

# A Laboratory Experiment of Knowledge Diffusion Dynamics

ANDREA MORONE<sup>\*</sup>, PIERGIUSEPPE MORONE<sup>\*\*</sup>, RICHARD TAYLOR<sup>\*\*\*</sup> ♦

**ABSTRACT:** This paper aims to study, by means of a laboratory experiment and a simulation model, some of the mechanisms which dominate the phenomenon of knowledge diffusion in the process that is called 'interactive learning'. We examine how knowledge spreads in different networks in which agents interact by word of mouth. We define a regular network, a randomly generated network and a small world network structured as graphs consisting of agents (vertices) and connections (edges), situated on a wrapped grid forming a lattice. The target of the paper is to identify the key factors which affect the speed and the distribution of knowledge diffusion. We will show how these factors can be classified as follow: (1) learning strategies adopted by heterogeneous agents; (2) network architecture within which the interaction takes place; (3) geographical distribution of agents and their relative initial levels of knowledge. We shall also attempt to single out the relative effect of each of the above factors.

**JEL classification:**

**Keywords:** Knowledge, Network, Small world, Experiment, Simulation.

---

\* University of Bari, Department of Economics, Via Camillo Rosalba 54, Bari, Italy.  
Email: [a.morone@dse.uniba.it](mailto:a.morone@dse.uniba.it)

\*\* University of Roma 'La Sapienza', Istituto di Economia e Finanza, Piazza A. Moro 5, Roma, Italia.  
Email: [piergiuseppe.morone@uniroma1.it](mailto:piergiuseppe.morone@uniroma1.it)

\*\*\* Manchester Metropolitan University, CPM – Centre for Policy Modelling, Manchester, England.  
E-mail: [richard@cfpm.org](mailto:richard@cfpm.org)

♦ Authors' names are ordered alphabetically

## **1. INTRODUCTION**

Several authors have recently attempted to study how knowledge diffuses by means of informal interactions. Both theoretical and empirical works addressed this topic with the aim of shedding some light on the complex mechanisms which regulate informal learning (Ellison and Fundenberg, 1993; Rauch, 1993; Bala and Goyal, 1995, 1998; Achemoglu and Angrist, 1999; Chwe, 2000). The growing interest in this field of research is due to the increasing importance of knowledge for economic growth and development, coupled with the recognition of the dominant role of informal interactions in producing and diffusing this knowledge. In fact, modern economy has often been described as knowledge-based, or as a learning economy, due to the central role that knowledge and learning play for economic development (OECD, 1996).

The aim of this paper is to present an original contribution to the debate on informal learning using a laboratory experiment designed to reproduce the complex learning dynamics that take place when people exchange knowledge by means of face-to-face interaction. The experiment is based upon the conceptual model of knowledge diffusion specified in our earlier work (Morone and Taylor, 2004a) and presented in section 3.1. In section four, the experimental findings are investigated by comparing these results with the outputs of agent-based simulation models calibrated upon the experimental set-up. What we are trying to achieve is to reproduce the core dynamics of the experiment in terms of the global behaviour of the system, whilst not in any way claiming that the underlying agent motivations are the same. The final part of this paper then further explores the behaviour of the model by simulating over much larger range of parameter settings than would be possible with laboratory methods, thereby extending the analysis of the influence of network factors upon knowledge diffusion patterns and contributing to the debate on informal knowledge diffusion.

## **2. LITERACY REVIEW**

The mechanisms which dominate informal processes of knowledge diffusion have been increasingly investigated by evolutionary economists as well as by game theorists and applied economists. Mechanisms of innovation diffusion (Clark, 1984 and Rogers, 1995) are often viewed as good examples of informal learning processes

because they tend to occur through interaction within geographical and other informal networks, involving social externalities. Several researchers have investigated the patterns through which different agents adopt new technologies by means of theoretical as well as simulation models. (Ellison and Fundenberg, 1993, 1995; Bala and Goyal, 1995, 1998). Another common way of modelling the mechanisms of social learning and technology diffusion makes use of evolutionary game theory (Chwe, 2000; Ellison, 1993, 2000; Anderlini and Ianni 1996; Berninghaus and Schwalbe, 1996; Goyal, 1996; Akerlof, 1997; Watts, 2001).

Along with the speed of new technologies' diffusion, several researchers have focused on the impact of peers' behaviour upon individual decisions in areas such as propensity to crime, use of drugs, school dropout and school attainments (Brock and Durlauf, 1995; Bénabou, 1993; Durlauf, 1996; Gleaser, Sacerdote and Scheinkman, 1996).<sup>1</sup> What all the studies considered so far have in common is their reliance on the idea that learning from neighbours occurs and that under certain conditions it leads to the desirable stable equilibrium. However, none of these studies go beyond a binary definition of learning.

Jovanovic and Rob (1989) proposed for the first time a model in which incremental improvements in knowledge were defined as a complex process of assembling different ideas by means of information exchange among heterogeneous agents. The new insight brought by the authors is that knowledge was defined as something more complex than a binary variable and that, therefore, growth of knowledge could be defined as an interactive process tightly linked to its diffusion.

Cowan and Jonard (1999) made a subsequent attempt to study the effects of incremental innovations and their diffusion within a network of heterogeneous agents. Knowledge in their model is considered as a vector of values and is exchanged via a simple process of barter exchange. Depending on the network structure, the authors found that there is a trade-off between the speed of knowledge diffusion and the variance of knowledge. In other words, there is a spectrum of states of the world, varying from a situation of high knowledge inequality and fast knowledge diffusion (i.e. small-world), to the opposed situation, more equal in terms of knowledge variance but less efficient in terms of knowledge diffusion.

---

<sup>1</sup> For a more detailed review see Morone and Taylor, 2004b.

Along the lines of these works, Morone and Taylor (2004b) defined a model in which agents exchanged knowledge exclusively by means of face-to-face interactions. The network structure was endogenous to the model and could vary over time. The authors showed how small-world networks emerged and coexisted with both a very unequal and a very equal diffusion of knowledge, different outcomes depending upon the initial conditions.

As already mentioned, the objective of this paper is to shed some light on informal learning by means of a laboratory experiment based upon this model of knowledge diffusion. Then, the characteristics of the model will be thoroughly investigated through agent-based simulation. Moreover the simulation model will be calibrated using experimental findings, aiming to reproduce the core dynamics of the experiment in terms of the global behaviour of the system. Through analysis of experimental data we identify learning strategies followed by the experimental subjects, and use this knowledge to inform agent design in the simulation model<sup>2</sup>.

### **3. THE MODEL**

The model presented in this paper is based on the above revised literature and aims at bringing, within this intellectual framework, new insights on the features which shape learning patterns. We shall argue that there are three fundamental factors which influence the speed and the distribution of knowledge diffusion within any closed network. Namely these factors are: (1) the learning strategies adopted by heterogeneous agents; (2) the network architecture within which the interaction takes place; (3) the geographical distribution of agents and their relative initial levels of knowledge.

In what follows we will single out the impact of each of these factors on learning dynamics by testing the model by means of a laboratory experiment. Subsequently, the experimental results will be compared with the outcomes of simulation modelling, using the same network sizes, architectures, geographical distributions, and initial levels of knowledge. In doing so we aim at replicating the experimental learning dynamics by means of simple behavioural rules followed by artificial agents.

---

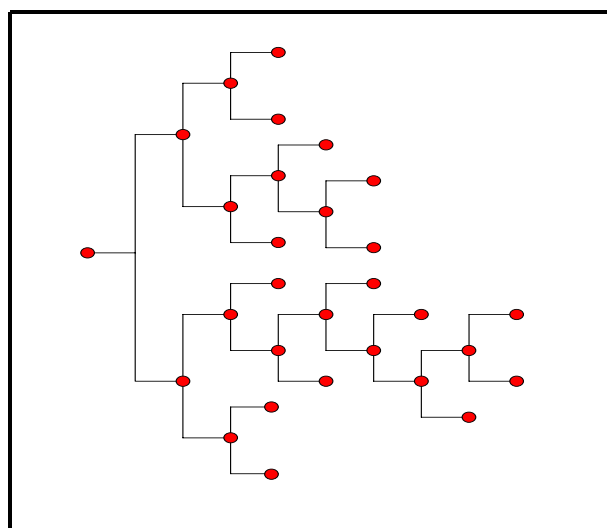
<sup>2</sup> We are not in any way claiming that the underlying agent motivations are the same as the experimental case.

The first objective of the simulation modelling therefore, is to test a number of artificial learning strategies to see how they compare with the experiments in terms of the efficiency and equality of diffusion. The second objective is to further explore the model by considering much larger multi-agent systems, and analyse these over hundreds of repetitions to remove any unwanted artefacts associated with particular configurations or initialisation routines. We will now discuss more in detail the overall framework of the model, the experimental setting and the simulation methodology.

### 3.1 *Defining Complex Learning Dynamics*

This paper is grounded on the idea that learning is a complex process (Jovanovic and Rob, 1989; Ancori *et al.*, 2000). More precisely, the theoretical foundation of the experiment relies on the works of Morone and Taylor (2004a,b), where agents exchange knowledge by means of face-to-face interactions, and every time there is a knowledge transfer, the knowledge is mastered through a backward process by which it is confronted and articulated with previous knowledge. This learning process is achieved through the introduction of a nonlinear complex cognitive structure. Basically, each agent is endowed with a cognitive map that resembles the structure of a tree in which each vertex (node) represents a piece of crude knowledge, and each edge (link) represents knowledge that agents have already mastered and learned how to use (see Figure 1).

**Figure 1.** Cognitive Map



Moving from left to right in the cognitive map, we advance from less to more specialised knowledge. This observation justifies the assumption that new nodes can only be activated if they are directly connected to active nodes. In other words, new knowledge can only be acquired if it can be integrated with the existing, accumulated knowledge in order to be used.

The model is stepped forward through simulated time, where each discrete time-step is known as a ‘cycle’. From one cycle to the next agent states may change according to rules specified in the model. These rules are executed simultaneously for each member of the agent population. Interaction amongst agents is based on the transmission of knowledge. Every cycle each agent has got the opportunity to initiate a learning activity by contacting one of her/his acquaintances in the social network (see section 3.3 for description of the network upon which interactions take place) as follows. An agent, whom we shall call A, contacts an acquaintance, B, from her/his list of acquaintances. She/he then selects a node of her/his cognitive map which she/he would like to learn (of course this node has to be pegged to already activated nodes). If the contacted agent possesses the bit of knowledge required, the player that initiated the contact acquires it and activates it in her/his cognitive map. Otherwise, he/she will remain at the former level of knowledge.

### **3.2 *Initialising the Cognitive Maps***

While constructing agents’ cognitive maps, it was assumed that everybody knew the first node (i.e. was endowed with the most elementary level of knowledge). The nodes in the second column were assigned subject to a specialisation process (i.e. each agent could specialise either in scientific knowledge, knowing the upper part of the tree, or in technical knowledge, knowing the lower part of the tree). The initial CM configuration was constructed as follows. Firstly, an integer indicating the depth of knowledge was selected at random from the interval between 1 and  $N_{cols}$ , where  $N_{cols}$  is the total number of columns in the CM. In all experiments reported here, we used a value of 6 for the integer  $N_{cols}$ .

Subsequently, another integer was chosen to define the kind (i.e. specialisation) of the knowledge possessed by each agent. This number was either a 1, indicating technical knowledge, or a 2, indicating scientific knowledge.

If the first integer took the maximum value  $N_{cols}$ , then the whole map was activated in the corresponding part of the tree. Otherwise, the nodes in all the columns

up to and including the column indexed by the first integer value were activated in the region of specialisation defined by the second integer.

Given the technique used to assign the initial level knowledge we could calculate an agent's initial probability of knowing each node of the CM as follows: each agent will know node 1 with probability 1, node 2 (3) with probability  $5/12$ . If he/she knows node 2 (3) then could know nodes 4 and 5 (6 and 7) in the third row. The probability of knowing these nodes is set equal to  $4/12$ . Following this reasoning every agent will know nodes 8, 9, 10 and 11 (12, 13, 14 and 15) with probability  $3/12$ ; nodes 16-23 (24-31) with probability  $2/12$  and nodes 32-47 (48-63) with probability  $1/12$ .

A consequence of this particular structure of the CM is that we can easily derive subjects' best node selection strategy. Each (rational) agent would try to learn the node with the highest probability of being known by any of his/her acquaintances. As observed above, scientific and technical knowledge, in the same column, have the same associated probabilities, so initially (i.e. at cycle 0) agents are indifferent between learning a node allocated in the upper or lower part of the tree. On the contrary, nodes allocated in different columns will have different probabilities, more precisely each node which appears in column 2 has a higher probability of each node in columns 3, 4, 5 and 6. Similarly nodes in column 3 have a higher associated probability than those in columns 4, 5 and 6 but a lower probability than nodes in column 2. Following this line of reasoning we can conclude that at cycle 0 each (rational) agent will try to learn one of the node positioned in the lowest empty (or partially empty) column. If this is the best strategy for all agents at cycle 0, this will raise the probabilities of the basic nodes of being known at cycle 1. Therefore also at cycle 1 the best learning strategy will be to learn one of the nodes positioned in the lowest empty (or partially empty) column. This reasoning can be extended to every subsequent cycle leading us to the conclusion that learning vertically will be a dominant strategy<sup>3</sup> for all agents and so it will lead to a Nash equilibrium.

However, this analysis does not take into consideration the fact that much of this information about the structure of the game (i.e. how the *CMs* are initialised) is not available to the agents. Furthermore, in our model players do not know the absolute levels of knowledge of others. The existence of different opportunities to learning

---

<sup>3</sup> Note that this is also a stable strategy because subject will not be influenced by the timing.

through acquaintance selection - contingent upon the geographical location in the network - significantly complicates the task of solving to find the most efficient strategy. It is reasonable to assume that the subjects may be able to surmise some of this through experience of playing the game. But it is not clear what is the ‘rational’ strategy that would lead to the best outcome in terms of individual gains. Furthermore, there are a very large number of potential strategies (for acquaintance selection) which would need to be tested in this kind of analysis.<sup>4</sup>

### 3.3 *Setting of the Experiment*

The aim of the “game” for each player is to increase his/her own level of knowledge. The cognitive map assigned to players resembles the graphical representation reported in figure 1. In figure 2 we reproduce a demonstrative cognitive map that each player might see displayed on his/her screen. Each tab is labelled between 1 and 63 and represents a node that can be acquired. Yes/No values indicate whether or not the tab is activated, i.e. whether that particular piece of knowledge is at present in the player’s cognitive map. As already discussed, in order to activate a new tab (i.e. to learn a new bit of knowledge) the acquired knowledge has to be linked to an already activated tab in the cognitive map.

The game is simultaneous and composed of  $N$  (=100) periods. The networks upon which interactions take place are constructed following the method first used by Watts and Strogatz (1998). Each agent is initially assigned a random, unique position in a one-dimensional wrapped grid (i.e. a ring). The social network is then created by connecting an agent with all other agents located within her/his *neighbourhood*. Social neighbourhoods are defined as the region on the grid that includes the adjacent cells falling within the agent’s visible range. We therefore specify a *regular network* shaped as a ring lattice within which each agent interacts only with a number ( $n$ ) of nearest neighbours.

---

<sup>4</sup> This should not surprise as it is in the nature of Evolutionary Economics a degree of complexity which suggests a close proximity to agent-based simulation techniques as a tool for investigation. “While neoclassical theory describes with precision and rigour a simple world that apparently does not exist” (Dopfer 2004), “evolutionary economics tries to analyse a complex reality of economic

**Figure 2.** Experimental definition of the cognitive map

Departing from this initial network structure we then constructed two other networks following the *re-wiring procedure* first introduced by Watts and Strogatz (1998). With a certain probability  $p$ , we allowed any individual to disconnect him/herself from the local connection and to re-connect that edge to a vertex chosen uniform randomly over the entire lattice. As the value of  $p$  changes, the structure of the network changes. By choosing a fairly small value of  $p$ , the lattice remains almost regular and highly clustered (i.e. with high degree of cliquishness), 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” (Cowan and Jonard, 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.

As suggested in the literature (Watts and Strogatz, 1998; Watts 1998) there is a region of  $p$  ( $0.01 \leq p \leq 0.1$ ) within which the network structure is that of a *small*

---

phenomena deeply intertwined with cognitive, institutional, organisational and political dimensions, i.e. a world of variety full of *complexity* and *informational fuzziness*” (Pyka and Ahrweiler, 2004).

*world*: i.e. one in which a short average path length ( $L$ ) and high degree of cliquishness ( $C$ ) coexist. This contrasts with the two extreme cases of  $p=0$  and  $p=1$  where the network structure becomes respectively that of a *regular network* and that of a *random network*. We defined our two new networks architectures by setting the  $p$  parameter to  $0.1^5$  and  $1$ . The experimental networks were then defined as follow: regular network ( $p=0$ ), small world network ( $p=0.1$ ) and random network ( $p=1$ ).

The huge appeal of small-world networks lies in the impact they are said to have on dynamical systems. Watts and Strogatz (1998) maintain the point that “models of dynamical systems with small-world coupling display enhanced signal propagation speed, computational power, and synchronizability.” Furthermore, the authors showed, using simplified model for the spread of an infectious disease, that “infectious diseases are predicted to spread much more easily and quickly in a small world” than in regular lattices (1998: 442). These findings have profound implications for many manmade and natural systems. For instance, in a transportation network, small-world architecture could improve the flow of people or goods through the network. In a network within which people are exchanging knowledge by means of face-to-face interactions, small-world connectivity might improve the ease with which ideas and knowledge diffuse through the system.

#### **4. EMPIRICAL FINDINGS: COMPARING EXPERIMENTAL AND SIMULATION RESULTS**

In this section we will present our experimental findings, we will then compare those results with simulation outcomes in order to investigate artificial learning strategies. By basing the agent strategies upon those identified through analysis of the experimental data we shall develop a simulated strategy able to replicate the learning patterns observed in the laboratory. In this way we will develop a sort of *history friendly* model (which we should call *experimentally-friendly* model), following the methodological approach first developed by evolutionary economists like Nelson and Winter (1982), Silverberg et al. (1988) Dosi et al. (1995).

---

<sup>5</sup> Given the small number of agents composing our network (this constrain was due to the experimental nature of the work) we chose the highest possible value to the rewiring probability to generate a small world network. We tested the network structure following the Watts and Strogatz (1998) criteria for identifying the network as small worlds:  $\mathcal{L}(t) \cong \mathcal{L}_{random}(t)$  and  $C(t) \gg C_{random}(t)$ .

After investigating the impact of different learning strategies upon knowledge flows, we will turn our attention to the specific role played by the network architecture. By means of simulation analysis we shall single out the impact of different network structures upon the learning dynamics and knowledge dispersion, comparing the performances of small world networks, regular networks and random networks. We will finally turn our attention to the role played by the geographical distribution of knowledge, pointing out the importance of ‘access to knowledge’ and ‘equal learning opportunities’ as an inequality-decreasing device.

#### ***4.1 Studying learning strategies***

The experimental study was carried out with a network of  $N=14$  players interacting in the experiment for a total of 100 cycles. On average to each agent were assigned four acquaintances. The experiment was initially conducted under the small world network architecture (the experiment was subsequently replicated with different network structures as it will be discussed in the following sections).

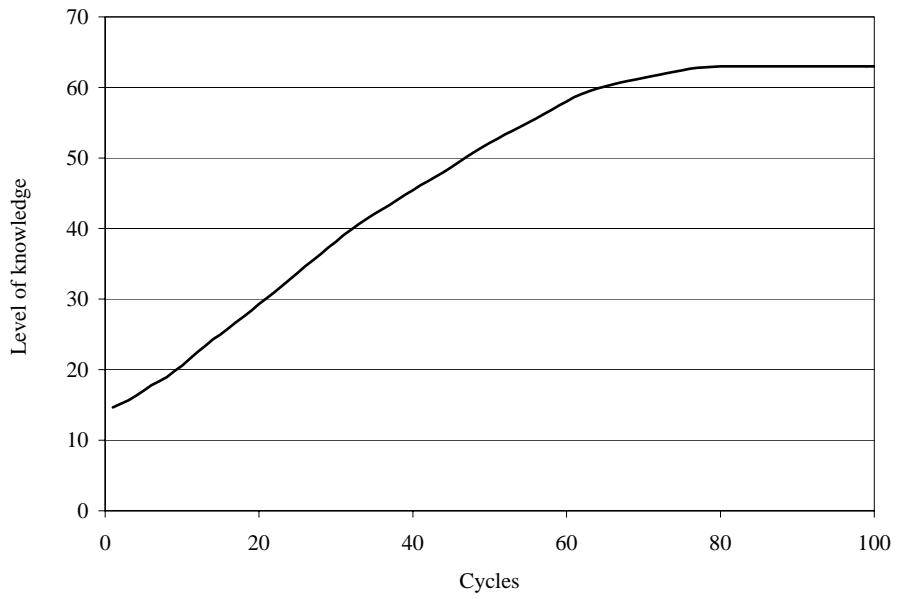
Small world phenomena are regarded by some researchers to only be worthy of investigation in the case of large, sparse networks. However, under experimental conditions it is not practical to undertake studies involving hundreds of participants. The largest laboratory available to us has the capacity to accommodate a maximum of forty subjects, however we decided to carry out an initial round of experiments with a much reduced number of subjects in order to expedite the research<sup>6</sup> and to test the assumption that large networks are required for small world analysis to have value.

As mentioned earlier, every cycle each player had the opportunity to contact one neighbour and ask for some knowledge. Independently on the network structure the experimental agents reached always the steady state within the time frame of the experiment. In figures 3 and 4 we present the learning dynamics (both in terms of learning speed and knowledge dispersion) for the first experiment.

---

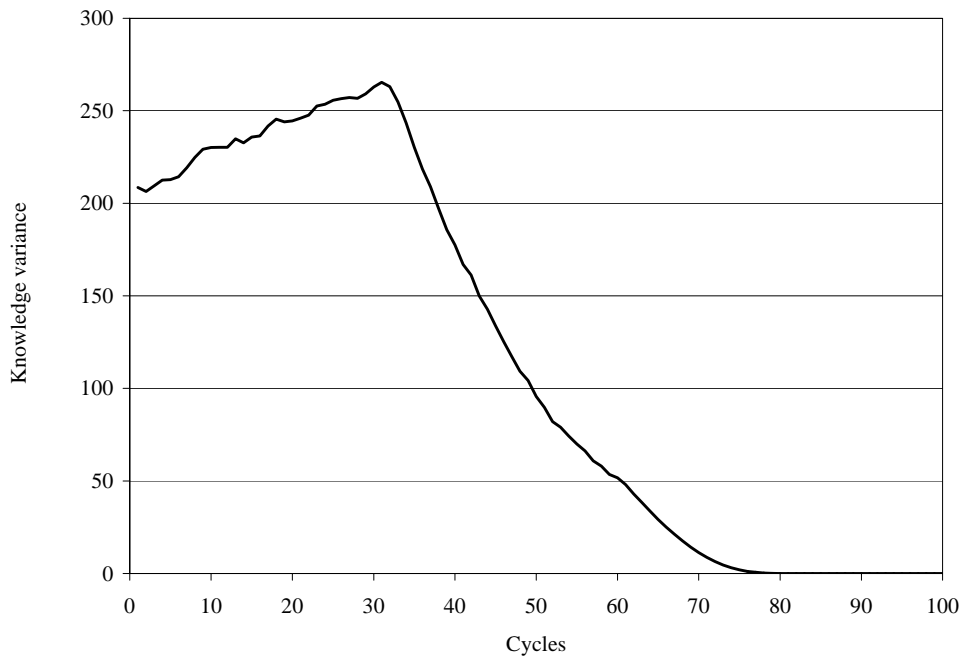
<sup>6</sup> We are in negotiation to scale up this investigation.

**Figure 3.** Small world network, learning dynamic  
(experimental results)



**Source:** Experiment results

**Figure 4.** Small world network, variance in learning dynamic  
(experimental results)



**Source:** Experiment results

As a first step we shall focus our analysis on the overall picture. As we can see all 14 players converge to the steady state in a time frame of less than 80 cycles, gaining on average 38 nodes in 79 cycles (hence learning almost at an average speed of 0.5 nodes per cycle). Moreover, the knowledge variance grows in the short term, showing an initially unequal learning process. After the first 30 cycles it starts converging towards the zero-variance steady state.

In order to elicit the learning strategies followed by experimental agents we both studied the cognitive maps of each player as well as used questionnaires asking to state how they decided the node to be learned and the acquaintance to contact. For the former we allowed a choice of three possibilities: random strategy, width-first strategy or depth-first strategy (we also provided an ‘other’ checkbox). As discussed in section 3.1, the depth-first strategy would entail a specialised learning process, whereas a width-first strategy would entail the acquisition of less specialised knowledge first.

**Figure 5.** Cognitive Map and learning strategies

1	2	4	8	16	32
				17	33
					34
			18	35	
				36	
	5	9	19	37	
			20	38	
			21	39	
		10	22	40	
			23	41	
	3	6	11	24	42
				25	43
				26	44
				27	45
				28	46
7		12	29	47	
			30	48	
			31	49	
		13	50		
			51		
14	52				
	53				
	54				
15	55				
	56				
	57				
			58		
			59		
			60		
			61		
			62		
			63		

An example could be useful to clarify the different learning strategies. Let us assume that a subject has the following cognitive map: {1, 2, 3, 4, 5} (see figure 5). In the first cycle she/he can choose the new bit of knowledge to be learned from a set of six possible nodes: {6, 7, 8, 9, 10, 11}.

We shall maintain that if she/he asks one of the following nodes {6, 7}, then she/he adopts a vertical strategy (hence, learning the least specialised knowledge first). On the contrary, if she/he tries to learn {8, 9, 10, 11}, then she/he adopts a horizontal strategy (hence, she/he is specialising).

Let us assume now that in the first cycle our subject asked and learned node 6; then, in cycle two her/his cognitive map will be updated as follows {1, 2, 3, 4, 5, 6} (see figure 6 below).

**Figure 6.** Cognitive Map and learning strategies

1	2	4	8	16	32	
			5	9	17	33
				6	10	18
		11			19	35
		3	6	12	20	36
				13	21	37
	7			14	22	38
				15	23	39
	7		7	14	24	40
					25	41
					26	42
	7	7	14	27	43	
				28	44	
29				45		
30				46		
31				47		
						48
			49			
			50			
			51			
			52			
			53			
			54			
			55			
			56			
			57			
			58			
			59			
			60			
			61			
			62			
			63			

In the second cycle she/he can try to learn one of the following nodes {7, 8, 9, 10, 11, 12, 13}. If she/he will try to acquire node 7, then she/he will be following again a vertical strategy. On the other hand, if she/he will ask one of the other possible nodes {8, 9, 10, 11, 12, 13}, then it will be argued that she/he is adopting a horizontal strategy.

Finally, we shall note that if our subject 's strategy shows some kind of 'time inconsistency' (i.e. it changes over time), then we will classify her/him as a *random* strategy follower.

It is interesting to compare the results obtained studying the learning strategy actually adopted by experimental players with the results obtained through a questionnaire where players were asked to declare the learning strategy they adopted.

**Table 1.** Different node selection strategies

	Questionnaires		CM investigation	
	share	percentage	share	percentage
<i>width-first strategy</i>	10	27.0%	24	57.1%
<i>depth-first strategy</i>	2	5.4%	5	11.9%
<i>random strategy</i>	25	67.6%	13	31.0%

**Source:** Experiment results

Interestingly enough, the figures obtained through direct investigation of players' *CMs* seem to support rather strongly the hypothesis that subjects try to learn in a vertical way: 57% of the subjects have always adopted a vertical strategy whereas only 12% of the players have consistently adopted a horizontal one. Confronting these results with those obtained through our questionnaire shows that most of the players who adopted a width-first strategy were unable to categorise their behaviours as such. Hence, we can conclude that they were unconsciously following a particular strategy.

In part based on this analysis we designed the artificial learning strategies for both node and acquaintance selection. These are summarised in table 2. Comparing the results of the questionnaires and the analytic study of *CMs* we have defined three strategies for acquaintance selection: "random" or zero-intelligence (ZI), intelligent "sequential" and one based on the adaptation of the relative "strength" of connections, which was shown to be efficient in previous works (Morone and Taylor, 2004a,b).

**Table 2.** Possible Learning Strategies

Node selection strategy	Acquaintance selection strategy
random strategy	random strategy
width first strategy	sequential strategy
depth first strategy	preferential model strategy

The sequential strategy requires some further clarification. Following this strategy, in the first cycle the agent would select an acquaintance to contact at random. Depending on the outcome of this interaction, the agent would select the same acquaintance in the following cycle if the interaction was successful. If the previous interaction was unsuccessful however, the agent would revert to random selection from among acquaintances.

We therefore have three possible strategies for each option, making a total of nine possible *combined strategies*. The following step was to set up a simulation that would provide results that could satisfactorily<sup>7</sup> replicate the experimental dynamics.

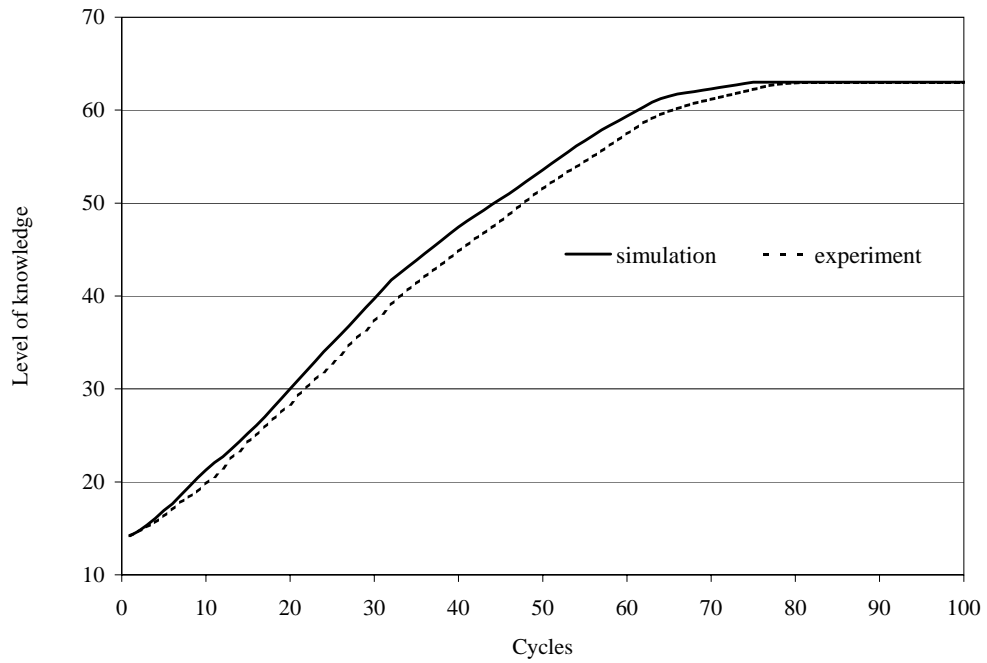
In testing the various combinations we started first with a population of zero-intelligence agents behaving randomly in respect of both action selections. This produced a very poor result in terms of efficiency of learning. Our second trial was to specify a strategy based on the heuristic of adaptive ‘sequential’ selection combined with width first node selection. As it emerges from figures 5 and 6 we were able to replicate the experimental results in a satisfactory way adopting a *width-first sequential strategy*. Both the mean and variance simulation series were statistically not different from the respective series obtained in the laboratory.

Clearly it was not a difficult task to design artificial learning strategies resulting in a very similar overall system performance in terms of knowledge flows. Whereas the width first strategy might be regarded as close to optimal (in the sense of a Nash equilibrium) the sequential strategy almost certainly is not (because it is based only on a memory of one cycle, and because behaviour is ZI after unsuccessful interactions). However, in order to answer this question definitively we need to explore a larger range of available strategies. A first step towards this would involve tabling the outcomes of the above 9 combined strategies, under different network configurations, in order to compare them in terms of knowledge flows and make comparisons with the experimental case.

---

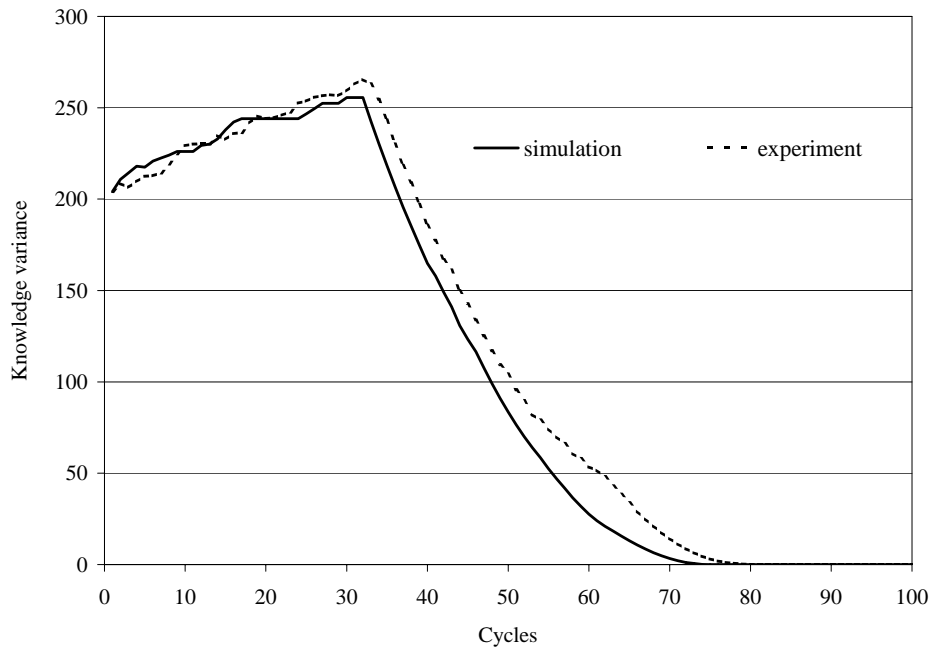
<sup>7</sup> By ‘satisfactorily’ we mean that the difference of the mean of the two series (i.e. experimental data and simulation data) is not statistically significant. This was tested by mince of a t-test.

**Figure 5.** Small world network, learning dynamic comparison  
(experimental results vs. simulation results)



Source: Experiment and simulation results

**Figure 6.** Small world network, variance in learning dynamic comparison  
(experimental results vs. simulation results)

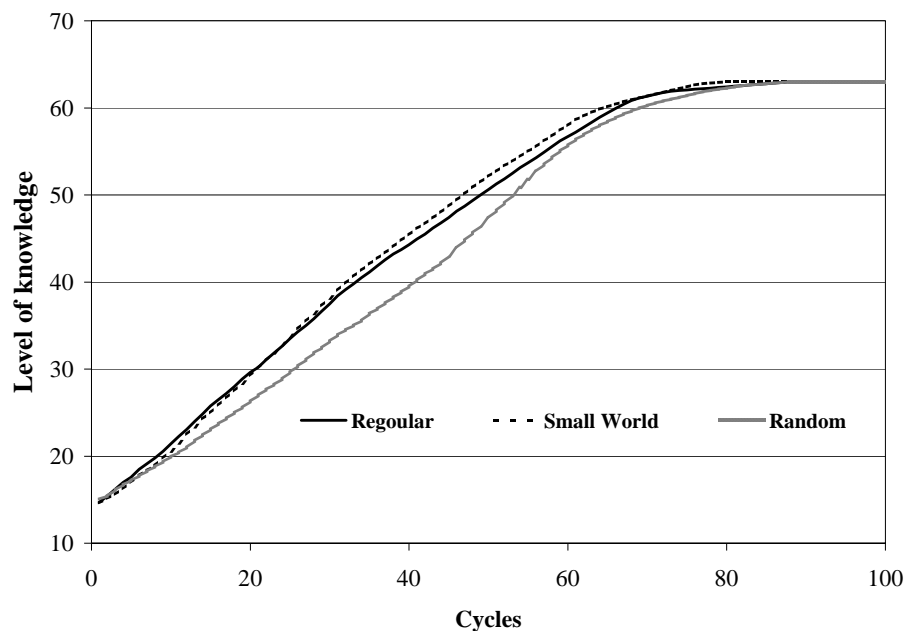


Source: Experiment and simulation results

#### 4.2 Studying network structure and geographical distribution

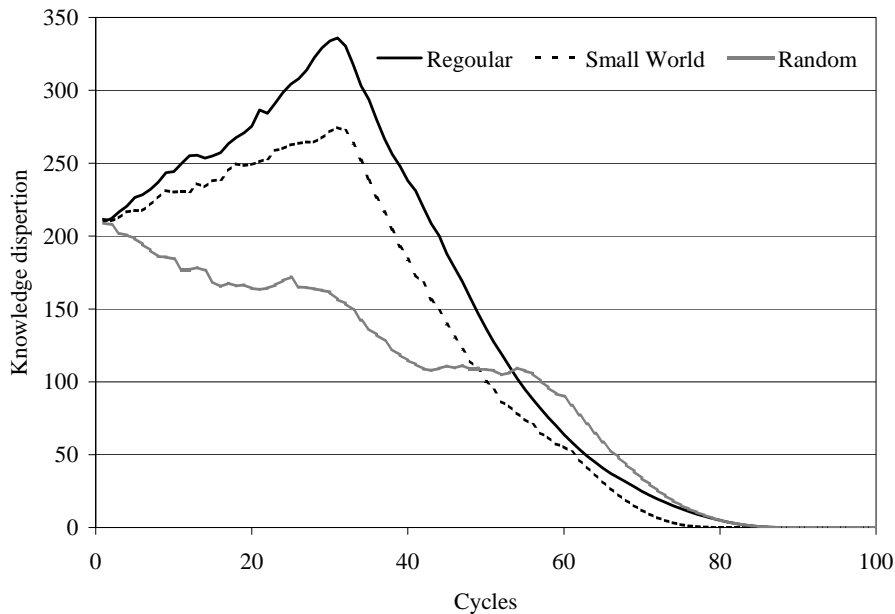
We can now start looking at the role played by network structure in the learning process. In order to do so we replicate the experiment using the same agents (endowed with the same cognitive maps) but allocating them on different network architectures. Tuning the  $p$  value as described in section 3.2 we constructed a random network ( $p=1$ ) and a regular network ( $p=0$ ). We then compared these results with those obtained in the first experiment. Looking at the learning pace (in figure 7) we could immediately observe that overall the three models performed quite similarly: in each case the system converged to a steady state within the first 90 cycles. Moreover, the small world network over-performed when compared with the other two networks. We also carried out simulations for the regular and random networks and then compared the three cases.

**Figure 7.** Different networks, learning dynamic comparison  
(experimental results)



Source: Experiment results

**Figure 8.** Different networks, variance in learning dynamic comparison  
(experimental results)



**Source:** Experiment results

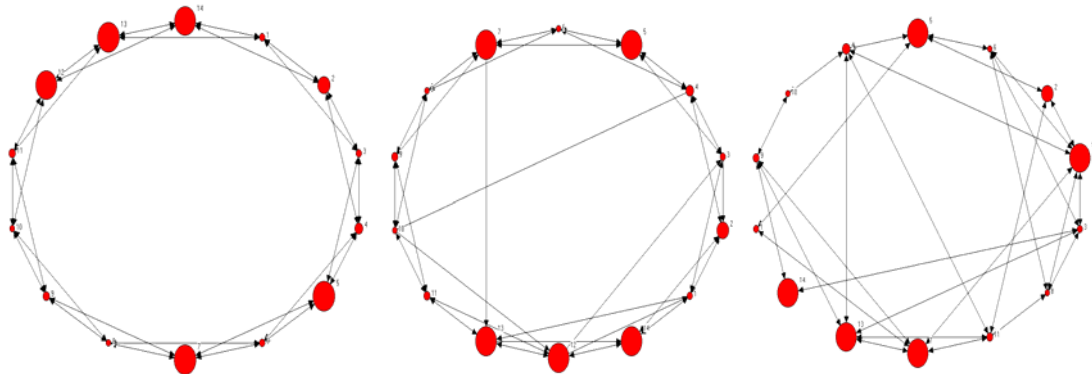
Interestingly enough the worst-performing system was the random network. This result appeared to be counterintuitive since the random network was also the system with the shortest average path length and therefore was expected to perform better than the regular network in terms of speed of knowledge diffusion. Looking more closely into the experimental results we observed that the random network was also the system performing better in terms of knowledge dispersion, displaying a short term converging pattern as opposed to small world and regular system (figure 8).

The answer to this apparently odd result lies in the different geographical distribution of agents obtained in the random network when with re-wiring probability  $p=1$  each original connection was interrupted to be reconnected randomly with another node in the graph. In fact, looking at the three network's structures we could clearly see how the most knowledgeable agents tend to be clustered together in the small world and in the regular network, whereas the same agents are quite disconnected in the random network. In light of these observations we could expect to see knowledgeable agents learning faster in the first two networks as opposed to the random one. This observation was corroborated when looking at individual learning patterns: the learning ability of the most knowledgeable agents in the random network was far less efficient than that of the same agents in the other two network structures.

This gap in the performance was due to the fewer learning opportunities available to knowledgeable agents in the random network.

**Figure 9.** Experiment network structures

Regular network, Small world network and Random network



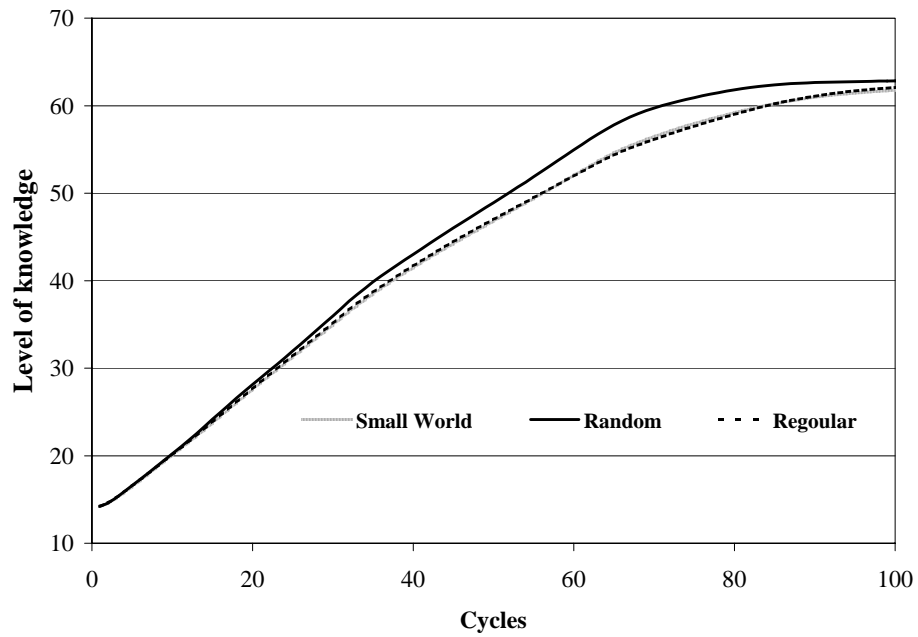
**Source:** Experiment and simulation results

This result leads us to conclude that learning dynamics are not just affected by the network structure and the learning strategy, but also (and perhaps mainly) by the *learning opportunities* provided to different agents in the network.

In order to test independently the effect upon learning dynamics of the network structure and of the geographical distribution of agents we run batches of 100 simulations for each network's structure always reallocating the agents in different ways. Then, we compute the average performance of each network hence clearing out the geographical effect.

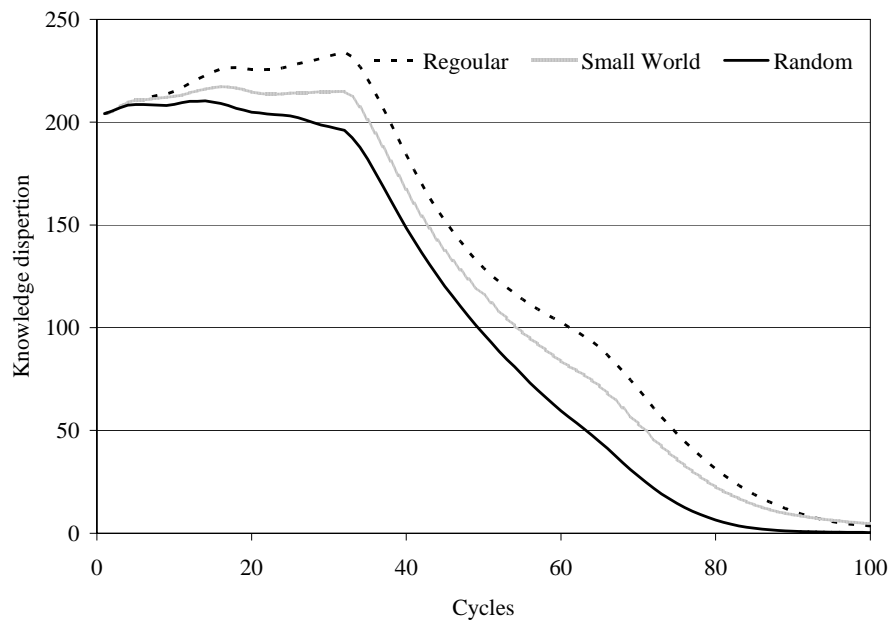
The results obtained in this way show, in fact, a different picture: the random network structure performs no longer the worst in terms of speed of learning, but actually appears to be the most efficient network. On the other hand, the relative performance of small world and regular network is very similar in terms of speed of convergence toward the long-run steady state.

**Figure 10.** Different networks, learning dynamic comparison  
(average over batches of simulation results)



**Source:** Simulation results

**Figure 11.** Different networks, variance in learning dynamic comparison  
(average over batches of simulation results)



**Source:** Simulation results

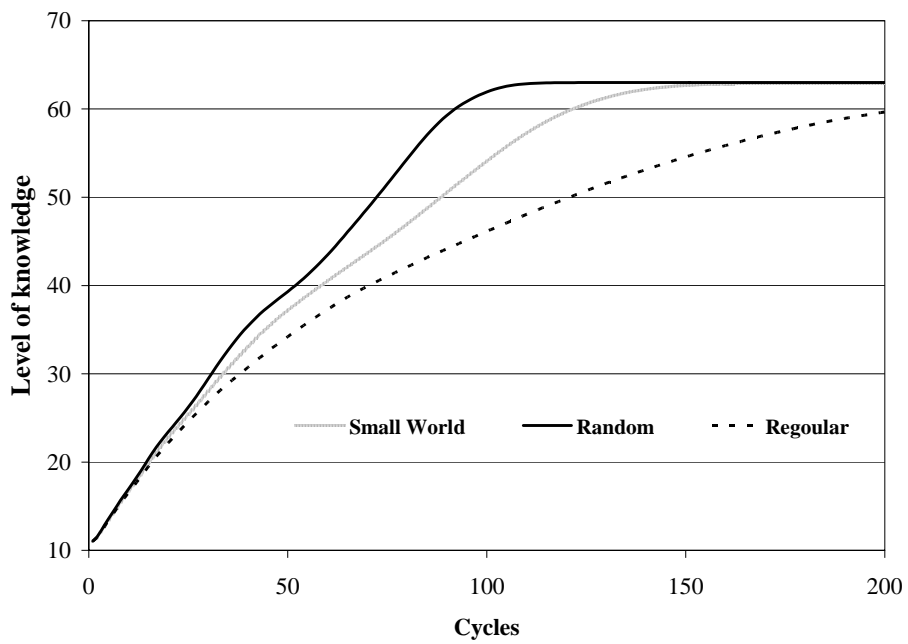
From a distributional point of view the random structure shows almost a constant decrease in the knowledge variance in the short-term transition towards the natural equilibrium of the model. This contrasts with small world and regular architectures, which display a general inequality increase over the short run.

What emerges from this analysis is that small world is not the most efficient nor the most unequal network as previously observed in the literature on knowledge diffusion (Cowan and Jonard, 1999). In our laboratory experiments, as well as in our simulations, the small world network was not the most unequal, and when we corrected for the particularities of the geographical distribution used in the laboratory experiments, we obtained the result, through simulations, that the random network was the most efficient.

A possible explanation for these contrasting results could be the different network size used in different simulations. In fact, our experiment, and subsequently our simulations, were based on a rather small network (i.e. 14 agents) compared to that of Cowan and Jonard (500 agents). According to Watts (1999) small world phenomena can only be observed in large and sparse systems. We can certainly say that in such small networks, the regular and the random would tend to look more similar in terms of sharing small  $L$  and high  $C$  values, but on the basis of this fact they should not be precluded, in our view, from the study of knowledge diffusion patterns within the small worlds framework. In fact we have found it useful to study small networks in the laboratory, since it has allowed us to easily trace the paths of agents' decision making and learning processes. Furthermore, these experiments, as well as the subsequent simulations, have shown that significant differences in outcomes can be attributed to changes in the network configuration, investigated here by consideration of the theory of small world networks.

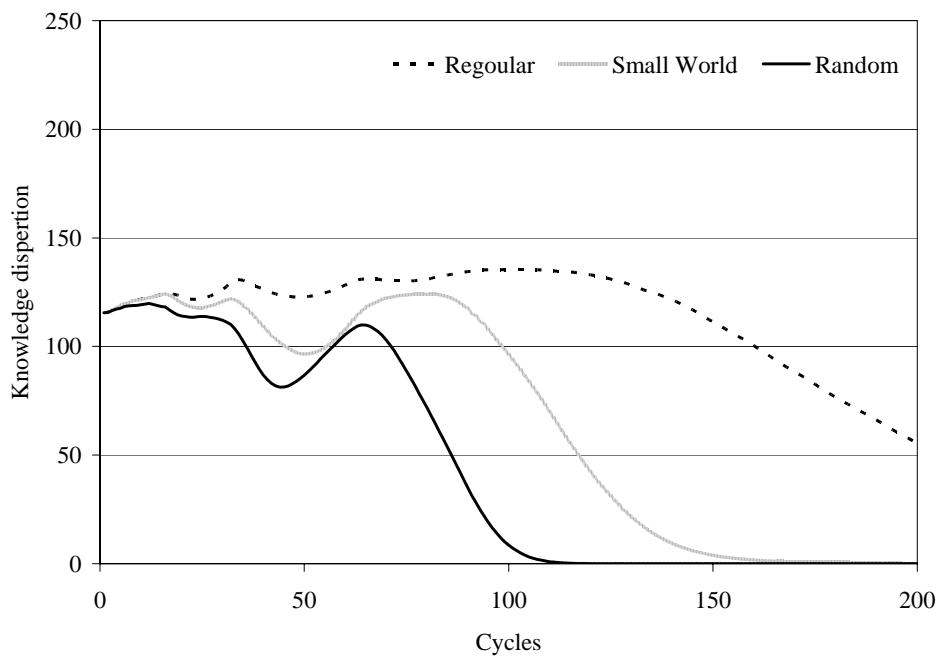
However, we wanted to investigate the impact of system scale upon this result. In order to test our new hypothesis, that *small world networks are neither the most efficient nor the most unequal*, we ran new batches of simulations, with subsequent increases in the size of the network. Nonetheless, much bigger networks (i.e. 100 and 500 agents) displayed similar results to those obtained with 14 agents: random networks being consistently the best performing in terms of speed of knowledge diffusion as well as in terms of knowledge distribution, followed by small world networks and finally by regular networks.

**Figure 12.** Different networks, learning dynamic comparison  
(100 agents' network, average over batches of simulation results)



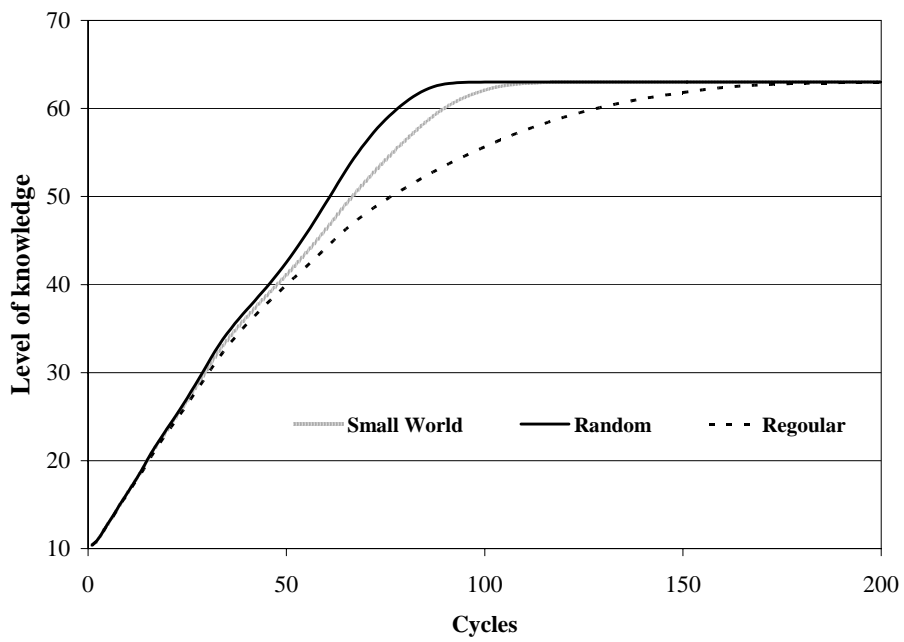
Source: Simulation results

**Figure 13.** Different networks, variance in learning dynamic comparison  
(100 agents' network, average over batches of simulation results)



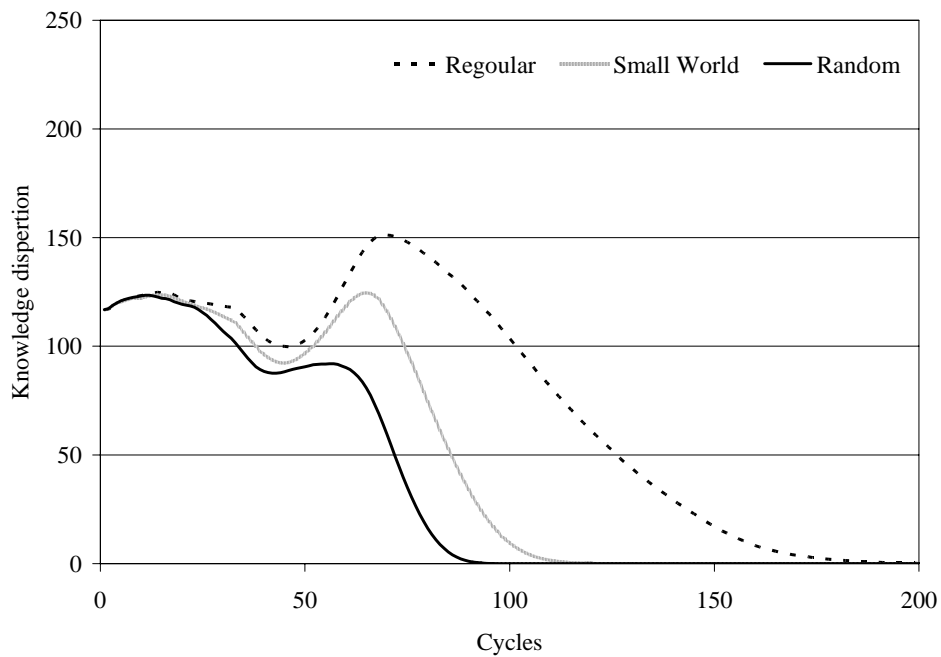
Source: Simulation results

**Figure 14.** Different networks, learning dynamic comparison  
(500 agents' network, average over batches of simulation results)



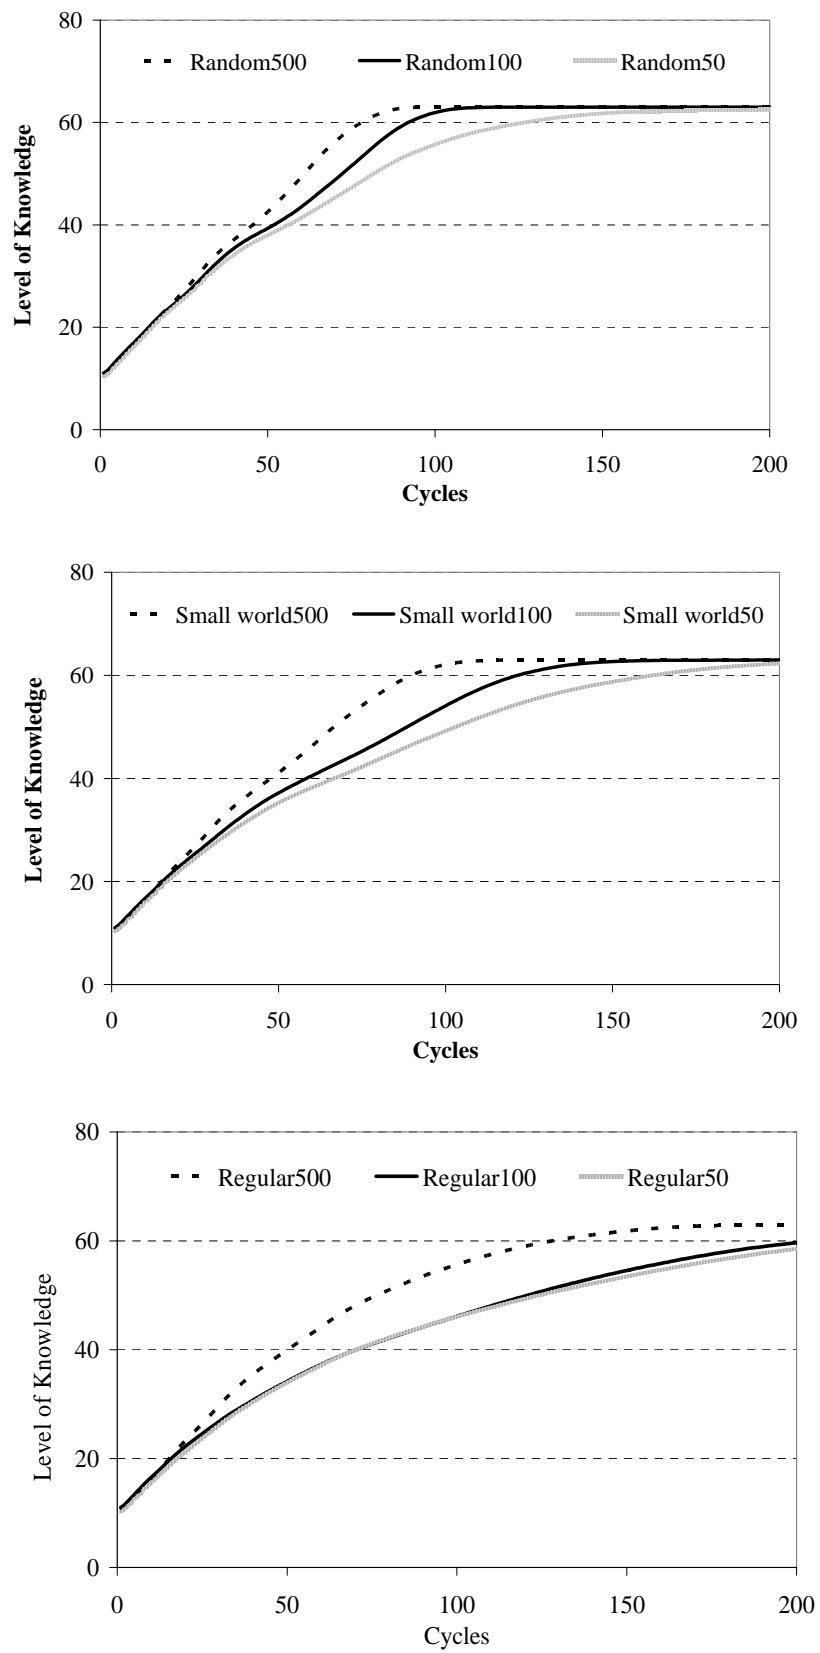
Source: Simulation results

**Figure 15.** Different networks, variance in learning dynamic comparison  
(500 agents' network, average over batches of simulation results)



Source: Simulation results

**Figure 16.** Comparing size of networks and speed of convergence (500, 100 and 50 agents' networks, average over batches of simulation results)



**Source:** Simulation results

Looking at figures 12 through 15 an interesting observation can be drawn: the larger the network the faster the converging process (i.e. the faster are knowledge flows). This result is quite interesting since it shows that, other things being equal, the system scale is a key determinant of relative speed of knowledge diffusion, larger networks being more efficient. This result is confirmed for each network architecture as figure 16 below shows.

As far as the variance of convergence is concerned there is no generic rule which can be drawn since small world simulations show a rather similar fluctuating pattern, whereas random and regular networks show fairly fuzzy fluctuating patterns which make any attempt of comparison not really informative.

## 5. CONCLUDING REMARKS

In this paper we investigated the dynamics of knowledge diffusion in a small population of agents by means of a laboratory experiment as well as by a simulation model. As stated in the introduction, the aim of the paper was to present an original contribution to the debate on informal learning, combining a laboratory experiment designed to reproduce complex learning dynamics with a simulation model able to reproduce the core dynamics of the experiment in terms of the global behaviour of the system. Departing from this analysis, we further explored the behaviour of the model by simulating over a much larger range of parameter settings than would be possible with laboratory methods, thereby extending the analysis of the influence of network factors upon knowledge diffusion patterns. While conducting this investigation, several elements arose as key determinants of flows dynamic within a closed network. Namely these factors were: (1) the learning strategies adopted by heterogeneous agents; (2) the network architecture within which the interaction took place; (3) the geographical distribution of agents and their relative initial levels of knowledge; (4) the network size.

Concerning the learning strategy, we were able to identify a set of nine possible *combined strategies* which would allow us to investigate the two actions that each player had to undertake while attempting to acquire knowledge. Any time an agent was trying to learn a new bit of knowledge she/he had to make two choices: first, decide what to learn (this decision was constrained to the set of possible nodes which could be articulated with previous knowledge); second, decide with whom to interact

(also this decision was bounded in the set of acquaintances in the agent's social network). The following step was to identify, through simulation, the *combined learning strategy* that would be sufficient to replicate the experimental dynamics. Carrying out this exercise, we were able to define a combined strategy which closely replicated experimental global behaviour of the system. Moreover, the *best combined learning strategy* was then showed to be the closest thing to an optimal strategy (i.e. one which leads to a Nash equilibrium).

As already mentioned, a second step in our investigation was examining those factors which affect knowledge flows. A preliminary result of our investigation showed that learning dynamics is heavily affected by the *learning opportunities* provided to each agent in the network. By *learning opportunities* we mean the chances each agent has got to interact with more knowledgeable agents. In other words, a particular geographical distribution of agents (endowed with different knowledge) could substantially affect learning dynamics. In order to test independently the effect upon learning dynamics of the network structure and of the geographical distribution of agents, we run batches of simulations for each network's structure, while keeping reallocating the agents in different ways. Then, we computed the average performance of each network hence clearing out the geographical effect. Once corrected for any possible geographical bias we could conclude that small world networks do perform better than regular networks, but consistently underperformed when compared with random networks. This finding contrasts with previous literature (Cowan and Jonard, 1999), which maintained that small world network is the most efficient (as well as the most unequal) system.

At this point in our research we do not know why these results differ so markedly. It suggests to us the need for further investigation into the effect of network configuration using the small world framework. A further result obtained in this investigation, which is to our knowledge an original finding, was that the bigger the network size, the faster the diffusion is. Interestingly enough this result was shown to be independent from the particular network architecture.

In conclusion, investigating the relative effect of each and every of the four factors, we can maintain that studying the nexus between structure and flows is a rather complex task which involves a large number of aspects that concur to define flows dynamics.

As a suggestion for further research we would propose to make an attempt to produce a clear taxonomy of all the factors which might affect knowledge flows that occur in social networks. It could be then interesting to divide these factors into some broad categories. As a pure exercise we could suggest *structural* and *individual*, the first one including all those factors directly referring to the network architecture and size, and the second one referring to individual decisions such as strategy decisions.

## REFERENCES

- Acemoglu, D. and J. Angrist (1999), "How Large Are the Social Returns to Educations? Evidence From Compulsory Schooling Low", *NBER, Working Paper*, 7444.
- Akerlof, G. (1997), "Social Distance and Social Decisions", *Econometrica*, 65: 1005-1027.
- Anderlini, L. and A. Ianni (1996), "Path Dependence and Learning from Neighbours", *Games and Economic Behaviour*, 13: 141-177.
- 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.
- Bala, V. and S. Goyal (1995), "A Theory of Learning with Heterogeneous Agents", *International Economic Review*, 36(2): 303-323.
- Bala, V. and S. Goyal (1998), "Learning from Neighbours", *Review of Economic Studies*, 65: 595-621.
- Bénabou, R. (1993), "Working of a City: Location Education and Production", *Quarterly Journal of Economics*, 58: 619-652.
- Berningaus, S. and U. Schwalbe (1996), "Evolution, Interaction, and Nash Equilibria", *Journal of Economic Behaviours and Organization*, 29: 57-85.
- Borgatti, S.P., M.G. Everett, and L.C. Freeman, (1999), *UCINET 5.0 Version 1.00. Natick: Analytic Technologies*.
- Brock, W. and S. Durlauf (1995), "Discrete Choice with Social InteractionI:Theory" *NBER Working Paper*, 5291.
- Chwe, M. S-Y. (2000), "Communication and Coordination in Social networks", *Review of Economic Studies*, 67: 1-16.
- Clark, G. (1984), *Innovation diffusion: Contemporary geographical approaches*, Geo Books: Norwich.
- Cowan, R. and N. Jonard (1999), "Network Structure and the Diffusion of Knowledge", *MERIT Working Papers*, 99-028.

- Dopfer, K. (2004), *The Evolutionary Foundations of Economics*, Cambridge University Press, 2004, forthcoming.
- Dosi, G., Marsili O., Orsenigo I. and Salvatore R. (1995), "Technological Regimes, Selection and Market Structures", *Small Business Economics*.pp. 411-436.
- Durlauf, S. (1996), "Neighborhood Feedbacks, Endogenous Stratification, and Income Inequality", in: W. Barnett, G. Gandolfo and C. Hillinger (ed), *Dynamic Disequilibrium Modelling*, Cambridge, Cambridge University Press.
- Ellison G. (1993), "Learning, Local Interaction, and Coordination", *Econometrica*, 61: 1047-1071.
- Ellison G. (2000), "Basins and Attraction, Long Run Stochastic Ability and the Speed of Step-by-Step Evolution", *Review of Economic Studies*, 67: 17-45.
- Ellison G. and D. Fudenberg (1993), "Rules of Thumb of Social Learning", *The Journal of Political Economy*, 101(4): 612-43.
- Glaeser, E. L., B. Sacerdote, and P. Scheinkman (1996), "Crime and Social Interactions", *Quarterly Journal of Economics*, 111: 507-548.
- Goyal S. (1996), "Interaction Structure and Social Change", *Journal of Institutional and Theoretical Economics*, 152: 472-494.
- Jovanovic, B. and R. Rob (1989), "The growth and diffusion of Knowledge", *Review of Economic Studies*, 56(198-9): 569-82.
- 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), "Learning Dynamics with Complex Cognition", in Pyka, A., Hanusch H. (eds.), *Applied Evolutionary Economics and the Knowledge-based Economy*, Edward Elgar, Cheltenham, 2004 (forthcoming).
- Nelson R. and Winter S. (1982), *An Evolutionary Theory of Economic Change*, Harvard University Press
- OECD (1996), *Employment and Growth in the Knowledge-Based Economy*, Paris.
- Pyka A. and P. Ahrweiler (2004) " Applied Evolutionary Economics and Social Simulation ", *Journal of Artificial Societies and Social Simulation*, 7:2
- Rauch, J. (1993), "Productivity Gains From Geographic Concentration of Human Capital: Evidence From the Cities", *Journal of Urban Economics*, 34: 380-400.

- Rogers E. M. (1995), *Diffusion of Innovation*, 4<sup>th</sup> edition, Free Press: New York.
- Silverberg, G., Dosi, G. and L. Orsenigo (1988), Innovation, Diversity and Diffusion: A Self-Organization Model, *The Economic Journal* 98, pp. 1032-1054.
- Watts, D. J. (1999), *Small Worlds*, Princeton University Press: Princeton.
- Watts, D. and S. Strogatz (1998), "Collective Dynamics of Small-World Networks", *Nature*, 393: 440-442.