Was Hayek an Ace? *

Nicolaas J. Vriend, *Queen Mary and Westfield College*, University of London, UK

May 1999

Abstract: In order to address the question whether Hayek might have been an Agent-based Computational Economist (ACE) *avant-la-lettre*, we consider an ACE model concerning the phenomenon of information contagion. Alongside increasing returns, network externalities, and information cascades, information contagion has been presented in the literature as an explanation for particular patterns of macrobehavior that may seem at odds with the underlying micromotives. But whereas these other explanations have been shown to have a proper microfoundation, information contagion has remained a phenomenon that seemed to occur only when certain ad hoc rules of thumb for individual behavior are assumed. We show how information-contagious behavior can emerge in a coevolutionary process of interacting adaptive agents, how this is related to various Hayekian themes, and how ACE research in general is an application of Hayek's methodological insights.

J.E.L. classification codes: B21, B31, B41, D11, D83, O33

Keywords: Hayek, Agent-based Computational Economics, Decentralized interaction, Distributed knowledge, Reinforcement learning, Information aggregation, Self-organization, Spontaneous order, Methodology

---

Queen Mary and Westfield College, University of London, Dept. of Economics, Mile End Road, London, E1 4NS, UK, <n.vriend@qmw.ac.uk>, http://www.qmw.ac.uk/~ugte173/.

* Comments by participants of the Liberty Fund symposium on *"The Legacy of F.A. Hayek"*, Freiburg, Germany, May 6-9, 1999, are gratefully acknowledged. The usual disclaimer applies.

1. Introduction

Hayek was without doubt one of the great minds of economics, and not only of economics. Obviously, this paper will not pretend to question him being an ace. The Ace in the title rather refers to Agent-based Computational Economist (ACE). This is a recently established field of economics, forming a subfield of the equally contemporary field of computational economics. As Tesfatsion, one of the pioneers of the field, puts it on the ACE Web-site:

> *"Agent-based computational economics (ACE) is roughly characterized as the computational study of economies modelled as evolving decentralized systems of autonomous interacting agents. A central concern of ACE researchers is to understand the apparently spontaneous formation of global regularities in economic processes, such as the unplanned coordination of trade in decentralized market economies that economists associate with Adam Smith's invisible hand. The challenge is to explain how these global regularities arise from the bottom up, through the repeated local interactions of autonomous agents channeled through socio-economic institutions, rather than from fictitious top-down coordination mechanisms such as imposed market clearing constraints or an assumption of single representative agents. ACE is thus a specialization to economics of the basic complex adaptive systems (CAS) paradigm."* (Tesfatsion [1998])

The descriptions used in this informal definition of ACE must look rather familiar to experts of Hayek. Hayek shares with ACE the belief that the economy needs to be understood from a bottom-up perspective. In this he stands out from both Keynesian macroeconomics and Walrasian general equilibrium theory (e.g., Debreu [1959]), which came to dominate the science of economics during Hayek's life. *"Keynes' theories will appear merely as the most prominent and influential instance of a general approach to philosophical justification of which seems to be highly questionable. Though with its reliance on apparently measurable magnitude it appears at first more scientific than the older micro-theory, it seems to me that it has achieved this pseudo-exactness at the price of disregarding the relationships which really govern the economic system"* (Hayek [1978], p. 289). Hayek, on the other hand, insisted on the need to consider a market economy as a truly decentralized system of interacting individual agents: *"( ) true individualism is the only theory which can claim to make the formation of spontaneous social products intelligible"* (Hayek, [1948b], p. 10), and Hayek seemed to share with Adam Smith the belief that an emergent, spontaneous order tends to be beneficial: *"true individualism believes ( ) that, if left free, men will often achieve more than individual reason could design or foresee"* (p. 11). Although the emergence of global regularities stands central on the ACE research agenda as well (see the Tesfatsion quote above), whether and when such spontaneous orders are beneficial is more an open research question in ACE.

Hayek's work was firmly rooted in the *"antirationalistic"* approach of the English individualism as known, for example, from Adam Smith's Invisible Hand metaphor (see, e.g., Hayek [1948b]). Now, some people have claimed that general equilibrium theory simply *'has finally proved mathematically* *what Smith argued two centuries ago'*. For example, the First Fundamental Theorem of Welfare Economics is considered to be *"a formal and very general confirmation of Adam Smith's asserted "invisible hand" property of the market"* (Mas-Colell et al. [1995], p. 549).[1] But we have argued elsewhere (Vriend [1994]), that general equilibrium models cannot be considered as ideal representations of decentralized economies because their main contribution has been to make Adam Smith's transcendental hand *visible* by imposing various centralized concepts. These models base their analysis of a decentralized economy upon the actions of autonomous agents, pretending to discard any kind of external determination of the behavior of the individual agents. But it turns out that in order to explain anything, resort must be made to many concepts and structures that transcend the level of the individual agents. In particular, the existence of an auctioneer, the division of time into meta- and real time, and the rules of the game in these models are in no way the result of the behavior of the individual agents. Moreover, the set of possible actions of the individual agents is predetermined by the rules and structure of the model, and each individual takes the structure of the model into account in calculating his choices, *anticipating* the equilibrium character of the overall outcome, which should, instead, be explained by their actions. How can an individual agent in these models understand that the economy will turn out to be in equilibrium; unless he is God, the man with the invisible hand, or the auctioneer?

Given the easily recognizable connections between Hayek and ACE it is no surprise that many Hayek experts seem interested in the recent complex systems literature. At the same time, however, most complexity and ACE researchers are hardly aware of Hayek's work at all. Every now and then somebody might mention that it would be interesting to have a closer look at Hayek's work, but that is about it. In this paper we will take up these suggestions. We present a concrete example of an ACE research project as a guide to address the question whether Hayek might have been an Agent-based Computational Economist (ACE) *avant-la-lettre* in great detail. Apart from the purely intellectual motivation for such a study, some of the underlying questions motivating this project are: How could Hayek's insights and theories help to understand ACE research? And what, if anything, could we learn from current ACE research about Hayek's work? Far from offering definite answers to these questions, this paper will suggest there might be some reasons to believe that a close encounter between Hayek and ACE has potential benefits that might work in both directions.

This paper is organized as follows. In section 2 we give a brief introductory sketch of the literature on complex adaptive systems. Section 3 presents our ACE model of the emergence of information contagion, and relates it to Hayek's work on the division of knowledge and information aggregation. In section 4 we put our model into a somewhat wider perspective by discussing some related literature on information contagion and social learning. Section 5 presents an analysis of the properties of the model, while section 6 concludes with a discussion, in particular of some methodological issues.

---

[1] This Theorem states that if a set of prices and a corresponding allocation of commodities constitute a competitive equilibrium then this allocation is Pareto optimal, which means that no agent can be made better off without making at least one other agent worse off.

2. Complex adaptive systems

As Tesfatsion explained, ACE is a specialization to economics of the basic complex adaptive systems paradigm. A *'complex system'* is a system consisting of a large number of relatively independent parts that are interconnected and interactive. Such a system is *'adaptive'*, if these parts are agents that change their actions as a result of events occurring in the process of interaction. Some examples of complex systems are biological systems, immune systems, brains, the weather, ecologies, and societies.[2] A decentralized economy, consisting as it does of a large number of locally interconnected and interacting rational agents who are all continually pursuing advantageous opportunities, is also an example of a complex adaptive system. Persuasive arguments why this is the case can be found in Hayek [1948a], written long before the term 'complex adaptive system' as such was known.[3] An essential characteristic of such systems is that their global properties cannot be derived simply from an examination of the individual components. Even when each individual agent is inherently simple, the behavior both of the system as a whole and the individual agents may become complex (see, e.g., Holland [1992], Langton [1989] or Kauffman [1993]).

A complex system is not the same as a chaotic system. In general, complex systems tend to evolve away from the extremes of, on one side, absolute order, and on the other side, what looks like complete randomness. Currently, the key theoretical concepts are self-organization (the formation of regularities in the patterns of interaction) and selection (through system constraints). Selection seems to act in many self-organizing systems to push the system continuously back to some *'boundary'* between order and chaos. Around this *'edge'*, these systems appear to carry out the most complex behavior and adapt most readily to *changing* environments. We will illustrate precisely this property later in the paper, but here we just want to stress the importance of this feature (see also Hayek [1948]).

As the interactions between the individual agents are in general nonlinear, from a mathematical point of view such systems are often intractable. And as far as economics is concerned, the available analytical apparatus borrowed from graph theory, statistical mechanics and the theory of interacting particle systems appears to be restrictive with respect to the economic contents of the models. It seems especially difficult to incorporate the essential fact that the interactions that take place between economic agents in reality, are not determined by their given position in a grid, graph or lattice, or by some kind of anonymous matching device. Most interactions in a decentralized economy are determined by economic agents who are themselves actively pursuing those interactions that are the most advantageous to them. Transactions do not take place in Walrasian central markets, or through anonymous random matching devices. Instead, market interactions depend in a crucial way on local

---

[2] For an introduction to the topic of complexity in a broad sense, see Waldrop [1992], Lewin [1993], Casti [1994], and Gell-Mann [1994].

[3] These arguments did not influence many formal economic models, in which direct interaction between the agents is usually absent. See Kirman [1994] for a recent survey of economic models with interacting agents.

knowledge of the identity of some potential trading partners. A market in a decentralized economy, then, is not a central place where a certain good is exchanged, nor is it the aggregate supply and demand of a good. In general, markets emerge as the result of locally-interacting individual agents pursuing advantageous contacts, i.e., they are self-organized.

Given the limitations of formal analyses, one often has to rely on computational methods. One approach is to model each individual agent and their interactions *explicitly*; in other words, to provide these agents (real robots, or simulated machines) with artificial *'flesh and blood'*. That *'bottom-up'* approach is also known as the artificial life, or *'alife'*, approach (Langton [1989]), and in economics nowadays as Agent-based Computational Economics (ACE). Modeling *homo oeconomicus* as a *'machine'* does not seem to pose any particular conceptual difficulties to economic theory (see, e.g., Lane [1992]). After all, as Lucas puts it, doing economics means *"programming robot imitations of people"* (in Klamer [1984], p.49). Although Lucas' statement was only meant as a metaphor, current computational capabilities suggest taking it literally, and to consider its usefulness for economic theory. What Lucas was referring to, of course, was that *homo oeconomicus* is a rather mechanical representation of real, live agents. The fundamental characteristic of *homo oeconomicus* is that he simply chooses (one of) the most preferred option(s) in his perceived opportunity set. In fact, *homo oeconomicus* is an *'opportunist'*; always doing the best he can (see also Vriend [1996]).

Making some simple assumptions concerning the preferences and technologies of the individual agents, this implies that the question of the modeling of the perceived opportunity sets of the individual agents becomes of central concern. During the process of interaction between the individual agents in a decentralized economy, perceived opportunities evolve; due either to a change in the underlying situation or to a change in the perception of this situation, i.e., learning. A general characteristic of agents living in the complexity of a *'large world'* is that they do not have a true, well-specified model to work with. That is, the agents' problem situation is ill-defined (see Arthur [1992] and [1994]). Hence, instead of basing their actions on deductive reasoning from universal truths, they are forced to use inductive reasoning. Such reasoning proceeds from the actual situation faced by an agent. In Savage's [1954] terminology, this is the *'cross that bridge when you meet it'* principle. In a large world individual agents are constantly searching, learning, and adapting to their environment.

In a decentralized economy, each individual's activities will in a certain way be *'involved'* in the activities and decisions of some other agents. Hence, each agent will have a different relevant *'environment'* for different kinds of activities, and these environments are influenced by the actions of other individual agents. In other words, while an individual agent is adapting to its environment, parts of the environment are adapting to him. In biology this phenomenon is known as *coevolution*. In principle, such a process might go on forever (see, e.g., Sigmund [1995]). Therefore, rather than analyzing whether an equilibrium exists for an economy with some *given* structure, in ACE one analyzes how structures and patterns *emerge* as regularities in the process of interaction of the individual agents. In other words, instead of 19th-century physics, it is biology or meteorology that provides the relevant metaphors for this approach to the study of decentralized economies.

3. The Agent-based Computational model

The basic choice problem we consider is that of a population of individual agents who each face a decision problem between two items with uncertain qualities. We can think of these two items as new products, technologies, services, financial gurus, or whatever binary choice agents might need to make frequently in every day life. The rationale to focus on the introduction of new items is explained by Hayek [1948e]: *"It is, perhaps, worth stressing that economic problems arise always and only in consequences of change. As long as things continue as before, or at least as they were expected to, there arise no new problems requiring a decision, no need to form a new plan"* (p. 82).

The only information the agents have is the choice plus the corresponding value experienced by a sample of other agents who had faced the same decision problem before them. This basic choice problem has been considered in the literature on information contagion, and in some literature on social learning. We will discuss some of this literature in section 4. Information contagion occurs when the probability that an individual agent chooses a certain option is a positive monotonic function of the number of agents who had made that choice already (independent from the actual values experienced by them). For any form of social learning to which this applies, this implies a diffusion process such that a certain choice once it starts being made by a certain number of people spreads quickly in a population. One basically unanswered question is, where such information contagious behavior comes from.

3.1 The basic problem situation

There is a population of 250 decision makers. In a given period they face a choice between two items that were previously unknown to them. Each new item i is characterized by the expected value of the utility it will generate, EV_i. These expected values are unknown to the individual agents. Given an expected value, EV_i, the value that a specific agent will actually experience from an item will be a random draw from a uniform distribution with support from EV_i-0.25 to EV_i+0.25. Hence, if a given item i is characterized by an EV of, say, 0.40, the actual utility levels experienced by the individual agents choosing this item will range from 0.15 to 0.65, with every utility level in this range equally likely to occur. Notice that we do not have any increasing real returns to scale of any form, no change in taste, endogenously determined utility depending on the number of adopters, nor are there complementarities or network externalities. Each individual agent's utility of a certain item is simply a random draw from the same uniform distribution characterized by the item's expected value EV. Figure 1 gives an example with two items with expected values EV_1 = 0.40, and EV_2 = 0.55.

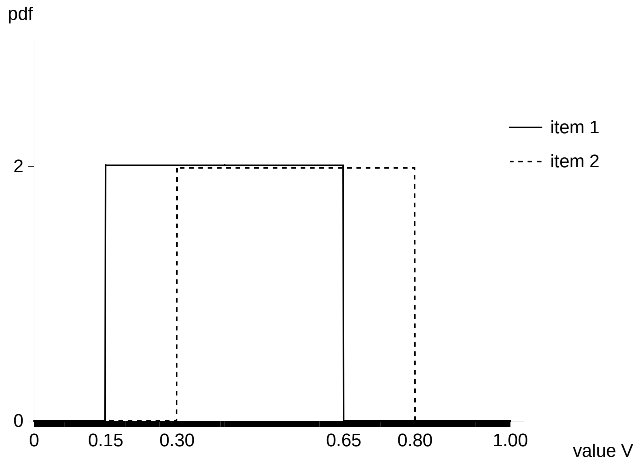

Fig. 1 Probability density function for two items

This might seem a simple statistical problem that could be solved by a central planner. However, as explained by Hayek [1948e], *“( ) the sort of knowledge with which I have been concerned is knowledge of the kind which by its nature cannot enter into statistics and therefore cannot be conveyed to any central authority in statistical form. The statistics which such a central authority would have to use would have to be arrived at precisely by abstracting from minor differences between the things, by lumping together, as resources of one kind, items which differ as regards location, quality, and other particulars, in a way which may be very significant for the specific decision. It follows from this that central planning based on statistical information by its nature cannot take direct account of these circumstances of time and place and that the central planner will have to find some way or other in which the decisions depending on them can be left to the "man on the spot"”* (p. 83).[4]

These ‘men on the spot’, then, face their choice problem sequentially, with the order of the agents being random. Although each individual agent himself has no experience with these two specific new items, he can draw six random samples from the people that have made already a decision before him. For each of the elements in his sample, he can observe the choice made, and the value actually experienced by the agent. Given this sample information, an agent makes a choice himself, and then the next agent in the queue makes his decision, until the end of the queue is reached. Before the first agent in the sequence makes his decision in a given period, we add six dummy agents. Three of these dummies choose one item, and the other three the other item. This 50-50 seeding prevents any bias at the start of a period. The reason to do this is that lock-in due to the choice of these very first dummy agents would be an uninteresting artifact.[5]

---

[4] Ellison & Fudenberg [1993] give some agricultural examples which illustrate this point.

[5] Notice that we take the number of observations sampled as given. However, as Streit & Wegner [1989] point out, there is a fundamental paradox here, one that we avoid. The value of information is unknown until you have got it. Agents can learn something without knowing whether this information constitutes a net gain. Hence, equalizing marginal costs and marginal revenue is impossible. In fact, there is a dual decision problem. First, to collect information and second, to make choices.

We can think of the sampling by the individual agents as a further reflection of the fact that information as to the utility of a specific item is largely local. Ellison & Fudenberg [1993] present some evidence suggesting *"( ) that farmers often distrust the information of central authorities and experts, and prefer to see how innovations work out in their neighborhood"* (p. 631). Therefore, each agent only samples in his own 'neighborhood', where a 'neighborhood' is not necessarily spatial but the subset of agents with whom he has contact, and that he believes to be relevant for the issue concerned. Although the random payoffs generated by a given item reflect Hayek's observations that *"every individual has some advantage over all others because he possesses unique information of which beneficial use might be made"* (Hayek [1948e], p. 80) of the item, and that *"(a)t any given moment the equipment of a particular firm is always largely determined by historical accident, and the problem is that it should make the best use of the given equipment ( )"* (Hayek [1948f], p. 101), at this point we can also think of the random component of the payoffs as 'noise' in the sampler's observations of the values actually generated.

The fact that different individual agents will have different samples implies that *"( ) we deal ( ) with a situation in which a number of persons are attempting to work out their separate plans, (and hence) we can no longer assume that the data are the same for all the planning minds"* (Hayek [1948f], p. 93). The fact that no individual agent observes all data concerning a new item reflects Hayek's [1948e] observation concerning *"( ) an essential part of the phenomena with which we have to deal: the unavoidable imperfection of man's knowledge and the consequent need for a process by which knowledge is constantly communicated and required"* (p. 91). The sampling of some observations models the fact that we are dealing with *"a process which necessarily involves continuous changes in the data for the different individuals. (T)he causal factor enters here in the form of the acquisition of new knowledge by the different individuals or of changes in their data brought about by the contacts between them"* (Hayek [1948f], p. 94).

In any case, as figure 1 illustrates, in general the information sampled will be far from conclusive to determine which of the two items has the greatest expected value. For example, a utility level of 0.6 experienced by a specific agent in a sample could have been generated by an item with an expected value of 0.35, but also by an item with an expected value of 0.85. Obviously, this uncertainty matters a great deal for an agent that needs to make such a decision. We assume that each agent has in mind a set of simple rules of thumb to choose an item, and that the propensity to use any of these rules may change over time as a result of an agent's experience in the use of these rules. Therefore, before we explain in detail the modeling of the decision making and learning by the individual agents, we need to clarify how the individual agents face a similar basic choice problem over and over again.

3.2 Choice dynamics

All individual agents face the same basic choice problem for 25,000 periods. In every period two new items arrive, and all agents sequentially face a choice between them, with the order of the agents being determined at random in every single period. The fact that we modeled the sampling in a given period as random is a short-cut to take into account that for every day-to-day decision an individual agent may have a different relevant ‘neighborhood’, and we do not want to impose any given, fixed structure on these neighborhoods.

As we explained in section 3.1, every item appearing is characterized by its expected value. This expected value itself, which is unknown to the agents, is also a random draw from a uniform distribution; this time with support from 0.25 to 0.75. Hence, the worst item that can ever appear has an expected value of 0.25 (generating values for individual agents between 0.0 and 0.5), and the best possible item has an expected value of 0.75 (yielding utility levels between 0.5 and 1.0). Obviously, the ranges of utility levels that can be generated by intermediate items overlap with each other, as shown in figure 1. Every 500th period, the expected values of the two items are identical (0.5). This identical expected value case serves as a useful benchmark to see how much information contagion and path-dependency has emerged. While we use this benchmark every 500th period, in all other periods the expected value of the two items will not be identical, with one of the two items being superior in a statistical sense.

Although we have not said much about individual decision making and learning yet, intuition might suggest that this must be a trivial problem. If we run the model for 25,000 periods, and if in every period (apart from the benchmark periods) one of the two items is superior, then, surely, eventually every agent will easily discover which item is better. However, matters are slightly more complicated. Every period, two new, unknown items appear, and each item is up for choice only once during the entire history. Hence, the learning concerns the general rules of behavior, and not the specific, particular items as such. The fact that the agents learn the usefulness of general rules of behavior, and not the value of specific items also implies that if an agent oversees a certain sample of prior adoptions by other agents he might choose item 1, whereas he might have chosen 2 if he were confronted with the same two items but a different sample of prior adoptions.

3.3 Individual decision making and learning

The individual agent’s decision making is modeled for each individual agent separately by means of a Classifier System. Figure 2 presents one such stylized Classifier System.

condition	action	strength
if	then	....

Fig. 2 Classifier System

A Classifier System consists of a set of rules, each rule consisting of a condition part (*'if ...'*), and an action part (*'then ...'*), plus to each rule attached a measure of its strength. The Classifier System does two things. First, it decides which of the rules will be the active rule in a given period. Hence, it checks the condition part, and all rules satisfying the *'if ...'* condition make a 'bid' as follows: bid = strength + ε, where ε is white noise. The rule with the highest 'bid' in this 'stochastic auction' wins the right to be active. Second, the Classifier System updates the strength s of a rule that has been active, and has generated a reward from the environment in a given period t-1, as follows: $s_t = s_{t-1} - c \cdot s_{t-1} + c \cdot reward_{t-1}$, where $0<c<1$. Hence, $\Delta s_t = c \cdot (reward_{t-1} - s_{t-1})$. In other words, as long as the reward generated by the rule in period t-1 is greater than its strength at t-1, its strength will increase. As a result, the strength of each rule converges to the weighted average of the rewards from the environment generated by that rule.[6] In our implementation of this model, the strengths of all rules are equal at the start.

Classifier Systems are a form of reinforcement learning. Reinforcement learning is related to multi-armed bandit problems, and is based on two principles. First, agents try actions. Second, actions that led to better outcomes in the past are more likely to be repeated in the future. There is a family of stochastic dynamic models of such individual behavior in the scientific literature, for which different backgrounds can be distinguished. The idea was first developed in the psychological literature. See especially Hull [1943], and Bush & Mosteller [1955], on which Cross [1983] is based. Much later, reinforcement learning was independently reinvented twice as a machine learning approach in computer science. See, e.g., Barto et al. [1983], and Sutton [1992] for a survey of an approach called reinforcement learning. The other reinforcement learning approach in computer science is known as Classifier Systems. See Holland [1975] for early ideas on this, or Holland et al. [1986] for a more

[6] We presented this specific learning model in Kirman & Vriend [1995].

elaborate treatment of the issue of induction in general. In the economics literature reinforcement learning became better known only recently through Roth & Erev [1995].

Interestingly, Hayek [1973] seems to have anticipated much of the insights of this literature, and in particular his ideas seem very close to those presented in Holland et al. [1986]. *"'Learning from experience' ( ) is a process not primarily of reasoning but ( ) of practices which have prevailed because they were successful"* (p. 18). And *"(w)hat we call understanding is in the last resort simply his capacity to respond to his environment with a pattern of actions that helps him to persist"* (p. 18). Hayek then goes on *"to use the conception of evolution ( ) as an explanation of the rise of rules of conduct"* (p. 24) instead of *"construct(ing) such rules by deduction from explicit premises"* (p. 21). Hence, Hayek advocates an inductive approach as worked out in great detail in Holland et al. [1986].

More in general, Hayek favored what is nowadays known as adaptive behavior, and which is usually linked to the concept of 'bounded rationality', something which Hayek [1948b] called 'antirationalistic'. *"The antirationalistic approach, which regards man not as a highly rational and intelligent but as a very irrational and fallible being, whose individual errors are corrected only in the course of a social process, and which aims at making the best of a very imperfect material, is probably the most characteristic feature of English individualism"* (p. 8/9). In Savage's [1954] terminology, the adaptive behavior implied by this bounded rationality is known as following the *'cross that bridge when you meet it'* principle, which is necessary when an agent is in a *'large world'*, as opposed to the *'small world'* to which Subjective Expected Utility theory applies. In a large world, the agent's situation is ill-defined in the sense that he does not have a well-specified model of his environment. Hence, instead of *deducing* optimal actions from universal truths, he will need to employ *inductive* reasoning, i.e., proceeding from the actual situation he faces. As Hayek [1973] puts it: *"Evolutionary rationalism () recognizes abstractions as the indispensable means of the mind which enable it to deal with a reality it cannot fully comprehend"* (p. 30). Hayek also saw this adaptive behavior in terms of 'if ... then ...' rules. *"Whenever a type of situation evokes in an individual a disposition towards a certain pattern of response, that basic relation which is described as 'abstract' is present"* (p.30). And abstractness is *"the basis of man's capacity to move successfully in a world very imperfectly known to him- an adaptation to his ignorance of most of the particular facts of his surroundings"* (p. 30).

It should be stressed that the CSs are not models of agents using only simple decision rules. Although each rule for itself in a CS is a simple rule, it is the *set* of rules that forms the link between actions and previous actions and outcomes, and it is not the individual rules that matter. As is well-known, such a representation of knowledge is not restrictive in any sense, and any program that can be written in a standard programming language can be implemented in a CS. That is, these systems are *'computationally complete'* (see Minsky [1967]). Hence, a CS may be thought to model the most complex and sophisticated human decision procedures, as well as the most simple. In other words, *any* decision can be modeled *'as if'* made by a CS. This is getting close to Friedman's as if argument. Hayek [1973] uses this variant: *"( ) the rules ( ) need not be rules which are 'known' to these* *elements; it is sufficient that the elements actually behave in a manner which can be described by such rules"* (p. 43). Hence, a CS is a minimal form of modeling learning, in the sense that we do not need to make many assumptions about the reasoning procedures actually followed by the agents. As Hayek put it: *"( ) we can make use of so much experience, not because we possess the experience, but because, without our knowing it, it has become incorporated in the schemata of thought which guide us"* (p. 30/31).

All these concepts (the use of simple rules of thumb, and in particular 'if ... then ...' rules, the evolution of the rules used, which is based on the experience with the use of these rules, in particular their degree of success, and the role of abstraction and schemata) are worked out in great detail in Holland et al. [1986] in relation to Classifier Systems. But one important difference between Hayek and more recent approaches should be noticed. When Hayek [1973] uses the evolutionary argument, what he has in mind is that *"( ) selection will operate as between societies of different types"* (p. 44). Rules of behavior emerge *"( ) often not because they conferred any recognizable benefit on the acting individual but because they increased the chances of survival of the group to which he belonged"* (p. 18). The group selection argument is nowadays largely out of favor in evolutionary theories. However, in the Classifier System literature, and the reinforcement learning literature in general, the evolutionary argument operates at the level of the rules of behavior or conduct, nowadays usually known as rules of thumb, themselves. That is, each individual agent considers a set of rules, and these rules compete with each other.[7]

[7] But obviously, the social element has not disappeared completely. Which rules are good depends on which rules other people follow as well. Hence, evolution also takes place at a social level. This is called coevolution: one individual's set of rules evolve in response to the rules used by other individuals, with the sets of all these individuals evolving at the same time.

rule	choice
1	highest average
2	highest average if at least 2
3	highest average if at least 3
4	highest minimum
5	highest minimum if at least 2
6	highest minimum if at least 3
7	highest maximum
8	highest maximum if at least 2
9	highest maximum if at least 3
10	‘Bayesian’
11	‘Bayesian’ if at least 2
12	‘Bayesian’ if at least 3
13	majority with margin of at least 1
14	majority with margin of at least 3
15	majority with margin of at least 5
16	follow last
17	follow last 2
18	follow last 3
19	random
21-38	opposite choice of 1-18

Table 1 Decision Rules

Table 1 summarizes the set of rules we actually used in our model. Obviously, rules 1 to 12 are only eligible if both items appear in the sample (the same applies to rules 21 to 32). Rules 10 to 12 measure both the average and the variation of the values in a sample. To illustrate that these rules of thumb compete with each other, and that given the six sample observations, different rules of thumb may lead to different product choices, consider the following example. If the choices in an agent’s sample are three times item 1, and three times item 2, with utility levels of 0.48, 0. 71, and 0.28 for item 1, and 0.41, 0.37, and 0.44 for item 2, then rule 1 (choose highest average) would point to item 1, while rule 4 (choose highest minimum) would lead to item 2. The relative importance of each rule of thumb in a decision makers decision process depends on the payoffs generated by these rules of thumb, such that rules that gave riser to higher payoffs are more likely to be used. As explained above, the agents continuously update their beliefs in this respect.[8] Besides through the white noise added to the ‘bids’ of the Classifier System (see above), the agents experiment through some kind of ‘trembling hand’, mistakenly picking the item they did not intend to with a given small probability.

[8] Obviously, a more general analysis, including also the issues of creativity and innovation, would allow for new rules of thumb to be generated (rules we perhaps could not even imagine right now). This could be modeled with a Genetic Algorithm combined with our Classifier System.

4. Some related literature on information contagion and social learning

Alongside increasing returns (Arthur [1989]), network externalities (Katz & Shapiro [1985] and [1986]), information cascades (Bikhchandani et al. [1992), and herding behavior (Banerjee [1992]), information contagion (Arthur & Lane [1991]) has been presented in the literature as an explanation for particular patterns of macrobehavior (for example, path-dependency and lock-in effects) that may seem at odds with the underlying micromotives. But whereas these other explanations have been shown to have a proper microfoundation (either related to changing productivity or changing preferences, or to Bayesian updating in the face of uncertainty), information contagion has remained a phenomenon that occurs only when certain ad hoc rules of thumb for individual behavior are assumed.[9]

Making use of the theory of generalized Polya urn schemes (see Hill et al. [1980]), Arthur & Lane [1991] show that in a population of naive Bayesian optimizers information contagion tends to drive the market to stable (but not complete) domination by a single product. This is caused by the fact that the more a product has been chosen already by others, the more likely it is to be in an individual agent's sample on which he has to base his choice, and hence the more precise his information concerning the true quality of the product. Dosi et al. [1994] show similar results for a given rule of thumb that is more directly imitative.

Narduzzo & Warglien [1996] carried out two experiments to test the empirical importance of information contagion. Between 50 and 170 experimental subjects were instructed that they faced the choice between two products of which the objective value is uncertain, while the same product can generate a different value for different subjects. The players made their choice sequentially, and each player received information concerning the choice and value generated of a random sample of six subjects that had already made their choice at that point. In fact, the values generated by the two products had the same uniform distribution on [0.25, 0.75] and in another treatment these values were drawn for both products from [0.60, 1.00]. They observed that path-dependent dynamics emerged and that lock-in of market shares occurred, with early accidental choices giving rise to a seemingly stable cumulative market shares with the prevalence of one product over the other. They, then, tried to find out which choice heuristics the players used. Therefore, they interviewed some players after the experiment and they did some 'thinking-aloud' protocol analysis. They found four basic choice heuristics: the mean rule (highest average), the min rule (highest minimum), the max rule (highest maximum), and the popularity rule (follow majority). They observe that these rules are not necessarily used in isolation, and that subjects may have changed their rule on the basis of their experience from one run to another. Also, there might be context-sensitivity, with different samples inducing the use of different rules.

---

[9] A difference between the information contagion literature on the one hand, and the literature on information cascades and herding behavior on the other hand, is that in the latter an agent does not observe the payoffs generated by other agents, but only their choices as such.

Lane & Vescovini [1996] analyzed the aggregation effects of these four rules, with the assumption that all subjects follow the same rule. They find that the mean and the min rule never produce path-dependent behavior. The popularity rule always generates path-dependent market shares, and the max rule only when the products are exactly identical. Hence, three of the four rules reported by the subjects do not generate path-dependence. So the question is, where did the observed path-dependence come from? Lane & Vescovini note that the emergence of path-dependence and information contagion is related to the mix of rules actually used in the population, and that an important question is how people change rules after they experience outcomes. Since learning is a coevolutionary process (while one agent is learning all other agents are learning simultaneously), these two points should be considered combined. That is what we do. We want to understand the process through which information contagion emerges. Narduzzo & Warglien's experiments are one-shot games, but the players must have faced very many analogous decision problems outside the laboratory. Why does information contagion emerge, what role does it play, and what are the effects (both with respect to individual players and the society as a whole)? Where does a configuration of rules used in a population that leads to path-dependence come from? Is it based on arbitrary, extremely boundedly rational behavior? Or is it reasonable to learn rules of behavior that give rise to information contagion?

Ellison & Fudenberg [1993] consider a closely related problem of the use of rules of thumb in a situation where agents need to choose between two items with unknown value, and where social learning takes place. They assume that players use exogenously specified, simple rules of thumb. One justification they give for this is that they do not consider fully Bayesian learning a realistic assumption, because it requires calculations that may be too complicated. In each period some fraction of the players have the opportunity to revise their choices. They only observe last period's payoffs and choices of all agents. They present some simple and some more complicated rules which are all some form of a popularity weighted choice of the highest average, and they derive the optimal weight of the popularity for some of these rules. However, since they assume rules of thumb that are exogenously given, a relevant question is whether it is likely that these optimal weights will actually be used. In support of a positive answer Ellison & Fudenberg note that they showed in a working paper that all agents using the optimal weights constitutes an equilibrium, but they stress that important extensions of their analysis are needed. The reason for this is that the precise specification of these rules supposes more sophistication of the agents than they find themselves compelling. Therefore, conjecturing that the optimal popularity weighing might emerge from an adaptive process, Ellison & Fudenberg explicitly ask for a complementary study. As they put it, *"(i)t would be interesting to complement these results with an analysis of a dynamic process by which players adjust their rules of thumb along with their choice of technology"* (p. 638).

Obviously, organizing an experiment with the sequential choice of a large number of players is quite complicated, even if it is for only one period. Hence, it is impossible to study the dynamics of such learning behavior with human subjects by the experimental method. Another problem with organizing experiments is that it is difficult to characterize individual behavior since the rules of thumb used by the players are not directly observable, while subjects' reports are not very reliable. One reason being, that each player might actually use very many possible rules, without being aware which rules he exactly uses at a certain moment. Hence, questionnaires tend to be rather inconclusive. Therefore, this seems an excellent case to use an ACE model. Not only does this allow us to analyze very long run dynamics, but we can also do an explicit analysis of the rules of thumb actually used.

The basic choice situation in our ACE model as described in section 3.1 follows closely the one used in the information contagion papers discussed in this section. The two main differences are the following. In the information contagion experiments by Narduzzo & Warglien [1996], the expected values of the two items are actually identical, although this was not known to the subjects. In our ACE model the two expected values are generally different in every period (apart from the benchmark periods). The reason to use these different expected values in all periods not being a multiple of 500 is that if the two products would have an identical performance distribution in every period, then there would no relation between the rules of thumb used and the payoffs generated. That is, since on average any of the two items is equally good, any rule of thumb is as good as any other rule, and hence there would be no selection pressure at all. Therefore, in our model in each cycle of 500 periods we have 499 periods in which the average performance of the two products is different, in order to give selection some bite, and then we check in the 500th period what the selection process has achieved by using the two products with identical performance as a benchmark case.

In Ellison & Fudenberg [1993] there are only two items, and the agents can revise their choice between these two specific items in every period. As explained above, in our model, agents face the choice between two specific items only once. Obviously, this could be easily modified by allowing agents to revise their choice once they have experienced it themselves (as in Ellison & Fudenberg [1993]). However, the fundamental problem, then, is still the same. On the basis of a small sample of observations the agents need to make a choice between two items. The way they do this is using some rule of thumb. Hence, the choice of rules of thumb will still be an essential issue, and a dynamic model with adaptive agents remains a promising approach to understand this.[10]

5. Analysis of the model

In order to analyze the properties of our ACE model, we run the model sketched in section 3 with 100 agents for 25,000 periods. The central question we want to analyze is: *"How can the combination of fragments of knowledge existing in different minds bring about results which, if they were to be brought about deliberately, would require a knowledge on the part of the directing mind which no single person can possess?"* (Hayek [1948c], p. 54). From an objective point of view, in almost every period one of the two items is superior. But as Hayek points out: *"The equilibrium relationships*

[10] Notice that the dynamics that Ellison & Fudenberg [1993] study concerns the choice of technology from period to period, and not the dynamics of the rules of thumb to choose technologies from period to period.

cannot be deduced merely from the objective facts, since the analysis of what people will do can start only from what is known to them" (p. 44). This leads to *"( ) the general question of why the subjective data to the different persons correspond to the objective facts. Our problem of knowledge here is just the existence of this correspondence ( )"* (p. 51/52).

Although knowledge is very much divided in our model, each agent has a sample of six observations, and these samples of the individual agents overlap. Related to this, Hayek [1948e] conjectures that *"( ) the whole acts as one market ( ) because their limited fields of vision sufficiently overlap so that through many intermediaries the relevant information is communicated to all"* (p. 86). Hence, some more specific questions we want to answer are the following. Do the agents through their spontaneous interaction learn to use rules of thumb that solves the division of knowledge problem? How do the market outcomes look like? And do we get path-dependence and lock-in effects?

5.1 Path-dependency and lock-in

We first focus on the benchmark periods in which the expected value of both items is 0.50, i.e., the periods that are a multiple of 500. We want to know how the market shares of the two items develop as we go down the sequence of 100 agents in a given period, and in particular we want know how this development changes over time as the agents learn which rules of thumb to use. Figure 3 shows some examples of a typical run: the development of the market share of one of the items in periods 1000, 10,000, and 23,000.[11]

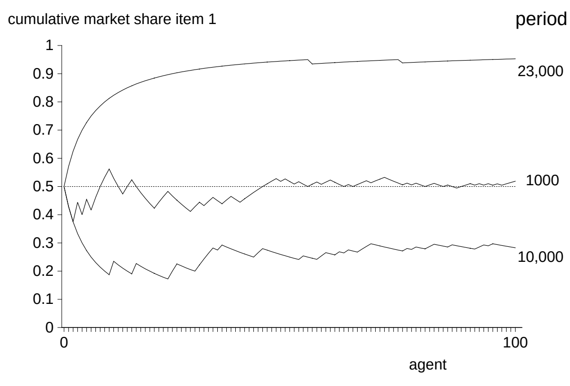

Fig. 3 Cumulative market shares

[11] We did various runs, but for presentational reasons most results here are given for one representative run only. One of the reasons being that any single run shows some typical fluctuations that are largely canceled out by averaging over various runs. Each sequence starts with a market share of 0.50 because of the initial choices by the six dummies. The market share of the other item is just one minus the share of the item shown, that is, the curve shown mirrored in the straight line at 0.50.

If there were no information externalities or contagion at all, every choice would be an independent decision, with each of the two items being equally likely to be chosen (as in these benchmark periods the two items were equally good), and the development of the market shares would be basically a random walk, zigzagging around a 0.50 market share. As we see in figure 3, this is what happens in period 1000, when the agents have had only little opportunity to learn, and they basically behave like 'zero intelligence' agents (see Gode & Sunder [1993]). If we showed this curve in period 1000 for different runs, or other benchmark periods towards the beginning of a run, we would get a series of different zigzag curves that stay all close to the 0.50 market share line.[12] The market share curve shown for period 10,000 looks very different. Just as in period 1000, we see some deviations from a 0.50 share early on, but unlike in period 1000, this time we see that the item that gets a smaller market share early on continues to lose ground for some time. Eventually, its share stabilizes at a level just below 30%. The rather smooth curve for period 23,000 shows the positive feedback effect even much stronger. We see one item right from the start increasing its market share continually until it dominates the market almost completely. Although the two items are identical in this period, the information contagion (as the agents happen to learn such behavior, as we will show below) leads to lock-in. Which of the two item gets to dominate is basically random, due to small historical events. That is, it is path-dependent.

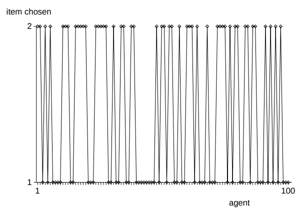

Fig. 4.a Individual choices, period 1000

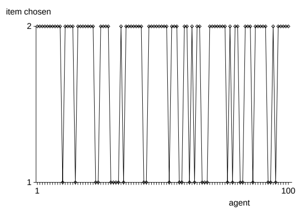

Fig. 4.b Individual choices, period 10,000

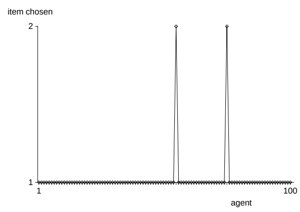

Fig. 4.c Individual choices, period 23,000

Figure 4 looks at the same phenomenon, the setting in of information contagion, focusing on the individual choices of all 100 agents. In figure 4.a we see an almost random sequence of choices in period 1000, shifting from one item to the other item all the time, and there is very little order, if any. In figure 4.b, showing the same for period 10,000, and in figure 4.c for period 23,000, we see an increasingly orderly pattern. In figure 4.c, although the two items are identical, item 2 does not seem very fashionable, with agent after agent choosing the item 1, and only some occasional deviations from the norm.[13]

[12] The fact that the zigzags appear to become smoother towards the end of the sequence is due to the fact that each additional decision maker carries less weight in the cumulative market share as we move down the line of 100 agents.

[13] If the fashion changed from one item to the other every now and then, this would show up as waves in figure 4.

As figures 3 and 4 show, the decentralized interaction of the individual agents leads to a situation in which almost all agents choosing the same item emerged as a spontaneous order, where "*(b)y 'order' we ( ) describe a state of affairs in which a multiplicity of elements of various kinds are so related to each other that we may learn from our acquaintance with some spatial or temporal part of the whole to form correct expectations concerning the rest, or at least expectations which have a good chance of proving correct*" (Hayek [1973], p. 36).

The behavior of a complex system is often said to be characterized by a *'2+2=5 effect'*, the system being more than the sum of its parts (see, e.g., Parker & Stacey [1994]). It might be that this term comes from the description of the behavior of simple non-linear dynamic functions, where the chaotic outcomes, going 'all over the place', seem profuse given the simple input specification. However, the striking feature of self-organized systems, as stressed by Hayek [1973], is not their *chaos* (that anything can happen), but quite on the contrary their *order* (that something very precise happens). It might also be that the *'2+2=5'* label stems from the observation that it is not possible to understand the behavior of a complex system by examining only its separate components. While this is correct, when it comes to self-organization the issue seems not so much that a complex system is *more* than the sum of its parts, but that it is *less* than the sum of its parts, with the difficulty arising because one cannot predict which of the possibilities will be realized by examining only the constituent parts. Hence, when it comes to describing the self-organizing aspect of a complex adaptive system, the emerging spontaneous order, it seems more useful to call it a *'2+2=3 effect'*.[14]

The market share curves and the individual choice curves shown suggest a simple story. As time goes on, the more information contagion develops, and the more the population gets self-organized, the more the development of market shares gets a particular pattern, with rather smooth curves concentrated in a relatively small space with either a very high or a very low market share. However, as we will show in a moment, matters are slightly more complicated. The spontaneous order emerging turns out to be far from absolute. In every run it takes some time before the information contagion starts giving rise to lock-in and path-dependency effects, but once the population gets self-organized this turns out to be not a monotonic process at all. Giving the period 10,000 or 23,000 curves for different runs, or giving these curves for different benchmark periods from the same run, would show market share curves going all over the place. Sometimes one item almost completely dominates the market, other times we see the fashion switching at some point from one item to the other, and sometimes this switching occurs so frequently that we get a zigzag curve similar to the one shown for period 1000. Hence, the curves seem to drift about in all directions, and the system moves all the time between almost complete order and almost complete disorder, but never stays at either of these. We will explain this phenomenon in section 5.2, but first we will illustrate this a little bit more by using different measures for what goes on in the market.

[14] We thank Martijn Huynen for this label.

Obviously, the final market share of an item is not exhaustively informative concerning the amount of lock-in generated. One change in fashion at the middle of the sequence would be sufficient to end up with 50-50 shares. Therefore, we take as a measure the average absolute distance between the cumulative market share curve (as shown in figure 3) and the straight line at 0.50, divided by 0.50. Basically, we take the relative size of the area between the market share curve and the straight line at 0.50, relative to the size of the rectangle defined by the axes and the 0.50 line as a measure of the path-dependence in the population's decisions. The more systematically the market stays away from a 50-50 distribution, the more lock-in we have. This measure, the lock-in rate, is a number between 0 and 1, and is shown in figure 5 for all benchmark periods, i.e., those that are a multiple of 500. As we see, lock-in ranges from low values around 0.1 at the beginning, and tends to get higher values as time goes on, up to about 0.80, but there remains a lot of variation all the time, with lock-in regularly falling back to the low initial values.

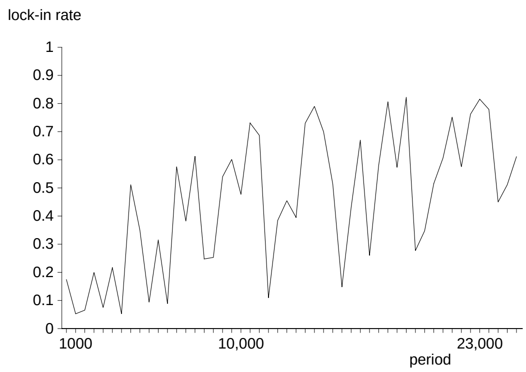

Fig. 5 Lock-in rates in benchmark periods

Another way to measure how much lock-in into one of the two items is present is the rate at which the choices of the agents switch from one to the other item in the benchmark periods. If each individual decision is taken independently, and the items are equally good, the expected switch rate is 0.50. Figure 6 shows the switch rates for each of the benchmark periods. As we see, the switch rate starts indeed around 0.50, and then comes down as time proceeds, but just as with the lock-in rates above, this goes with a lot of fluctuations. The switch rate regularly comes down to values close to 0, implying a very orderly state in which every agent chooses the same item, but almost equally regularly the switch rate jumps back to levels close to 0.50, the maximum disorder, as if all agents choose randomly.

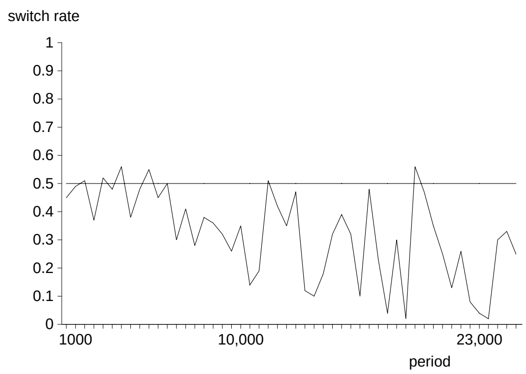

Fig. 6 Switch rates in benchmark periods

5.2 Performance over time

In the benchmark periods analyzed in section 5.1, the two items were always identical. In those periods, any item was as good as the other item, and hence any decision rule was as good as any other decision rule. We used those periods to see how much information contagious behavior the agents had acquired during the periods in between the benchmark periods, periods in which the two items were generally not identical. Before we analyze the behavior of the individual agents, we first want to see what the effects of the learning of the agents is on the overall outcomes for the society.

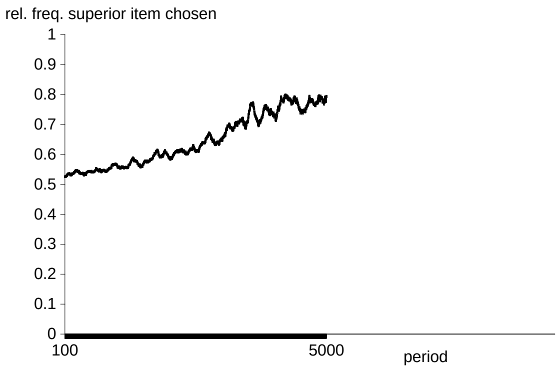

Fig. 7.a Mov. avg. performance, periods 100-5000

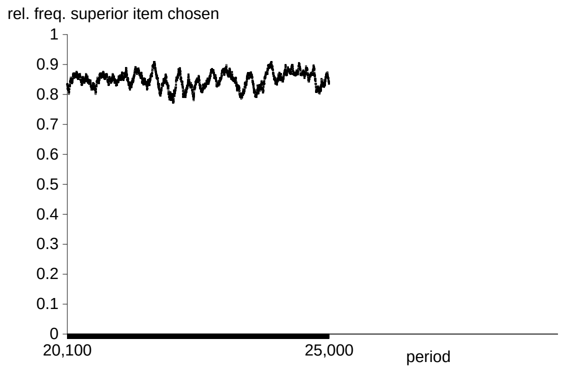

Fig. 7.b Mov. avg. performance, periods 20,100-25,000

Figure 7.a shows the 100-period moving average of the relative frequency with which the best of the two items is picked in each of the first 5000 periods, and figure 7.b does the same for the periods 20,100 to 25,000. The relative frequency with which the superior item is chosen, i.e., the final cumulative market share of the superior item in a given period, is a good measure of social efficiency.

At the start, with people making almost random choices, about 50% of the agents pick the correct item. This average frequency increases over time, and in figure 7.b it has reached a level of about 86%, without much further increase suggested by the trend.

The increase in performance as such takes place in part because the agents start using better rules, but in part this is also because they start aggregating information. We will analyze this in more detail below. That we do not reach higher efficiency levels is related to the fact that often the two items have an expected performance that is extremely close. In fact, on average the expected performance for the worst item turns out to be 0.42, and for the best item 0.58. Hence, in many periods the difference in expected performance is close to zero.

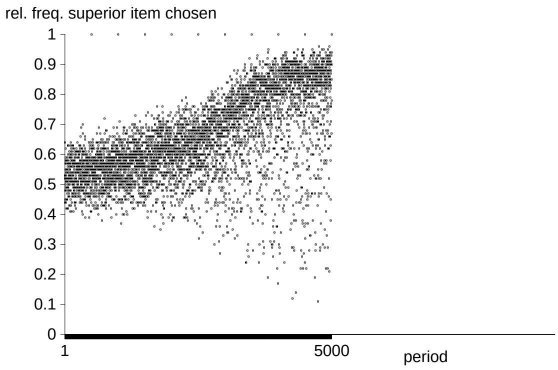

Fig. 8.a Performance, periods 1-5000

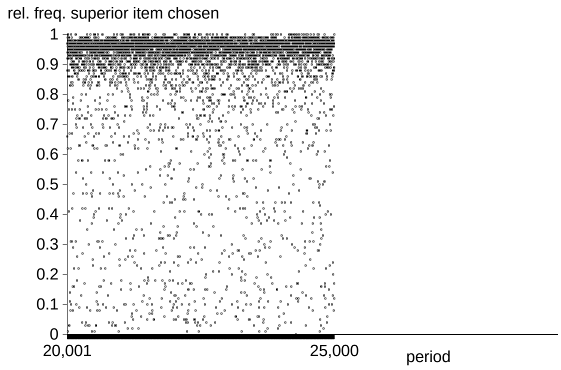

Fig. 8.b Performance, periods 20,001-25,000

Instead of the 100-period moving average of the performance, figure 8 shows every single period, and reveals that underlying these moving average performances something interesting is happening. Notice that every 500 periods the items are equally good, and hence everybody makes the right choice. More interesting is the observation that while the (moving) average performance goes up, the spread increases as well. In the beginning, in every period about 50% of the agents choose the correct item. Sometimes this is a little bit lower, and sometimes a little bit higher, but never very much so. For some time performances never exceed the 35 to 70% band. But as times goes on, and average performance goes up, occasionally periods occur in which only 30% of the agents pick the superior item. Later on there are periods with just 15% choosing correctly, and eventually (after about 5000 periods) it sometimes even happens that almost nobody recognizes which is the best item. All the time, though, the moving average of the performance shows an upward trend. As noticed above, in part the spread in performance is due to the fact that on average the expected performances of the two items is rather close, occasionally leading many people to the wrong choice. But the frequency with which the expected performances are close to each other (making mistakes likely) does not change over time. Hence, the change in spread over time that we observe is due to the adaptive behavior of the agents. As they learn, they improve their average performance, but occasionally this leads to disasters, with almost everybody choosing the wrong item.

As mentioned above, there are basically two things going on during the history presented. First, the agents learn to use better rules as such, i.e., the rules that lead to higher utility levels because they are better at recognizing which is the superior item on the basis of six sample observations. The dynamics are in part the result of this evolution of the rules being used. Second, as a given agent learns which rules to use, this influences his choices, which has two effects. First, it gives the choosing agent a certain utility. But at the same time, there is also an *externality*, as the choice of the given agent is added to the information pool on which the choices of future agents will be based. Hence, there is interaction between the agents' choices. As one agent learns and changes his behavior, other agents are learning as well, partly in response to this. This is a coevolutionary process; the rules that an agent uses evolving in response to the evolution of other agents' rules. Why is this information externality relevant? That is, why would an agent want to adapt his rules in response to what other agents do?

This is because of the phenomenon of information aggregation. On the one hand, as Arthur & Lane [1991] argued, an increase in the number of observations concerning a certain item leads to more reliable estimates of its value, and hence makes the item more attractive for risk-averse agents. On the other hand, there is even more direct information aggregation. Suppose every agent bases his choice on the six observations he has sampled, i.e., on the choice of item and the payoffs actually generated for those six agents. Now, consider a rule that tells an agent to follow the choice of the majority in his sample. This rule does not consider the actual payoffs generated by the six agents in the sample. But if each of the six agents in the sample had considered the payoffs in their samples of six, then the 'follow the majority' rule implicitly considers six times six or 36 sample payoffs instead of only six. That is, the rule aggregates the information of the agents in the sample.

To analyze the effects of these two factors that explain the increase in the average performance (i.e., the learning as such, and the information externality) we did the following experiment, which excludes the information externality. The basic choice situation in this variant of the model is the same as above. But this time every agent, when making his choice, does not observe what other agents did before him, nor the payoffs they realized. Instead, when an agent's turn comes, he can six times randomly choose an item himself, and observe the payoffs generated.[15] Hence, the only difference with the standard model is that there is no interaction between the agents, hence no information externality, and thus no possibility of information aggregation.

[15] Notice that to follow the previous setup closely, and to keep matters simple we do not consider the important issue of what the optimal sampling strategy would be.

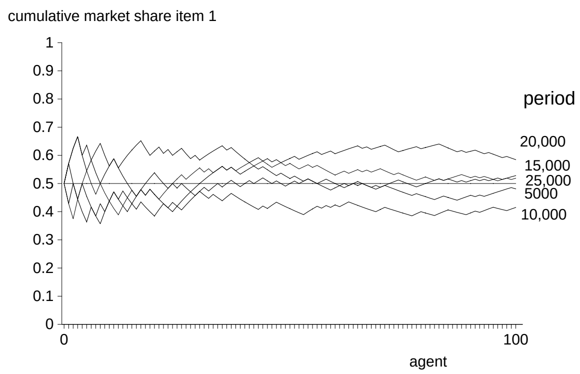

Fig. 9 Cumulative market shares (variant)

Figure 9 shows the cumulative market shares for the periods 5000, 10,000, 15,000, 20,000, and 25,000. As we see, there is no lock-in or path-dependency, which also applies to all benchmark periods not shown. The cumulative market share always stays around 0.50. This was to be expected, because these were all benchmark periods in which the two items are identical, and all agents make their choice independently. Since there is no information externality, we get no path-dependence, and hence no lock-in.

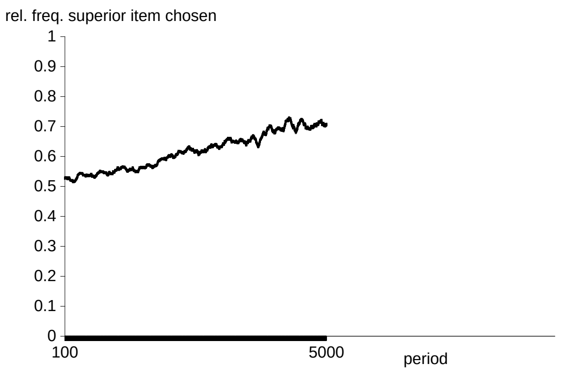

Fig. 10.a Mov. avg. performance, periods 100-5000 (variant)

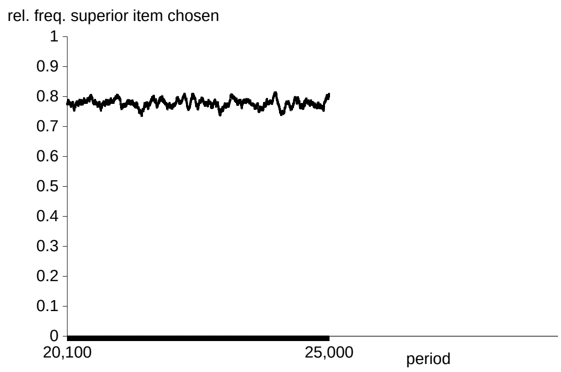

Fig. 10.b Mov. avg. performance, periods 20,001-25,000 (variant)

Although there is no information externality here, of course, the agents still learn which rules are more likely to pick the superior item on the basis of a sample of six observations. Figure 10.a shows the 100-period moving average of the relative frequency with which the best of the two items is picked in each of the first 5000 given periods, and figure 10.b does the same for the periods 20,100 to 25,000. We observe that the performance, starting from a level of 0.50 which even random choice would achieve, increase to a level of about 0.77. That is, the agents do learn to improve their performance by using the better rules, but they stay below the average performance in the standard version, when it reached a level of about 0.86. In other words, taking advantage of the information externality by aggregating knowledge, the agents succeed in winning another 10% in performance in the standard version.[16]

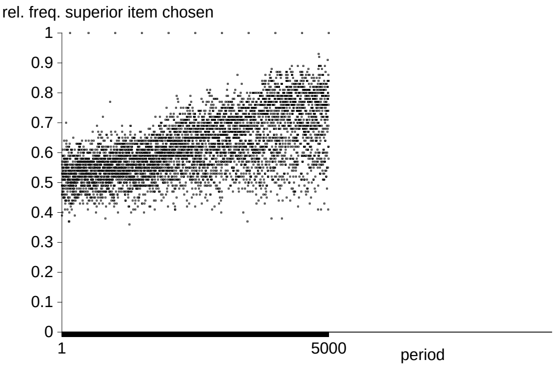

Fig. 11.a Performance, periods 1-5000 (variant)

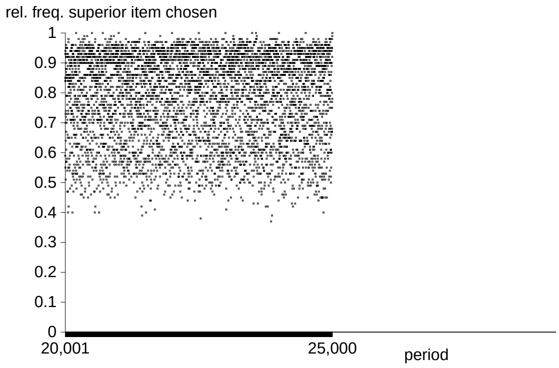

Fig. 11.b Performance, periods 20,001-25,000 (variant)

Figure 11 shows every single period of the graphs underlying figure 10. As we see, the performance tends to rise, but, apart from the benchmark periods, it almost never gets close to 1, and there are also no disasters. In the most unfortunate periods, still about 40% of the agents succeed in choosing the superior item. This illustrates at the same time the advantage and the limits of the information aggregation. By aggregation the agents succeed in reaching very high performance levels in many periods, higher than they could ever achieve on their own. But when agents aggregate information (e.g., following the majority rule), an agents wastes some information as well, since he does not use the information concerning the actual payoffs realized by the six people in his sample.

As we explained above, if a single agent aggregates information he implicitly uses six times six, i.e., 36, observations instead of only the six in his own sample. But if each of the six agents in his sample would also be aggregating information, they would each implicitly use 36 observations, and hence our single agent would be using six times 36, i.e., 218, observations. Hence, the more agents use aggregating rules, the more aggregation of knowledge occurs. But when too many agents aggregate information, too many agents waste their own information. At some point a tiny little bit of knowledge starts getting aggregated *ad absurdum*. In some sense, the agents start aggregating ignorance instead of knowledge.

[16] The average performance in the standard version in the periods 24,001 to 25,000 is 0.861, and in the variant without externalities this is 0.774, implying a difference of 11.2%.

5.3 The individual decision rules

One of the advantages of an ACE approach is that we, as modelers, know for each single period which of the two items is superior. Hence, for each single decision to be made by any of the agents, given his sample of six observations, we can check for each of the 39 rules whether it would have picked the superior item. Obviously, the individual agents do not obtain this information. They only try one rule of behavior in every period, and observe the payoffs they generate doing so. Figure 12 shows the time series of the relative frequencies that a given rule would have picked the superior item.[17] In fact, this graph shows the relative frequency that a given rule belongs to an agent's best-reply correspondence.

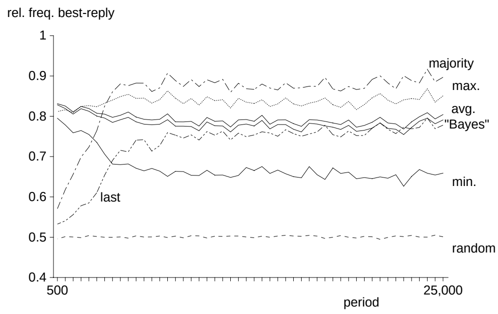

Fig. 12 Specific rules as best-replies

As we see, picking an item at random (rule 'random') leads to the superior item in about 50% of the cases, and this remains constant over time. Looking at just one other agent, and imitating whatever he picked (rule 'last') starts close to 50%, but as other agents learn to make better choices, the performance of this rule increases considerably, and gets close to the rules that choose the highest maximum in the sample (rule 'average'), or the highest maximum compensated for the variance (rule '"Bayesian"'). This increase in performance applies even much more to the rule that says to follow the majority of the six observations sampled (rule 'majority'). This rule, which does not use any of the available information concerning the utility levels obtained by the six agents sampled, at some point starts beating all other rules.[18] Two other rules stand out. The rule that chooses the highest minimum (rule 'minimum') deteriorates over time. The explanation for this is that, implicitly, it does the opposite of information aggregation. It favors the item that is the least chosen, because the more

[17] These frequencies are normalized for eligibility, since, as we explained above, in some cases the *'if ...'* part of a rule is not satisfied. Each observation concerns one cycle of 500 periods (from one benchmark period to the next). For presentational reasons we only show the rules 1, 4, 7, 10, 13, 16, and 19 (see table 1).

[18] Arthur & Lane [1991] argue that lock-in coming out of the simple imitation of other people is not interesting, but what makes it interesting here is that the simple imitative behavior itself is an emergent phenomenon.

an item is chosen the more likely it is that some observation will be in the lower part of the distribution, and hence be the lowest minimum in the sample. Exactly the opposite applies to the rule that chooses the highest maximum (rule 'maximum'). The more an item is chosen, the more likely it is it will provide the highest maximum in the sample.

The important thing to notice here is that the degree to which a given rule is objectively good changes over time as a result of the other agents are changing the rules they use. To show how the effect of the information externality makes it a coevolutionary process, i.e., the agents adapting to each others' adaptation to each other ..., figure 13 presents the frequencies with which the individual rules form part of an agent's best-response correspondence in the variant in which there are no information externalities. As we see, these frequencies remain constant over time, apart from some random noise. The only thing the agents need to learn is to figure out which of these rules are most often good in a given situation. Obviously, for different situations different rules might be best. But which rule is good for a given sample configuration does not change over time. This is very much unlike figure 12, in which the learning of the agents influences in turn what the other agents have to learn.[19]

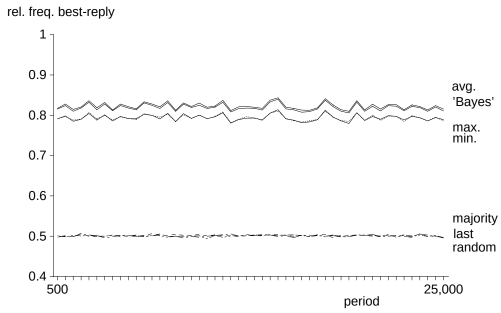

Fig. 13 Specific rules as best-replies (variant)

To conclude our analysis of the model, could it be that the famous QWERTY lock-in has less to do with network externalities and other real payoff matters than with information contagion? After all, with the current technology, and most people using a personal computer, switching a keyboard layout is relatively easy. It is true that it requires a little bit of personal investment (time and effort to change the layout itself, plus some re-training), but if individual agents knew an alternative keyboard were superior, that would be no obstacle. The only problem seems that individual agents do not know whether it worth choosing an alternative keyboard layout, and generating own sample observations by trying various different keyboard layouts is rather costly. Hence, an individual agent needs to base his

[19] Another advantage of an ACE approach is that we can do a precise analysis of what each individual agent learns. In particular, there is no reason to assume that every agent learns exactly the same, but we will not pursue this here.

decision on the choices made by other people, and as our ACE model demonstrates, it might be that it is the emergence of information contagious behavior that leads to a QWERTY lock-in.

6. Discussion

Hayek greatly contributed to keeping the tradition alive, and extending it considerably, of studying the economy as a self-organizing system (a tradition going back at least to Adam Smith), against the dominating influence of both the Walrasian general equilibrium theory and the Keynesian macroeconomics. In doing so, he presented us with important research questions, some of which have been taken up in this study in which individuals have to make repeatedly a choice between two previously unknown items while they can rely only on some information from previous adopters.

One of the particular contributions of Hayek is his view, first, that through a self-organizing process the economy will overcome the problem of the division of knowledge, and second, that in aggregating information this will lead to the socially desirable outcomes. Our ACE model basically confirms Hayek's vision, but one of the benefits from this ACE approach is that we see that matters are slightly more complicated than perhaps expected by Hayek.

First, our ACE model exhibited self-organization, and the emergence of spontaneous orders, in which typically most agents choose the same, superior item, and indeed we observed that *"such an order will utilize the separate knowledge of all its several members, without this knowledge ever being concentrated in a singe mind ( )"* (Hayek [1973], p. 41). But it turned out that this was not a simple monotonic process from disorder to order until the solution had been reached, with a happy ending. Instead, the system was continually moving back and forth between order and disorder. That is, the self-organization is a continuing, ongoing story, in which the emerging order unravels time and again.

Second, the emerging spontaneous order was beneficial. That is, on average. But along with the improved average performance we also saw an increase in both the number and degree of disasters. We explained how there is a tension between generating knowledge and aggregating knowledge. If enough knowledge is generated by the individual agents, aggregation leads to good outcomes, but if everybody would merely aggregate over and over again a tiny little bit of knowledge, this might lead occasionally to very bad outcomes for the society. In fact, it is this which keeps the self-organizing process from being a monotonic one. If it were monotonic, we would get stuck with only disasters.

Hence, we can address Hayek's [1948c, 1948e] optimistic conjecture that the overlap of the knowledge of the individual agents will aggregate the divided knowledge in such a way that the objectively optimal solution will emerge. As our ACE model shows, pointing to occasional disasters (QWERTY?, VHS?) is not sufficient to reject this hypothesis, as the average performance did improve considerably. This was the result of the emergence of information contagious behavior of the individual agents. Information contagion is a way to aggregate distributed knowledge in society, allowing the individuals and the society to achieve higher performance levels. This came together with path-dependent lock-in effects, but in some sense, the remarkable thing is not so much the emergence of these effects, but the fact that this was related to a high performance.

Although Hayek did not have the ACE tools available,[20] he had not only a view of the economy as a complex system, but his explanation for the emergence of beneficial spontaneous orders also relied on a view of the role of the individual agents that seems very much congruent with the one used in the ACE approach, which focuses on bounded rationality and rule-based behavior. For one reason or another almost every paper in economics that uses the concept of bounded rationality refers to Simon (see e.g., Simon [1955], [1957], [1959] or [1976]). But Hayek [1948b] asserts that viewing human agents as inherently boundedly rational has a much longer tradition. Moreover, Hayek's view of the importance of inductive behavior, with the emergence of simple rules of behavior on the basis of the outcomes experienced in the use of such rules, fits almost perfectly with much more recent approaches in computer science and ACE (see, in particular, Holland et al. [1986]). And although Hayek used a group selection argument, which is largely out of favor nowadays, our ACE example illustrates that the basic idea of Hayek can be retained, with selection operating on a set of ideas instead of groups. In our ACE analysis the beneficial information aggregation, and more in particular the best rule being the one to follow the majority, did not emerge because of an evolutionary process working through group selection, nor did it come through a selection of individuals. It arose through a coevolutionary process, the simultaneous evolution of rules of behavior used by the individual agents.

And apart from envisaging the tools used in recent ACE research, Hayek also correctly foresaw that it is not the use of simple rules of thumb as such that matters, but also the fact that this usage is based all the time on the experience of the outcomes thus generated. Only the current state-of-the-art allows us to model these matters precisely. Old fashioned simulations based on fixed rules of thumb as done in the fifties, sixties or seventies would not work. For example, the information aggregation, and in particular the rule to follow the majority *emerge*. If we had specified *a priori* that the individual agents follow the majority rule then we would have stayed at a performance level of 0.50. Also, when the majority rule emerges as a good rule, this does not imply that everybody should follow it. If they did, then the performance would fall back again to 0.50. Hence, what matters is also the precise configuration of rules used in the population. And the continuously changing configurations that emerged turned out to lead to both a high performance level, and information contagion with path-dependent lock-in.

In all excitement about the computer output in ACE research, methodological issues are sometimes a little bit neglected. Nevertheless, they are quite important. Not because ACE presents innovative methodological insights, but because it turns out to be a very advantageous way to actually apply the

[20] Littlechild [1986] even asserts that Austrian models in general have not been represented mathematically.

abstract methodological insights of Hayek and others. Moreover, acknowledging this, in turn, helps to put ACE research in the right perspective, facilitating a fruitful interpretation of its output as well.

Social theory is the explanation of social phenomena, where explanation is the brain's modeling of the complex events to be explained. As Weimer [1982] puts it: *"explanation is modeling"* (p. 271). According to Hayek [1948d], what we do is, *"we construct hypothetical models in an attempt to reproduce the patterns of social relationships which we know in the world around us"* (p. 68). In contrast to the natural sciences, in the social sciences *"(e)xperimentation is impossible, and we have therefore no knowledge of definite regularities in the complex phenomena in the same sense as we have in the natural sciences"* (Hayek [1948g], p. 126). That is, as Weimer [1982] explains, in the natural sciences there is the ability to simplify and control a situation to the extent that it can be repeated, either under identical conditions or those that we choose to vary systematically, such that we can isolate and identify the definite regularities in observed phenomena. However, *"(t)he empirical research in complex social phenomena consists in the construction of situations in which we demonstrate to ourselves that we can produce "facts" of which we are already well aware. Our demonstrations "test" our theoretical models only in the sense already noted; they compare the consistency of our theoretical model with an analogical knowledge of social phenomena, but they neither confirm nor refute them in any logical sense"* (Weimer [1982], p. 252). And Hayek [1948d]: *"The theory itself, the mental scheme for the interpretation, can never be "verified" but only tested for its consistency. It may be irrelevant because the conditions to which it refers never occur; or it may prove inadequate because it does not take account of a sufficient number of conditions. But it can no more be disproved by facts than can logic or mathematics"* (p. 73). And *"( ) a simple theory of phenomena which are in their nature complex ( ) is probably merely of necessity false ( )"* (Hayek [1967b], p. 28).

This idea that it is not possible to test the truth of a social theory, and that the best we can aim for is doing consistency checks, resembles Friedman's [1953] *'as if'* argument. *"We may not be able directly to confirm that the causal mechanism determining the phenomenon in question is the same as that of our model. But we know that, if the mechanism is the same, the observed structures must be capable of showing some kinds of action and unable to show others; and if, and so long as, the observed phenomena keep within the range of possibilities indicated as possible, that is so long as our expectations derived from the model are not contradicted, there is good reason to regard the model as exhibiting the principle at work in the more complex phenomenon. ( ) Our conclusions and predictions will also refer only to some properties of the resulting phenomenon, in other words, to a kind of phenomenon rather than to a particular event"* (Hayek [1967a], p. 15).

Hence, Hayek advocates not only an 'as if' argument, but he also argues that we can hope to explain at best general principles, or stylized facts. Although he did not discuss ACE models as such, Hayek [1982] seems to use the same argument concerning the degree of explanation that we can achieve also with respect to ACE models of complex social phenomena. *"Assume I could construct a rat - that is, a mechanical model that can do all a rat does. ( ) To be a really true model, it would* *clearly have to do also a great many things we could not predict, even though we know precisely how the mechanism we have built works. It would both occasionally have to respond to external stimuli in a manner that we cannot predict, but also have to act "spontaneously" in response to internal processes that we cannot observe. The reason for our inability to predict, in spite of our precise knowledge of the mechanism that moves it, would be that our mind is not capable of perceiving and digesting, in the same manner as the mechanical rat does, all the particular stimuli that operate upon it and all the processes of classification that proceed in it. The only means by which we could achieve predictions would be to build a computer that imitates all that the mechanism of the rat performs; or, in other words, to build another rat identical in structure with the first one and making it live from the beginning in exactly the same environmental conditions, so it would perceive an learn exactly what the first rat does. That is, in order to understand what a rat will do and why it does it, we would have to become another rat"* (Hayek [1982]. p. 292/293; emphasis added).

Perhaps it is useful to stress that Hayek is here arguing in *favor* of building a *"really true model"* of a rat (or ACE models, for that matter). The skeptical part of his remarks is related to the fact that he explicitly uses this illustration to justify his contention concerning the *"absolute limit to our powers of explanation"* (p. 292). In exactly the same way, ACE models are abstractions from reality, and not aimed at replicating reality. Hence the term *'simulation'* to describe ACE models might create confusion, since ACE models do not try to simulate reality as such, but only to understand some general phenomena, the stylized facts. As Kirman & Vriend [in press] explain: *"We will not try to build a model fitting all aspects of the real world for the following reasons. First, every model is by definition an abstraction. If enough data can be collected, statistical testing will reject any model. Second, when modeling by building artificial worlds, one might get a very good fit without gaining understanding. There exist economic simulation models with more than 10,000 variables. At some point it might be that one mainly succeeds in building a copy of the real world, about which we have the same degree of understanding as about the real world. Therefore, we will only consider specific questions concerning the stylized facts of the real market that appear remarkable or important. We will try to build a minimal model that generates, and with which to test those stylized facts. This might suggest ways to understand, or not, those phenomena. This understanding is of the same type as with formal mathematical models. The question is whether we might consider the real world to be working 'as if' it were like our model".*

Hence, social theory is explanation, and explanation means modeling. Models can be presented in various forms. They could be either purely verbal or quantitative, where quantitative models, in turn, might be a set of mathematical expressions or a set of computer instructions. That is, a computer program as used in ACE *constitutes* a model. And since explanation *is* modeling, and this is what social theory is all about, an ACE computer program as such *is* social theory. The only, and essential, reason to execute the computer program is to carry out the consistency checks; both with respect to reality and with what one anticipated the model to produce.[21] A possible advantage of quantitative models in general might be that it can be analyzed more precisely. That is, the consistency checks can be done more carefully. Such a consistency analysis can be a formal, mathematical analysis, or it could be a numerical analysis. The great advantage of ACE models, then, is that one can do much more extensive and elaborate consistency checks. In some sense, ACE allows us to do experiments now also in the social sciences. But notice that Hayek's distinction between the natural sciences and the social sciences remains still valid, because the experiments made possible by ACE concern the level of the *model*, and not reality itself.

To conclude, it would probably be presumptuous to judge whether Hayek might have been an ACE, but in any case it seems clear that ACE is social theory in a Hayekian tradition. And therefore a further exchange of insights would seem fruitful.

References

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

Arthur, W.B. (1992). *On Learning and Adaptation in the Economy* (Working Paper 92-07-038) Santa Fe Institute.

Arthur, W.B., & Lane, D.A. (1991). *Information Constriction and Information Contagion* (Working Paper 91-05-026) Santa Fe Institute.

Banerjee, A.V. (1992). A Simple Model of Herd Behavior. *Quarterly Journal of Economics*, *107*, 797-817.

Barto, A.G., Sutton, R.S., & Anderson, C.W. (1983). Neuronlike Adaptive Elements That Can Solve Difficult Learning Control Problems. *IEEE Transactions on Systems, Man, and Cybernetics*, *13, No. 5*, 834-846.

Bikhchandani, S., Hirshleifer, D., & Welch, I. (1992). A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades. *Journal of Political Economy*, *100*, 992-1026.

Bush, R.R., & Mosteller, F. (1955). *Stochastic Models for Learning*. New York: Wiley.

Casti, J.L. (1994). *Complexification: Explaining a Paradoxical World through the Science of Surprise*. New York: HarperCollins.

Cross, J.G. (1983). *A Theory of Adaptive Economic Behavior*. Cambridge: Cambridge University Press.

Debreu, G. (1959). *Theory of Value: An Axiomatic Analysis of Economic Equilibrium*. New York: Wiley.

Dosi, G., Ermoliev, Y., & Kaniovski, Y. (1994). Generalized Urn Schemes and Technological Dynamics. *Journal of Mathematical Economics*, *23*, 1-19.

Ellison, G., & Fudenberg, D. (1993). Rules of Thumb for Social Learning. *Journal of Political Economy*, *101*, 612-643.

Friedman, M. (1953). The Methodology of Positive Economics. In M. Friedman, *Essays in Positive Economics* (pp. 3-43). Chicago: University of Chicago Press.

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

Gode, D.K., & Sunder, S. (1993). Allocative Efficiency of Markets with Zero-Intelligence Traders: Markets as a Partial Substitute for Individual Rationality. *Journal of Political Economy*, *101*, 119-137.

---

[21] Notice that this implies that the following view of ACE research is definitely not correct. Reality leads to *facts* (which need to be explained), and a computer simulation produces *artifacts* (which, in turn, need to be explained). The alleged objective, then, would be to show that the explanation for the facts could be the same as the explanation for the artifacts, but achieving this is meaningless anyway because the facts produced by the computer program are inferior and subordinate anyway as they concern only *artifacts*. As we explained, the output of an ACE computer program are not artifacts to be explained. The computer program itself is the model that explains the social facts.

Hayek, F.A. (1948a). *Individualism and Economic Order*. Chicago: University of Chicago Press.

Hayek, F.A. (1948b). Individualism: True and False (reprint). In F.A. Hayek (Ed.), *Individualism and Economic Order* (pp. 1-32). Chicago: University of Chicago Press.

Hayek, F.A. (1948c). Economics and Knowledge (reprint). In F.A. Hayek (Ed.), *Individualism and Economic Order* (pp. 33-56). Chicago: University of Chicago Press.

Hayek, F.A. (1948d). The Facts of the Social Sciences (reprint). In F.A. Hayek (Ed.), *Individualism and Economic Order* (pp. 57-76). Chicago: University of Chicago Press.

Hayek, F.A. (1948e). The Use of Knowledge in Society (reprint). In F.A. Hayek (Ed.), *Individualism and Economic Order* (pp. 77-91). Chicago: University of Chicago Press.

Hayek, F.A. (1948f). The Meaning of Competition (reprint). In F.A. Hayek (Ed.), *Individualism and Economic Order* (pp. 92-106). Chicago: University of Chicago Press.

Hayek, F.A. (1948g). Socialist Calculation I: The Nature and History of the Problem (reprint). In F.A. Hayek (Ed.), *Individualism and Economic Order* (pp. 119-147). Chicago: University of Chicago Press.

Hayek, F.A. (1967a). Degrees of Explanation (reprint). In F.A. Hayek (Ed.), *Studies in Philosophy, Politics and Economics* (pp. 3-21). Chicago: University of Chicago Press.

Hayek, F.A. (1967b). The Theory of Complex Phenomena (reprint). In F.A. Hayek (Ed.), *Studies in Philosophy, Politics and Economics* (pp. 22-42). Chicago: University of Chicago Press.

Hayek, F.A. (1973). *Law, Legislation, and Liberty* (Vol. 1, Rules and Order). Chicago: University of Chicago Press.

Hayek, F.A. (1978). *New Studies in Philosophy, Politics, Economics, and the History of Ideas*. Chicago: University of Chicago Press.

Hayek, F.A. (1982). The Sensory Order after 25 Years. In W.B. Weimer & D.S. Palermo (Eds.), *Cognition and the Symbolic Processes, Volume 2* (pp. 287-293). Hillsdale, New Jersey: Lawrence Erlbaum.

Hill, B.M., Lane, D.A., & Sudderth, W.D. (1980). A Strong Law for Some Generalized Urn Processes. *Annals of Probability*, *8*, 214-226.

Holland, J.H. (1975). *Adaptation in Natural and Artificial Systems*. Ann Arbor: University of Michigan Press.

Holland, J.H. (1992). *Adaptation in Natural and Artificial Systems. An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence* (2nd ed.). Cambridge, MA: MIT Press.

Holland, J.H., Holyoak, K.J., Nisbett, R.E., & Thagard, P.R. (1986). *Induction: Processes of Inference, Learning, and Discovery*. Cambridge, MA: MIT Press.

Hull, C.L. (1943). *Principles of Behavior*. New York: Appleton-Century-Crofts.

Katz, M.L., & Shapiro, C. (1994). Systems Competition and Network Effects. *Journal of Economic Perspectives*, *8*, 93-116.

Katz, M., & Shapiro, C. (1985). Network Externalities, Competition and Compatibility. *American Economic Review*, *75*, 424-440.

Katz, M., & Shapiro, C. (1986). Technology Adoption in the Presence of Network Externalities. *Journal of Political Economy*, *94*, 822-841.

Kauffman, S.A. (1993). *The Origins of Order. Self-Organization and Selection in Evolution*. New York: Oxford University Press.

Kirman, A.P. (1994). *Economies with Interacting Agents* (Working Paper 94-05-030) Santa Fe Institute.

Kirman, A.P., & Vriend, N.J. (1995). *Evolving Market Structure: A Model of Price Dispersion and Loyalty* (mimeo) Paper presented at the Workshop Economic ALife, Santa Fe Institute.

Kirman, A.P., & Vriend, N.J. (in press). A Model of the Marseille Fish Market with Adaptive Agents. In D. Delli Gatti, M. Gallegati & A.P. Kirman (Eds.), *Market Structure, Aggregation & Heterogeneity*. Cambridge: Cambridge University Press.

Klamer, A. (1984). *The New Classical Macroeconomics. Conversations with New Classical Economists and their Opponents*. Brighton: Wheatsheaf.

Lane, D.A. (1992). *Artificial Worlds and Economics* (Working Paper 92-09-048) Santa Fe Institute.

Lane, D., & Vescovini, R. (1996). Decision Rules and Market Share: Aggregation in an Information Contagion Model. *Industrial and Corporate Change*, *5*, 127-146.

Langton, C.G. (1989). Introduction. In C.G. Langton (Ed.), *Artificial Life* (Vol. VI, Santa Fe Institute Studies in the Sciences of Complexity). Redwood City, CA: Addison-Wesley.

Lewin, R. (1993). *Complexity. Life at the Edge of Chaos*. London: Phoenix.

Littlechild, S.C. (1986). Three Types of Market Process. In R.N. Langlois (Ed.), *Economics as a Process*. Cambridge: Cambridge University Press.

Mas-Colell, A., Whinston, M.D., & Green, J.R. (1995). *Microeconomic Theory*. Oxford: Oxford University Press.

Minsky, M.L. (1967). *Computation: Finite and Infinite Machines*. Englewood Cliffs, NJ: Prentice-Hall.

Narduzzo, A., & Warglien, M. (1996). Learning from the Experience of Others: An Experiment on Information Contagion. *Industrial and Corporate Change*, *5*, 113-126.

Parker, D., & Stacey, R. (1994). *Chaos, Management and Economics*. London: Institute for Economic Affairs.

Roth, A.E., & Erev, I. (1995). Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term. *Games and Economic Behavior*, *8*, 164-212.

Savage, L.J. (1954). *The Foundations of Statistics*. New York: Wiley.

Sigmund, K. (1995). Darwin's "Circles of Complexity": Assembling Ecological Communities. *Complexity, 1, No. 1*, 40-44.

Simon, H.A. (1955). A Behavioral Model of Rational Choice. *Quarterly Journal of Economics*, *69*, 99-118.

Simon, H.A. (1957). *Administrative Behavior*. New York: Free Press.

Simon, H.A. (1959). Theories of Decision-Making in Economics and Behavioral Science. *American Economic Review*, *49*, 253-283.

Simon, H.A. (1976). From Substantive to Procedural Rationality. In S.J. Latsis (Ed.), *Method and Appraisal in Economics* (pp. 129-148). Cambridge: Cambridge University Press.

Streit, M.E., & Wegner, G. (1989). Information, Echange et Catallaxie. Réflexions sur les concept-clés de la théorie évolutionniste du marché. *Journal des Economistes et des Etudes Humaines, 1, No. 1*, 3-24.

Sutton, R.S. (1992). Introduction: The Challenge of Reinforcement Learning. *Machine Learning*, *8, Nos. 3/4*, 225-227.

Tesfatsion, L. (1998). *Agent-based Computational Economics* (http://www.iastate.edu/tesfatsi/ace.htm).

Vriend, N.J. (1994). A New Perspective on Decentralized Trade. *Economie Appliquée*, *47, No. 4*, 5-22.

Vriend, N.J. (1996). Rational Behavior and Economic Theory. *Journal of Economic Behavior and Organization, 29, No. 2*, 263-285.

Waldrop, M.M. (1992). *Complexity. The Emerging Science at the Edge of Order and Chaos*. London: Viking.

Weimer, W.B. (1982). Problems of Complex Phenomena: An Introduction to the Theoretical Psychology of The Sensory Order. In W.B. Weimer & D.S. Palermo (Eds.), *Cognition and the Symbolic Processes, Volume 2* (pp. 241-285). Hillsdale, New Jersey: Lawrence Erlbaum.