and high degree of cliquishness coexist. Making an analogy between the social network described above and a friendship network, one can see the average path length as the number of friendship in the shortest chain connecting two agents, and the cliquishness as the extend to which friends of one agent are also friends of each other. The figure below shows a graphic representation of this dynamic:

Figure 1: Transition from regular to random structure

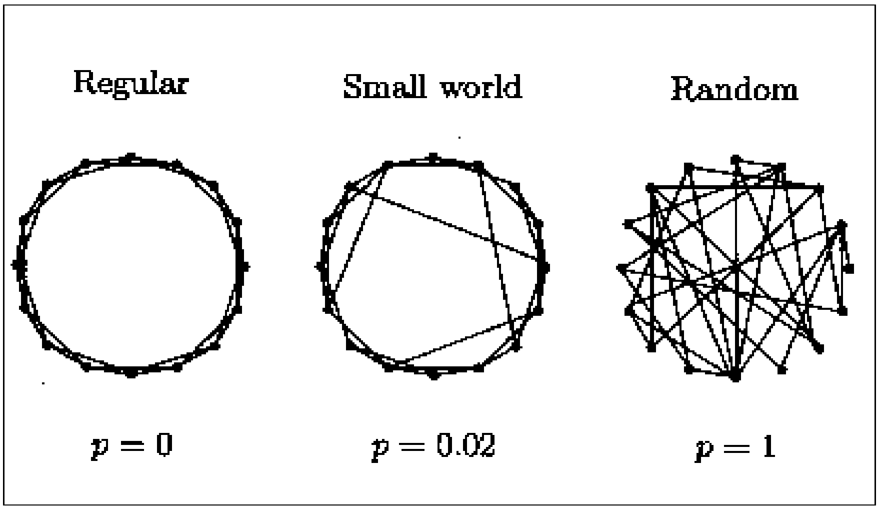

Source: C&J, 1999.

A simulation exercise developed by C&J on the basis of this model leads to some interesting results. The authors have found that the long-run average level of knowledge is a non-monotonic function of *p*, with its absolute peak corresponding to values of *p* around 0.06. These values are clearly in the small-world region (which varies between $p = 0.01$ and $p = 0.1$ approximately). From this first result the small-world situation appears to be the most efficient in terms of long-run knowledge diffusion. The picture changes when the authors consider the long-run knowledge variance. In this case the small-world region is the region in which the variance is the highest. In other words knowledge is distributed in the less equal way. To sum up, the small-world region is the place in which “the structural properties of network of agents’ relationships are at the same time producing the best overall performance in terms of how much knowledge diffuses through the system, and the worst overall performance when homogeneity of the allocation is considered” (C&J, 1999: 11). This pattern is generated by the fact that “in a small world, agents with direct connection to experts relatively rapidly become expert themselves, thus forming new (possibly distant) high-knowledge clusters. The high cliquishness of the small world ensures that agents near experts, and new second-generation experts also have high knowledge levels. This keeps them distinct from regions that are not closely connected with experts” (C&J, 1999: 12).

3. The knowledge diffusion model

We identify four main limitations of the C&J model and we will try to tackle them in this section. First limitation: there is no attempt to identify the process by which a social network may become a small world, or to consider situations where the nature of the environment may change due to changes in agent behaviour or changes in the network structure. In the C&J model the small-world configuration is imposed by changing the initial conditions (i.e. the value of *p*). This allows the authors to define average path length and cliquishness as a function of *p* but leaves us without a clear understanding of the relation between different values of *p* and the real world.

Second limitation: it is difficult to identify this variable *p* in real-world situations, where non-local connections certainly do exist but it is very hard to say what we could measure to find a value for *p*. If we don't have a good definition of *p* in the real world, nor do we then have a way to validate the network structure of the model. We should not forget that, just because it has small world properties (within a certain range of *p*), does not necessarily mean that it has a plausible structure.

A third limitation is the assumption that agents interact on the basis of complete information on the level of knowledge of their acquaintances. It is not in the spirit of the knowledge generation and diffusion field to assume such *a priori* knowledge of what agents need to learn. Search, trial and error and mistakes are very important in this field, meaning that decisions under uncertainty are a major part of the game[1].

A fourth and more complex limitation, which will not however be addressed in this paper, is the adoption of an oversimplifying notion of knowledge. Considering knowledge as a number (or as a vector of numbers) is indeed a convenient simplification, but restricts our understanding of the complex structure of knowledge generation and diffusion.

We try to address the first three limitations by introducing a social network structure which is emergent as a consequence of interactions, and which, we argue, specifies a more plausible mechanism for forming connections between distant agents, and in which individuals build internal models of the expected gain from interactions with their acquaintances. The key result of C&J is that the small world region has high mean and variance. These findings should be tested by developing alternative model specifications and making the same statistical analysis of the results. The

[1] We are grateful to Jorge Katz for pointing out this limitation of the C&J type of model.

alternative model described in the next section, will address some of those limitations of the C&J model, and will likewise be concerned with the efficiency and the equality of the diffusion process.

3.1 The model specifications

We will now consider in greater detail the specification of the model. We start by considering the definition of networks and neighbours. We assume a population of N agents. Each agent is initially assigned a random position in the grid, and interacts with its closest neighbours. The initial social network is created by connecting an agent with all other agents located within the neighbourhood of the agent. The local environment of the agent is called the "neighbourhood" of the agent and it is defined as the region on the grid that includes those cells adjacent in the four cardinal directions (north, south, east and west) and within the agent's visible range. The global environment is the entire grid of cells and is shaped as a wrapped grid forming a torus. We specify a wrapped grid so that there are no edge effects - where we might have different behaviour due to the boundaries because some agents will have a smaller local neighbourhood and will be less able to make connections to other agents (i.e. we would have peripheral agents and central agents). The graph associated with our network can be defined following the general definition of graph: 'a graph G consists of a nonempty set of elements, called vertices, and a list of unordered pairs of these elements called edges' (Wilson and Watkins, 1990). In our simulation vertices correspond to agents and edges are agents' connections. Formally we have G (I, Γ), where $I = \{1, . . ., N\}$ is the set of agents, and $\Gamma = \{\Gamma(i), i \in I\}$ gives the list of agents to which each agent is connected. This can also be written $\Gamma(x) = \{y \in I \setminus \{x\} \mid d(x, y) \leq v \}$, where $d(x, y)$ is the length of the shortest path from agent x to agent y (i.e. the path which requires the shortest number of intermediate links to connect agent x to agent y), and v (visibility) is the number of cells in each direction which are considered to be within the agent's spectrum. Intuitively, Γ_x (we will use this notation rather than $\Gamma(x)$ from now on) defines the neighbourhood of the agent (vertex) x.

Not all the cells of the grid are occupied by agents, and those occupied are occupied by only one agent. A neighbourhood can be defined as the adjacent cells on the grid, and there are normally two kinds: the von Neumann neighbourhood, which includes adjacent cells in the four cardinal directions (north, south, east, west), and the

Moore neighbourhood, which includes adjacent cells in both the cardinal and diagonal directions (the surrounding eight cells). In our model we will adopt the von Neumann definition of neighbourhood. The "lack of diagonal vision is a form of imperfect information and function to bound the agents' *rationality*" (Epstein and Axtell, 1996: 24).

Each agent has an initial, arbitrarily chosen, level of knowledge *k*. *K* is an integer chosen at random from a given range. Agents may make acquaintances (we call this form of interaction 'communication interaction') with other agents who belong to the same neighbourhood, and with distant agents. Initial acquaintances are the immediate neighbours (which are those within the visible spectrum). Subsequently, an agent can learn of the existence of other agents through interactions with its acquaintances (i.e. it can be introduced to acquaintances of its acquaintances). Therefore the number of acquaintances change over time. In this way we introduce a dynamic element in the model since we do not assume that the number of acquaintances is constant over time. Having defined Γ_x as the number of initial acquaintances of agent *x* (or *first generation connections*), we define $\varphi_{x,t}$ as the number of acquaintances of the acquaintances at time *t* (or *next generation connections*). Now we can define the total number of acquaintances at a certain moment of time as:

$$\Phi_{x,T} = \Gamma_x + \sum_{t=1}^{T} \varphi_{x,t} \ . \qquad (1)$$

Each acquaintance has an associated strength, $\tau \in (0,1)$, which is a measure of the strength of the relationship from the agent to its acquaintance. Note that this model is not constrained to have symmetry of relationships between agents: in general, more prestigious agents (with higher levels of knowledge) will be the object of strong relationships with more peripheral agents (with lower levels of knowledge), which may be unreciprocated or reciprocated only weakly.

The unit of time we define in our model is called the *cycle*. In each cycle, all individuals are sorted into a random order, and then each is permitted to interact with one acquaintance. There are two types of interaction, called the "communication interaction" and the "diffusion interaction". During the communication interaction, information about the existence of other individuals may be transmitted as follows: if the selected acquaintance is connected to other individuals of which the agent is not aware, then a connection is made from the agent to the acquaintance of the acquaintance. If there is more than one acquaintance of the acquaintance, then just one of them is chosen at random.

An agent will have preferences for interactions with acquaintances with which it has strong relations, because of the way that the acquaintance is selected for interaction. Agent *y* will be selected for interaction with agent *x* with probability given by:

$$p^x(y) = \frac{\tau_y^x}{\sum_{l=1}^{\Phi} \tau_l^x},$$

(2)

Agent interaction is based on the transmission of knowledge. However, in this model, agents' interaction is not based on the assumption that each agent has, at any moment of time, full information about other agents' knowledge levels. Rather, each agent will build internal models of the likelihood of having gain interactions. We define a "gain interaction" as the case where an agent increases its knowledge level as the result of an interaction. Each time an interaction occurs, the strength of the relationship τ_i, where $i = \{1, \ldots, \Phi\}$, is adjusted (we drop for simplicity the index of the agent and use it only when strictly necessary):

$$\tau_{i,t} = \tau_{t-1} + \varepsilon \quad where\ \varepsilon = \begin{cases} 1 & \text{if the interaction is a gain interaction;} \\ 0 & \text{otherwise.} \end{cases}$$

(3)

In this way the agent will develop preferences for selecting for an interaction with acquaintances with which it has previously experienced gain interactions. This is what we mean when we state that the agent builds internal models of preference.

The gain from the diffusion interaction is calculated in the following way: first calculate the distance in knowledge between him/herself and the acquaintance:

$$\delta_t = k_t^y - k_t^x,$$

(4)

then calculate the actual gain:

$$G = \frac{k_{t-1}^y}{\delta_t}.$$

(5)

(NB since we are working with integer values of knowledge, the gain *G* is rounded down to the nearest integer value).

The new value for the individual's knowledge is calculated as:

$$k_{new} = k_{old} + G \qquad (6)$$

The strengthening of relationships serves to make interactions between the same agents more likely in subsequent periods. This measure of strength of relationship is used largely in the (network analysis) calculation of prestige, or the extent to which the agent has relationships directed towards it compared with the relationships it directs towards others. This rule, the knowledge diffusion rule, is based on the assumption that agents with similar levels of knowledge are more likely to have gain interactions and therefore will be more likely to interact. This fact is theoretically supported by the literature on the 'epistemic community' or 'community of practice'[2].

[2] See, for instance, Wenger (1998).

4. Model Parameters

As anticipated earlier the focus of the simulation is to observe the knowledge distribution dynamics and the network properties of the model. To tackle the distribution question we calculate the average level of knowledge in the network after any cycle (a cycle correspond to a time period):

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

and the variance in knowledge distribution:

$$\sigma^2(t)=\frac{1}{N}\sum\bar{\mu}_i^2(t)-\bar{\mu}^2(t).$$

To study the network properties we calculate the average path length:

$$L(t)=\frac{1}{N}\sum_{x=1}^{N}\sum_{x\neq y}\frac{d(x,y)}{N-1};$$

and the average cliquishness:

$$C(t)=\frac{1}{N}\sum_{x=1}^{N}\sum_{y,z=1}^{\Phi}\frac{X(y,z)}{|\Phi_x|(|\Phi_x|-1)/2},$$

where $X(y, z) = 1$ if y and z are connected at time t (no matter whether the connection is a first generation or next generation connection), and $X(y, z) = 0$ otherwise.

We consider a small population of 40 agents allocated randomly over a grid composed of 100 cells. Each agent has a visibility ν = 2, meaning that each agent can see the two cells situated in the four cardinal directions (north, south, east and west)[3]. The initial knowledge k can take any natural value between 0 and 99, in the first case analysed, and between 50 and 99, in the second case. In fact, we run two simulations[4], one in which the initial knowledge gap (i.e. knowledge variance) is relatively large, and a second case in which the knowledge gap has shrunken by one-half. Knowledge is randomly distributed among agents.

The target of the simulation exercise is to test the hypothesis that a small world implies high mean and variance in the knowledge distribution. As anticipated in the second section, this was the major result obtained by C&J model. We will try to find whether this result is consistent with our results, which will be presented and discussed in the following section.

[3] Since agents do not occupy all cells, there is a critical density problem. Running several simulations, given a constant visibility ν = 2 and a constant grid dimension of 100 cells, we have detected this level to be around 20 agents: below this number of agents the results tend to be significantly different, while above this threshold we observe a significant stability in the results.

5. Simulation Results

After letting agents interact for thirty cycles (which implies a large number of interactions, since during each cycle all agents have a chance for interaction) we observe significant changes in the knowledge distribution: both variance and mean increase over the whole period, suggesting a polarization of knowledge distribution and an increase in the knowledge gap. The two figures below show these dynamics: in Figure 2 we plot the variance in knowledge against time and we observe σ^2 growing at a stable rate over the whole period. Plotting μ against time we observe a similar pattern: the average knowledge of agents increase from an initial level of nearly 55 to almost 70.

Figure 2: Change in knowledge variance. Initial knowledge range [0, 99]

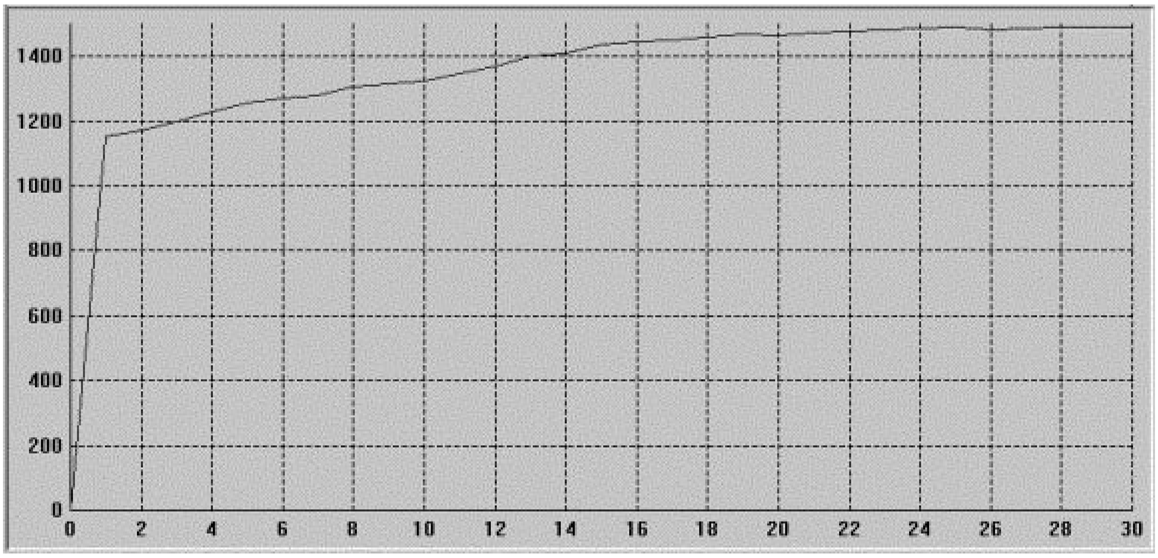

Source: Simulation results

Figure 3: Change in the average level of knowledge. Initial knowledge range

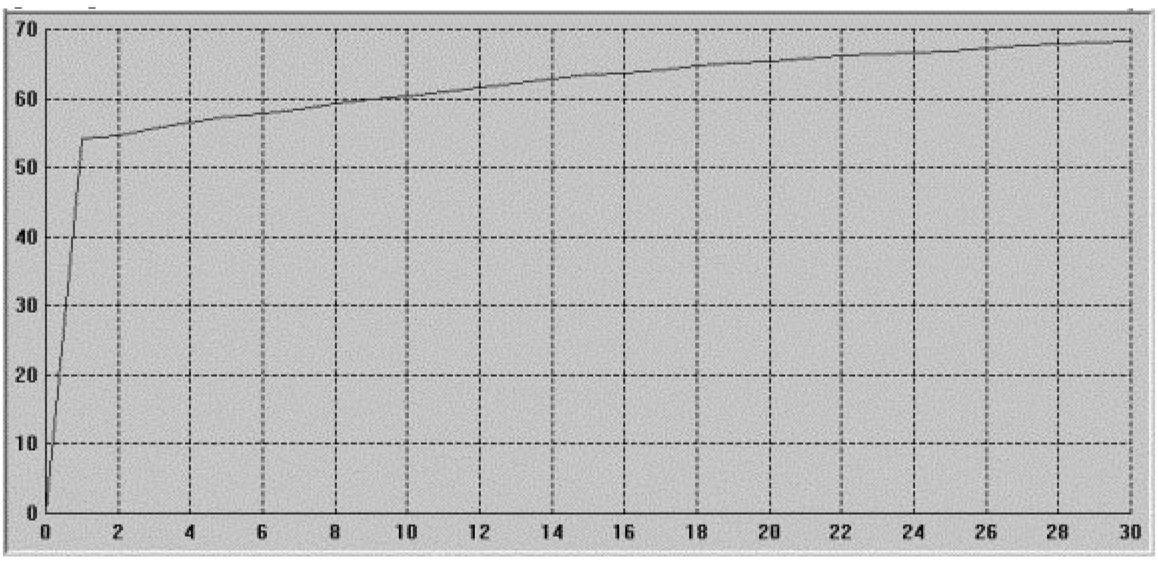

Source: Simulation results

Investigating the individual performance of each agent over time we can divide the whole population into three groups. The first group consists of 'fast catching-up agents', which grow very fast, reaching the highest level of knowledge possible (which is the initial level of knowledge of the most knowledgeable agent, since in the model there is no individual learning). A second group of 'slow catching-up agents', which over time increase substantially their level of knowledge display constant increases in knowledge, but take much longer to complete the catching-up process. Finally, there is a third group of 'unable to catch-up agents', which are completely unable to reach high levels of knowledge and after a few cycles are unable to have any positive interaction. This can be observed by plotting the individual's level of knowledge against the time cycles:

Figure 4: Individuals' levels of knowledge. Initial knowledge range [0, 99]

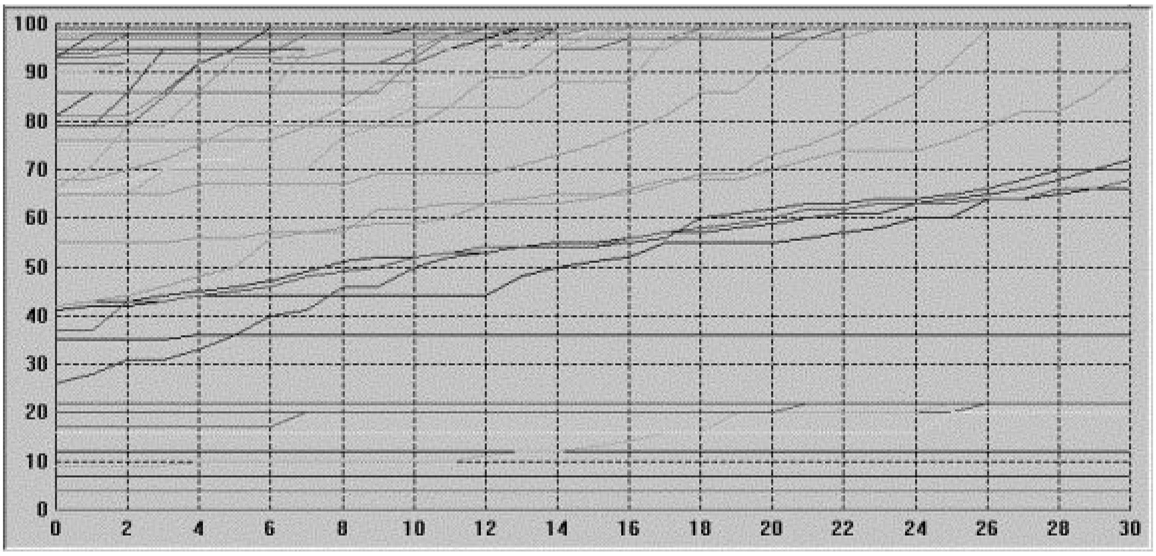

Source: Simulation results

From Figure 4 it can be observed how all those agents (except two cases) who have an initial level of knowledge lower than 40 are unable to catch-up. Within this group, though, there is a local catching-up process which reduces the variance within the group. From these findings we can conclude that this simulation shows a diverging path in the knowledge distribution, generating an overall increase in the knowledge gap.

A completely different picture emerges if, keeping all other parameters unchanged, we reduce the initial knowledge variance. Taking $k \in [50, 99]$ we observe a constant reduction in the variance and a constant increase in the mean. Both these indicators reach the highest possible values in less than 30 cycles:

Figure 5: Change in knowledge variance. Initial knowledge range [50, 99]

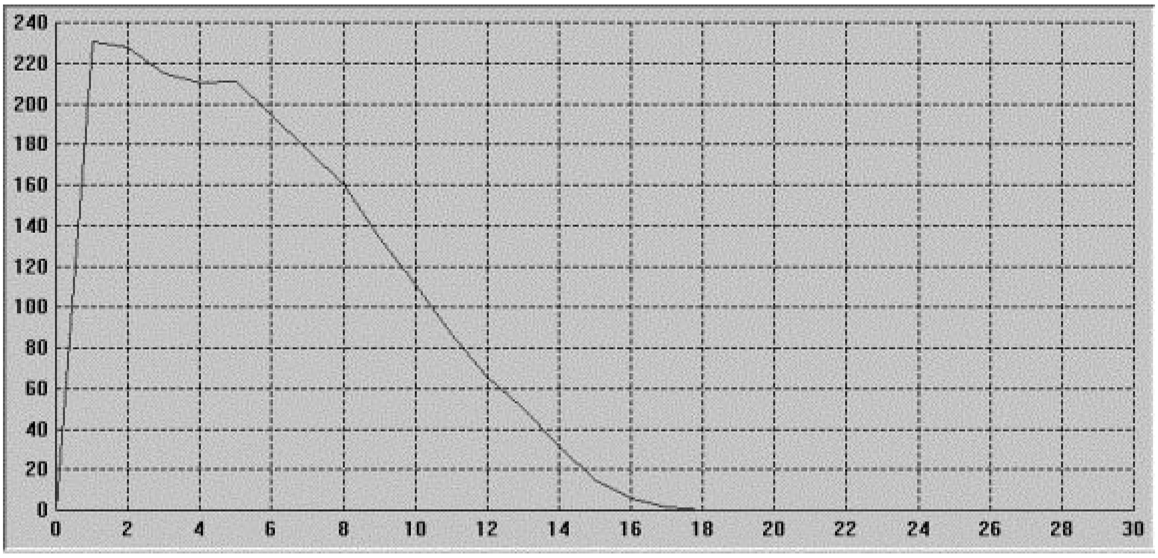

Source: Simulation results

Figure 6: Change in the average level of knowledge. Initial knowledge range

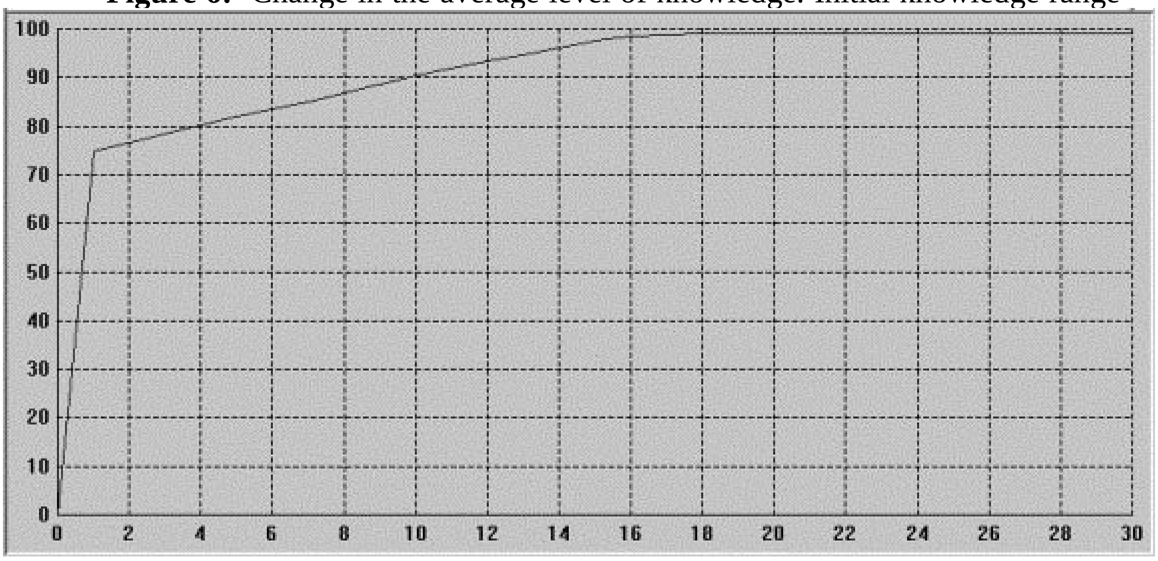

Source: Simulation results

In Figure 5 above the variance goes to zero after 19 cycles, which means that all agents have reached the highest possible level of knowledge. In this case interaction is generating a fast converging path and all agents reach the highest level of knowledge possible.

Figure 7 below shows how convergence occurs. Agents with highest level of knowledge are those who earlier converge to the highest level of knowledge possible. There is not a one-to-one correspondence between the initial position and the speed of convergence, since the pace of convergence will also depend upon the surrounding environment (i.e. if you are surrounded by agents with lower level of knowledge or much higher level of knowledge you have little margin for learning; on the other hand if your neighbours have similar level of knowledge there are more chances for learning).

Figure 7: Individuals' levels of knowledge. Initial knowledge range [50, 99]

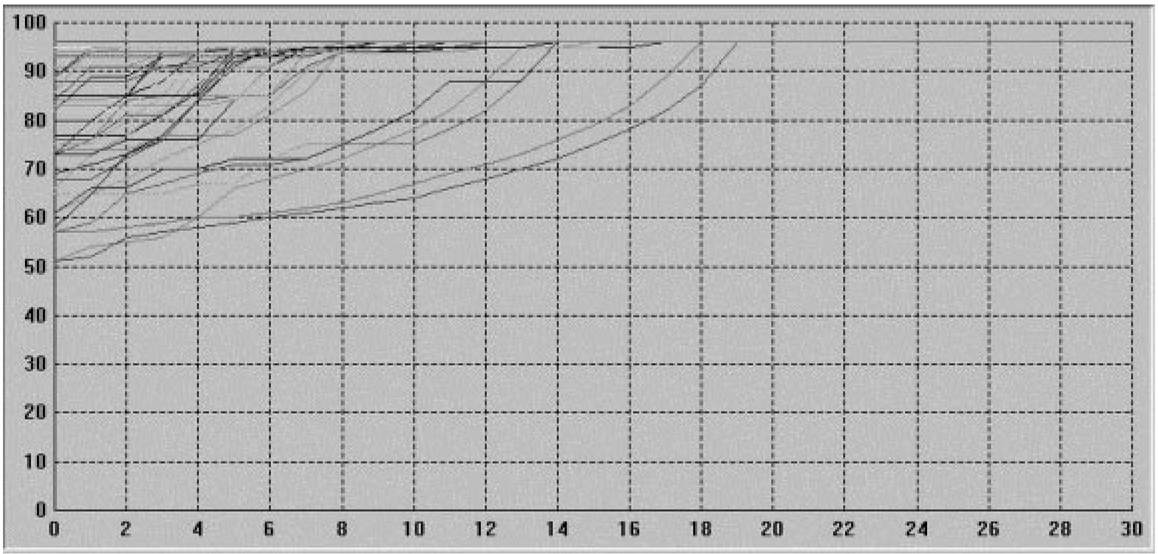

Source: Simulation results

When, we consider the network properties of the model we discover that in both cases the final network, which emerges from agent interactions, is a small world, meaning that it has a short average path length and high degree of cliquishness.

Table 1. Small world calculation results

		C	*L*
First simulation	First cycle	0.1364706	16.457471
	Last cycle	0.9093333	2.3229885
Second simulation	First cycle	0.0893333	26.637931
	Last cycle	0.9376667	2.3011494

Source: Simulation results

In the table above we present the results of the small world properties calculations. For each simulation we have calculated average path length and cliquishness for the first and the last cycle. Looking at this figure we observe, in both cases, a sharp increase in *C* and a significant decrease in *L*, generating a final situation in which short average path length and high degree of cliquishness coexist. These findings, in line with our predictions, clearly suggest that in both cases the small world emerges as a result of agents' interaction.

6. Conclusions

The model presented in this paper aimed to investigate the dynamics of knowledge diffusion and the network properties of a small population of agents. The network was divided into several neighbourhoods, and agents were allowed to interact with their acquaintances (neighbours) through face-to-face interaction. We wanted to observe if small world properties could emerge by allowing agents to acquire new acquaintances over time. The basic interaction between two acquaintances gave each agent the chance to enrich their knowledge level and to acquire new acquaintances. The results of this model were consistent with our expectations. We observed, in different simulations, the emergence of small world properties after allowing for some interactions. Changing the initial knowledge gap between agents we observed, though, very different knowledge diffusion dynamics: in a first case in which the knowledge gap was relatively big the population divided into three groups, a first group of knowledgeable agents, a second group of catching-up agents, and a third group of marginalised agents. Compressing the initial knowledge gap the picture changed drastically: all agents manage to catch-up, reaching the highest possible level of knowledge within relatively few cycles.

An important implication of these results is that small world dynamics are not necessarily associated with high variance in knowledge. It can well be the case that small world dynamics generate an efficient and equal process of knowledge diffusion. From our model it follows that the key variable to discriminate between convergence and divergence in knowledge diffusion is the initial level of knowledge variance: if initially agents are endowed with extremely different levels of knowledge, even though they live in a small world, they will not be able to converge towards the highest level of knowledge. Some agents will quickly reach the highest level of knowledge while some others (those initially endowed with low level of knowledge) will stay ignorant forever. On the other hand, if the initial knowledge gap is reasonably small (i.e. the group is reasonably homogeneous) the process of knowledge diffusion will quickly converge to the highest possible level of knowledge. In our understanding this result suggests that small world properties can facilitate the equal diffusion of knowledge only if some ‘barriers to communication’ are initially removed. If this is the case, the small world properties speed up the catching-up process, otherwise a diverging path will emerge.

References

Ancori, B., A. Bureth, et al. (2000), 'The Economics of Knowledge: the Debate About Codification and Tacit knowledge', *Industrial and Corporate Change*, 9, (2): 255-287.

Arrow, K. (1962), 'The Economic Implications of Learning-by-Doing', *Review of Economic Studies*, 29, : 155-173.

Cowan, R., P. A. David, et al. (2000), 'The Explicit Economics of Knowledge Codification and Tacitness', *Industrial and Corporate Change*, 9, (2): 211-253.

Cowan, R. and N. Jonard (1999), 'Network Structure and the Diffusion of Knowledge', *MERIT Working Papers*, 99-028.

Cowan, R. and N. Jonard (2000), 'The Dynamics of Collective Invention', *MERIT Working Papers*, 00-018.

Epstain, J. and R. Axtell (1996), *Growing Artificial Societies* The MIT Press: Cambridge, Massachusetts & London, England.

Granovetter, M. (1973), 'The Strength of Weak Ties', *American Journal of Sociology*, 78: 1360-1380.

Milgram, S. (1967), 'The small world problem', *Psychol. Today*, 2: 60-67.

Nonaka, I. and N. Takeuchi (1995), *The Knowledge Creating Company*, Oxford University Press: Oxford.

Polanyi, M. (1962), *Personal Knowledge*, Routledge: London.

Rosenberg, N. (1982), *Inside the Black Box: Technology and Economics*, Cambridge University Press.

Saviotti, P. P. (!998), 'On the Dynamic of Appropriability, of Tacit and Codified Knowledge', *Research Policy*, 26: 843-856.

Steinmueller, W. E. (2000), 'Will New Information and Communication Technologies Improve the 'Codification' of Knowledge?', *Industrial and Corporate Change*, 9: 361-376.

UNDP (1999), *Human Development Report 1999*, Oxford University Press: New York, Oxford.

Watts, D. J. (1999), *Small Worlds*, Princeton University Press: Princeton.

Wenger, E. (1998), 'Communities of Practice; Learning and Social System', *System Thinker*, June.

Wilson, R. J. and J. J. Watkins (1990), *Graphs: An Introductory Approach*, Wiley: New York.

World Bank (1999), *World Development Report 1998/1999*, Oxford University Press: New York.