Simulating I-Space (SIS):
An Agent-based Approach to Modeling Knowledge Flows[*]

MAX BOISOT
IN3, Universitat Oberta de Catalunya, Av. Tibidabo, 39-43, 08035 Barcelona, Catalonia, Spain
and Sol C. Snider Entrepreneurial Research Center, Vance Hall, Wharton School, University of Pennsylvania, PA 19104-6374, USA.
E-mail: boisot@attglobal.net

AGUSTÍ CANALS
IN3, Universitat Oberta de Catalunya, Av. Tibidabo, 39-43, 08035 Barcelona, Catalonia, Spain
and Sol C. Snider Entrepreneurial Research Center, Vance Hall, Wharton School, University of Pennsylvania, PA 19104-6374, USA.
E-mail: acanalsp@uoc.edu

IAN MACMILLAN
Sol C. Snider Entrepreneurial Research Center, Vance Hall, Wharton School, University of Pennsylvania, Philadelphia, PA 19104-6374, USA.

Word count: 7433

---

[*] Previous versions of this paper were presented at the Agent-Based Simulation Conference in Montpellier and at the 2003 EGOS Colloquium in Copenhagen.

Keywords

Information, Knowledge Management, Agent-Based Simulation, Model Validation, Schumpeterian Economy.

Abstract

Knowledge management often generates theories that are too general or abstract to be easily testable. In some cases, simulation modeling can help. In this paper we develop an agent-based simulation model derived from a conceptual framework, the Information Space or I-Space and use it to explore the differences between a neoclassical and a Schumpeterian information environment. After introducing the knowledge management issues involved, we first briefly present the conceptual framework. This is followed by a presentation of the agent-based model and a development of a number of hypotheses designed to help validate the model. We then present a number of model runs designed to test these hypotheses in a preliminary way. We find broad support for the hypotheses and conclude that the model exhibits enough consistency to warrant further development.

1. Introduction

Knowledge has always played a key role in economic processes, although economists themselves have tended to think of it as having only a support role (Mirowski, 2002; Boisot, 1995). In the last few years, however, knowledge has come to be acknowledged as perhaps *the* decisive economic resource for twenty-first century firms. This belated recognition of the importance of knowledge has given rise to a whole new set of practices aimed at improving the way that knowledge is created, used and transferred within and between organizations. These practices have been grouped under the label of knowledge management.

Although much knowledge management thinking can trace its origins to more established disciplines such as economics, sociology, philosophy or psychology (Nonaka and Takeuchi, 1995; Davenport and Prusak, 1998; Prusak, 2001), it still lacks theoretical foundations that it can call its own and therefore remains largely a practice rather than a full-fledged intellectual discipline. The new economy, however, needs its theoreticians, and, perhaps for this reason, the management of knowledge has recently generated a growing interest among academics to become one of the most dynamic among the sub-fields of management (Spender, 2002; Grant, 2002). The immediate theoretical need is for a deeper understanding of the information processes that drive agent behaviours in markets and organizations (Canals, 2002).

One of the main obstacles that stand in the way of knowledge management becoming a full-fledged intellectual discipline is the difficulty of building solid bridges between theory and empirical practice. Theoretical knowledge management models necessarily show a high degree of both complexity and abstraction that makes it difficult to operationalize them and to derive from them empirically testable hypotheses. And what is currently available empirically is either case-based and/or anecdotal with only tenuous links to theory (Boisot, 1995; 1998; Choo and Bontis, 2002).

The methodological challenge facing the fledgling field of knowledge management is compounded by a deep change that is taking place in our understanding of human organizations as objects of study. Briefly summarized, we are coming to view such organizations as instances of *complex adaptive systems* (CAS) (Kauffman, 1995; Holland, 1992; Gell-Mann, 1994). These systems exhibit emergent behaviours, that is, behaviours that can neither be reduced to their constituent components, nor predicted on the basis of how these components interact with each other. Such behaviours are nonlinear and in human organizations they reflect a complex interplay of feedback and feed-forward effects, compounded by the ability of human agents to forge alternative *representations* of their situations (Conte, 1999). In short, many of the emergent properties displayed by human organizations result from the fact that they are made up of knowledgeable agents capable of forging distinctive cultures and value systems for themselves (Axelrod and Cohen, 1999).

It is this capacity to generate representations of states of the world and of their own place in them that distinguish social systems such as organizations from other complex adaptive systems that we can find in nature. The agents that we take to make up such systems - these can vary according to our chosen level of analysis to cover individual human beings, firms, tribes, and other social groupings - are able to recognize emergent higher order structures, reason about them, and take them into account when formulating appropriate actions. It is thus representational capacity – ie, a capacity for having a certain kind of *knowledge*[1] – that gives rise to "second-order emergence" and that distinguishes social systems from other complex systems[2] (Gilbert and Troitzsch, 1999).

One of the main methodological challenges we meet in dealing with complex adaptive systems like human organizations resides in the difficulty of deriving predictive schemes from high-level theories

---

[1] Representational knowledge is but one kind of knowledge. *Embodied* knowledge is another (Clark, 1997).

[2] We are not excluding certain evolved animal societies from having such knowledge. It may reside either in the brains of individual agents – this might be the case with primates - or in the interactions taking place at the societal level – the case of certain insect societies (Resnick, 1994).

(Gilbert and Troitzsch, 1999). The kind of "middle-range" theorizing initially proposed by Merton (1968) can only take us part of the way since most non-linear systems cannot be understood analytically. There is seldom a set of equations that we can use to predict the future state of the system. For this reason, we resort to computational methods and attempt to apprehend the system by simulating its relevant features. In the case of social systems, of course, we can hardly expect to derive detailed predictions from simulation models – for reasons just given, their sheer complexity usually makes that impossible - but in running a simulation we are often able to gain some insights into how a given system works and to then generate hypotheses that can subsequently be tested empirically (Carley and Hill, 2001).

Simulation methods are progressively being applied to organizational phenomena, giving rise to fruitful new approaches to organization science like Computational Organizational Theory (Carley and Prietula, 1994; Lomi and Larsen, 2001). We believe that the same simulation methods could also usefully be applied in the new field of knowledge management. In this paper, we present a simulation model, Simulated I-Space, or *SIS*, that implements a conceptual framework for the analysis of information flows, Boisot's Information-Space or *I-Space* (Boisot, 1995; 1998; Boisot and Child, 1996; 1999). The usefulness of a general conceptual framework such as the I-Space depends on the theorizing that it gives rise to as well as to what these theories can predict. The under-determination of theory by facts, however – a key insight of the Duhem-Quine thesis (Duhem, 1914; Quine, 1953)– often makes it hard to derive robust predictions from such theorizing. Could a simulation approach bring a general conceptual framework such as the I-Space closer to the real world?

In what follows we first briefly outline the main features of the conceptual framework, I-Space (section 2), and of the SIS simulation model that it gave rise to (Section 3). In section 4, we initiate the validation of SIS by developing simple and testable hypotheses concerning the way that knowledge might flow respectively in a neoclassical and in a Schumpeterian regime. In section 5, we present the runs relevant to these hypotheses that were derived from the simulation model. We discuss these in section 6 and a conclusion follows in section 7.

2. The I-Space

As a conceptual framework, the I-Space is built on a simple and intuitively plausible premise: structured knowledge flows more readily and extensively than unstructured knowledge. If the highly situated and tacit knowledge of the Zen master, for example, is only accessible to a small number of disciples through prolonged face-to-face interactions, the abstract, symbolic knowledge of prices and quantities manipulated by a bond trader is often available to a global market in a matter of seconds. Human knowledge is built up through the twin processes of discrimination and association (Thelen and Smith, 1994). Framing these as information processes, the I-Space takes information structuring as being achieved through two cognitive activities: codification and abstraction.

Codification articulates and helps to distinguish from each other the categories that we draw upon to make sense of our world. The degree to which any given phenomenon is codified can be measured by the amount of data processing required to categorize it. Generally speaking, the more complex or the vaguer a phenomenon or the categories that we draw upon to apprehend it – ie, the less codified it is - the greater the data processing effort that we will be called upon to make.[3]

Abstraction, by treating things that are different as if they were the same (Dretske, 1981), reduces the number of categories that we need to draw upon in order to apprehend a phenomenon. When two categories exhibit a high degree of association – ie, they are highly correlated – one can stand in lieu of the other. The fewer the categories that we need to draw upon to make sense of phenomena, the more abstract our experience of them.

---

[3] The way that we measure codification bears more than a passing ressemblance to the way that Kolmogorov or Chaitin measure complexity (Kolmogorov, 1965; Chaitin, 1974).

Codification and abstraction work in tandem. Codification facilitates the associations required to achieve abstraction, and abstraction, in turn, by keeping the number of categories needed down to a minimum, reduces the data processing load associated with the act of categorization. Taken together, they constitute joint cognitive strategies for economizing on data processing. The result is more and usually better structured data. Better-structured data, in turn, by reducing encoding, transmission, and decoding efforts, facilitates and speeds up the diffusion of knowledge within a given a given population of agents while economizing on communicative resources.

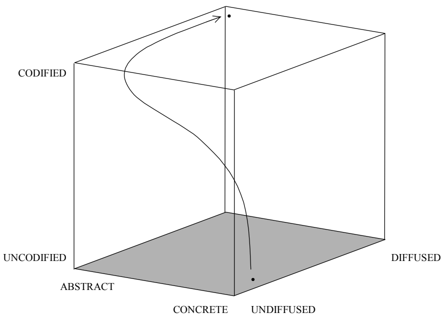

Figure 1: The Diffusion Curve in the I-Space

The relationship between the codification, abstraction and diffusion of knowledge is illustrated by the diffusion curve of Figure 1. The figure tells us that the more codified and abstract a given message, the larger the population of data processing agents that it can be diffused to in a given time period. Such agents might be individual human beings, but they might also be aggregates of these such as small groups, departments, or whole organizations, such as firms. All that is required to establish the candidacy of an agent is firstly, an ability to receive, process and transmit data to other agents within a population, and secondly, a capacity for unified agency

Codification, abstraction, and diffusion, make up only one part of a social learning process. Knowledge that is diffused within a target population must also get absorbed by that population and then get applied in specific situations. When applied, such knowledge may not fit in well with existing schema and may trigger a search for adjustments and adaptations – what Piaget described as a process of assimilation and accommodation (Piaget, 1967) and we shall refer to as scanning. The social learning process that we have just described forms a cycle in the I-Space – the Social Learning Cycle or *SLC* - that is illustrated in Figure 2. It is made up of six steps: scanning, codification, abstraction, diffusion, absorption, and impacting. that are outlined in Table 1. Many different shapes of cycle are possible in the I-Space, reflecting both the obstacles and the incentives to the learning process. Where learning leads to the creation of new knowledge, however, we hypothesize that the cycle will move broadly in the direction indicated by the figure.

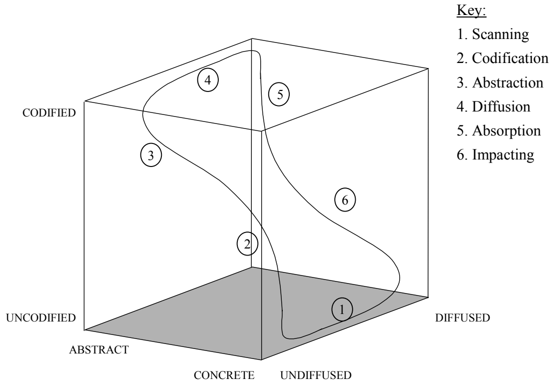

Figure 2: The Social Learning Cycle in the I-Space

Six Phases of SLC
1. **Scanning**
Identifying threats and opportunities in generally available but often fuzzy data – ie, weak signals. Scanning patterns such data into unique or idiosyncratic insights that then become the possession of individuals or small groups. Scanning may be very rapid when the data is well codified and abstract and very slow and random when the data is uncodified and context-specific
2. **Problem Solving**
The process of giving structure and coherence to such insights – ie., codifying them. In this phase they are given a definite shape and much of the uncertainty initially associated with them is eliminated. Problem-solving initiated in the uncodified region of the I-Space is often both risky and conflict-laden.
3. **Abstraction**

Generalizing the application of newly codified insights to a wider range of situations. This involves reducing them to their most essential features – ie., conceptualizing them. Problem-solving and abstraction often work in tandem.
4. **Diffusion**
Sharing the newly created insights with a target population. The diffusion of well codified and abstract data to a large population will be technically less problematic than that of data which is uncodified and context-specific. Only a sharing of context by sender and receiver can speed up the diffusion of uncodified data; the probability of a shared context is inversely achieving proportional to population size.
5. **Absorption**
Applying the new codified insights to different situations in a "learning by doing" or a "learning by using" fashion. Over time, such codified insights come to acquire a penumbra of uncodified knowledge which helps to guide their application in particular circumstances.
6. **Impacting**
The embedding of abstract knowledge in concrete practices. The embedding can take place in artifacts, technical or organizational rules, or in behavioural practices. Absorption and impact often work in tandem.

Table 1: The Six Phases of the Social Learning Cycle

In moving around an SLC, an agent incurs both costs and risks. There is no guarantee that the cycle can be completed. How, then, does an agent extract enough value from its learning processes to compensate for the efforts and risks incurred? If we take the term value in its economic sense, then it must involve a mixture of utility and scarcity (Walras, 1874). In the I-Space, utility is achieved by moving information up the space towards higher levels of codification and abstraction. Codification and abstraction together economize on data processing and transmission resources while increasing the reliability and generalizability of information so created. Scarcity, by contrast, is achieved by keeping the knowledge assets created located towards the left hand side of the diffusion curve – clearly, the scarcity of information will be inversely related to the number of people who possess it. Here we encounter a difficulty which is unique to knowledge goods. As indicated in figure 3, maximum value is achieved in the I-Space at point MV, that is, at the point where codification and abstraction are at a maximum and where diffusion is at a minimum. Yet, as can be seen from the diffusion curve, this is a point at which the forces of diffusion are also at a maximum. The point is therefore unstable and a cost must therefore be incurred – ie, patenting, secrecy, etc - to prevent diffusion taking place.

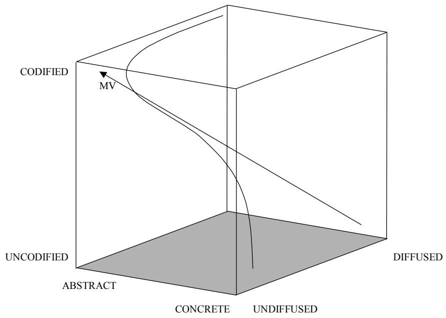

Figure 3: Maximum Value (MV) in the I-Space

With a knowledge good, then, and in contrast to the case of a purely physical good, utility and scarcity are inversely related. The greater the utility achieved, the more difficult it becomes to secure the scarcity necessary to extract full value from the good in question. The paradoxical nature of value in the case of an information good can be dealt with in two ways:

1. By hoarding – this strategy build upon the diffusion dynamics suggested by Figure 1. It is assumed that all the potential economic returns offered by a given knowledge asset over and above a normal accounting rate of return – ie, its economic *rent* - will be exhausted by the time it has diffused to the population as a whole. The strategy then consists of blocking or slowing down its diffusion in order to maintain its economic rent at some positive level. Since this is a strategy that is driven by equilibrium thinking that we associate with neoclassical economics, we label it Neoclassical learning or an *N-learning* strategy (Boisot, 1998).

2. By sharing – this strategy builds upon the learning dynamics suggested by Figure 2. It is assumed that new knowledge assets - and hence new value - are built up in the absorption, impacting and scanning phases of the SLC, that the creation of these new knowledge assets is in some way dependent on the knowledge assets currently being diffused, and that the value of these new knowledge assets will be greater than that lost through the erosion of scarcity brought about by diffusion of existing knowledge

assets. The strategy then consists of moving around the SLC faster than competitors in order to secure first-mover advantages in the creation of new knowledge and the destruction of existing knowledge. We label such a strategy Schumpeterian learning or an *S-Learning* strategy (Boisot, 1998).

It is clear that N-learning strategies focus on preserving existing knowledge whereas S-learning strategies focus on challenging or destroying it through the creation of new knowledge – ie, through innovation (Nelson and Winter, 1982). In section 4 we will develop hypotheses associated with these two strategies that could be explored via our simulation model – SIS - and subsequently subjected to empirical testing. We now turn to a description of the simulation model itself.

3. SIS: An Agent-based Simulation Model

Model Architecture

SIS is a multi-agent simulation characterized by mixture of competition and collaboration. Individual agents aim at nothing more than surviving, and although at this stage in the simulation's development they have no learning capacity – that is, they have no memory - the game as a whole displays elements of evolutionary behaviour. Profits are a means of survival and if agents run out of money they are 'cropped' from the simulation – ie, they are selected out. They can, however, also exit the simulation while they are still ahead.

How does SIS implement the concepts of the I-Space? The I-Space is a conceptual framework for analyzing the nature of information flows between agents as a function of how far such information has been structured through processes of codification and abstraction. Structuring information facilitates its flow and such flows, in turn, facilitate the exchange and diffusion of knowledge assets. Where given types of exchange are recurrent, they will form transactional patterns that can be institutionalized (Boisot, 1995). In SIS, we focus on the creation and exchange of knowledge assets *tout court* without concerning ourselves with the phenomenon of recurrence. In later versions of the model, recurrence will become our central concern.

SIS is populated with agents that carry knowledge assets in their heads. Each of these knowledge assets has a location in the I-Space that changes over time as a function of codification, abstraction, and diffusion processes as well as of how agents decide to exploit them. These, for example, have the possibility of exchanging their knowledge assets in whole or in part with other agents through different types of dealing arrangements. Knowledge assets, however, in addition to having their value eroded by uncontrolled diffusion processes, can also grow obsolete over time.

Natural selection is at work in SIS at two levels. Agents survive by making good use of their knowledge assets. They can make use of these assets directly to earn revenue – economic rents - or they can make indirect use of these assets by entering into trades with other agents who will then use them directly. Agents that fail to make direct or indirect use of their knowledge assets in a timely fashion are eventually likely to get selected out of the simulation – ie, to get “cropped”. On the other hand, knowledge assets, somewhat like Dawkins’ “memes” (Dawkins, 1999), “colonize” the heads of agents and survive by inhabiting the heads of as many agents as possible. If they fail to occupy at least one agent’s head, they also die out.

Existing agents have the option of quitting the game while they are ahead and before they are cropped. Conversely, new agents can be drawn into the game if the environment becomes sufficiently rich in opportunities. Here, the rate of new entries is based on mean revenues generated by the game in any given period. Both entry and exit rates are based on the difference in mean revenues between two periods. The rate of entry and exit is a parameter that is set at the beginning of the simulation and related to percentage changes in mean revenue.

SIS has three model components: 1) an agent component that specifies different agent characteristics; 2) an knowledge asset component that specifies the different ways that agents can invest in developing their knowledge assets; 3) an agent interaction component that specifies the different ways that agents can interact with each other. In what follows, we discuss each model component in turn. We start with agent characteristics and then describe the knowledge assets component. This is followed by a brief discussion of the agent interaction component.

1) The Agent Component:

SIS operates through a population of agents that make up the diffusion dimension of the I-Space. In the model as developed, agents are intended to represent organizations – firms or other types of information-driven organizations – within an industrial sector. It would be quite feasible, with suitable parameter settings, to have the agents represent individual employees within a firm and hence to simulate the behaviour of such agents within a single organizations. It would also be possible to have an individual agent representing the behaviour of a strategic business unit. Conversely, one could run SIS above the firm level and simulate knowledge flows within a population of industries, having each industry represented by an agent[4].

As we have already seen, agents can enter or exit SIS according to circumstances and can also be cropped from the simulation if their performance falls below a certain threshold. Agent entry and exit is an important source of variation within the simulation. Clearly, the agent population that is located along the diffusion dimension of the I-Space will vary in size at different moments in the simulation.

Agents aim to survive within the simulation and to maximize their wealth over the periods of the simulation. Agent wealth is expressed both in terms of *money* and in terms of *knowledge* and is taken to be the sum of revenue streams and of revenue-generating knowledge assets. Wealth expressed in terms of

---

[4] This assumes that an industry is sufficiently organized internally to exhibit the property of agency. The assumption may not hold in the case of fragmented industries.

money builds up a *financial fund*. Wealth expressed in knowledge terms builds up an *experience fund*. This latter fund constitutes an intangible asset that we can associate with the agent's intellectual capital. It is non-fungible. Agents modify their wealth either by changing the location of their knowledge assets in the I-Space, and hence altering their revenue-generating potential, or by trading in these assets with other agents thereby enlarging or shrinking their asset base. The details of how this is done are given under the heading of 'agent interaction'.

From their financial and experience funds, agents draw budgets that are used respectively for meetings and for investing in knowledge assets. Money that is not spend gets put back into the relevant fund and accumulates. The agent's preference for drawing from one type of fund or for another – ie, its bias either for acquiring its knowledge through meeting others or for acquiring it through creating some for itself - is set at the beginning of the game for all agents. The part of total revenues that go to one fund or the other can be switched with a toggle.

2) <u>The Knowledge Asset Component</u>:

In SIS, knowledge assets are represented in network form. A knowledge network consists of a collection of elements and of relations between elements. Depending on the application envisaged, the elements could consist of concepts, facts, or more aggregated knowledge domains. The relations between the elements, in turn, could consist of rules of inference, associations, or more complex links between knowledge domains. We shall refer to the elements of the network as *nodes* and to the relations between elements as *links*. Nodes and links can be combined with certain probabilities[5]. A knowledge asset, then, can either be a node or a link between two nodes. Each node and each link varies in how far it has been codified, made abstract, or has been diffused to other agents. Thus each node and link has a unique location in the I-Space that determines its value to the agent and hence its revenue-generating potential. The more codified and abstract a knowledge asset the greater its utility and hence the greater its value.

---

[5] Probability linkages only indirectly addresses the issue of *coherence* between elements. Our knowledge representations have but limited realism.

Likewise, the less diffused a knowledge asset, the scarcer it is and hence, again, the greater its value. Agents can enhance the value of their knowledge assets – and hence their revenue-generating potential – in two ways: 1) by investments in the Social Learning Cycle (SLC) that offer the possibility of changing the location of knowledge assets in the I-Space; 2) by combining nodes and links into networks that under certain circumstances can be *nested* and in this way building up more complex knowledge assets. The different locations in the I-Space thus have different revenue multipliers applied to them to reflect their different degrees of utility and scarcity. The proneness of the asset to diffusion or to obsolescence also varies with its location in the I-Space. The details of how this works are given in the appendix.

3) The Agent Interaction Component:

Agents meet each other throughout the simulation with a frequency that can be varied. When they meet they can ignore each other or they can attempt to engage in different types of transactions. In the second case, they need to be able to inspect each other's knowledge assets in order to establish whether a transaction is worth pursuing. Having established that it is, they can either: 1) engage in straight buying as selling of knowledge assets; 2) license other agents to use their knowledge assets; 3) enter into a joint-venture with another agent by creating a new agent that is jointly owned; 4) acquire another agent and convert it into a wholly-owned subsidiary; 5) merge with another agent, thus reducing the number of agents in the simulation. The cost of inspections and of agent interactions will be an inverse function of how codified and abstract the knowledge assets of interacting agents turn out to be.

SIS in the I-Space

The way that the simulation model maps onto the I-Space architecture is indicated by the flowchart of Figure 4.[6]

---

[6] Information on the detailed components of the model are can be sent to any scholar interested.

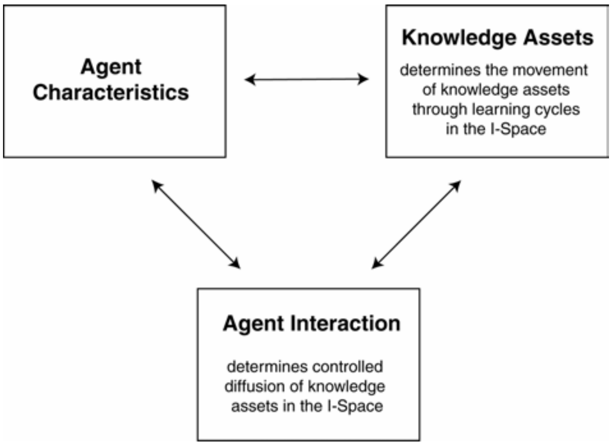

Figure 4: SIS architecture and the I-Space

The SIS parameters make it possible to fine-tune the simulation in order to model a variety of diverse situations with a considerable degree of flexibility, thus improving its grounding in the real world compared to other simulations. However, before it can be used to model specific cases, we need to assess the simulation's overall validity (Gilbert and Troitzsch, 1999). We need to ask how well it implements the basic assumptions that underpin our conceptual framework, the I-Space. We initiate the process in the following sections.

4. Validating SIS: Some Hypotheses

We now initiate the validation of SIS by developing a number of simple hypotheses relating to the general nature of N-learning and S-learning and testing these in a preliminary way. We draw on the work of Nelson and Winter (1982) as well as Schumpeter and standard micro-economics textbooks in order to compare two types of stylised economy: an economy run according to neoclassical prescriptions and one run on Schumpeterian lines[7]. These economies are not intended to be descriptively accurate. Rather, they reflect distinct theoretical orientations to be found in the literature. Like Nelson and Winter, we focus on

---

[7] Nelson and Winter (1982) term their evolutionary approach "neo-Schumpeterian" (p. 39), arguing that it supports a neo-Schumpeterian perspective.

industries rather than individual firms, that is, we work with populations of firms whose knowledge bases make up products which can substitute for each other (Porter, 1980).

The neoclassical paradigm, and particularly its hard core, the theory of general equilibrium - such as, ie, Arrow-Debreu - is empirically devoid of content (Blaug, 1992). The prescriptions of this paradigm are not thought to be suitable for a dynamic economy, one in which the constraints of friction and time operate. What we need do as we move from a neoclassical to a Schumpeterian economy, therefore, is to place an increasing number of such real world constaints on economic agents.

In contrast to Nelson and Winter, natural selection is built into the programming architecture of our simulation. However, our agents have no memory and make their decisions at random. Individual agents, therefore, cannot learn. Only the agent population as a whole can be said to learn, and this, only in the sense that random decisions by agents allow some to survive and others not. We must therefore emphasize that, its evolutionary orientation notwithstanding, not much discernable evolution takes place within the time frame in which our model operates.

We now briefly outline the features of our two economies

The Neoclassical Economy:

The first welfare theorem of the Arrow-Debreu model – that a competitive equilibrium allocation is Pareto-efficient – assumes that all economic agents have the same information about all economic variables (Arrow and Hahn, 1971; Debreu, 1959; Postlethwaite, 1989). In general equilibrium, markets clear (Bray, 1990). General equilibrium theory constitutes the core of the neo-classical research program (Weintraub, 1985; in Blaug, 1992, p. 168). However, the theory makes no predictions and all modern work of the Arrow-Debreu variety of the theory has been confined to 'existence theorems' – ie, theorems that state the conditions under which a general equilibrium system has a unique solution.

The perfect foresight assumption made by General Equilibrium theory (as qualified in Nelson and Winter, 1982, p. 8) requires full knowledge of existing and future market conditions. Learning, taken as the acquisition of knowledge over time, is therefore unecessary. Although perfect foresight is an equilibrium concept (Hicks, 1939; Bray, 1990), the information costs associated with reducing uncertainty or with transactions do not figure in the Arrow-Debreu General Equilibrium model (Hahn, 1984; in Blaug, 1992, p. 168). Resources are taken to be perfectly mobile and the knowledge resources that go into products or production processes, therefore, diffuse rapidly – if not instantaneously - between players. For this reason, perhaps, neoclassical economics has no theory of entrepreneurship and neither can nor has the need to create new knowledge (Baumol, 1968; in Nelson and Winter, 1982, p. 32).

In a neoclassical economic regime, competition is impersonal and atomistic. There is no rivalry leading to interaction – competitive or otherwise – between players (Fergusson, 1969). There are thus no meetings between industry players, no mergers no acquisitions, no joint ventures. There is no licensing of knowledge either. It is a requirement of neoclassical theorizing that agents remain small relative to the size of the market.

The above suggests that a neoclassical economy requires no activation of either the knowledge asset component or the agent interaction component of our simulation model. But our model cannot fully meet the neoclassical requirement of a perfectly homogeneous product with no differentiation, since the agents in the model have differing initial knowledge endowments, and, by implication differentiated product offerings. Nevertheless, rapid diffusion erodes the differentiation of their respective knowledge assets. And since, by assumption, agents cannot create new knowledge for themselves, differentiation can only erode over time. Although our model is based on the heterogeneity of firm endowments in knowledge assets, we can set our parameters to rapidly attain equilibrium conditions. In equilibrium, firms meet the no profit condition, only making an accounting profit equal to the rate of return in perfectly competitive industries.

The Schumpeterian Economy:

Schumpeter argued that innovation was constantly upsetting the move towards equilibrium and generating volatility. There is thus no assumption that equilibrium is the natural state towards which economic activity moves. The actions of time and change are now admitted, and, by implication, so is the creation of new knowledge. Technical externalities associated with new knowledge creation correspond to increasing returns to scale and hence to market failure (Arthur, 1989). Firms are represented as having certain capabilities that are a source of competitive advantage and that are modified by the creation of new knowledge, by the acquisition of knowledge in the market, and by random events. In a Schumpeterian economy, natural selection operates to weed out the poor performers (Nelson and Winter, 1982). Creative destruction is therefore at work in such an economy.

Schumpeter (1934), and then later J.M.Clark (1955) and J.K.Galbraith (1952), all pointed out that firms employ R&D as a competitive weapon (Nelson and Winter, 1982). New knowledge is also created by individual entrepreneurs who then use it to enter the market. Their entry, however, is often bunched, thus becoming a source of market boom and bust - one reason that economic behaviour is cyclical (Schumpeter, 1934). According to Schumpeter, economic growth and development is driven by the creation of new knowledge, which he describes as “new combinations”. In our simulation model, new combinations are the direct result of the *creation* of knowledge assets and the indirect result of the *trade* in knowledge assets - that is, from the activation of both the knowledge asset component and the agent interaction component of the simulation model.

In our simulation, we model only the *knowledge* on which firms base their output decisions. The positive returns earned by agents measure the economic rents their knowledge assets secure over and above an accounting profit. Clearly, given the diversity of initial agent endowments in knowledge, the simulation does not start in equilibrium, although it may move towards equilibrium if the ease of imitation erodes firm-specific advantage faster than it can be reconstituted by new knowledge creation. Recall that profits from innovation and new knowledge creation are disequilibrium phenomena (Schumpeter, 1934; Nelson and Winter, 1982).

From N-Learning to S-Learning

N-learning strategies reflect equilibrium thinking and the strategic concern of firms to avoid finding themselves in an equilibrium situation. They assume that knowledge gets codified, abstracted, and diffused. Although the structuring of knowledge is exogenously given, its diffusion is taken to be instantaneous. S-learning strategies, by contrast, reflect far-from-equilibrium thinking (Prigogine and Stengers, 1984). They assume that in addition to getting codified, abstracted, and diffused, knowledge gets absorbed, and impacted, and, in giving rise to scanning activities, it becomes a source of new knowledge and hence of competitive advantage. N-learning strategies express belief in a stabilizing neoclassical order whereas S-learning strategies express belief in a destabilizing Schumpeterian order, one characterized by “gales of creative destruction” (Schumpeter, 1934).

Our two stylized economies, the neoclassical and the Schumpeterian, represent different levels of activation of the SIS model components. The lowest level of activation models is what we have labeled the N-Learning condition. Here neither the knowledge creation component of the model nor the agent interaction component is activated. Only the agent characteristic component is used at this level. We associate these settings with a neoclassical economy, one in which new knowledge creation is absent as are the possibilities of collaboration between agents. We are thus dealing with the static case of pure competition that we shall label *the neoclassical case*. In a knowledge-based economy this situation is of little practical interest except as a baseline case. We do not expect to see anything very dramatic occurring in the baseline case. The highest level of activation models what we have labeled the S-Learning condition. Here all model components are brought into play. Both new knowledge creation and collaborative agent interactions are now allowed. We associate these settings with the Schumpeterian condition, one in which learning through time generate waves of creative destruction (Schumpeter, 1934). We label this *the Schumpeterian case*.

Neither of our chosen cases has any pretension to realism. They each model theories that have been held about limited aspects of economic behaviour and therefore find only distant echoes in the real world. Each of them, however, would be expected to exhibit the "logic" of its situation under simulated conditions. Like Nelson and Winter, we start from a case rooted in economic orthodoxy. Like them, we contrast it with a Schumpeterian condition, albeit – in contrast to Nelson and Winter – one derived from an agent-based approach. Since we cannot give them unambiguous theoretical or empirical referents, we do not model cases that are intermediate between the neoclassical and the Schumpeterian cases.

Each agent in SIS stands for a firm that competes or collaborates with other firms in its sector. The sector itself comprises an initial set of 20 agents. As the simulation progresses, new agents can be created through different mechanisms: in the neoclassical case, they can only be created from scratch; in the Schumpeterian case they can also be created as subsidiaries from other agents or as a consequence of mergers or through joint ventures. Each agent starts with an initial endowment of knowledge assets. In the neoclassical case, new knowledge can only enter the simulation with the entry of new agents. In the Schumpeterian case, new knowledge in the form of nodes and links may also be created during the simulation runs; this new knowledge will vary in its degree of codification and abstraction as well as in its level of complexity. In the neoclassical case, the diffusion of knowledge only takes place uncontrollably as a result of diffusion decay – ie, as a result of random diffusion forces. In the Schumpeterian case, the diffusion of knowledge will also take place either as a consequence of different kinds of interaction among agents.

Hypotheses

In this preliminary validation exercise, our aims are modest. We aim to model some tractable aspects of economic theory rather than some intractable aspect of the real world. If our hypotheses are corroborated - that is, if our results are consistent with some the claims of neoclassical and Schumpeterian theories - we will have taken a first step towards validating SIS. The seven hypotheses that we put forward are all derived from the literature that we have cited above. They compare the simulation outputs that we would expect from a neoclassical economic order with those that we would expect from a neoclassical one:

H1 Given the comparative stability of the neoclassical case in equilibrium as compared to the Schumpeterian regime of creative destruction, we would expect the average number of agents – ie, of market players - in the former case to be more stable than in the latter one.

H2 In a neoclassical case approaching equilibrium, when only an accounting profit can be earned, fewer agents have any incentive to enter the industry than in the Schumpeterian case.

H3 Fewer agents will get cropped in a neoclassical regime close to equilibrium conditions than under a Schumpeterian regime of creative destruction.

H4 But in a neoclassical regime close to equilibrium, more agents would have an incentive to leave the industry given the lack of economic rents available (ie, the lack of a knowledge-based competitive advantage) as evidenced by declining revenue rates.

H5. In the neoclassical case, since all agents are taken to share the same knowledge instantaneously and technical conditions are *given*, the only way that new knowledge enters the market is through new entrants. In the Schumpeterian case, by contrast, new knowledge creation by agents is admitted. This leads to the hypothesis that a larger number of new knowledge assets will be created in the Schumpeterian case than in the neoclassical one.

H6. This hypothesis is an indirect consequence of H5. Since new knowledge that is not instantly diffused constitutes a source of rents for agents, we would predict that revenues per agent would be higher in the Schumpeterian case than in the neoclassical one.

H7. This hypothesis is a consequence of H2, H3, H4, and H 6. With more agents entering the market, fewer exiting, and more revenues being earned per agent in the Schumpeterian than in the neoclassical case, then, in spite of more agents being cropped in the former case, we would expect a larger volume of revenue being available to agents in the first case than in the second.

Our seven hypothese are summarized in Table 2 below.

HYPOTHESES			
H	**Variable**	**Neoclassical Case**	**Schumpeterian Case**
H_1	**Agent numbers**	More stable	Less stable
H_2	**Agent entries**	Less entries	More entries
H_3	**Agents cropped**	Less agents cropped	More agents cropped
H_4	**Agent exits**	More exits	Less exits
H_5	**Total knowledge generated**	Less knowledge assets created	More knowledge assets created
H_6	**Revenues per agent**	Less revenues per agent	More revenues per agent
H_7	**Total revenues**	Less total revenues	More total revenues

Table 2: Seven Hypotheses

5. Results

We performed 200 runs for each of our four cases and used the average values for interpretation. The number of runs proved sufficient to show significant differences at the 95% level of confidence between our two cases. Each run terminated after 200 periods. In this exercise, we first focused on the number of agents that participate in the simulation, the number of agents that entered the game, the number of agents that were cropped from the game and the number that exited the game voluntarily. We then looked at the number of new knowledge assets – nodes and links - that were created during the game – these offer us a proxy measure of the overall social benefits of the economic order in which they were created. Finally, we looked at the total revenues generated as well as the revenue per agent. This gave us a measure of the price society had to pay the agents for the knowledge that they generated.

In what follows, we present the results of model runs, firstly in summary form in Table 3, and then secondly in graphic form that gives a better indication of how, in each case, the simulation evolved over time. In Table 3, we present the mean and standard deviation of relevant variables taken at three points in time - period 0, period 100, and period 199 - for each of the two cases modeled. We then plot the evolution of some of the variables in a graph, indicating the mean and the 95% interval of confidence for each period.

PERIOD		Number of agents	Accum. entries	Accum. crops	Accum. exits	Accum. K-assets creation	Rev. per agent	Accum. total revenues
Neoclassical case								
0	Mean	20,00	0,00	0,00	0,00	0,00	2,17	43,32
	S.D.	*0,00*	*0,00*	*0,00*	*0,00*	*0,00*	*0,27*	*5,32*
100	Mean	18,69	0,00	0,00	1,33	0,00	1,71	3.835,56
	S.D.	*1,19*	*0,00*	*0,00*	*1,22*	*0,00*	*0,17*	*589,49*
199	Mean	16,86	0,00	0,00	3,14	0,00	1,40	6.519,64
	S.D.	*1,82*	*0,00*	*0,00*	*1,88*	*0,00*	*0,12*	*926,31*
Schumpeterian case								
0	Mean	20,02	0,00	0,00	0,00	1,27	0,70	14,02
	S.D.	*0,12*	*0,00*	*0,00*	*0,00*	*1,05*	*0,11*	*2,17*
100	Mean	14,46	136,74	290,09	0,04	358,63	2,43	5.455,00
	S.D.	*3,13*	*12,96*	*39,40*	*0,21*	*53,76*	*0,48*	*732,63*
199	Mean	15,22	296,83	630,25	0,04	870,73	1,30	7.790,29
	S.D.	*4,42*	*22,41*	*48,29*	*0,21*	*79,48*	*0,31*	*1.165,00*

Table 3: Mean and Standard Deviation of variables at representative periods for each case.

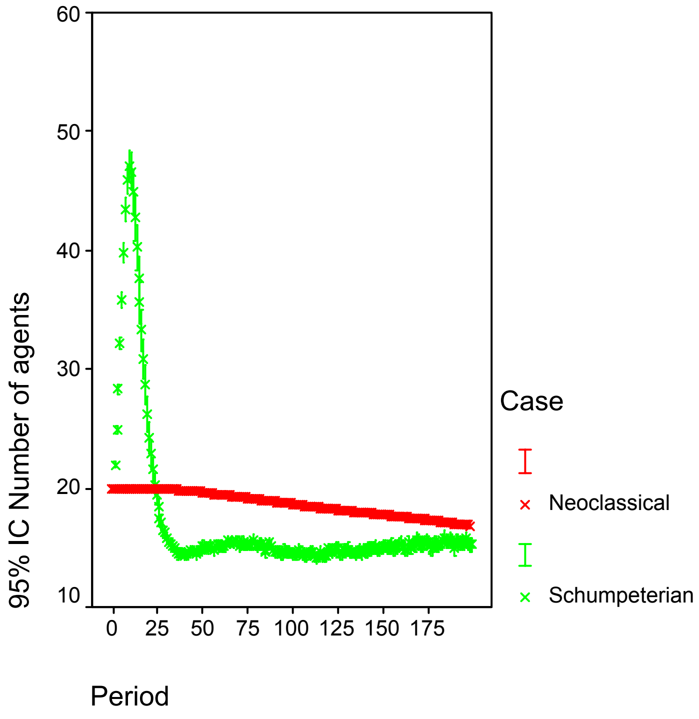

Figure 5: Evolution of the total number of agents (95% interval of confidence)

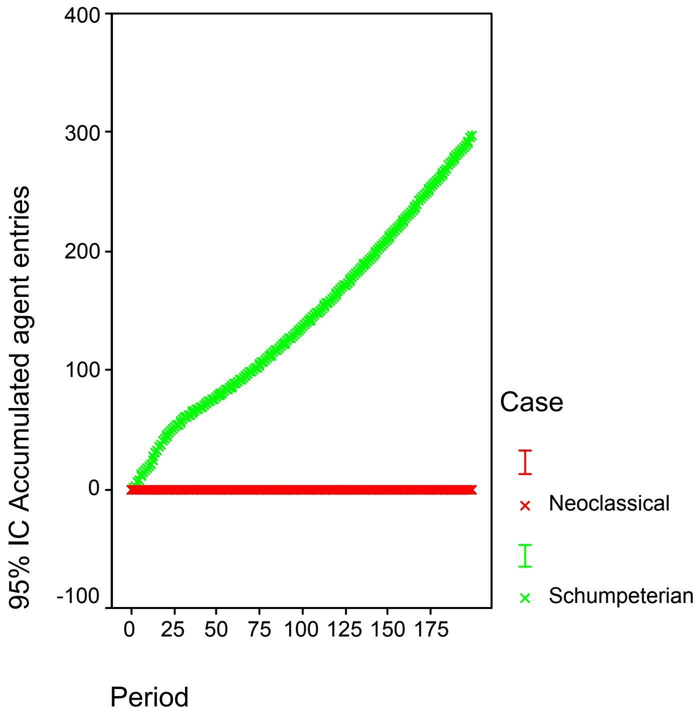

Figure 6: Accumulated entries of new agents (95% interval of confidence)

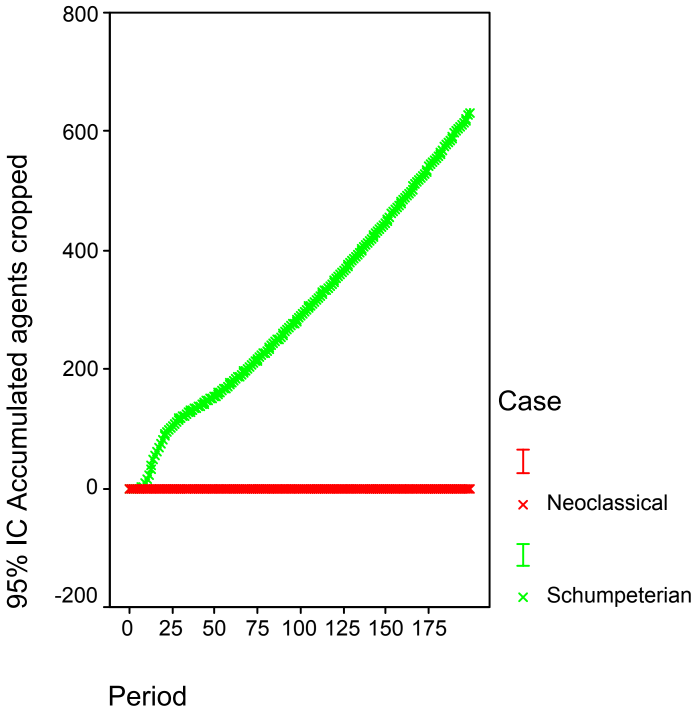

Figure 7: Accumulated number of agents cropped (95% interval of confidence)

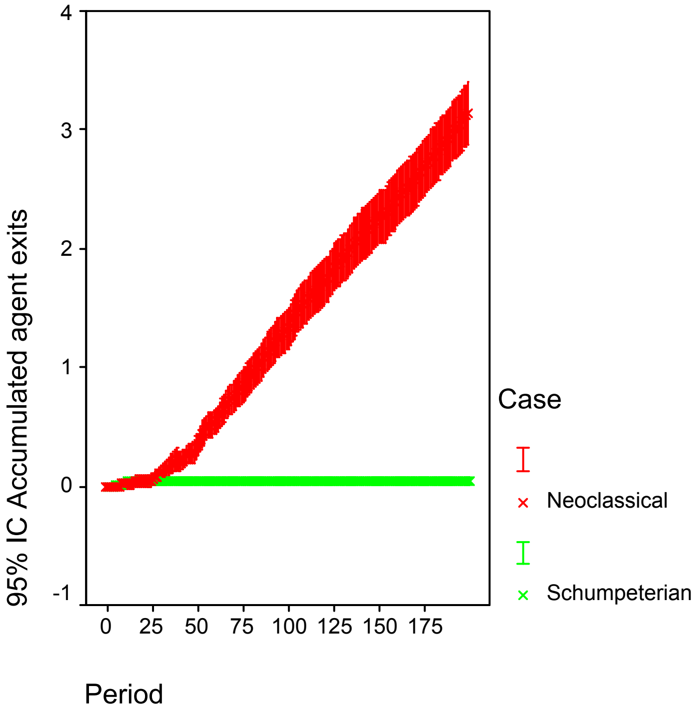

Figure 8: Accumulated number of exits (95% interval of confidence)

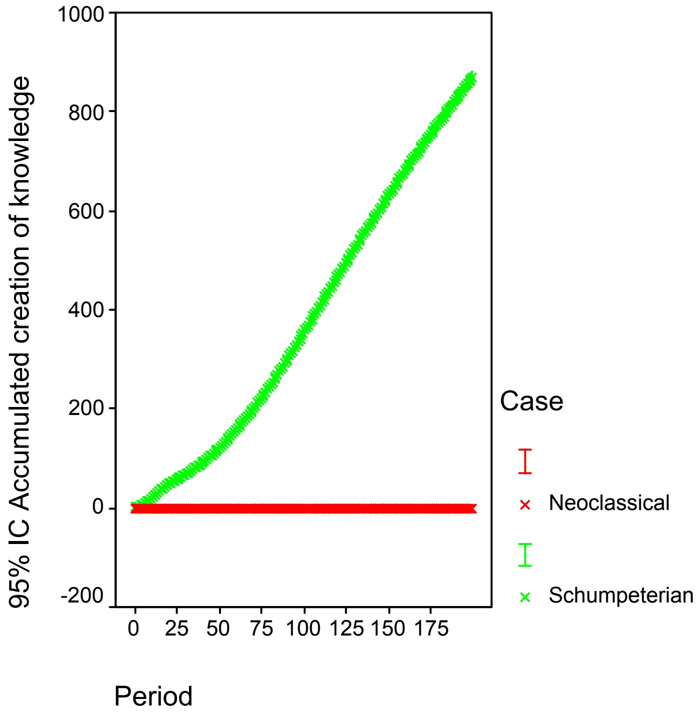

Figure 9: Accumulated number of knowledge assets created (95% interval of confidence)

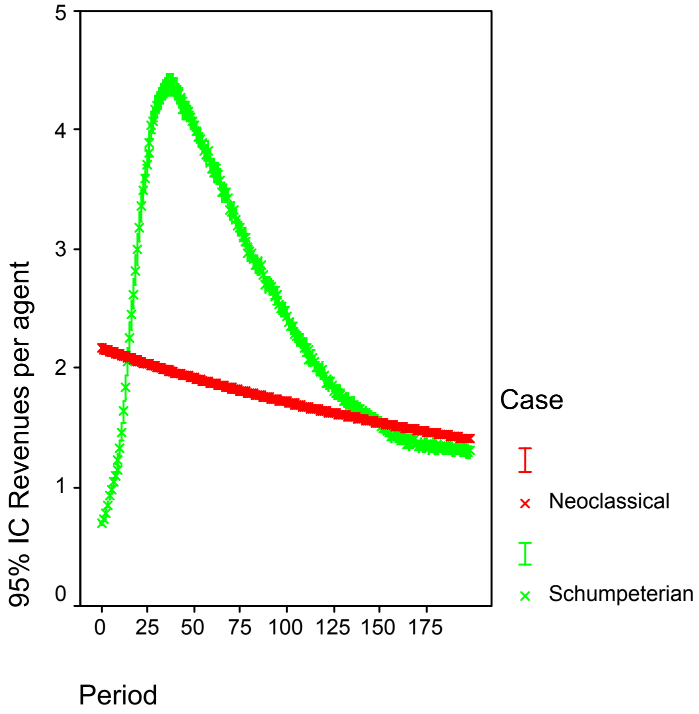

Figure 10: Average revenues per agent (95% interval of confidence)

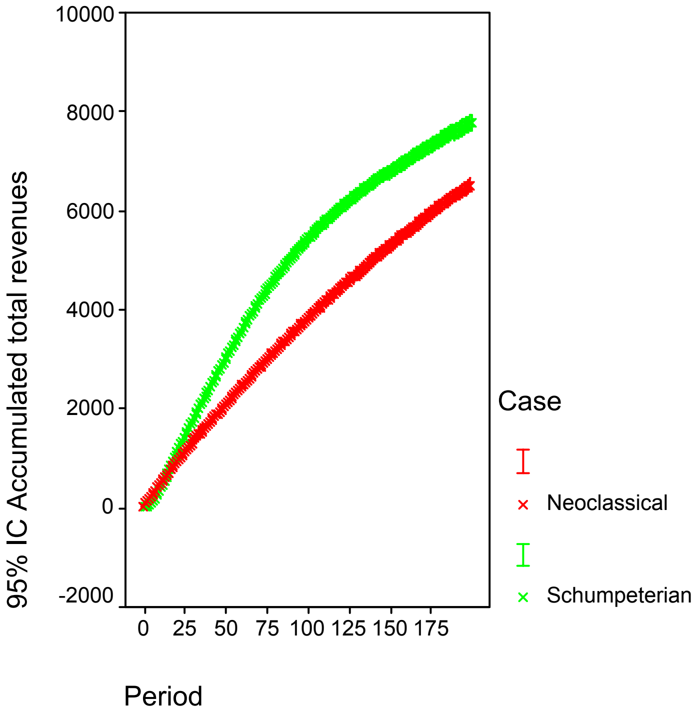

Figure 11: Accumulated total revenues (95% interval of confidence)

6. Discussion

In Table 2, we put forward seven hypotheses designed to see whether the simulation model could meaningfully distinguish between neoclassical and Schumpeterian economic processes as they manifest themselves in the I-Space. How strongly do the results of the simulation runs support or challenge these hypotheses? What do the results tell us about the conceptual framework or about the simulation model?

Applying the outputs of Table 3 to the hypotheses set out in Table 2 gives us the assessment of Table 4. We observe that all of the hypotheses that we put forward hold, thus supporting the claim that our model in a simple way can distinguish between neoclassical and Schumpeterian scenarios as manifested in knowledge-related processes.

HYPOTHESES ASSESSMENT				
H	**Variable**	**Neoclassical Case**	**Schumpeterian Case**	**Accepted or rejected**
H_1	**Agent numbers**	More stable	Less stable	Accepted
H_2	**Agent entries**	Less entries	More entries	Accepted
H_3	**Agents cropped**	Less agents cropped	More agents cropped	Accepted
H_4	**Agent exits**	More exits	Less exits	Accepted
H_5	**Total knowledge generated**	Less knowledge assets created	More knowledge assets created	Accepted
H_6	**Revenues per agent**	Less revenues per agent	More revenues per agent	Accepted
H_7	**Total revenues**	Less total revenues	More total revenues	Accepted

Table 4: Hypotheses assessment

7. Conclusions

This paper describes the first results of a research project that aims to use simulation modeling to generate empirically testable hypotheses concerning the knowledge-based behaviour of economic agents – individuals, firms, etc. It has to be recognized that our results are not particularly strong yet - suggestive rather than conclusive. We are still in the model validation phase of our research, and even here, only at the beginning. The next step in the validation process will be to examine how agents fare under different regimes of knowledge diffusion and obsolescence. This would allow us to relate different kinds of intellectual property rights regimes to different rates of technical change.

Although we would expect stronger support from a more refined model, we nevertheless, feel that the results that we have presented here are promising enough to justify further development of the simulation model. As we move beyond validation, we hope to gradually get closer to the real world, modeling real cases and exploring the fine grain of knowledge-based agent behaviour. It might be possible, for example, to model particular classes of agents and to endow them with memory. This would allow us to mix evolutionary and developmental processes – ie, phylogenetic and ontogenetic learning processes. In the first type of process, learning only takes place at the level of a population of agents. In the second type of process, learning can also take place at the level of the individual agent. Beyond that we would like to explore the structure of recurrent transactions between agents. What could it teach us about the nature of institutionalization? Under what circumstances will recurrent transactions look like market processes and under what circumstances will hierarchical transactions be favoured? Finally, we would like to be able to give labels to different kinds of knowledge assets in the simulation model. With suitable specification, we might then be able to simulate knowledge flows at the industry level. Our hope is that this kind of simulation-based work will help to place knowledge management on a sound theoretical footing.

REFERENCES

Arrow, K.J. and Hahn, F.H. (1971) *General Competitive Analysis*, San Francisco, CA.: Holden-Day, 1971.

Arthur, W.B. (1989) 'Competing Technologies, Increasing Returns, and Lock-In by Historical Events'. *Economic Journal* **99**, 394, 116-131.

Axelrod, R. and Cohen, M.D. (1999) *Harnessing Complexity: Organizational Implications of a Scientific Frontier*, New York, NY: The Free Press, 1999.

Baumol, W.J. (1968) 'Entrepreneurship in Economic Theory'. *American Economic Review* **58**, 2, 64-71.

Blaug, M. (1992) *The Methodology of Economics: Or How Economists Explain*, Cambridge, UK.: Cambridge University Press, 1992.

Boisot, M.H. (1995) *Information Space: A Framework for Learning in Organizations, Institutions and Culture*, London: Routledge, 1995.

Boisot, M.H. (1998) *Knowledge Assets: Securing Competitive Advantage in the Information Economy*, New York: Oxford University Press, 1998.

Boisot, M.H. and Child, J. (1996) 'From Fiefs to Clans and Network Capitalism: Explaining China's

Emergent Economic Order'. *Administrative Science Quarterly* **41**, 600-28.

Boisot, M.H. and Child, J. (1999) 'Organizations as Adaptive Systems in Complex Environments: The Case of China'. *Organization Science* **10**, 3, 237-252.

Bray, M. (1990) 'Perfect Foresight'. In: Eatwell, J., Millgate, M. and Newman, P., (Eds.) *Utility and Probability: The New Palgrave*, London, UK.: MacMillan, 1990.

Canals, A. (2002) 'Quo vadis, KM? La complexitat com a nou paradigma per a la gestió del coneixement'. IN3-UOC Working Paper Series, WP02-005. [http://www.uoc.edu/in3/dt/20000/index.html]

Carley, K.M. and Hill, V. (2001) 'Structural Change and Learning Within Organizations'. In: Lomi, A. and Larsen, E.R., (Eds.) *Dynamics of Organizations: Computational Modeling and Organizational Theories*, pp. 63-92. Menlo Park, CA.: AAAI Press/The MIT Press, 2001.

Carley, K.M. and Prietula, M.J. (1994) 'ACTS Theory: Extending the Model of Bounded Rationality'. In: Carley, K.M. and Prietula, M.J., (Eds.) *Computational Organizational Theory*, Hillsdale, N.J.: Lawrence Erlbaum, 1994.

Chaitin, G.J. (1974) 'Information-Theoretic Computational Complexity'. *IEEE Transactions on Information Theory* **20**, 1, 10-15.

Choo, C.W. and Bontis, N. (2002) *The Strategic Management of Intellectual Capital and Organizational Knowledge*, New York, NY.: Oxford University Press, 2002.

Clark, A. (1997) *Being There: Putting Brain, Body, and World Together Again*, Cambridge, MA.: The MIT Press, 1997.

Clark, J.M. (1955) 'Competition: Static Models and Dynamic Aspects'. *American Economic Review* **45**, 2, 450-462.

Conte, R. (1999) 'Social Intelligence Among Autonomous Agents'. *Computational & Mathematical Organization Theory* **5**, 3, 203-228.

Davenport, T.H. and Prusak, L. (1998) *Working Knowledge: How Organizations Manage What They Know*, Boston, MA.: Harvard Business School Press, 1998.

Dawkins, R. (1999) *The Extended Phenotype: The Long Reach of the Gene*, Oxford, UK.: Oxford University Press, 1982.

Debreu, G. (1959) *Theory of Value*, New York: Wiley, 1959.

Dretske, F.I. (1981) *Knowledge and the Flow of Information*, Stanford: CSLI publications, 1999.

Duhem, P. (1914) *La Théorie Physique: Son Object, Sa Structure*, Paris: Rivière et Cie., 1914.

Fergusson, C. (1969) *Microeconomic Theory*, Homewood, IL.: Richard Irwin, 1969.

Galbraith, J.K. (1952) *American Capitalism: The Concept of Countervailing Power*, Boston, MA.: Houghton Miffin, 1952.

Gell-Mann, M. (1994) *The Quark and the Jaguar* , New York: Freeman, 1994.

Gilbert, G.N. and Troitzsch, K.G. (1999) *Simulation for the Social Scientist*, London : Open University

Press, 1999.

Grant, R. (2002) 'The Knowledge-Based View of the Firm'. In: Choo, C.W. and Bontis, N., (Eds.) *The Strategic Management of Intellectual Capital and Organizational Knowledge*, Oxford, UK.: Oxford University Press, 2002.

Hahn, F.H. (1984) *Equilibrium and Macroeconomics*, Oxford, UK.: Basil Blackwell, 1984.

Hicks, J. (1939) *Value and Capital*, Oxford: The Clarendon Press, 1939.

Holland, J.H. (1992) *Adaptation in Natural and Artificial Systems*, Cambridge, MA.: MIT Press, 1975.

Kauffman, S. (1995) *At Home in the Universe: The Search for Laws of Complexity*, London: Penguin, 1996.

Kolmogorov, A.N. (1965) 'Three Approaches to the Quantitative Definition of Information'. *Problems in Information Transmissions* **1**, 3-11.

Lomi, A. and Larsen, E.R. (2001) *Dynamics of Organizations: Computational Modeling and Organization Theories*, Menlo Park, CA.: AAAI Press / The MIT Press, 2001.

Merton, R.K. (1968) *Social Theory and Social Structure*, New York, NY.: The Free Press, 1968.

Mirowski, P. (2002) *Machine Dreams: Economics Becomes a Cyborg Science*, Cambridge, UK.: Cambridge University Press, 2002.

Nelson, R.R. and Winter, S.G. (1982) *An Evolutionary Theory of Economic Change*, Cambridge, MA.: Belknap Press of Harvard University Press, 1982.

Nonaka, I. and Takeuchi, H. (1995) *The Knowledge-Creating Company*, New York: Oxford University Press, 1995.

Piaget, J. (1967) *Biologie Et Connaissance: Essai Sur Les Relations Entre Les Regulations Organiques Et Les Processus Cognitifs*, Paris: Gallimard, 1967.

Porter, M.E. (1980) *Competitive Strategy: Techniques for Analysing Industries and Competitors*, New York, NY.: The Free Press, 1980.

Postlethwaite, A. (1989) 'Asymmetric Information'. In: Eatwell, J., Millgate, M. and Newman, P., (Eds.) *Allocation, Information and Markets: The New Palgrave*, London, UK.: MacMillan, 1989.

Prigogine, I. and Stengers, I. (1984) *Order Out of Chaos: Man's New Dialogue With Nature*, Toronto: Bantam Books, 1984.

Prusak, L. (2001) 'Where did knowledge management come from?'. *IBM Systems Journal* **40**, 4, 1002-1007.

Quine, W.V.O. (1953) 'The Dogmas of Empiricism'. In: Quine, W.V.O., (Ed.) *From a Logical Point of View: 9 Logico Philosophical Essays*, pp. 20-46. Cambridge, MA.: Harvard University Press, 1961.

Resnick, M. (1994) *Turtles, Termites, and Traffic Jams: Explorations in Massively Parallel Microworlds*, Cambridge, MA.: The MIT Press, 1994.

Schumpeter, J.A. (1934) *The Theory of Economic Development: An Enquiry into Profits, Capital, Credit,* *Interest and the Business Cycle*, New Brunswick, NJ.: Transaction Publishers, 1983.

Spender, J.-C. (2002) 'Knowledge Management, Uncertainty, and an Emergent Theory of the Firm'. In: Choo, C.W. and Bontis, N., (Eds.) *The Strategic Management of Intellectual Capital and Organizational Knowledge*, pp. 149-162. New York, NY.: Oxford University Press, 2002.

Thelen, E. and Smith, L.B. (1994) *A Dynamic Systems Approach to the Development of Cognition and Action* , Cambridge, MA.: The MIT Press, 1994.

Walras, L. (1874) *Elements of Pure Economics or the Theory of Social Wealth*, Philadelphia, PA.: Orion Editions, 1984.

Weintraub, R. (1985) *General Equilibrium Analysis: Studies in Appraisal*, Cambridge, UK.: Cambridge University Press, 1985.