MODELING KNOWLEDGE-BASED ECONOMIC PROCESSES: A SIMULATION APPROACH

Max Boisot

Universitat Oberta de Catalunya, Barcelona, Catalonia, Spain.

Sol Snider Center for Entrepreneurial Research, The Wharton School, University of Pennsylvania, Philadelphia, USA.

e-mail: *boisot@attglobal.net*

Agustí Canals

Universitat Oberta de Catalunya, Barcelona, Catalonia, Spain.

e-mail: *acanalsp@uoc.edu*

Av. Tibidabo, 39-43
08035 Barcelona – Catalonia, Spain
Tel: +34 932 537 594
Fax: +34 934 176 495

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 four hypotheses designed to help validate the model. We then present a number of model runs designed to test these four hypotheses. We find relative 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 (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.

One of the main obstacles that stand in the way of knowledge management becoming a full-fledge 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 (Nonaka and Takeuchi, 1995; Davenport and Prusak, 1998).

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 feedforward effects, compounded by the ability of human agents to forge alternative representations of their situations. 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 knowledge – that gives rise to

"second-order emergence" and that distinguishes social systems from other complex systems[1] (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 (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 numerical 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).

Computational Organizational Theory applies simulation methods to organizational phenomena (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 – Smart Assets - inspired by a conceptual model of information flows, Boisot's Information-Space or I-Space. (Boisot, 1995; 1998; Boisot and Child, 1996; 1999). The usefulness of a general conceptual model such as the I-Space depends on what it can predict. The under-determination of theory by facts, however – a key insight of the Duhem-Quine thesis (Duhem, 1914; Quine, 1953)– makes it hard to derive robust predictions from general 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 Smart Assets simulation model inspired by the model (Section 3). In section 4, we initiate the validation of Smart Assets by developing a simple and testable hypothesis concerning the way that knowledge might flow respectively in a neoclassical and in a Schumpeterian regime as well as in intermediary cases. 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 develops a simple, intuitively plausible premise: structured knowledge flows more readily and extensively than unstructured knowledge. 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.

[1] 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).

Codification articulates 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.[2]

Abstraction reduces the number of categories that we need to draw upon 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.

Codification and abstraction work in tandem. Codification facilitates the associations required to achieve abstraction and abstraction in turn reduces the data processing load associated with the act of categorization. Taken together, they constitute joint strategies for economising 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 while economizing on communicative resources.

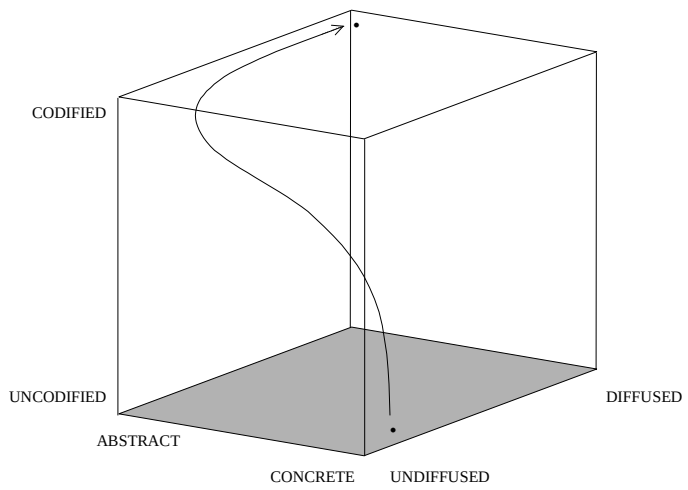

Figure 1: The Diffusion Curve in the I-Space

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 that it can be diffused to in a given time period. 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

[2] 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).

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. 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.

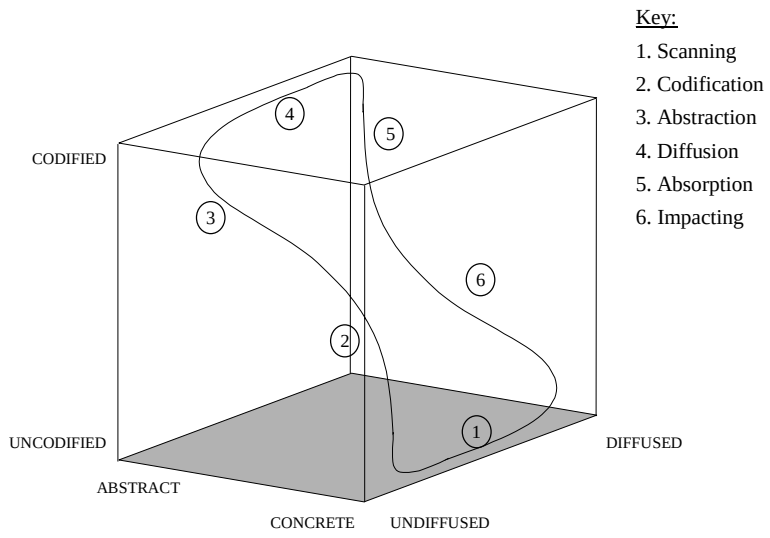

Figure 2: The Social Learning Cycle in the I-Space

In the I-Space, utility is achieved by moving up the space towards higher levels of codification and abstraction. Scarcity is achieved by keeping the knowledge assets created located towards the left hand side of the diffusion curve. 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.

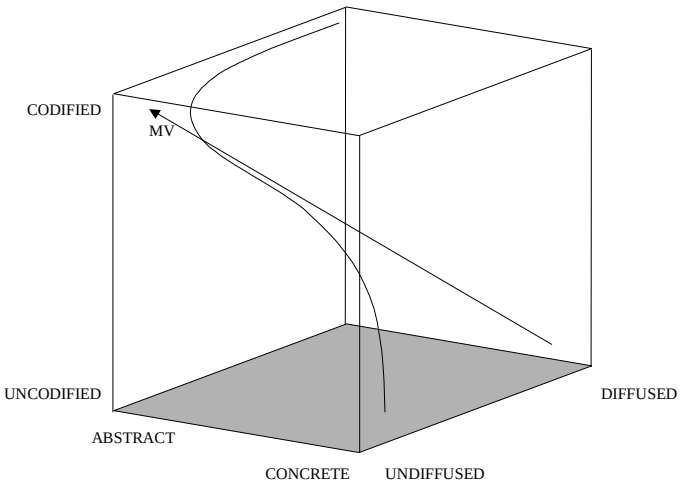

Figure 3: Maximum Value (MV) in the I-Space

The paradoxical nature of value in the case of an information good can be dealt with in two ways:

1. Hoarding – this strategy build upon the dynamics suggested by Figure 1. It is assumed that all potential value of a given information good is exhausted by the time it has diffused to the population as a whole. The strategy then consists of blocking or slowing down its diffusion. We label such strategies Neoclassical learning or N-learning strategies (Boisot, 1998).

2. Sharing – this strategy builds upon the dynamics suggested by Figure 2. It is assumed that new value can be generated by the absorption, impacting and scanning phases of the SLC and that this new value will be greater than that lost through the erosion of scarcity brought about by diffusion processes. The strategy then consists of moving around the SLC faster than competitors in order to secure first-mover advantages. We label such strategies Schumpeterian learning or S-Learning strategies (Boisot, 1998).

It is clear that N-learning strategies focus on preserving existing knowledge whereas S-learning strategies focus on challenging or destroying it. In section 4 we will develop these two strategies into hypotheses that could be explored via a simulation model – Smart Assets - and subsequently subjected to empirical testing. We now turn to a description of the simulation model itself.

3. SMART ASSETS: AN AGENT-BASED SIMULATION OF KNOWLEDGE FLOWS

3.1 Model Architecture

How does Smart Assets implement and embody 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 flows have been structured through processes of codification and abstraction. Such flows, over time, give rise to the creation and exchange of knowledge assets.

Smart Assets 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 diffusion processes as well as of what agents decide to do with them. These have the possibility of exchanging their knowledge assets in whole or in part with other agents through different types of dealing arrangements. Knowledge assets can also grow obsolete over time.

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, entry is based on mean revenues generated by the game in any given period. Entry and exit 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 for every percentage change in mean revenue. Change of rate of entry and exit is a function of percentage change in mean revenue.

Smart Assets has three model components: an agent component that specifies agent characteristics; a knowledge asset component that specifies the different ways that agents can invest in developing their knowledge assets, and an agent interaction component that specifies the different ways that agents can interact with each other.

i) Agents

Smart Assets operates through a number 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.

Agents can enter or exit Smart Assets 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 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. From their funds, agents draw budgets for meetings and for investing in knowledge assets.

ii) Knowledge Assets

In Smart Assets, knowledge assets are represented in network form. A knowledge network consists of a collection of elements and of relations between elements. 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[3]. 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. 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 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.

iii) Agent Interactions

Agents meet each other throughout the game and the frequency of encounters between agents can be varied. 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 a function of how codified and abstract the knowledge assets of interacting agents turn out to be.

3.2 Smart Assets in the I-Space

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

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

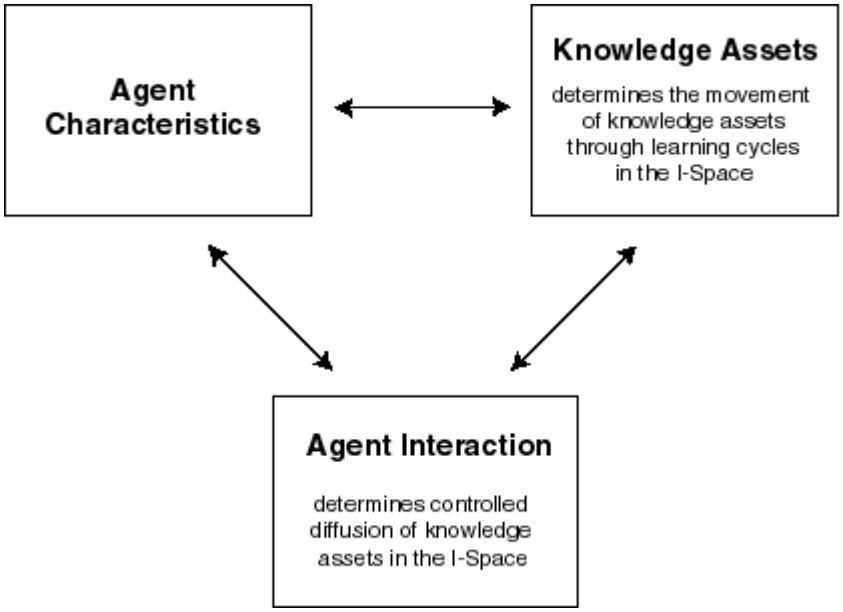

Figure 4: Smart Assets architecture

The number of Smart Assets parameters – about 65 variables can be parametized in the model - makes it possible to tune the simulation in order to build models representing diverse situations with a considerable degree of flexibility and 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 general validity of the simulation (Gilbert and Troitzsch, 1999). This will be our aim in the following sections.

4. VALIDATING SMART ASSETS: SOME HYPOTHESES

We now initiate the validation of Smart Assets by developing some simple hypotheses concerning the nature of N-learning and S-learning against which it can be tested.

We model the evolution in time of four distinct information modes, each representing a distinct economic dynamic. Each agent in Smart Assets 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: from scratch, 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. New knowledge in the form of nodes and links may 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. The diffusion of this knowledge will take place either as a consequence of different kinds of interaction among agents as well as of random diffusion forces.

4.1 From N-Learning to S-Learning

Our four information modes represent different levels of activation of the Smart Assets model components. The lowest level of activation models 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 agricultural case. In a knowledge-based economy this situation is of little interest except as a baseline case. It could be associated with a Neolithic agricultural order. Different agents may possess different agricultural technologies, but in the absence of agent interaction or of knowledge creation there is little change in these assets over time.

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 consumer electronics case.

Two intermediate cases are also examined. In the first, no knowledge creation is possible although agent interaction is allowed. Agents can thus trade or exchange knowledge assets, but they do not invest in the creation of new ones. If new knowledge assets appear, it is through the action of new entrants into the market. Such collaborative "non-competitive" relations between agents, coupled with a conservative attitude to new knowledge creation, are characteristic of many of the professions. We therefore label this the professional cartel case. In the second intermediate case, new knowledge creation is admitted but agent interaction is not. We take this to describe an anti-collusive environment characteristic of an anti-trust economic order. We label this the anti-trust case.

Our four cases span the continuum between the pure N-Learning situation and the pure S-Learning case:

Case 1: The agricultural case - purely neoclassical.

Case 2: The professional cartel case – intermediate between neoclassical and Schumpeterian.

Case 3: The anti-trust case– intermediate between neoclassical and Schumpeterian.

Case 4: The consumer electronics industry - the Schumpeterian case.

None of these four cases has any great 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. In effect, in certain respects our approach parallels that adopted by Nelson and Winter in their book, An Evolutionary Theory of Economic Change (Nelson and Winter, 1982). Like them we start from a case rooted in economic orthodoxy. Like them, we arrive at a Schumpeterian condition, albeit – in contrast to Nelson and Winter – one that is derived from an agent-based approach.

4.2 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. Agent-based modeling typically does not allow one to make detailed predictions; the processes modeled are too complex for that. At best one can predict certain gross features of the processes and this is what will be attempted here. If our hypotheses are corroborated - that is, if our results are consistent with the theoretical claims of some of the current economic and managerial theories that have demonstrated a measure of empirical validity – we will have taken a first step towards validating Smart Assets. Our hyptotheses are summarized in Table 2.

In case 1, where we switch off the knowledge creation and agent interaction components of Smart Assets, we end up having something quite similar to a neoclassical economic setting. In neoclassical economic theory, knowledge is considered a public good instantly and costlessly accessible to all players – ie, knowledge diffuses instantly. Under such circumstances, there is no need for knowledge sharing interactions between agents. Furthermore, given that knowledge is a public good, knowledge creation generates no appropriable value for a firm engaged in research activities. Such an activity therefore is pointless from a commercial perspective. Since Smart Assets only models the generation of rents earned from the differential possession of knowledge assets, one should expect the agricultural case to result in fairly uneventful behaviours, punctuated only by the random entry of new agents (together with their knowledge assets) into the simulation. In this case, as in the pure neoclassical economy, time almost comes to a halt. We take this point to derive our first set of hypotheses. These are given in the top left-hand quadrant of Table 2.

After activating the knowledge creation and agent interaction components of Smart Asset, we would expect that Case 4, the consumer electronics case, would reflect a situation in which both innovation and inter-agent collaboration is important – ie, the Schumpeterian settings. Recall that this is a setting characterized by “gales of creative destruction”. This gives point gives us our second set of hypotheses as given in the bottom right-hand quadrant of Table 2. Hypotheses for the two intermediate cases, 3 and 4, are given respectively in the top right-hand quadrant and the bottom left-hand quadrant of Table 2.

	No agent interaction			Agents interact		
No knowledge creation	**Case 1** ***Agricultural - Neoclassical***			**Case 2** ***Professional Cartel***		
		Hypotheses			Hypotheses	
	Agent numbers	Most stable		Agent numbers	Less stable than case 1	More stable than case 4
	Agent entries	Minimum		Agent entries	Less than case 1	Less than case 4
	Agents cropped	Minimum		Agents cropped	More than case 1	Less than case 4
	Agent exits	Minimum		Agent exits	More exits than case 1	
	Total knowledge generated	Minimum		Total knowledge generated	More than case 1	
	Revenues per agent	Less than case 3	More than case 4	Revenues per agent	Less than case 3	More than case 4
	Total revenues	Less than case 3	More than case 4	Total revenues	Less than case 3	More than case 4
Knowledge created	**Case 3** ***Anti-trust***			**Case 4** ***Cons. electronics - Schumpeterian***		
		Hypotheses			Hypotheses	
	Agent numbers	Less stable cf case 1		Agent numbers	Least stable	
	Agent entries	More than case 1	Less than case 4	Agent entries	Maximum	
	Agents cropped	More than case 1	Less than case 4	Agents cropped	Maximum	
	Agent exits	More than case 1	Less than case 4	Agent exits	Maximum	
	Total knowledge generated	More than case 1	Less than case 4	Total knowledge generated	Maximum	
	Revenues per agent	Maximum		Revenues per agent	Minimum	
	Total revenues	Maximum		Total revenues	Minimum	

Table 1: Hypotheses

4.3 Procedure

We perform 200 runs for each of our four cases and use the average values for interpretation. The number of runs has proved sufficient to show significant differences at the 95% level of confidence between the four cases. Each run terminates after 200 periods. In this exercise, we first focus 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 look at the number of new knowledge assets – nodes and links - that were created during the game – these will offer us a proxy measure of the overall social benefits of the economic order in which they were created. Finally, we look at the total revenues generated as well as the revenue per agent. This gives us a measure of what society has to pay the players for the knowledge generated.

5. RESULTS

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 four cases modeled. We then plot the evolution of each variable in a graph, indicating the mean and the 95% interval of confidence for each period.

Case	Period		Agent Number	Agent Entries	Agents Cropped	Agent Exits	Knowledge Assets	Revenue per Agent	Total Agent Rev.
1	**0**	Mean	20,00	0,00	0,00	0,00	0,00	0,66	-786,89
		S.D.	0,00	0,00	0,00	0,00	0,00	0,09	1,80
	100	Mean	21,27	1,95	0,00	0,68	0,00	0,75	736,48
		S.D.	3,43	3,54	0,00	0,96	0,00	0,12	381,35
	199	Mean	20,50	3,85	0,00	3,35	0,00	0,82	2.380,01
		S.D.	3,68	4,33	0,00	2,07	0,00	0,15	803,18
2	**0**	Mean	20,01	0,00	0,00	0,00	0,00	0,67	-786,62
		S.D.	0,07	0,00	0,00	0,00	0,00	0,10	2,02
	100	Mean	16,41	115,35	250,73	0,03	0,00	3,64	5.343,19
		S.D.	2,89	11,91	52,01	0,17	0,00	0,46	475,80
	199	Mean	15,97	221,70	509,07	0,03	0,00	2,59	10.990,61
		S.D.	2,97	18,26	61,85	0,17	0,00	0,29	561,95
3	**0**	Mean	20,00	0,00	0,00	0,00	1,07	0,69	-786,29
		S.D.	0,00	0,00	0,00	0,00	1,00	0,10	1,94
	100	Mean	35,14	26,67	0,00	16,49	353,02	1,61	4.322,83
		S.D.	6,47	4,88	0,00	4,72	106,60	0,26	642,44
	199	Mean	91,67	52,12	13,11	43,58	1.909,83	0,71	10.158,02
		S.D.	23,51	10,01	13,71	6,08	720,34	0,17	1.457,09
4	**0**	Mean	20,02	0,00	0,00	0,00	1,27	0,70	-785,98
		S.D.	0,12	0,00	0,00	0,00	1,05	0,11	2,17
	100	Mean	14,46	136,74	290,09	0,05	358,63	2,43	5.278,29
		S.D.	3,13	12,96	39,40	0,21	53,76	0,48	536,12
	199	Mean	15,22	296,84	630,25	0,05	870,73	1,30	8.157,10
		S.D.	4,42	22,41	48,29	0,21	79,48	0,31	835,48

Table 2: Mean and Standard Deviation of variables at representative periods for each case.

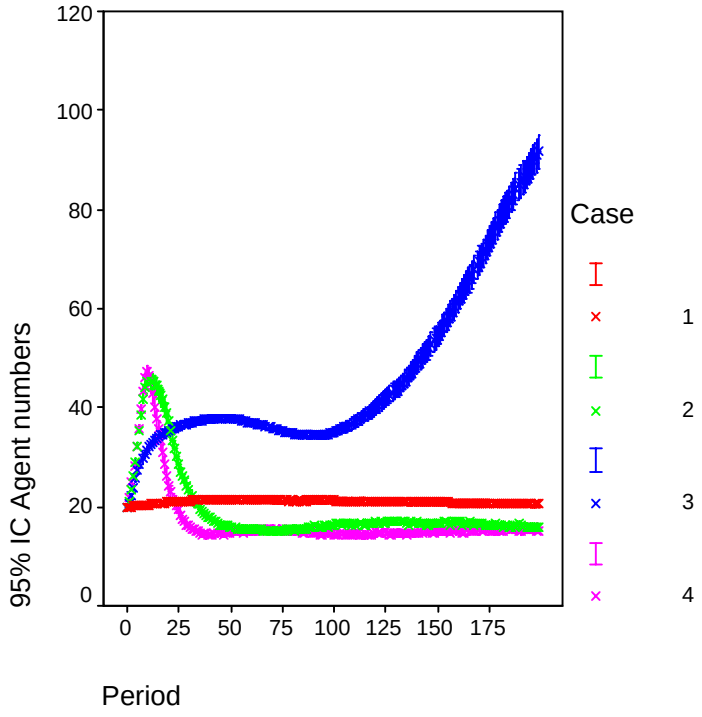

Figure 5: Evolution of the total number of agents in the four cases (95% interval of confidence)

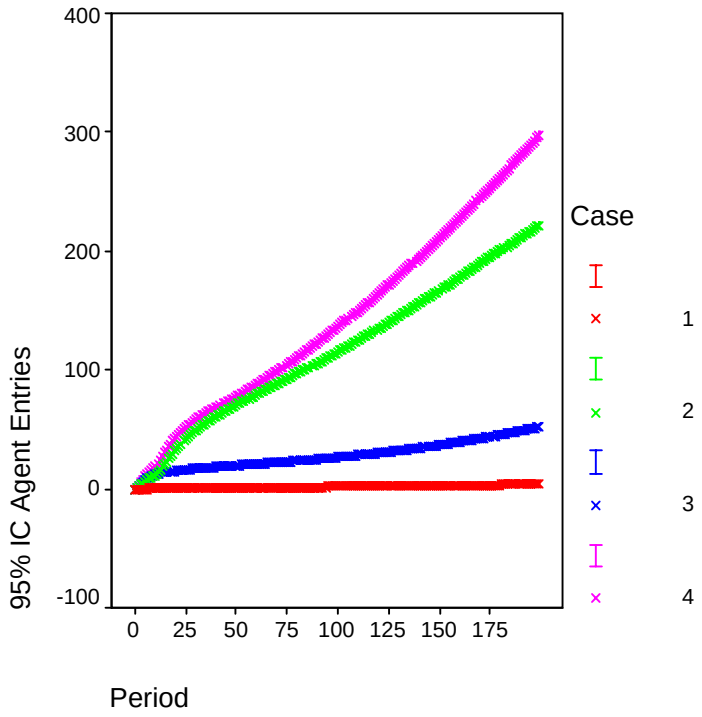

Figure 6: Evolution of the total number of agent entries in the four cases (95% interval of confidence)

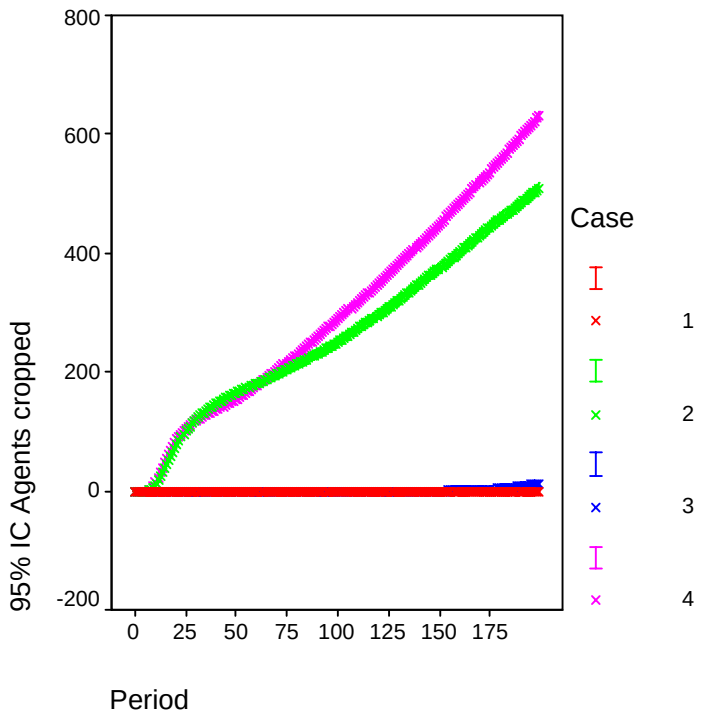

Figure 7: Evolution of the total number of agent cropped in the four cases (95% interval of confidence)

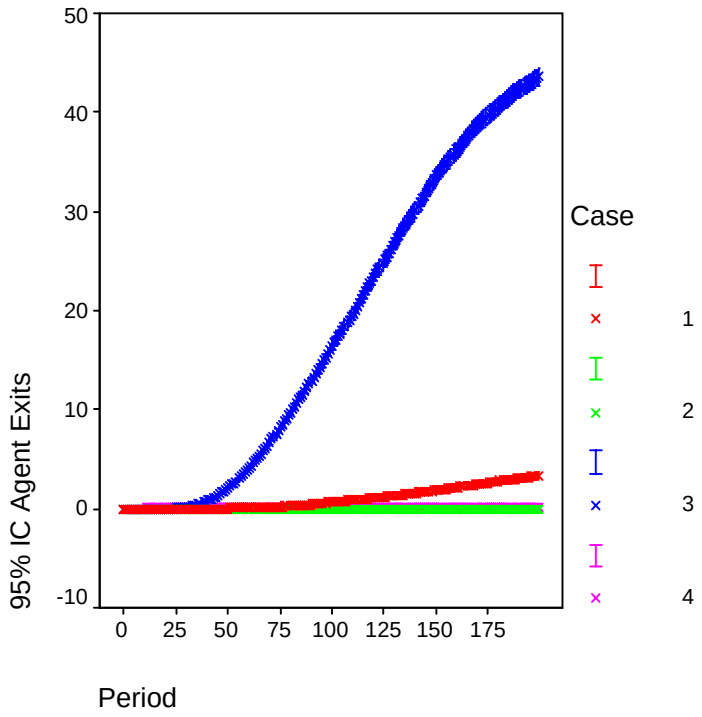

Figure 8: Evolution of the total number of agent exits in the four cases (95% interval of confidence)

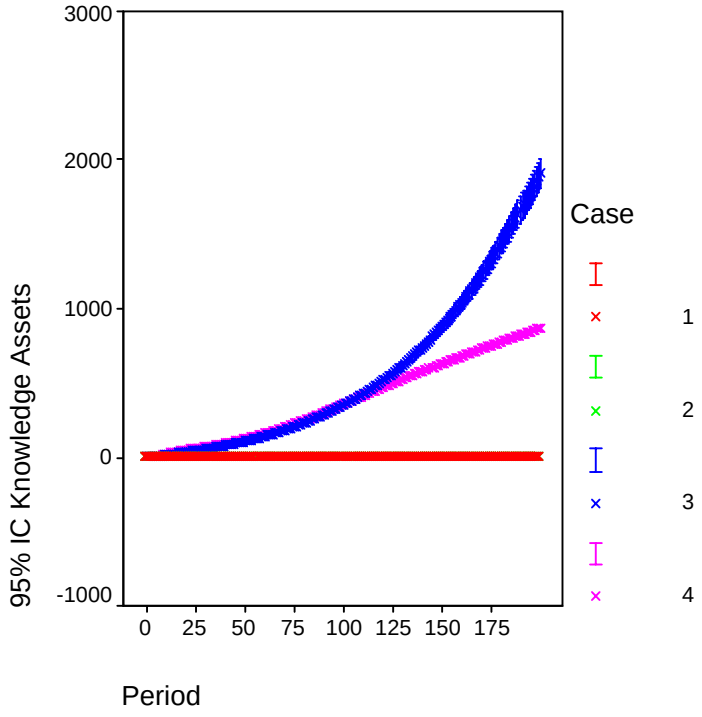

Figure 9: Total number of knowledge assets created (95% interval of confidence)

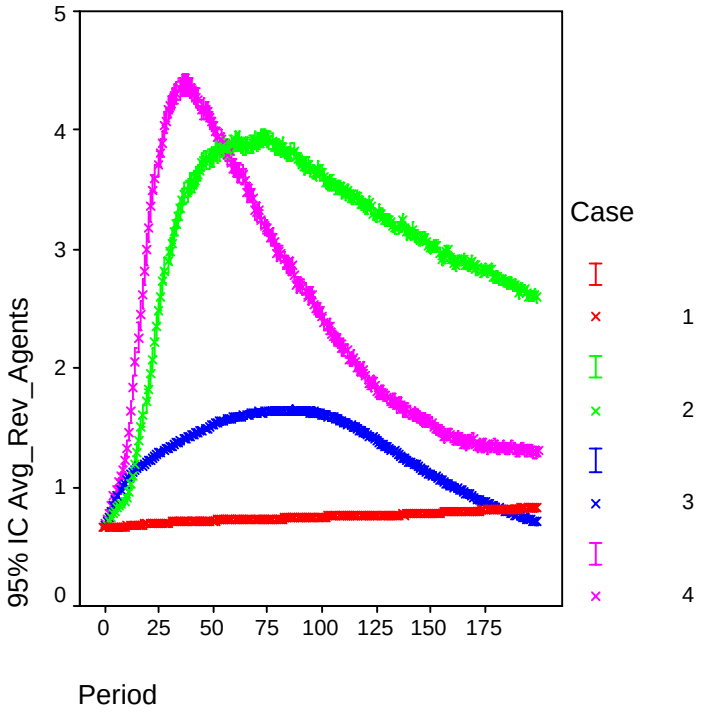

Figure 10: Evolution of the average revenue per agent in the four cases (95% interval of confidence)

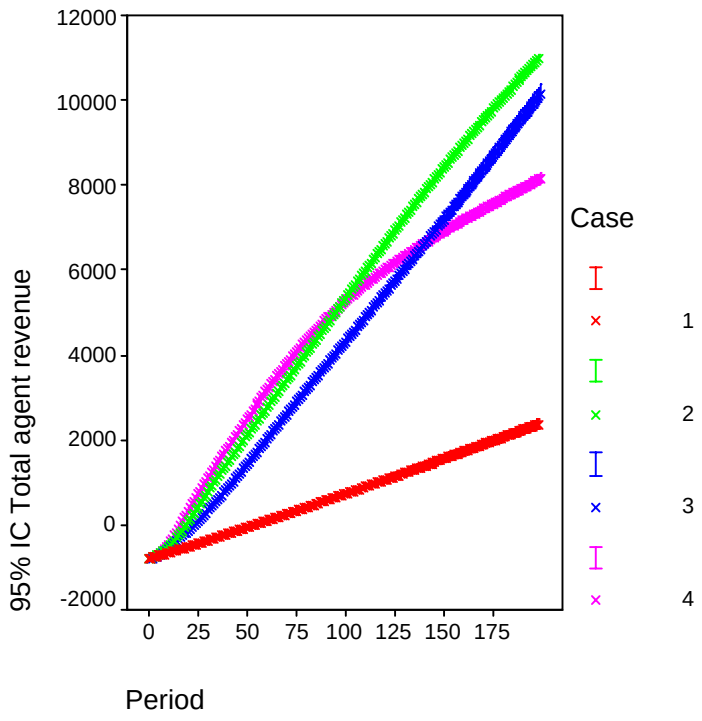

Figure 11: Evolution of the total agent revenues in the four cases (95% interval of confidence)

6. DISCUSSION

In Table 2, we put forward four hypotheses designed to see whether the simulation model could meaningfully distinguish between Neoclassical and Schumpeterian processes as they manifest themselves in the I-Space. How strongly do the results of the simulation runs support or challenge these hypotheses? What does this 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.

	No agent interaction					Agents interact				
No knowledge creation	**Case 1** ***Agricultural - Neoclassical***					**Case 2** ***Professional Cartel***				
		Hypotheses					Hypotheses			
	Agent numbers	Most stable	✓			Agent numbers	Less stable than case 1	✓	More stable than case 4	✓
	Agent entries	Minimum	✓			Agent entries	Less than case 1	✗	Less than case 4	✓
	Agents cropped	Minimum	✓			Agents cropped	More than case 1	✓	Less than case 4	~
	Agent exits	Minimum	✗			Agent exits	More exits than case 1	✗		
	Total knowledge generated	Minimum	✓			Total knowledge generated	More than case 1	✗		
	Revenues per agent	Less than case 3	~	More than case 4	✗	Revenues per agent	Less than case 3	✗	More than case 4	✓
	Total revenues	Less than case 3	✓	More than case 4	✗	Total revenues	Less than case 3	✗	More than case 4	~
Knowledge created	**Case 3** ***Anti-trust***					**Case 4** ***Cons. electronics - Schumpeterian***				
		Hypotheses					Hypotheses			
	Agent numbers	Less stable cf case 1	✓			Agent numbers	Least stable	✓		
	Agent entries	More than case 1	✓	Less than case 4	✓	Agent entries	Maximum	✓		
	Agents cropped	More than case 1	✓	Less than case 4	✓	Agents cropped	Maximum	✓		
	Agent exits	More than case 1	✓	Less than case 4	✗	Agent exits	Maximum	✗		
	Total knowledge generated	More than case 1	✓	Less than case 4	✗	Total knowledge generated	Maximum	~		
	Revenues per agent	Maximum	✗			Revenues per agent	Minimum	✗		
	Total revenues	Maximum	✗			Total revenues	Minimum	✗		

Table 3: Hypotheses assessment

At the most general level, we observe that the four cases exhibit markedly different behaviours in the simulation runs. At the very least, therefore, the simulation can distinguish between them. The key issue is whether the difference between cases is consistent with our hypotheses. Taking now each hypothesis in turn:

Agent numbers: Our hypotheses hold. Although the behavior of cases 2 and 4 are quite similar, case 4 can be considered the more unstable of the two because it reaches a higher maximum and a lower minimum, and shows steeper slopes.

Agent entries: Our hypotheses hold, except for case 2, which, as compared with case 1, shows more agent entries. One possible reason for this is that case 2 also has a lot of agents getting cropped. The number of agents in case 2 remains more or less constant after period 50, but with a high rate of agent turnover (agent cropping and entry). In case 2, as a consequence of

the high level of cropping, there is substantious revenue available for the agents that remain, and this attracts new entrants. By contrast, case 1 is much more stable. A similar balance is maintained, but with a much lower rate of renewal.

Agents cropped: Here, our hypotheses hold, except for the one comparing case 2 to case 4. And even in here, the hypothesis holds for most of the simulation, since case 4 is ahead in the cumulative number of agents cropped for all periods except those between roughly periods 25 and 75.

Agent exits: Here our results do not corroborate the hypothesis that the number of exits will be maximum in case 4 and minimum in case 1. We did not reckon with the fact that many exits will happen when agents are wealthy, something that is less likely to happen on a large scale in a Schumpeterian setting. The picture would change if we limited our focus to agent cropping.

Total knowledge generated: Here some of the hypotheses hold, and others don't. It is not the case that in case 2, for instance, *more* knowledge is created than in case 1. In both cases the creation of new knowledge assets is zero – for that is what we required of the system through our setting of the parameters. "New assets created" counts only those assets actually created by agents currently active in the simulation and not those that get introduced into the simulation via the entry of new agents. Another hypotheses that fails is the one that states that in case 3 the creation of knowledge will be at a lower level than in case 4. In case 3 there is in fact a lot of knowledge creation. One possible reason for this is the high number of agents created in that case. With more agents in the game, the probability of creation of agents increases.

Revenues per agent: Clearly, our hypotheses do not hold for this variable. The reason could be that perhaps when we think of revenues in the Neoclassical or Schumpeterian cases we have in mind all possible revenues – ie, those derived both from knowledge assets and from other sources. In our simulation, however, we only take into account those revenues derived from knowledge assets.

Total revenues: Our hypotheses do not hold well for this variable either. Since the two hypotheses are fairly closely related this might not be too surprising. The explanation might be the same in both cases.

In sum, we see that the simulation model achieves a moderate ability to discriminate between the four cases. It experiences the greatest difficulty with case 2, the professional cartel, but does quite well with case 3, the anti-trust case. And although we would expect the clearest overall contrast to be between case 1, the neoclassical case, and case 4, the Schumpeterian one, the model fares only modestly well in distinguishing between these two cases. We conclude that although the simulation model is capable of achieving some measure of distinction between the information conditions relating to the neoclassical case and those relating to the Schumpeterian one, both the model and the hypotheses would need further refinement to make the case fully convincing. Our achievement, therefore, is a modest one: our hypotheses, after all, are not very ambitious. Yet the overall results suggest that the model has achieved enough internal coherence to justify its further development.

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. We are still in the model validation phase of our research, and even here, only at the beginning. We would expect better support from a more refined model. Nevertheless, we 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 gradually expect to get closer to the real world, modeling real cases and exploring the fine grain of knowledge-based agent behaviour. The next step will be to model particular classes of agents and endow them with memory. This will 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 can 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 work will help to place knowledge management on a sound theoreticall footing.

REFERENCES

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

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.

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.

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

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

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.

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.

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.

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