X Seminario Latino-Iberoamericano de Gestión Tecnológica ALTEC 2003
"Conocimiento, Innovación y Competitividad: Los Desafíos de la Globalización"

Small World Dynamics and the Process of Knowledge Diffusion. The Case of the Metropolitan Area of Greater Santiago de Chile*

Piergiuseppe Morone
SPRU – Science and Technology Policy Research
University of Sussex, Mantell Building
P.Morone@sussex.ac.uk

Richard Taylor
CPM – Centre for Policy Modelling
Manchester Metropolitan University
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 'interactive learning'. We examine how knowledge spreads in a network in which agents have 'face-to-face' learning interactions. We define a social network structured as a graph consisting of agents (vertices) and connections (edges) and situated on a grid which resembles the geographical characteristics of the metropolitan area of Greater Santiago de Chile. The target of this simulation is to test whether knowledge diffuses homogeneously or whether it follows some biased path generating geographical divergence between a core area and a periphery. We also investigate the efficiency of our 'preference' model of agent decision-making and show that this system evolves towards a small-world type network.

Keywords: Agent-based, Chile, 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.

Small World Dynamics and the Process of Knowledge Diffusion. The Case of the Metropolitan Area of Greater Santiago de Chile

That the advance of knowledge lies at the core of modern growth process is more than an inference from the growth accounts. It is a perception enforced by well over a century of common experience. Economists have therefore applied themselves to learning more about the ways that practical knowledge is gained and exploited. A new outlook has developed and spread. It is not yet well defined.

Moses Abramovitz,
Thinking about growth,
Cambridge University Press, 1989.

1. INTRODUCTION

The nexus linking education and personal income is well established in the labour economic literature. It has been thoroughly investigated both theoretically and empirically and its main conclusion is that investments in human capital have a positive impact on one's future income. The learning process through which labour force builds up human capital has then been divided into schooling and on-the-job experience.

In this work, following Morone (2001), we argue that the process of human capital creation should be decomposed into two distinct learning processes: *individual learning* and *interactive learning*. The first is basically the learning process which takes place at school or through training courses, while the second, neglected until now, is the learning process which takes place through simply every-day interactions which each person has with her/his acquaintances.

The focus of this work is precisely this second learning process. We will study how interactive learning affects the knowledge creation in a middle-income country like Chile. We decided to focus on the Chilean case because of its particularly biased wage distribution and its exponential returns to schooling.[1]

The main target of this work is to model one possible type of interactive learning known as 'face-to-face' learning. Departing from previous works on knowledge diffusion we aim to develop a model which takes into consideration the complexity of the process of knowledge acquisition. In doing so we define a cognitive structure for each agent (cognitive map) and use this structure to define the process of knowledge acquisition.

The paper is structured as follows: section 2 revises briefly the existing literature on informal knowledge diffusion. In section 3 we present a formal model which investigates how knowledge spreads among agents situated on a grid composed of several neighbourhoods. Section 4 develops the concept of Cognitive Map (*CM*) and the processes of knowledge diffusion and accumulation in the presence of a complex cognitive structure.

Then section 5 introduces the case study and shows the results of the simulation exercise. Finally, Section 6 concludes the paper.

2. REVIEWING LITERATURE ON INTERACTIVE LEARNING

The learning process which takes place simply by interacting within a neighbourhood can be considered as a social externality. For instance, the diffusion of a new technology is a positive externality, which is generated by the informal interaction of entrepreneurs. Moreover, learning from neighbours can represent a negative externality. Consider, for instance, the problem of children in a poor neighbourhood choosing whether to pursue higher education or drop out of school. If in such an environment the most popular decision is dropping out, this would represent a negative externality for those children facing the choice subject to social influence.

Most studies found that social learning plays a positive role in determining the long-term equilibrium. Ellison and Fudenberg (1993) developed a model in which agents consider the experience of neighbours in deciding which of two technologies to use. Their work is structured around two simple models of the learning environment. First, they consider a homogeneous population which reacts identically to the two technologies, with one technology being superior to the other; subsequently they introduced heterogeneous agents which consider the two technologies in different ways. In the first case the issue is whether the better technology will be adopted, while in the second case the question is whether the two technologies will be adopted by the appropriate players. In both environments agents use exogenously specified 'rules of thumb' that ignore historical data but might incorporate a tendency to use the more popular technology. Under this condition, the outcome of their work suggests that "even very naïve learning rules can lead to quite efficient long-run social states", but adjustment can be slow when a superior technology is first introduced (Ellison and Fudenberg, 1993: 637).

In a subsequent paper (1995), the authors focused their attention on the patterns of information exchanges, studying the way in which word-of-mouth communication contributes to aggregate information of individual agents. They defined a nonstrategic environment composed of homogeneous agents which face the decision of choosing one of two products. Their findings show how "despite the naïve play of individuals, this type of information flow may lead to efficient learning on the social level, and that social learning is often most efficient when communication between agents is fairly limited" (Ellison and Fudenberg, 1995: 120).

Bala and Goyal (1995) analysed social learning in an organisational environment in which agents have a time-limited life experience and heterogeneous beliefs. Departing from the original work of Kiefer (1989), they consider the case of a monopolistic firm attempting to calculate its true demand curve. “The learning process – adopted by the authors – is then one in which the sequence of managers learn about the demand curve through their own actions as well as the experience of earlier managers” (Bala and Goyal, 1995: 303). In this case again, learning from others augments the probability of converging to the set of ex-post optimal actions.

Subsequently, Bala and Goyal (1998) investigated the relation between the structure of social networks and learning processes in a world where agents learn from currently available social information, as well as from past experiences (as opposed to the previous works of Ellison and Fudenberg). Their findings show that the structure of the neighbourhood has important implications for the likelihood of adopting new technologies, for the coexistence of different practices, and for the temporal and spatial patters of diffusion in a society. More precisely, they showed how neighbourhood structures characterised by the presence of locally independent agents (i.e. agents with non-overlapping neighbourhoods) generally facilitate social learning.

A common way of modelling the mechanisms of social learning and technology diffusion makes use of evolutionary game theory. Several authors examined local interaction games[2] in which each person’s payoff depends on the actions of his/her neighbours. Most of these studies pointed out that local interaction might result in the diffusion of personal behaviours in certain dynamic systems. In a recent work, Morris (2000) extended this finding, linking the occurrence of social learning (that he calls contagion) to some qualitative properties of the interaction system such as cohesion, neighbour growth, and uniformity. Chwe (2000) modelled social learning as dependent on both social structure and individual incentives. In this way he obtained a model that he called a ‘local information game’ as, he argued, “locality is represented by information and not necessarily by payoffs […]. Local interaction games make sense for local coordination, such as keeping our street clean; for ‘big’ coordinations such as political changes, informational locality is more appropriate” (Chwe, 2000: 11).

Along with the speed of new technologies’ diffusion, several researchers have focused on the impact of peers’ behaviour upon individual decisions in several areas such as propensity to crime, use of drugs, school dropout and school attainments. For instance, the *role model* provided by good students might generate positive spillovers in the sense that peer students may imitate successful behaviours. Nonetheless, "when social interactions act as strategic complementarities between agents, multiple equilibria may occur in absence of any coordination mechanisms" (Brock and Durlauf, 1995: 1). The school-based interaction was studied by Benabou (1993) in his model of occupational and residential choices. The author shows that when the cost of individual education is a decreasing function of investment decision of one's neighbours, neighbourhoods can exhibit multiple equilibria in the steady-state level of human capital.

Durlauf (1996) shows how these spillovers can have significant distributional effects generating stubborn inequality in the long run. Income inequality "emerges through the interaction of positive feedback structure between members of a common neighbourhood with the tendency of families to stratify themselves endogenously into economically homogeneous neighbourhoods" (Durlauf, 1996: 505).

Positive feedbacks means, in the words of Durlauf, "that the income distribution of neighbourhoods will strongly affect the future economic status of children within the community" (Durlauf, 1996: 505). A similar argument was raised by Gleaser, Sacerdote and Scheinkman (1996), while arguing that social interactions can explain large differences in community crime rates.

What all the studies considered so far have in common is the fact that learning from neighbours occurs and that under certain conditions it leads to (highly unequal) multiple equilibria. Yet, none of these studies goes beyond a binary definition of learning. If there is a new technology, agents can either learn of its existence, and hence adopt it, or stay in their initial state of ignorance. Pupils that live in 'good' neighbourhoods have good chances of imitating good *role models*, while pupils which live in 'bad' neighbourhoods are more likely to imitate bad *role models* and hence end-up in perpetuating the initial differences. In other words, learning is defined and modelled as an imitational process deprived of its real complexity. There are in the economic literature few attempts to model informal knowledge diffusion in a more complex way. We will now examine these works.

Jovanovic and Rob (1989) presented a model in which knowledge was defined as a public good which is exchanged through pairwise meetings. The authors define incremental improvements in knowledge as a complex process of assembling different ideas by means of information exchange by heterogeneous agents. The new insight brought by Jovanovic and Rob is that knowledge is defined as something more complex than a binary variable and, therefore, that growth of knowledge can 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 diffusions over 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.

Along the lines of Cowan and Jonard (1999), Morone and Taylor (2001) defined a model in which agents exchange knowledge 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, depending upon the initial conditions. The model presented in the following section could be considered an extension of Morone and Taylor (2001) as it departs from similar assumptions but develops a more complex learning structure.

3. MODEL DESCRIPTION

We assume a population of N agents distributed over a grid that resembles the geographical configuration of the metropolitan area of Greater Santiago de Chile. The grid is divided into 34 portions, each corresponding to a defined district of Santiago and thus having different dimensions and different population densities. Each agent is initially assigned a district and then allocated to a cell at random within that district. However, agents do not occupy all the cells of the grid, and each occupied cell can have only one agent. Neighbourhoods, which may contain cells from several districts, are constructed according to the von Neumann definition: those cells adjacent in the four cardinal directions and within the agent's visible range. The initial *local-network* is created by connecting an agent with all other agents located within her/his neighbourhood. We also define a *cyber-network* as the ideal network connecting all those agents which have access to the Internet. The *cyber-network* generates a second system which has no geographical dimension but connects agents located in far distant regions through ICT support. As opposed to the *local-network*, all agents who have access to the Internet are neighbours within the *cyber-network* (i.e. the visibility is equal to the size of the system). Each agent has a list of acquaintances which includes the members of the *local-network* and the *cyber-network*. Data, which is used to initialise the model, is a sub-sample of the 1998 edition of the *Encuesta de Ocupación y Desocupación,* and will be described more fully later.

Our geographical network is expressed in graph notation following the standard definition: "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 the agents' connections. Formally, we have G (I, Γ), where $I = \{1,\ldots, N\}$ 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 $\Gamma(x) = \{(y\in I \setminus \{x\} \mid d(x, y) \leq \nu) \cup (y\in \omega)\}$, 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), ν (visibility), as already mentioned, is the number of cells in each direction which are considered to be within the agent's spectrum, and ω defines the *cyber-space*, which by definition encompasses all those agents endowed with ICT facilities (i.e. computers connected to the Internet). Intuitively, Γ_x (we will use this notation rather than $\Gamma(x)$ from now on) defines the neighbourhood of the agent (vertex) x.

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. Each interaction is aimed at diffusing knowledge. Each agent is endowed with a *cognitive map* (*CM*), which contains information on the level and the kind of knowledge possessed by her/him. The structure of the *CM* is that of a tree, where each node corresponds to a bit of potential knowledge and each edge corresponds to acquired knowledge. We will return to the *cognitive map* in the next section.

Initial acquaintances in the *local-network* are the immediate neighbours (i.e. those within the visible spectrum). Subsequently, an agent can learn of the existence of other agents through interactions with her/his acquaintances (i.e. she/he can be introduced to the acquaintances of her/his acquaintances). A new connection is made if the selected acquaintance is connected to another individual of whom she/he is not aware, where that individual was the one the acquaintance interacted with in the last cycle (this would tend to avoid the situation where it is introduced to one that has a low preference and not considered to be a good choice). Moreover, agents can stop interacting with some of their acquaintances if the connection does not tend to result in gain interactions and is therefore no longer useful.[3] Therefore the number of acquaintances changes over time, but does not necessarily increase over time. In this way we introduce a dynamic element into the network structure.

Having defined Γ_x as the set of initial acquaintances of agent x (or *first generation connections*), we define $\varphi_{x,t}$ as the set of acquaintances of the acquaintances at time t (or *next generation connections*), and the individual $m_t \in \varphi_{x,t}$ who is added at each t. We also define $\vartheta_{x,t}$ as the set of acquaintances dropped at time t and the individual $n_t \in \vartheta_{x,t}$ who is dropped at each t. Now we can define the total set of acquaintances for individual x at time $t=T$ as:

$$\Phi_{x,T} = \left(\Gamma_x \cup \varphi_{x,T}\right) / \vartheta_{x,T} \tag{1}$$

Each acquaintance has an associated strength, $\tau \in (0.05,1)$, which is a measure of the strength of the relationship from the agent to her/his acquaintance. Note that this model is not constrained to have symmetry of relationships between agents.

We define a rule governing how an agent chooses an acquaintance to interact with. In doing so, we make the assumption that an agent prefers interacting with acquaintances with whom she/he has strong relations. Agent y will be selected for interaction with agent x with probability given by:[4]

$$p^x(y) = \frac{\tau_y^x}{\sum_{i \in \Phi} \tau_i^x}, \tag{2}$$

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

$$\tau_{i,t} = \varepsilon \tau_{i,t-1} - \beta \qquad (3)$$

$$where \begin{cases} \varepsilon = 1 \quad \text{and} \quad \beta = 0 \quad \text{if the interaction is a gain interaction;} \\ \varepsilon = 0.6 \quad \text{and} \quad \beta = 0 \quad \text{if the interaction is not a gain ineraction;} \\ \varepsilon = 1 \quad \text{and} \quad \beta = 0.05 \quad \text{if an agent is not selected for interaction.} \end{cases}$$

As already mentioned, τ_i is bounded between 0.05 and 1. Whenever the τ_i attached to any acquaintance reaches the lower threshold of 0.05, the acquaintance is dropped from the acquaintances list. In other words, we assume agent memory decay in our model. However, if the acquaintance is a neighbour (i.e. an acquaintance asserted at the beginning of the simulation), then she/he is not dropped (i.e. they can reach the lower threshold of preference, but are not dropped from our acquaintance list). This is due to the fact that initial acquaintances are geographical neighbours with whom we keep meeting unless we move to a different neighbourhood (an option which is not considered in our simulation model).

Equations (2) and (3) together define our model of 'preferential acquaintance selection' where agents gather, over time, information on the opportunities for learning interactions, which they use to develop internal models of interaction preference. We suggest this process enables agents to make better choices of acquaintances and learn more efficiently.

In the following section we describe how the (knowledge) *diffusion interaction* takes place. In doing so we will describe in greater detail how the *cognitive map* of each agent is constructed and how this 'structured' knowledge is important in defining the model of interactive learning.

4. Cognitive Map: Implications for Knowledge Diffusion

One of the main limitations of the simulation models which aim to formalise the process of knowledge diffusion as an incremental process (Cowan and Jonard, 2000; Morone and Taylor, 2001) is the oversimplifying assumption that knowledge is accumulated as a stockpile (i.e. a vector of cardinal numbers indicating the level of knowledge). The roots of this problem are to be found in the distinction between economics of information and economics of knowledge. As pointed out by Ancori *et al.* (2000), the economics of knowledge differs from the economics of information in the sense that knowledge is no longer assimilated to the accumulation of information in a stockpile. The distinction between these two concepts has been repeatedly ignored by a certain branch of the economic literature (economics of information), which does not consider the cognitive structure that agents use to elaborate knowledge.

Following this distinction, Ancori e*t al.* (2000) develop an appreciative model of great interest to us: according to the authors, new knowledge is acquired "by a backward process through which the new knowledge is confronted and articulated with previous experience. […] the appropriation of crude knowledge – i.e. its integration in one's cognitive context – is not the result of a transmission, but rather the result of a re-engineering process" (Ancori *et al.*, 2000: 267). What the recipient agent is basically doing is de-codifying the knowledge received in order to be able to position it in her/his own cognitive map.

We can think of the cognitive map as a tree in which each vertex (node) represents a piece of crude knowledge and each edge (link) represents knowledge that we have already mastered and learned how to use.

Figure 1. Cognitive Map

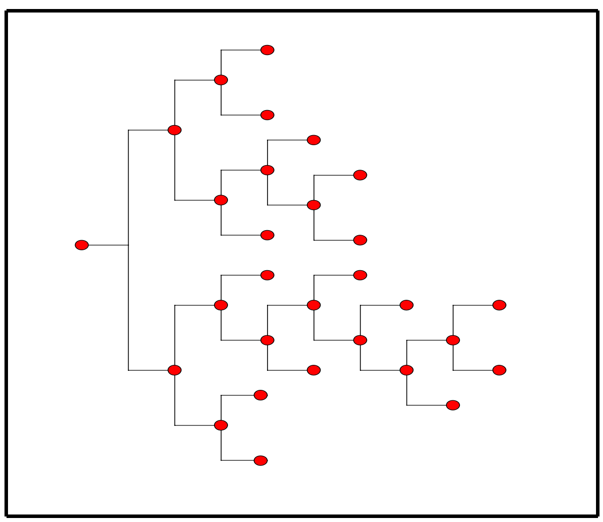

In the graphical representation, figure 1 above, we present an example cognitive map. This shows mastered knowledge as active nodes on the map (the coloured nodes), while all the other possible nodes which would complete the tree represent knowledge that at present is not in our cognitive map but could be activated through individual as well as interactive learning.

As assumed by Ancori *et al*., knowledge demands knowledge in order to be acquired; hence, in order to activate a new node it would have to be directly connected to active (coloured) nodes. Moving from left to right in the cognitive map we move from less to more specialised knowledge. This observation justifies the assumption that new nodes can only be activated (i.e. new knowledge can be acquired) if they are directly connected to active nodes.

Each agent is initially endowed with a different cognitive map, which depends upon her/his level and kind of education (measured as years of schooling and kind of school attended). Each column corresponds to a higher level of education.[5] Moreover, once education becomes specialised[6] the cognitive map will develop only in certain areas, showing more pieces of knowledge in those areas in which the agent is knowledgeable.

It is worth nothing here that in our model there is no knowledge creation, hence it is bounded to reach a stationary state in which no more knowledge exchange is possible (i.e. the *CM* being bounded in the number of columns and, therefore, in the number of *knowledge nodes*). In an ideal situation (i.e. characterised by a fully integrated system without any excluded agent or cluster of agents), the stationary state would be reached when all the agents have acquired all the possible knowledge (i.e. all cognitive maps are saturated). Nonetheless, the presence of isolated agents might generate a multiple equilibria situation. As a final note, it is important to mention that in the initial situation most knowledgeable agents will necessarily have completely saturated *CM* in one of the specialised areas.

The cognitive map can be formally described as: $CM(X, \mathrm{N})$ where X is the set of the whole possible knowledge available (i.e. the set of vertices), and N identifies the pieces of knowledge activated (i.e. edges of the graph).

We will now explain how the process of knowledge diffusion takes place. An agent, whom we shall call A, contacts an acquaintance, B, in accordance with equation (2). Once the contact has been established the algorithm compares the two cognitive maps subtracting the cognitive map of A from that of B. This can produce one of two possible results:[7]

$$CM_A(X, \mathrm{N}) - CM_B(X, \mathrm{N}) \begin{cases} = \varnothing \\ \\ \neq \varnothing \end{cases} \tag{4}$$

If the difference between the two sets is a non-empty set there is possibility for interaction; if not, agent A will have no interest in interacting with agent B as there is no possible gain.

We present an example that will clarify the issue. The two maps in figure 2 represent the cognitive maps of agent A and an acquaintance, agent B. Now, let us assume that agent A contacts agent B. If we calculate the distance between the two maps we get $CM_A(X,\mathrm{N}) - CM_B(X,\mathrm{N}) \neq \varnothing$ (this can be clearly observed in figure 3 below).

Figure 2. Comparing two Cognitive Maps

Agent *A* Agent *B*

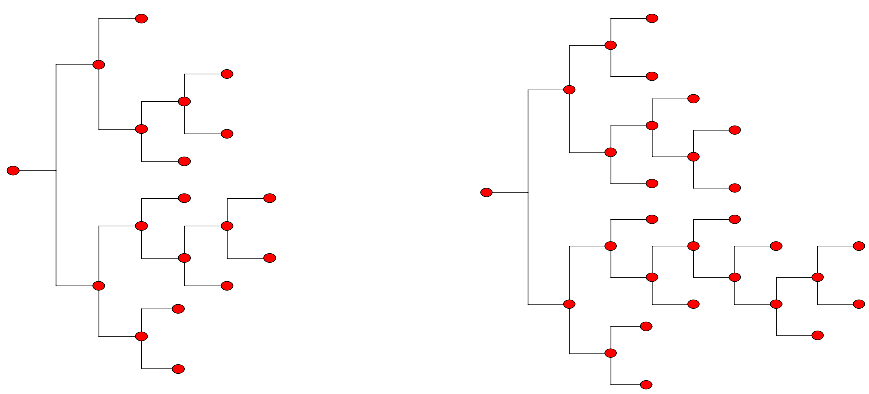

The first map below reproduces the difference between the two *CMs*. Once we have identified this difference, we need to identify the possible learning region where knowledge can be gained (i.e. additional nodes can be activated). To do so we recall the requirement that new knowledge has to be pegged to already existing knowledge, and thus we can cross out several of the coloured nodes in the first diagram. We conclude that the only knowledge that agent A can learn from agent B is that connected to activated nodes.

Figure 3. Knowledge interaction and the occurrence of learning

$CM_A - CM_B$ *Learning region*

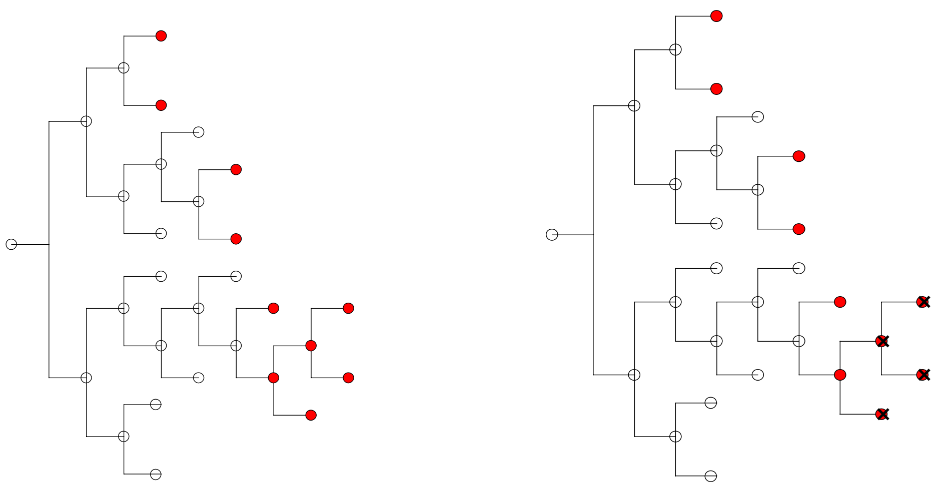

Defining the nodes of the learning region as *Ω*, then the actual learning can be expressed as $p\Omega$, where p represents the percentage of nodes of the learning region that will be activated as a consequence of the interaction. In other words, the agent that has started the interaction will activate (learn) p percent of the nodes, selected randomly (rounding always to the highest integer in the case of decimal numbers) from the *learning region*.[8] Since the number of nodes increases exponentially for each additional year of schooling, it implies that the higher the level of education of the interacting agents, the higher will be the learning opportunity. This mechanism reflects the idea that the 'absorptive capacity'[9] of each agent is a positive function of her/his level of education.

In conclusion, the learning process that we have defined in this model, using a cognitive map, extends the idea developed in Morone (2001) where the author defined the level of knowledge of each agent as a function of two main variables:

$$K_i = E_i + I_i \tag{5}$$

where E_i was the level of education obtained by individual i through a formal process of individual learning, and I_i was the level of education of individual i obtained through the process of informal interactive learning (face-to-face interaction).

The latter process of learning was then defined as a function of three variables: agent i's absorptive capacity (ψ_i); the degree of connectivity of the network (n) within which agents interact (θ_n); and the average level of education of other agents $\left(\frac{\sum_{j \in n} E_j}{N-1}\right)$. The knowledge function was then defined as:

$$K_i = E_i + I_i\left(\psi_i,\ \theta_n,\ \frac{\sum_{j \in n} E_j}{N-1}\right) \tag{6}$$

The model developed in this paper aims to explicate, through a simulation exercise, the basic idea expressed in equation (6). We will now briefly introduce the simulation setting. Then, we will describe the dataset used to calibrate the simulation. Finally, we will present the simulation results and their implications.

5. SIMULATION EXPERIMENTS

We performed simulation experiments with a population of 232 agents allocated over a grid that resembles the map of Greater Santiago de Chile. The relative populations of each of the 34 *comuna* (districts) composing the grid are respected, defining a grid with different levels of density over the whole map. The grid is composed of a total of approximately 2000 cells and has an approximate overall density of one agent per 8 cells. Each agent has a visibility $\nu = 2$, meaning that she/he can see the two cells situated in the four cardinal directions (North, South, East and West). The initial knowledge is defined by a cognitive map constructed according to a sample drawn from *Encuesta de Ocupación y Desocupación* regarding years of schooling and kind of schooling, as described above.

The basic objective of the simulation experiments was to test whether knowledge diffuses homogeneously throughout the population or whether it follows some biased path generating divergence and inequality. In particular we shall investigate differences amongst the modelled districts of Santiago and amongst core and periphery areas of the geographical grid. We also examine the efficiency of the model of ‘preferential acquaintance selection’ with which the agents are endowed. We investigate the impact on both individual learning and on system-level knowledge diffusion by comparing with the case where there is no preferential selection.

The model was programmed in the Strictly Declarative Modelling Language (SDML) developed at the CPM (Wallis & Moss, 1994) to support the modelling of social processes with multi-agent systems. The results were analysed using the graphical output capabilities of the SDML platform and the network analysis software toolkit UCINET 5.0 (Borgatti, Everett, & Freeman, 1999).

5.1 Data Description

As already mentioned, the data used for the simulation is a sample of 232 agents obtained from the 1998 edition of the *Encuesta de Ocupación y Desocupación* of the University of Chile. This survey has been performed every year since 1957 and it covers a sample of approximately 4000 households (with approximately 11000 individual observations) located in the Greater Santiago area.[10] The data collected covers: demographic information; occupational information; educational information; and wage information. The 1998 survey includes additional information on effective years of working experience; number of repeated school years; kind of school attended (characterised according to the funding system[11] and the rural/urban location); parents' level of education; religion; weight; height; and finally, use (or not) of a computer at work.

The variables that are of interest to our simulation experiment are the following: district of residence, years of schooling, kind of schooling, and use of computers at work. We will use these variables to distribute agents over the geographical grid, to build the *CM* of each agent and to construct the *cyber-network*. We consider male and female agents aged between 14 and 65. We chose this range of age as it appears to be the best representative sample of active population engaged in both interactive learning and individual learning.[12] We present below some summarising tables that describe the main statistics of our data.

Table 1. Original database

Variable	Num. of Obs.	Mean	Std. Dev.	Min.	Max.
District	7732	18.64602	10.28054	1	34
Schooling	7732	10.90184	3.892597	0	19
Kind	7732	2.504268	1.405629	0	6
Computer	7732	0.146275	0.353404	0	1

Table 2. Simulation sample database

Variable	Num. of Obs.	Mean	Std. Dev.	Min.	Max.
District	232	18.78017	10.45332	1	34
Schooling	232	11.15517	3.758058	0	18
Kind	232	2.577586	1.409005	0	6
Computer	232	0.1422414	0.350052	0	1

As expected, the two distributions are very similar, and in particular all the variables considered have similar standard deviations. In the following two graphs we plot the (kernel) density functions of the schooling variable for both databases (the original survey and the simulation sample).

Figure 4. Original database

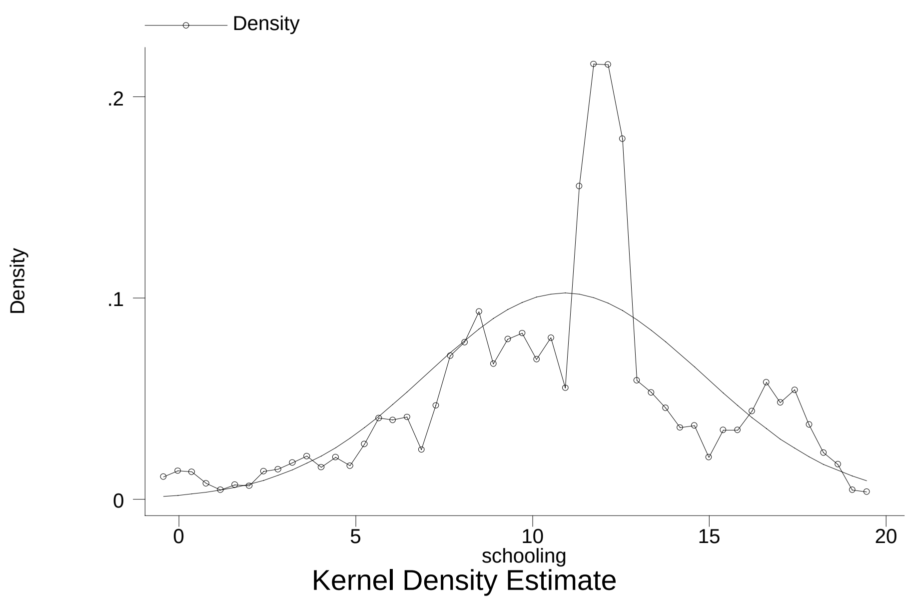

Figure 5. Sample database

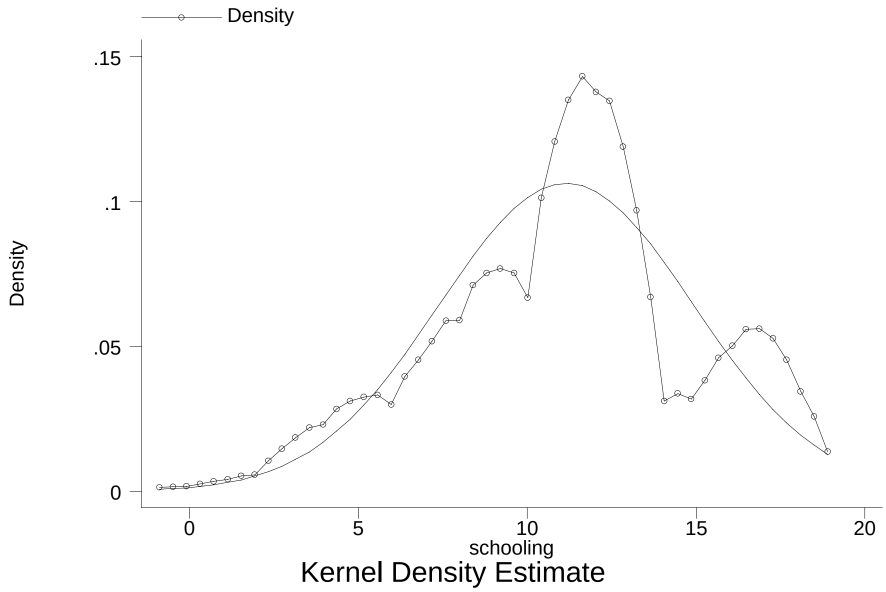

No significant differences can be detected between the two distributions. Both series have one absolute maximum and two local maxima respectively at the right and the left of the absolute maximum. This confirms the fact that the sub-sample is representative of the original sample. A more detailed analysis of the two datasets is presented in the annex.

5.2 Results and Interpretation

After letting agents interact for two hundred cycles we observe significant changes in the knowledge patterns. Indeed the average number of activated nodes increases substantially over the whole period, showing that the system is becoming more knowledgeable.[13] This is not surprising since agents can only increase their knowledge level (i.e. there is no loss of knowledge).

Figure 6. Changes in the average level of knowledge

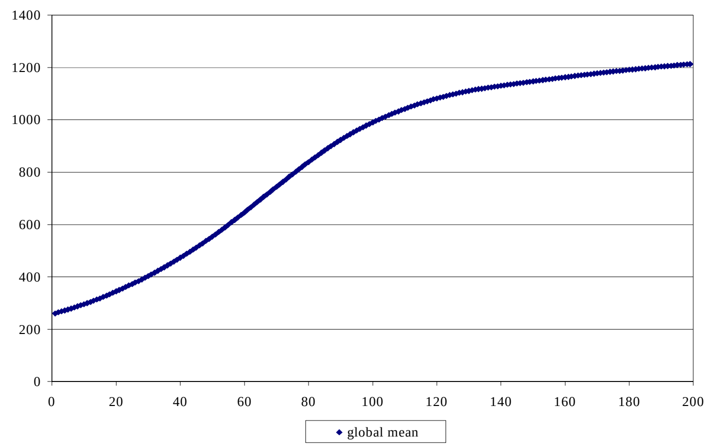

Source: Simulation results

Plotting μ we can observe how first the average number of activated nodes grows very fast, then after approximately 80-100 cycles there is a turning point (inflection point), after which μ grows more slowly. Finally, after about 150-160 cycles, the mean curve starts levelling off, meaning that the system is reaching a stable equilibrium.[14]

To explain these results, we first consider the dynamics underlying the structure of knowledge in the model. The number of fully saturated agents increases over time, and as agents approach this state they have a reduced potential for learning, i.e. the learning region becomes smaller. However, on the other hand, in the early stages of the simulation this region tends to widen in the *CM* of the majority of agents, giving the potential for greater gains. In addition, agents will have increased opportunities to gain from interactions as *CMs* become more heterogeneous. For example, two agents with identical schooling will not be able to gain from an interaction in cycle 0, whereas later in the simulation they most likely will experience a small gain. This begs the question: to what extent is the observed increase in knowledge due to the widening of the learning region (i.e. the structure of the *CM*), and to what extent is it due to agents making better choices for interaction (i.e. the preferential model of acquaintance selection). In the following scenario we shall test the hypothesis that the preference model for choosing among acquaintances is an efficient mechanism. To help us analyse these differences, we introduce another measurement, the ‘percentage of gain interactions’, which tells us how successful the learning is on an individual level. As shown in figure 7, our model of preferential selection appears to be very accurate since the percentage of gain interactions increases very sharply in the first 20 cycles and is constantly higher than 90% for more than 100 cycles. Then it starts decreasing when the system approaches a stable equilibrium.

A different picture emerges when we run the same simulation without the preferential model.[15] The number of gainful interactions is constantly smaller. Over the first 140 cycles, the simulation with preferential model is approximately 25% more efficient. Moreover, the efficiency gap (i.e. the distance in percentage gain interactions) grows bigger over the whole time span. After 140 cycles, when the model starts converging towards a stationary equilibrium, the speed at which the efficiency gap grows accelerates. In the last cycle the simulation with preferential model becomes about 70% more efficient.

The increase in the efficiency gap between the two simulations is due to the fact that agents keep meeting new acquaintances and - as the number of acquaintances grows bigger - the chances to select the right agent with whom to interact decrease. Moreover, this mechanism is accelerated after 140 cycles, as the *CMs* start becoming saturated: when the learning opportunities start decreasing (i.e. the learning regions become smaller) that is the moment when agents really need to know with whom it is worth interacting. Figure 7 below shows the differences between the two models so far discussed.

Figure 7. Percentage of gainful interactions. Studying the efficiency of the system

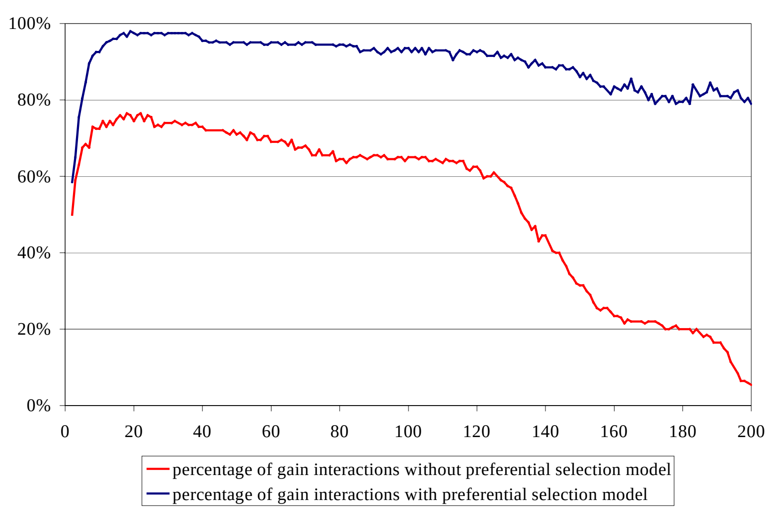

Source: Simulation results

Interestingly, but not unexpectedly, most agents endowed with ICT facilities will prefer to interact with cyber-acquaintances rather than with their regular (geographical) acquaintances. As illustrated in Figure 8 below, the interactions among cyber-network members are almost always gainful, while this is not the case for interactions which take place within the same district.

Figure 8. Percentage of gainful interactions within districts and cyber-network

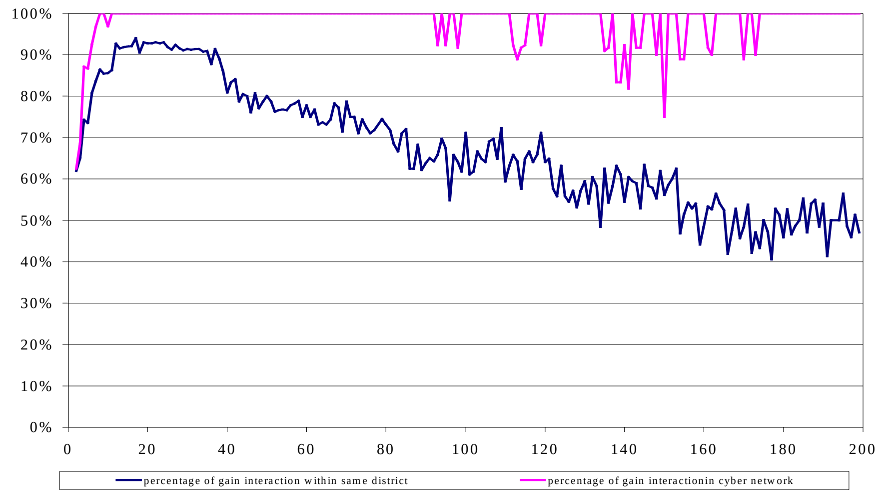

Source: Simulation results

In order to investigate the distributional changes of knowledge we will start our analysis by looking at the Kernel density functions of knowledge at different moments in time. This preliminary analysis will provide a clear picture of distributional changes over time, and pave the way to a more thorough analysis of knowledge distribution.

Figure 9. Kernel density functions of knowledge and normal distribution

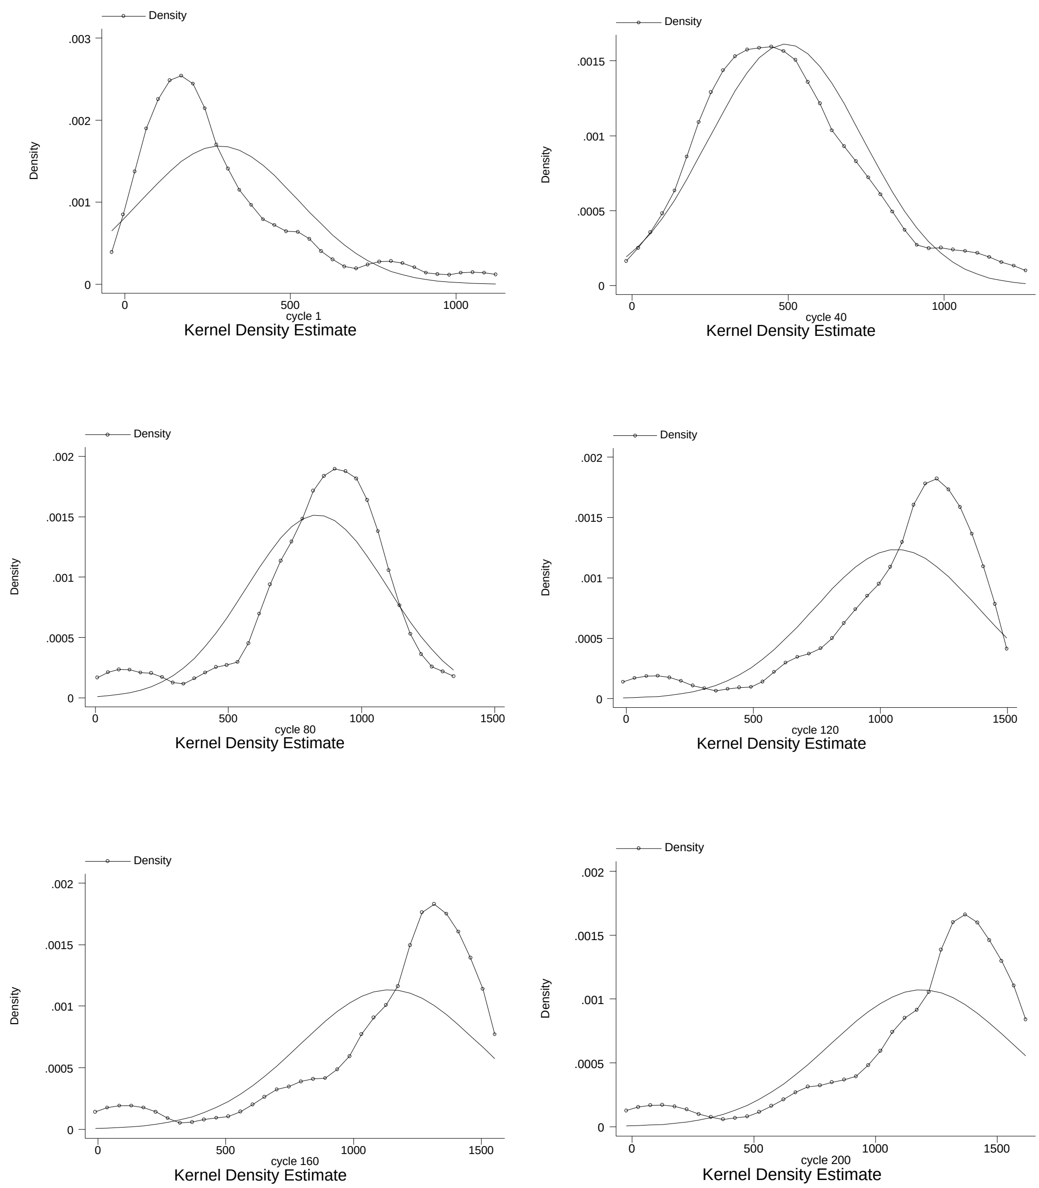

Source: Simulation results

As we can see from the figure above, there is a clear trend in the distributional changes: initially, there is a large majority of agents with a low level of knowledge (the mean of the distribution is in the left hand side of the first Cartesian diagram). Over time, the mean of the distribution shifts from left to right: after 40 cycles knowledge is almost normally distributed, and then it becomes skewed to the right hand side. This distributional dynamic shows clearly that the society has changed from one in which there is a majority of low knowledgeable people and an ‘intellectual oligarchy’, to a more mature society with a majority of knowledgeable people and some ‘pockets of ignorance’. In other words, the final knowledge distribution is the mirror image of the initial one. To our understanding this change has to be interpreted as a substantial improvement in terms of knowledge distribution. This intuition is indeed corroborated by the inequality index.[16]

In the figure below we report the Gini index calculated over the 200 cycles of the simulation. We can observe how knowledge inequality constantly decreases over the first 100 cycles to then stabilise.

Figure 10. Knowledge Inequality Dynamics (Gini Index)

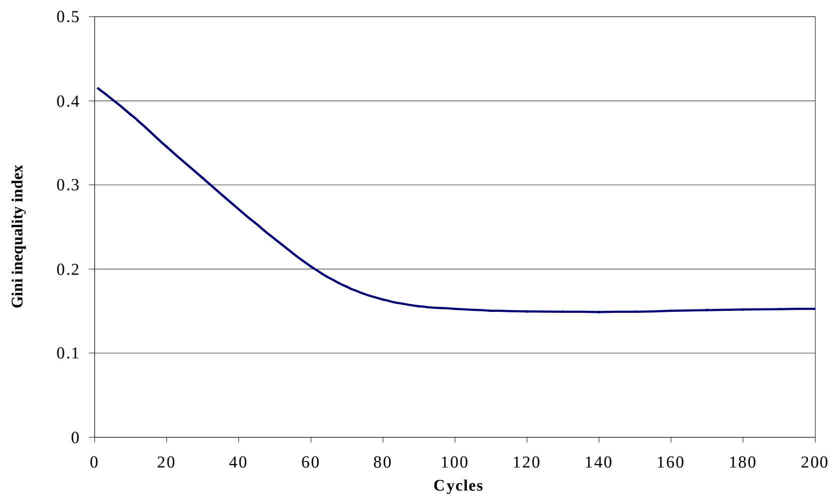

Source: Simulation results

So, overall the society has quickly become less unequal as well as more knowledgeable. This should be interpreted as a substantial improvement. Nonetheless, it will be interesting to further investigate the occurrence of some kind of an exclusion mechanism which generates the observed ‘pockets of ignorance’. In order to do so we start by looking at *poverty measures* which, in this context, should be interpreted as indexes of knowledge deprivation. We calculate the first two Foster, Greer and Thorbecke indices, (i.e. the head count ratio and the poverty gap ratio) with the *ignorance line*[17] computed as half the median of the average level of knowledge. The difference between the two measures is that the first one is just the ratio of the number of people having a knowledge level below the *ignorance line* over the total population. Whereas the poverty gap ratio is intended to provide some estimate of the average knowledge shortfall from the *ignorance line* and thus indicates the severity of the knowledge gap rather than just the number of people with a poor level of knowledge.

Figure 11. Poverty measures applied to knowledge

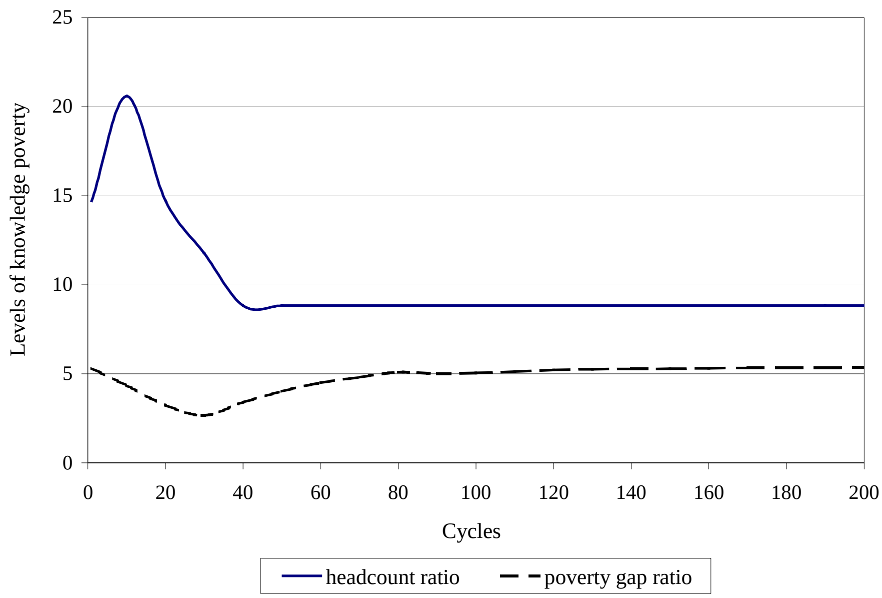

Source: Simulation results

The figure above shows how the headcount ratio first increases and then decreases. This implies that over the first 20 cycles the number of agents with the level of knowledge smaller than half the median is increasing. Nonetheless, those subjects are learning and getting closer to the *ignorance line*. We can infer this looking at the poverty gap ratio (or knowledge gap ratio) which, over the same time-span, is decreasing. The situation reverts afterwards: now the number of agents with the level of knowledge smaller than half the median is decreasing but there is a group of agents which is falling behind, being trapped in an exclusion mechanism. In other words, the knowledge gap between this group of agents and the rest of the population grows bigger over time.

The analysis of these results becomes more revealing if we look at the simulation outcome disaggregated at district level. The geographical dimension helps in understanding the mechanisms at work that generate the observed exclusion mechanism.

When we look at the mean and variance dynamics disaggregated at district level (see figure 9) we can draw two conclusions. First, some initially backward districts are able to catch up with the most knowledgeable ones, while others are marginalised from the learning process. Interestingly, the catching up process is not strongly correlated with the initial condition, meaning that some districts which start with a very low level of knowledge are able to reach, on average, very high levels of knowledge. At the same time, other districts which start from similar levels of initial knowledge are unable to catch up.

Figure 12. Changes in knowledge average and variance by districts

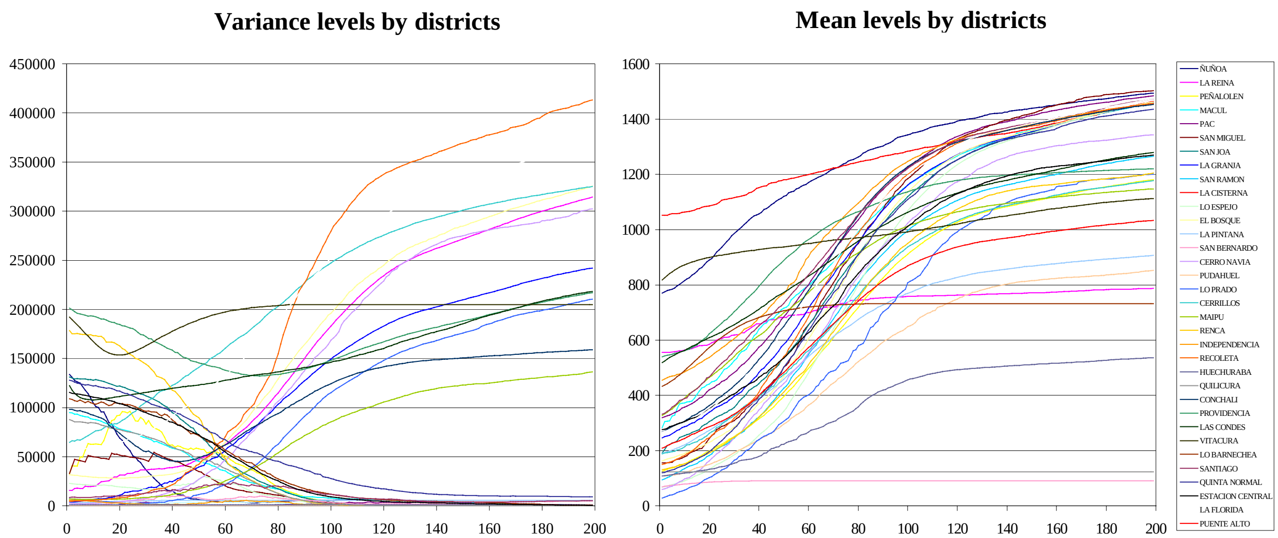

Source: Simulation results

The second conclusion is obtained when looking at the variance distribution by district. Some districts become less and less equal while others converge towards a more even equilibrium. In this case too the initial condition does not determine the twofold dynamic.

In order to understand which are the driving forces that cause only some backward districts to catch up, we need to investigate the actual geographical distribution of backward areas in the map of Greater Santiago. The four pictures reported in figure 13 below show the map of the Greater Santiago region divided into 34 districts.[18]

Figure 13. Changes in knowledge mean by districts

Initial situation

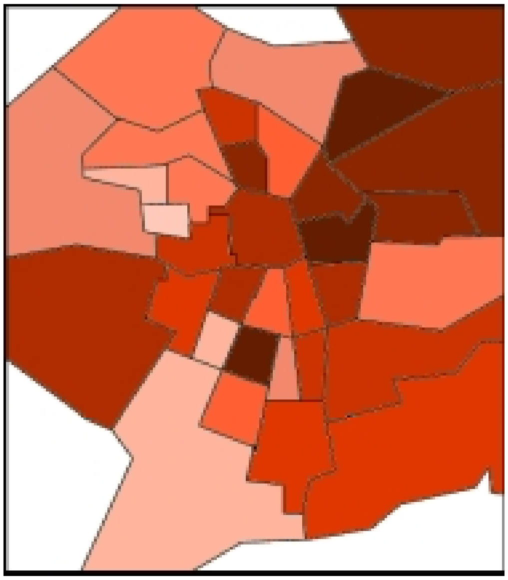

Knowledge distribution after 50 cycles

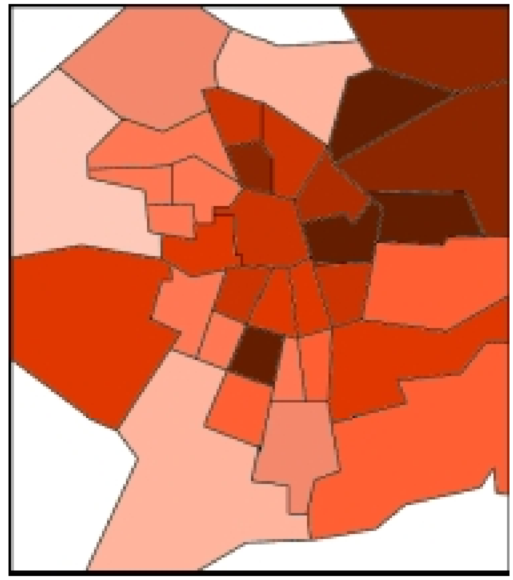

Knowledge distribution after 100 cycles

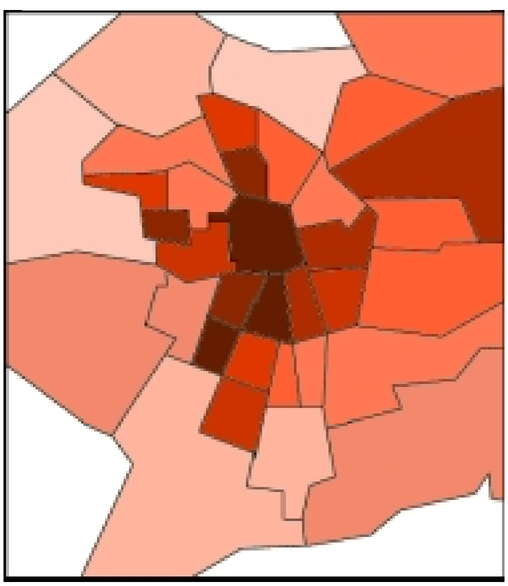

Knowledge distribution after 150 cycles

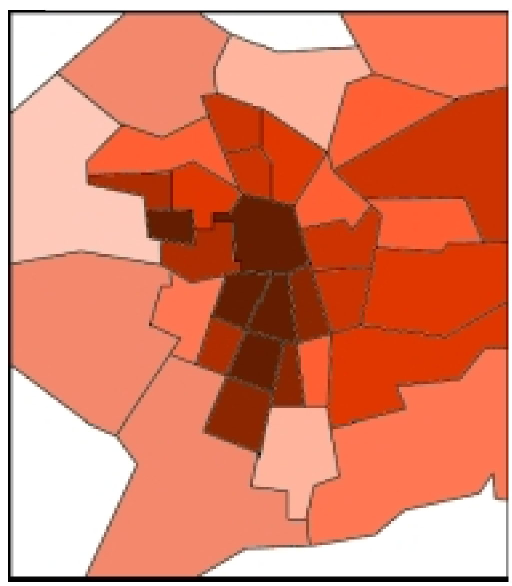

Source: Simulation results

The white areas adjoining the border of the map are not included in the Santiago district. The changes in the colours show the relative changes in average level of knowledge at different stages of the simulation (absolute changes are showed in figure 12).

It is clear from the first picture that initially, the most backward districts are those situated at the western periphery of the map. In general, central districts have an initially higher mean level of knowledge. There are a number of districts in the semi-periphery also having low mean. On average, the periphery has a slightly better initial condition than the semi-periphery (this is due to the fact that we also include in the periphery rich residential neighbourhoods such as *Las Condes*, *Lo Barnechea* and *Vitacura* located north-east in the map). The situation changes over time, showing how semi-peripheral districts are able to almost catch up with central districts, while far-peripheral districts are marginalised from knowledge growth.

To complete our analysis, we have divided the whole region into three macro districts,[19] (i.e. centre, semi-periphery, and periphery) and calculated the average level of knowledge for each macro district over time. In figure 14 below we can see the core-periphery dynamic more clearly. It emerges clearly how peripheral districts are those unable to catch up (in spite of a better initial condition with respect to semi-peripheral districts). We can therefore conclude that the geographical dimension plays a crucial role in determining the long-term equilibrium.

A further insight into this dynamic is provided by the changes in the variance calculated at macro district level (i.e. mean of district variance). The periphery and semi-periphery are the areas in which the variance grows more considerably over the long period. On the other hand, the average level of the intra-district variance decreases in the case of the central districts. Comparing with the earlier figure 12, this points out that the distribution of knowledge becomes more unequal amongst districts as well as within them. Hence, less educated people who live in peripheral districts will be marginalised with respect to those who live in central districts, as well as from the most educated people within the peripheral macro-district.

Figure 14. Centre-Periphery dynamics (Average level of knowledge and Average change in knowledge variance by districts)

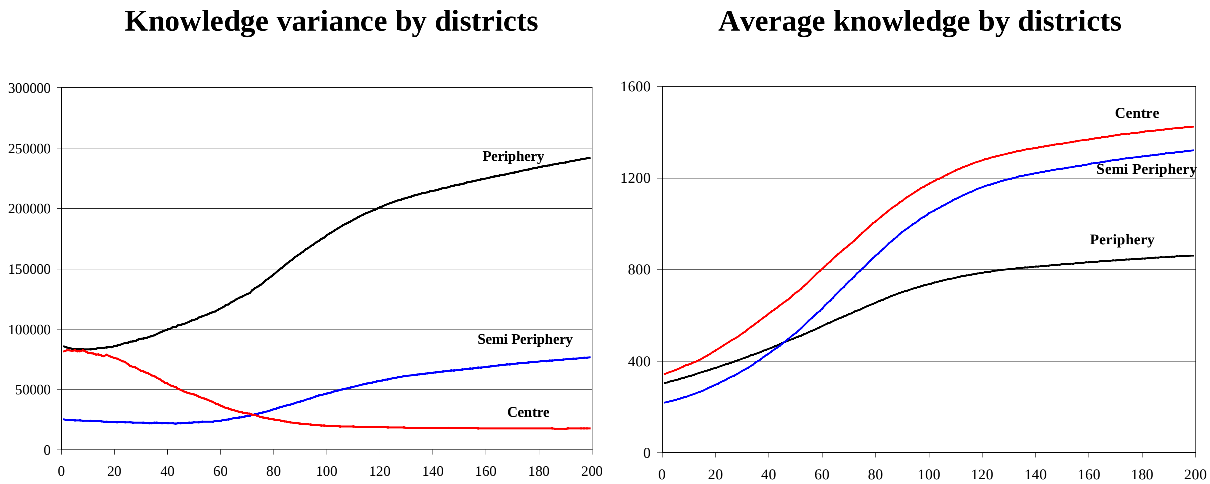

Source: Simulation results

From the results obtained so far we can conclude that initial conditions do not matter for those people who live in central or semi-peripheral neighbourhoods, while they do matter for those living in peripheral areas. An intuitive explanation of this result is that if a person had not had the opportunity to acquire knowledge through individual learning but does have the chance to interact with many people, then she/he will be able to cover, in the long run, the initial gap through interactive learning. On the other hand, if a person has a low initial level of knowledge and she/he lives in a peripheral neighbourhood, then she/he will be marginalised from society, being actually caught in a *poverty trap* mechanism.

5.3 SMALL WORLD CALCULATIONS

As discussed earlier, the simulation aims to observe the knowledge distribution dynamics as well as the network properties of the model. To tackle the issue of network properties, we follow the same methodology used in Morone and Taylor (2001). 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.

For calculating the average path length and cliquishness of a random network, we shall use the same approximation as Watts and Strogatz (1998) that $L_{random}(t) \cong \ln N / \ln n$ and $C_{random}(t) \cong n/N$, where n is the average number of connections of each agent and N is the total number of agents. The criteria for identifying the network as small worlds are that $L(t) \cong L_{random}(t)$ and $C(t) >> C_{random}(t)$.

We shall compare our dynamic network with a random one at different stages in the simulation to show whether or not the small worlds architecture is emerging in our system. Since the number of connections in our network is not constant (due to the mechanism by which agents make acquaintances of their acquaintances), in order to make an appropriate comparison we need to construct the random network with an equivalent number of connections. We shall, therefore, construct several random networks, one for comparison at each stage of interest of the simulation. The average path length and cliquishness will be calculated at both stages in the simulation. If, when comparisons are made with the random network, we find that the Watts-Strogatz criteria are observed, this will be evidence to suggest that a small worlds network structure is emergent from our model.

In Table 3 below we present our results for the small world calculations. Looking at this figure we observe that L (average path length) is substantially stationary, while on the other hand, C (average cliquishness) first decries and then increases again. The initial situation is definitely a small world with $L(t) \cong L_{random}(t)$ and $C(t) >> C_{random}(t)$ (i.e. where short average path length and high degree of cliquishness coexist). Then, after fifty cycles, the architecture of the network changes slightly (with a cliquishness which decreases by more than 60%). After another 50 cycles, the network evolves again into a small world and remains stable until the end of the simulation.

Table 3. Small world calculation results

		C	*L*
Simulation Results	Cycle 1	0.4396	2.1727
	Cycle 50	0.1425	2.3082
	Cycle 100	0.2661	2.9281
	Cycle 150	0.2853	3.2699
	Cycle 200	0.2578	2.9669
Random Network	Cycle 1	0.0388	2.4789
	Cycle 50	0.0345	2.6193
	Cycle 100	0.0345	2.6193
	Cycle 150	0.0302	2.7991
	Cycle 200	0.0302	2.7991

Source: Simulation results

From the results obtained we can conclude that the network structure of the model evolves over time. This dynamic process is particularly interesting since it shows that the network evolves into a very efficient network structure (i.e. small world), by means of an individual learning process, which takes place through the model of preferential acquaintance selection described above. In other words, what is efficient for the individual is also efficient for the network as a whole.

6. CONCLUSION

In this paper we have modelled the dynamic process of knowledge diffusion, where knowledge is structured non-linearly as a c*ognitive map*. The process through which agents interact allows each agent to learn from her/his acquaintances and to increase her/his knowledge in the learning region of the *CM*. Whether or not the interaction is gainful will depend upon the precise structure of that knowledge. The simulation was calibrated on the Greater Santiago region, using a sub-sample of the original database obtained from the *Encuesta de Ocupación y Desocupación* collected by the University of Chile for the year 1998.

We have also studied the network properties of the model over the whole simulation. We observed how the network architecture of the model changed over time, showing that the model evolved into a very efficient network structure (i.e. small-world) by means of an individual learning process, which takes place through the model of preferential acquaintance selection.

The results concerning the knowledge diffusion process are very interesting: in presence of high level of (knowledge) inequality there is a high risk of exclusion for those people initially endowed with low level of education. Moreover, we studied the spatial dimension of the exclusion process, finding that the *ignorance trap* mechanism is more likely to take place if an initial situation of low level of knowledge is coupled with geographical exclusion. In other words, those people who start with a high level of individual learning (i.e. schooling) will always be able to escape from the *ignorance trap* mechanism, while more backward people might be trapped if their low level of knowledge is cumulated with geographical exclusion. To summarise, each agent can find her/himself in one of the following four situations:

Figure 15. Possible long-term equilibria

	Centre and Semi-Periphery	Periphery
Low level of schooling	Long term equilibrium: **Converge towards high equilibrium**	Long term equilibrium: **Diverge towards low equilibrium**
High level of schooling	Long term equilibrium: **Converge towards high equilibrium**	Long term equilibrium: **Converge towards high equilibrium**

Source: Simulation results

These findings are extremely important from a policy prescription perspective. In the light of the new findings, two possible policy actions could avoid the occurrence of an *ignorance trap*: the policy maker could increase the level of education of more backward and marginalised peoples, and/or reduce the geographical gap between centre and periphery. This latter policy could be implemented through the development of infrastructure bridging the centre-periphery distance, as well as through the development of ICT. In other words, the *exclusion risk* could be minimised through the development of a more comprehensive cyber-network, so that also peripheral agents will have the same opportunity to interact with central and semi-peripheral agents.

ANNEX

Table A1 Schooling distribution, original sample

Schooling	Freq.	Percent	Cum.
0	144	1.86	1.86
1	26	0.34	2.2
2	68	0.88	3.08
3	134	1.73	4.81
4	159	2.06	6.87
5	127	1.64	8.51
6	390	5.04	13.55
7	157	2.03	15.58
8	724	9.36	24.95
9	553	7.15	32.1
10	625	8.08	40.18
11	480	6.21	46.39
12	2155	27.87	74.26
13	415	5.37	79.63
14	312	4.04	83.67
15	205	2.65	86.32
16	289	3.74	90.05
17	491	6.35	96.4
18	231	2.99	99.39
19	47	0.61	100
Total	7732	100	

Table A2 Schooling distribution, simulation sample

Schooling	Freq.	Percent	Cum.
0	1	0.43	0.43
2	2	0.86	1.29
3	2	0.86	2.16
4	9	3.88	6.03
5	5	2.16	8.19
6	10	4.31	12.5
7	6	2.59	15.09
8	20	8.62	23.71
9	16	6.9	30.6
10	20	8.62	39.22
11	11	4.74	43.97
12	68	29.31	73.28
13	12	5.17	78.45
14	6	2.59	81.03
15	5	2.16	83.19
16	12	5.17	88.36
17	19	8.19	96.55
18	8	3.45	100
Total	232	100	

Table A3 Geographical distribution, original sample

District	Freq.	Percent	Cum.
1	253	3.27	3.27
2	98	1.27	4.54
3	377	4.88	9.42
4	195	2.52	11.94
5	199	2.57	14.51
6	86	1.11	15.62
7	198	2.56	18.18
8	176	2.28	20.46
9	177	2.29	22.75
10	145	1.88	24.62
11	211	2.73	27.35
12	290	3.75	31.1
13	305	3.94	35.05
14	270	3.49	38.54
15	242	3.13	41.67
16	283	3.66	45.33
17	153	1.98	47.31
18	138	1.78	49.09
19	515	6.66	55.76
20	156	2.02	57.77
21	86	1.11	58.89
22	261	3.38	62.26
23	73	0.94	63.2
24	103	1.33	64.54
25	208	2.69	67.23
26	129	1.67	68.9
27	466	6.03	74.92
28	151	1.95	76.88
29	77	1	77.87
30	307	3.97	81.84
31	180	2.33	84.17
32	202	2.61	86.78
33	529	6.84	93.62
34	493	6.38	100
Total	7732	100	

Table A4 Geographical distribution, simulation sample

District	Freq.	Percent	Cum.
1	5	2.16	2.16
2	2	0.86	3.02
3	20	8.62	11.64
4	3	1.29	12.93
5	8	3.45	16.38
6	2	0.86	17.24
7	5	2.16	19.4
8	6	2.59	21.98
9	7	3.02	25
10	1	0.43	25.43
11	8	3.45	28.88
12	9	3.88	32.76
13	6	2.59	35.34
14	6	2.59	37.93
15	5	2.16	40.09
16	6	2.59	42.67
17	1	0.43	43.1
18	5	2.16	45.26
19	15	6.47	51.72
20	11	4.74	56.47
21	3	1.29	57.76
22	5	2.16	59.91
23	3	1.29	61.21
24	4	1.72	62.93
25	8	3.45	66.38
26	4	1.72	68.1
27	15	6.47	74.57
28	6	2.59	77.16
29	3	1.29	78.45
30	8	3.45	81.9
31	6	2.59	84.48
32	6	2.59	87.07
33	13	5.6	92.67
34	17	7.33	100
Total	232	100	

Table A5 Distribution of kind of school, original sample

Kind	Freq.	Percent	Cum.
0	139	1.8	1.8
1	1775	22.96	24.75
2	2995	38.74	63.49
3	1029	13.31	76.8
4	474	6.13	82.93
5	1305	16.88	99.81
6	15	0.19	100
Total	7732	100	

Table A6 Distribution of kind of school, simulation sample

Kind	Freq.	Percent	Cum.
0	1	0.43	0.43
1	53	22.84	23.28
2	89	38.36	61.64
3	33	14.22	75.86
4	13	5.6	81.47
5	42	18.1	99.57
6	1	0.43	100
Total	232	100	

Table A7 Distribution of use of computer at work, original sample

Computer	Freq.	Percent	Cum.
0	6601	85.37	85.37
1	1131	14.63	100
Total	7732	100	

Table A8 Distribution of use of computer at work, simulation sample

Computer	Freq.	Percent	Cum.
0	199	85.78	85.78
1	33	14.22	100
Total	232	100	

Figure A1. Map of Greater Santiago Region and District Codes

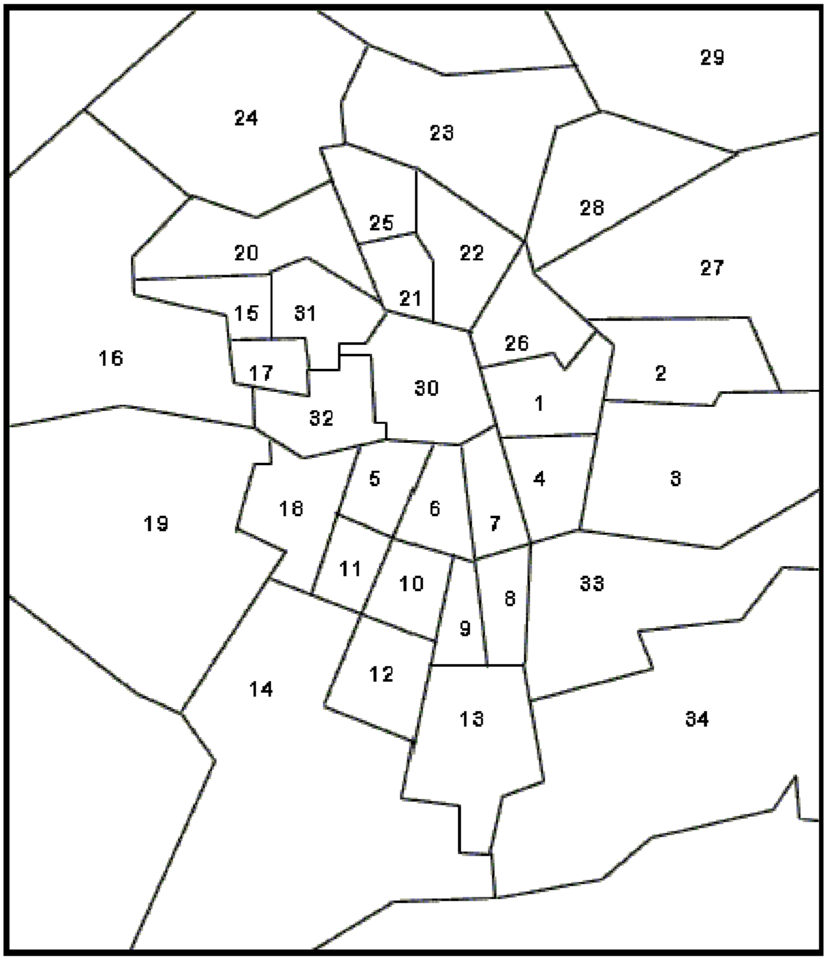

District	Code
ÑUÑOA	1
LA REINA	2
PEÑALOLEN	3
MACUL	4
PAC	5
SAN MIGUEL	6
SAN JOA	7
LA GRANJA	8
SAN RAMON	9
LA CISTERNA	10
LO ESPEJO	11
EL BOSQUE	12
LA PINTANA	13
SAN BERNARDO	14
CERRO NAVIA	15
PUDAHUEL	16
LO PRADO	17
CERRILLOS	18
MAIPU	19
RENCA	20
INDEPENDENCIA	21
RECOLETA	22
HUECHURABA	23
QUILICURA	24
CONCHALI	25
PROVIDENCIA	26
LAS CONDES	27
VITACURA	28
LO BARNECHEA	29
SANTIAGO	30
QUINTA NORMAL	31
ESTACION CENTRAL	32
LA FLORIDA	33
PUENTE ALTO	34

Macro-districts Composition:

Centre: 1, 4, 5, 6, 7, 21, 26, 30, 31, 32;

Semi-periphery: 8, 9, 10, 11, 12, 13, 15, 17, 18, 20, 22, 25;

Periphery: 2, 3, 14, 16, 19, 23, 24, 27, 28, 29, 33, 34.

Notes

[1] See among others: Beyer 2000; Morone 2001.
[2] See, among others, Ellison (1993, 2000), Anderlini and Ianni (1996), Berninghaus and Schwalbe (1996), Goyal (1996), Akerlof (1997), Watts (2001).
[3] An intuitive example would be that of a student who stops interacting with her/his teacher when she/he feels that she/he has nothing to learn from the interaction, or when her/his knowledge surpasses that of her/his teacher. Another example could be that of two equally educated agents which might stop communicating when they become specialised in two very different areas, so no shared interest remains.
[4] In this way we assume that agents are constrained by 'bounded rationality' in the sense that they respond to utility signals. This does not mean that they maximise utility (Katz, 2001).
[5] Moving from left to right, each pair of years of schooling corresponds to a complete column activated. We set the initial level of education (0 to 2 years of schooling) as three full columns activated (with the third column containing four nodes), then additional pairs of years of schooling correspond to an additional column in the cognitive map.
[6] After the eighth year of schooling students choose to specialise either in humanistic-scientific general education or technical-professional education.
[7] We define the cognitive map only as a function of X and N because at this stage we are not interested in the depth of knowledge.
[8] In the simulation model p is set equal to 0.1.
[9] We refer explicitly to the work of Cohen and Levinthal (1989) on returns from R&D. The concept of individual absorptive capacity has already been developed in Morone (2001).
[10] This region, called Región Metropolitana, includes 42% of the economically active national population (see: Bravo et all 1999:7).
[11] Chilean schools can be entirely funded by the government, partially funded by the government or entirely private.
[12] Agents younger than 14 years old are probably engaged mainly in individual learning as schooling represent a major part of their life, while people over 65 are probably less likely to be involved in knowledge exchange processes (because they are already too wise!).
[13] Given the structure of knowledge expressed by the cognitive map, we calculate changes in knowledge based on the total number of activated nodes for each agent.
[14] Longer simulations show how the curve completely levels off, meaning that the system reaches a long-run stable equilibrium.
[15] The only change was in the preferential selection: the value τ_i defined in equation (3) was not updated so it remained equal to 1. No acquaintances were dropped. Apart from that the two simulations here compared are identical.
[16] We studied distributional changes by means of inequality and poverty measures used typically in income distribution analysis. We did so because we are interested in studying the distributional properties of the model as well as its network and efficiency properties.

[17] We call ignorance line what is usually called poverty line.
[18] A detailed map with the district codes is reproduced in the annex.
[19] A detailed description of each macro-district composition is reproduced in the annex.

Reference

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.

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.

Anderlini, L. and A. Ianni (1996), "Path Dependence and Learning from Neighbours", *Games and Economic Behaviour*, 13: 141-177.

Appleton, S. and A. Balihuta (1996), "Education and agricultural productivity: evidence from Uganda", *Journal of International Development*, 8(3): 415-444.

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

Ashenfelter, O. (1991), "How Convincing is the Evidence Linking Education and Income?", *Industrial Relations Sections, Working Paper*, 292.

Ashenfelter, O. and A. Krueger (1992), "Estimates of the Economic Returns to Schooling From a New Sample of Twins", *NBER, Working Paper*, 4143.

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.

Becker, G. S. (1964), *Human capital, a theory and empirical analysis, with special reference to education*, New York: NBER.

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, & L.C. Freeman, (1999), *UCINET 5.0 Version 1.00. Natick: Analytic Technologies*.

Breschi, S. and F. Lissoni (2003), "Mobility and Social Networks: Localised Knowledge Spillovers Revisited", mimeo, paper presented at the DRUID *Academy Winter 2003 Conference*.

Brock, W. and S. Durlauf (1995), "Discrete Choice with Social InteractionI:Theory" *NBER Working Paper*, 5291.

Card, D. (2001), "Estimating the Return to Schooling: Progress on Some Persistent Econometric Problems", *Econometrica*, 69(5): 1127-60.

Case, A. and L. Katz (1991), "The Company you Keep: The Effects of Family and Neighbourhood on Disadvantaged Families", *NBER Working Paper*, 3705.

Chwe, M. S-Y. (2000), "Communication and Coordination in Social networks", *Review of Economic Studies*, 67: 1-16.

Cohen, W. M. and D. A. Levinthal (1989), "Innovation and Learning: The Two Faces of R&D", *The Economic Journal*, 99(September): 569-596.

Coleman, J. (1966), *Equality of Education Opportunities*, Washington DC: US Government Printing Office.

Cornes R., and T. Sandler, "The theory of externalities, public goods and club goods", Cambridge University Press, Cambridge MA, 1996.

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.

Dougherty, C. R. S., and E. Jimenez (1991), "The Specification of Earning Functions: Tests and Implications", *Economics of Education Review*, 10(2): 85-98.

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.

Ellison G. and D. Fudenberg (1995), "Word-of-Mouth Communication and Social Learning", *Quarterly Journal of Economics*, 110(1): 93-125.

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

Goethals, G., G. Winston and D. Zimmerman (1999), "Students Educating Students: the Emerging Role of Peer Effects in Higher Education", in M. E. Devlin and J. W. Meyerson,(eds.), *Future Forums: 1999 Papers*, New Haven, Conn. Forum Publishing.

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.

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

Haveman, R. and B. Wolf (1994), *Succeeding Generations: on the Effect of Investments in Children*, Russell Sage Foundation, New York.

Jovanovic, B. and R. Rob (1989), "The growth and diffusion of Knowledge", Review of Economic Studies, 56(198-9): 569-82.

Katz, J. (2001), 'Structural Reforms and Technological Behaviour. The Sources and Nature of Technological Change in Latin America in the 1990s', *Research Policy*, 30: 1-19.

Kiefer, N. (1989), "A Value Function Arising in Economics of Information", *Journal of Economic Dynamics and Control*, 7: 331-337.

Lucas, R. (1988), "On the Mechanics of Economic Development", *Journal of Monetary Economics*, 22: 3-42.

Mansfield, E. (1968), *Economics of Technological Changes*, Norton, New York.

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

Mincer, J. (1974), *Schooling, Experience and Earnings,* New York: Columbia University Press.

Morone P. (2001), "The Two Faces of Knowledge Diffusion: the Chilean Case", *Universidad de Chile*, *Departamento de Economia*, *Working Paper*, 178.

Morone, P. and R. Taylor, (2001), "Knowledge Diffusion Dynamics and Network Properties of Face-to-Face Interactions", mimeo, University of Sussex.

Morone, P. and R. Taylor, (2003), "Learning Dynamics with Complex Cognition", mimeo, University of Roma 'La Sapienza'.

Morris, S. (2000), "Contagion", *Review of Economic Studies*, 67: 57-68.

Murnane, R., J. B. Willett and F. Levy (1995), "The Growing Importance of Cognitive Skills in Wage Determination", *Review of Economics and Statistics*, 77: 251-66.

Nelson, R. and E. Phelps (1966), "Investment in Humans, Technological Diffusion and Economic Growth", *American Economic Review*, 56 (1/2): 69-75.

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

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

Rauch, J. (1993), "Productivity Gains From Geographic Concentration of Human Capital: Evidence From the Cities", *Journal of Urban Economics*, 34: 380-400.

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

Ryan, B. and N. Gross (1943), "The diffusion of hybrid seed corn in two Iowa communities", *Rural Sociology*, 8(1): 15-24.

Saviotti, P. P. (1998), 'On the Dynamic of Appropriability, of Tacit and Codified Knowledge', *Research Policy*, 26: 843-856.

Slicher van Bath, B. H. (1963), *The Agrarian History of Europe, A.D. 500-1850*, London: Arnold.

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

Summers, A. and B. Wolfe (1977), "Do School Make a Difference? ", *American Economic Review*, 67(4): 649-52.

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

Wallis, S., & S. Moss, (1994), *Efficient Forward Chaining for Declarative Rules in a Multi-Agent Modelling Language*, available online at: http://www.cpm.mmu.ac.uk/cpmreps

Watts, A. (2001), "A Dynamic Model of Network Formation", *Games and Economic Behaviours*, 34: 331-341.

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

Watts, D. and S. Strogatz (1998), Collective Dynamics of Small-World Networks, Letters to Nature, 393, June: 440-442.

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.