Working Paper

Speculation, Heterogeneity and Learning: A Model of Exchange Rate Dynamics

Luigi Marengo
Department of Economics, University Trento
Hélène Tordjman
Department of Economics, University of Paris XIII, and CREA, Ecole Polytechnique, France

WP-95-17
February 1995

IIASA International Institute for Applied Systems Analysis □ A-2361 Laxenburg □ Austria
Telephone: +43 2236 807 □ Fax: +43 2236 71313 □ E-Mail: info@iiasa.ac.at

Speculation, Heterogeneity and Learning: A Model of Exchange Rate Dynamics

Luigi Marengo
Department of Economics, University Trento
Hélène Tordjman
Department of Economics, University of Paris XIII, and CREA, Ecole Polytechnique, France

WP-95-17
February 1995

We thank the participants of the seminar "Evolutionary Approaches of Finance" held in Santa Fe in 1993, and especially Brian Arthur; the participants of the seminar "Confiance, apprentissage et anticipation économique" held in Compiegne in 1995; seminars at CREI (Paris XIII), CREA (Ecole Polytechnique), LEREP (Toulouse I), and André Orléan for valuable comments and suggestions on a previous draft. Support to the research by the G. Lurcy Foundation; the Center for Research Management, UC, Berkeley; IIASA, and the Italian Ministry of University and Research ("MURST" 40%) is gratefully acknowledged.

IIASA International Institute for Applied Systems Analysis □ A-2361 Laxenburg □ Austria
Telephone: +43 2236 807 □ Fax: +43 2236 71313 □ E-Mail: info@iiasa.ac.at

Preface

The research project on *Systems Analysis of Technological and Economic Dynamics* at IIASA is concerned with modeling technological and organisational change; the broader economic developments that are associated with technological change, both as cause and effect; the processes by which economic agents – first of all, business firms – acquire and develop the capabilities to generate, imitate and adopt technological and organisational innovations; and the aggregate dynamics – at the levels of single industries and whole economies – engendered by the interactions among agents which are heterogeneous in their innovative abilities, behavioural rules and expectations. The central purpose is to develop stronger theory and better modeling techniques. However, the basic philosophy is that such theoretical and modeling work is most fruitful when attention is paid to the known empirical details of the phenomena the work aims to address: therefore, a considerable effort is put into a better understanding of the 'stylized facts' concerning corporate organisation routines and strategy; industrial evolution and the 'demography' of firms; patterns of macroeconomic growth and trade.

From a modeling perspective, over the last decade considerable progress has been made on various techniques of dynamic modeling. Some of this work has employed ordinary differential and difference equations, and some of it stochastic equations. A number of efforts have taken advantage of the growing power of simulation techniques. Others have employed more traditional mathematics. As a result of this theoretical work, the toolkit for modeling technological and economic dynamics is significantly richer than it was a decade ago.

During the same period, there have been major advances in the empirical understanding. There are now many more detailed technological histories available. Much more is known about the similarities and differences of technical advance in different fields and industries and there is some understanding of the key variables that lie behind those differences. A number of studies have provided rich information about how industry structure co-evolves with technology. In addition to empirical work at the technology or sector level, the last decade has also seen a great deal of empirical research on productivity growth and measured technical advance at the level of whole economies. A considerable body of empirical research now exists on the facts that seem associated with different rates of productivity growth across the range of nations, with the dynamics of convergence and divergence in the levels and rates of growth of income, with the diverse national institutional arrangements in which technological change is embedded.

As a result of this recent empirical work, the questions that successful theory and useful modeling techniques ought to address now are much more clearly defined. The theoretical work has often been undertaken in appreciation of certain stylized facts that needed to be explained. The list of these 'facts' is indeed very long, ranging from the microeconomic evidence concerning for example dynamic increasing returns in learning activities or the persistence of particular sets of problem-solving routines within business firms; the industry-level evidence on entry, exit and size-distributions – approximately log-normal – all the way to the evidence regarding the time-series properties of major economic aggregates. However, the connection between the theoretical work and the empirical phenomena has so far not been very close. The philosophy of this project is that the chances of developing powerful new theory and useful new analytical techniques can be greatly enhanced by performing the work in an environment where scholars who understand the empirical phenomena provide questions and challenges for the theorists and their work.

In particular, the project is meant to pursue an 'evolutionary' interpretation of technological and economic dynamics modeling, first, the processes by which individual agents and organisations learn, search, adapt; second, the economic analogues of 'natural selection' by which interactive environments – often markets – winnow out a population whose members have different attributes and behavioural traits; and, third, the collective emergence of statistical patterns, regularities and higher-level structures as the aggregate outcomes of the two former processes.

Together with a group of researchers located permanently at IIASA, the project coordinates multiple research efforts undertaken in several institutions around the world, organises workshops and provides a venue of scientific discussion among scholars working on evolutionary modeling, computer simulation and non-linear dynamical systems.

The research focuses upon the following three major areas:

1. Learning Processes and Organisational Competence.
2. Technological and Industrial Dynamics
3. Innovation, Competition and Macrodynamics

1/Introduction

Speculative phenomena have for a long time fascinated a wide range of observers, from historians and psychologists to economists. However, their approaches are quite different. Historians and psychologists emphasize the recurrence of seemingly irrational individual and collective behaviours, speculative crisis being an example of collective manias and panics driven by rumours and other epidemics-like effects. On the contrary, most economists view speculation as the outcome of rational economic behaviours. In this latter view, rational agents engage in speculative trade because asymmetric information or differentiate risk aversion can lead them to think they can realize 'unexpected' profits. Both purposes are compatible with the Rational Expectation Hypothesis (REH). This approach is well-exemplified in e.g. Hirschleifer (1975), Figlewski (1978), and in rational bubbles models[1]. An alternative hypothesis to the same effect is that, if individuals were to be non rational and trade on the ground of "wrong" beliefs, they would be eliminated from the market via a process akin to natural selection (Friedman (1953)).

The theory of rational speculation has been challenged both by theoretical paradoxes and empirical puzzles. Concerning the latter, one should mention among others a) the persistence of predictable profit opportunities (positive autocorrelation of expected returns on the stock market, persistent bias in the forward discount on foreign exchange market, e.g. Bilson (1981)), b) statistical 'anomalies' of the price series, as leptokurtosis and volatility clustering (Friedman & Vandersteel (1982) and Baillie & MacMahon (1989)), and c) the micro-evidence of systematic biases in the way individuals form their expectations, thus leading to question the adequacy of the Rational Expectation Hypothesis. In particular, one observes contagion effects (Shiller (1989)) and threshold effects related to nominal values of exchange rates, (De Grauwe & Decupere (1992)) ; systematic biases in the formation of individual beliefs (Ito (1990) and Camerer (1987)) ; and heterogeneity of the interpretative models agents use to process information (Frankel & Froot (1987) and Froot & Ito (1989)). More generally, empirical studies of speculative episodes point to the central role of average opinion, of 'market psychology', in price dynamics. As shown by the adjectives usually employed to qualify its 'mood', e.g. "tense, feverish, depressed, optimistic...", the market is considered as an entity, endowed with a personality of its own (Arthur (1992)). Another general feature of the formation of beliefs on financial markets is the seemingly pervasiveness of tacit knowledge in the predictive rules used by agents. This tacit knowledge is often referred to by market operators as 'intuition', 'gut-feeling', 'common sense' (see the answers to

[1]There is a vast literature on rational bubbles; see the seminal contributions of Flood & Garber (1980) and Blanchard & Watson (1982), and for a survey, Rosser (1991).

Shiller's questionnaire (1989)) and other similar expressions falsely conveying the idea of a gift rather than that of a skill acquired through a learning process.

From a theoretical standpoint, "rational speculation" has to face even more serious problems.

First, as shown by recent research on speculative bubbles and sunspots equilibria, rational expectation models cannot generally rule out multiplicity of possible equilibria. In principle, convergence to any one of them may occur via two mechanisms: 'irrational' agents can be eliminated from the market by a selection process or they can learn to become 'rational' -whatever that means-, or combinations of the two. However, up to now no general result of convergence and stability has been achieved, leaving us with a basic indeterminacy concerning the aggregate outcome of out-of-equilibrium speculative dynamics (see for instance Bray (1982) and De Long *et al.* (1990))[1].

Second, if information is costly and equilibrium prices accurately reflects all available information, nobody will have an incentive to gather it because they can infer it from the observation of market prices. But then, as Grossman & Stiglitz (1980) are pointing out, what information will be revealed by prices if nobody is informed?

Third, if by whatever mechanism the economy converges to an equilibrium, there is no reason why trade should continue. This 'no trade' paradox is basically due to the assumption that individuals are homogeneous (which is implicit in rational expectations models). Being rational, they all share the same model of the world -which is the "true" model- and cannot form beliefs different enough to justify betting against each other. Hence, they have no incentive to trade and rational speculation is impossible (Tirole (1982) and Milgrom &Stokey (1982)).

This 'no trade' result addresses a fundamental question at the core of economic theory. Indeed, homogeneity of agents is crucial in that it allows (via the 'representative agent' technique) a simple aggregation of microbehaviours. If all agents behave in the same way, rationally pursuing their self-interest, the collective behaviour can be infered from the observation of the 'modal' individual behaviour. In this respect, the homogeneity assumption can be considered a kind of shortcut to surrogate A. Smith's Invisible Hand. Economic theory faces a dilemma: by maintaining a strong version of individual rationality as a foundation, it cannot explain satisfactorily why so much trade occurs on e.g. financial markets (as a reminder, the daily volume of transaction on the foreign exchange market alone is close to $ 1000 billions). Putting it differently, it appears to be

[1]For discussions of the various issues related to the mechanisms of convergence to a REE, see also Frydman & Phelps (Eds) (1983), and the other papers of the special issue of the *Journal of Economic Theory*, (1982).

impossible to keep together at the same time rationality -in the standard economists' sense-, speculation and equilibrium.

The purpose of this paper is to outline an alternative approach to speculative phenomena based on micro-heterogeneity and imperfect, adaptive, rationality. Speculation is seen as a disequilibrium dynamics generated by the interaction of less-than-rational heterogeneous agents that can be subject to contagious euphoria and panic, but that are also able to learn and modify their behaviour through time. The basic idea on which the representation of these learning processes is built is that of "mental models" developed, in different perspective, by Johnson-Laird (1983), Lakoff (1987), and Holland, Holyoak, Nisbett & Thagard (1986). Once we allow for agents to have different and evolving models of the world, the collective outcome of the market can only be studied by specifying the mechanisms by which agents interact. Learning processes and interaction mechanisms will be sketched out in section II. Section III presents a "pure speculation" model grounded on these hypotheses, and the simulation results are discussed in section IV. Section V highlights the major conclusions and the research task ahead.

II Speculation in disequilibrium

Pure speculation: the Beauty Contest metaphor

Speculative behaviour is generally defined as a transaction with the purpose to realize a capital gain (e.g. buying low today in order to sell high tomorrow, see Keynes (1936) and Kaldor (1939)). Forming an expectation about future prices is thus the main activity of a speculator.

If speculators are rational and know the fundamental value of the asset or the currency they are trading, they will buy when the price is inferior to its fundamental value and sell otherwise. By doing so, they contribute to push the price back to that fundamental value: speculation cannot be destabilizing (Friedman (1953)). However, this argument is valid only if all agents are rational, have the same model of equilibrium prices and share the same beliefs about the automatic return to these equilibrium values (and this is common knowledge). But if individuals don't share the same representations about equilibrium prices, no such stabilizing mechanism is likely to exist. It is one of the basic conjectures of this work that a speculative market where agents are truly heterogeneous can go through different regimes and experience very rich dynamics.

In a situation of 'true' uncertainty, in the Keynesian or Austrian sense, i.e. in a non transparent and constantly evolving environment, individuals cannot know the "true" model of the world[1]. They form hypotheses on its functioning, try, make mistake, learn and adapt. Their strategies can sometimes be locally optimal, but, because of the non stationarity of the environment (if only because of learning), this harmony cannot last infinitely. With such experience-dependent processes of beliefs formation, individuals will not in general share the same representation of the world. They are heterogeneous, and that is their principal incentive to transact: each agent believes that his model is "better" than those of the others, and that there is an opportunity of profit even in the absence of new information. Thus, a speculative market is ultimately a market where individual models of the world confront each others, i.e. where the price is determined by the interaction of agents trying to guess the average opinion in order to beat the market.

Such dynamics of cross-expectations has been highlighted by Keynes (1936), who was comparing in his famous parable the functioning of financial markets to a beauty contest where the participants had to choose the nicest face among some photographs of faces, the winner being the one whose choice was closest to the average choice. As the variable agents have to predict is the result of their collective activity, the dynamics of this type of market is highly self-referential. If agents are imperfectly rational and heterogeneous, they can only try to *imagine* what the others are thinking, and the collective dynamics results from these cross expectations. As just said, the main conjecture of this work is that such a formation of beliefs can lead to the emergence of unexpected aggregate outcomes exhibiting complicated dynamics, even without shocks on the fundamental variables.

In order to explore this conjecture, we built a model of 'pure speculation', without fundamentals, where the price results from the interaction of less-than-rational and heterogeneous agents. Because agents are persistently and unpredictably different in what they think and what they do, no simple procedure of aggregation exists, and the collective outcome of the system can be known only by specifying the mechanisms by which they interact. Studying this type of systems thus involves the development of appropriate theoretical tools, accounting explicitly for individual learning processes. These systems cannot generally be solved analytically and have to be studied by simulation. A methodology of simulation referred to as 'artificial life' has been recently developed, mainly in biology and cognitive and computer sciences, which allows to study the behaviour of such 'complex' systems[2]. An artificial world is, to paraphrase Lane (1992), a computer implementable system composed of a) micro-units interacting in an environment and b) an aggregate dynamics *emerging* from these interactions. The micro-

[1] See for instance Dosi & Egidi (1991) and Arthur (1992).
[2]See Langton et al (1991) and, for applications of this approach in economics, Lane (1992).

units are generally evolving through time via processes of learning and adaptation to their environment, thus permanently introducing novelty in the system: the dynamics of such a 'complex adaptive system' is open-ended (in the sense that unequivocal asymptotic outcomes can hardly be predicted).

Following this methodology, we built a model of an "artificial speculative market", where price dynamics can be analyzed under alternative hypotheses concerning the micro-units and the market mechanisms by which they interact[1]. Let us present our main hypotheses.

Category formation and mental models

If we abandon the hypothesis that agents know the 'true' model of the world, we have to account for the mechanisms by which individual representations emerge and evolve. Following recent developments in cognitive sciences, we start from the assumption that individuals are building "mental models of domains" (Johnson-Laird (1991), p. 2), i.e. sets of representations of the world linked together, whose "structure corresponds to the way in which human beings conceive the structure of the world" (*ibidem*). A mental model is a hierarchical structure (a system) of concepts and categories that can be "manipulated to produce expectations about the environment" (Holland *et al.* (1986), p. 12). The basic units constituting the model, namely, the categories, are not alone conveying meaning; as relations between these 'units' are reproducing in some ways the perceived structure of the world, the specific organization between categories is itself meaningful. Hence, this organization depends upon the content, the semantics of the categories and not, as it is the case in probabilistic approaches of induction, of formal syntactic rules[2]. Such a "semantic point of view" has fundamental implications about properties of knowledge accumulation, as it is able to explain why people make only a small part of all the possible inferences that can be imagined. Indeed, as the structure of the model is meaningful in itself, new information that is consistent with it will be more easily integrated in the model, whereas dissonant information will tend to be ignored[3]. Thus, inductive inferences are partly constrained by existing knowledge.

Two broad approaches of category formation are usually distinguished (Lakoff (1987)). The 'objective' approach defines categories by the set of properties they have in common: one has to be able to list exhaustively all the characteristics of an object in order to form a

[1]The first work of this type was done by Arthur, Holland, LeBaron, Palmer & Taylor at the Santa Fe Institute; Beltratti & Margarita (1992) have built another of such "artificial financial market".

[2]On approaches of induction in terms of mental models, see also Johnson-Laird (1983) and Lakoff (1987). On these points, see also Keynes (1921), Bell, Raiffa & Tversky (Eds.) (1988), and Gigerenzer *et al.* (1989).

[3]This last point is consistent with empirical evidence about the general phenomenon of cognitive dissonance, e.g. Aaronson (1972).

representation of it. Symmetrically, if one of these characteristics has changed, one is unable to recognize that it is the same object. Following the path-breaking work of E. Rosch, another approach has been developed, grounded on the recognition that even though the world is varied, uncertain, and constantly evolving, individuals are still able to form representations of it, understanding diversity and adapting to changes. This approach suggests that a category is formed by a 'prototype' assorted of defaults, i.e. of possible exceptions that one can accept without contradicting the category. Such categories can be built even if the world is imperfectly understood and changing all the time: by modifying and refining this network of defaults, one can adapt to diversity and novelty.

In this view, a mental model is organized in a hierarchy of defaults, i.e. in sets of categories of different levels (general/specific) linked together by these defaults[1]. The mechanisms by which these models are formed are mainly inductive ones. "Induction tries to find regularity and coherence behind the observations. Its most conspicuous instruments are generalization, specialization, analogy. Tentative generalization starts from an effort to understand the observed facts; it is based on analogy, and tested by further special cases" (Polya (1945), p. 117). When a category is contradicted by an observation, a default will be added to take into account this new information. This is an operation of specialization. If other informations repeatedly contradict the category, at some point, the individual may want to form a more general category that will not be falsified by observations[2]. Finally, people may use a model or part of a model developed in one domain to build a representation of another domain that is perceived as similar to the first one. This last operation can be seen as an analogy.

Another important issue concerns the place of intuition (or 'gut-feeling', as some traders would put it, e.g. Shiller (1989)). This issue is of course never addressed by rational choice theory since there is no place for it in, say, Savage's axiomatisation of individual choice. March & Simon (1993) distinguish between two complementary types of logic of action, a deliberative one and a tacit one. The deliberative, or explicit one (the "logic of consequences" in March & Simon's terms) consists in manipulating different models and comparing their respective predictions, i.e. what philosophers call *gedankenexperiment*. The tacit mode (the "logic of appropriateness") relates to pattern-recognition and ruled-based behaviour, and is a "skill in recognizing those things that have become familiar through past experience" (March & Simon (1993), p. 16). In terms of mental models,

[1]For more details on default hierarchies, see Johnson-Laird (1983) and Holland *et al.* (1986).
[2]This can give rise to some kind of "threshold effects". When exactly will thresholds trigger changes is another complicated matter. Studies of cognitive dissonance mentioned above show us that individuals seem to have very different levels of resistance to contradictory facts.

"things that have become familiar through past experience" are embedded in the specific organization of the categories, i.e. in the structure of a mental model, and intuition results from an implicit inferential process, done 'inside' one model (instead of comparing the outcomes of different models).

Finally, it should be emphasized that mental models can be of different type, depending upon the purpose for which they are developed. Some are descriptive, others include causality relationships and/or normative aspects, as needed by the agent with respect to the task he has to perform, the goal he wants to achieve or the problem he tries to solve. Knowledge accumulation is thus contingent to the activity of human beings, oriented by their roles and functions in the environment (or society at large): it is "embodied" (Lakoff (1987)) or "embedded" (Granovetter (1985)) in broader social and cultural forms[1].

Learning dynamics

A model can be seen as an hypothesis about the world, leading to predictions. In case of 'false' predictions, the hierarchical structure of the model allows for the identification of the part of the model responsible for the error. The impression of a gap between the world and the model one has of it becomes more acute. Depending of the amplitude of the gap, i.e. of the place of the defaults in the hierarchy, the model can either be modified or replaced by another one. If the falsified defaults are at the top of the hierarchy, the model will have to be abandoned, and some kind of exploration process will start. As noted before, induction is constrained by existing knowledge and contingent to the specific type of activity. A new hypothesis can be formed by recombining parts of the old model, and learning can be seen as a "procedure that discovers a way in which to combine old functions so as to create new ones" (Johnson-Laird (1983), p. 143). Often, repeated failures of the old model will have highlighted some of the 'weaknesses' of the previous representation, and alternative hypotheses may have started to be formed. It is also possible to think of alternative hypotheses permanently coexisting in one's mind, especially when confronted to a rapidly changing environment

Such a distinction between a kind of 'linear' accumulation of knowledge inside one model and the discontinuities of knowledge's evolution related to changes of models is close to the one made by Kuhn between 'normal' and 'revolutionary' science, as well as to March's 'exploitation' and 'exploration' dichotomy (1991). Indeed, one of the difficulties in studying learning mechanisms is to represent this tension between the exploitation of existing knowledge and the exploration of novel directions. As Holland

---

[1]Thus, induction is not only constrained by existing knowledge, but also oriented by individual's roles and objectives, thus accounting for the fact that, while making inferences, we do not explore all the 'space of possible worlds'.

(1986) puts it: "because of the uncertainty attached to all induction, the process has to be conducted in such a way that the system can absorb new rules, tentatives, without destroying abilities developed in familiar situations" (p. 594).

Learning dynamics results from a double interaction. First, categories interact to generate 'higher level' models, organized so as to represent the perceived structure of the world. Second, individuals interact with their environment, constantly receiving and processing new information. These pieces of information may reinforce some categories and contradict others, provoking a change in representations, i.e. a modification of the models' structure. Hence, such learning mechanisms imply a constant transformation of individual representations, a continuous process of adaptation and local adjustment to non-stationary environments.

The learning dynamics generated by these mechanisms has the following characteristics:
i) Learning is adaptive, because its 'directions' depend upon existing knowledge, upon path-dependent perceptions of changes in the environment and upon individuals goals.
ii) Changes in the environment are perceived and interpreted through already existing models, and in turn contribute to modify the models themselves. A model which is often and successfully used will tend to be reinforced because of its familiar usefulness, and contradictory facts that do not falsify its 'core' representations will tend to be ignored. Thus, this is an aspect of the interaction between the agent and its environment which might entail an implicit positive feedback leading to phenomenon of 'cognitive lock-in'[1].
iii) Finally, when categories at the highest level of the hierarchy are falsified, the model has to be abandoned and replaced while a subset of knowledge will be relegated to the background of the memory. This corresponds to a discontinuity in the accumulation of knowledge.

Such an approach to learning processes is quite far from the usual representation of learning in economics, most often formalized through bayesian probability updating procedures. Nonetheless, we think it can lead to an interesting representation of microbehaviours on speculative markets. First, as 'history matters', agents are supposed to have heterogeneous models of the world and are thus ready to take bets against each other. Second, because of its foundations upon prototypical categories and default hierarchies, this kind of approach can give an account of learning even in uncertain and non-stationary environments. Third, diverse mental models easily allow for both routine-type and deliberative actions. Fourth, the persistent tension and arbitrage between exploitation of existing knowledge and exploration of new domains is able to account for

[1] As mentioned, this type of dynamics is highly path-dependant, and to some extent irreversible (e.g. Arthur *et al.* (1987), David (1988), Arthur (1989) and (1992).

discontinuities, threshold effects and other 'peculiarities' in learning dynamics well documented in empirical studies. This basic approach shall inspire the modelling of the behaviour of our 'artificial agents'.
Before turning to market interaction, let us add a last remark. We do not claim of course that human cognition is strictly isomorphic to the representation of mental models that follows. But we do indeed claim that models of adaptation and search such as the one presented here help in understanding some basic stylized facts of microbehaviours (at least with regards to financial markets) and might even capture some generic properties of the dynamics of representations of agents facing uncertain environments.

The interaction of heterogeneous agents
When agents on a market are heterogeneous, price formation can be understood only by specifying the way they interact. "Knowing the norms, preferences, motives and beliefs of participants to collective behaviour can, in most cases, give necessary but not sufficient conditions in explaining the aggregate result; one needs to add a model giving an account of the interaction and of the aggregation of these preferences" (Granovetter (1978), p. 1421). To study the formation of a speculative price thus requires to specify the mechanisms by which agents communicate and trade, as well as the way their representations are modified through these interactions. Indeed, price is determined by the collective behaviour of the market. When the environment changes, individuals must modify their models of the world and their decision rules, i.e. must adapt. We have here a circular process, from the environment to individual behaviours, from the latter to collective behaviour, and back to the environment.

As Lesourne (1991) suggests, a distinction can be made between two functions of a market, namely an organizing and a creative one. The first concerns all institutional arrangements which are governing exchanges of goods and information (the way agents meet and trade, the bargaining rules leading to the effective price at which goods will be traded etc...). The institutional characteristics governing these interactions (what is exchanged? following what kind of procedures?) influence the mode of price formation. But "interaction between agents on a market can also endogenously generate a whole array of institutions" (Lesourne, *ibidem*, p. 20). It is the market's creative function. On a speculative market, interaction does not really create institutions, but dominant representations, norms of judgement, conventions. Two processes play a fundamental role in this respect, i.e. selection and learning. Selection ("what kind of behaviour brings positive payoffs?") determines the composition of the market, i.e. the type of strategies and models which are dominating the market at a certain time, and the price formation[1].

[1]The links selection and interaction entertain with each other are still to be explored. This is one of the numerous point a satisfactory evolutionist theory should address, but it will not be done here.

Learning, by modifying representations, expectations and decision rules, introduces a dynamics on the *nature* of the behaviours which are to become prevalent in the market at a certain time. It is creating the variety on which selection operates. As, contrary to most models of the evolutionary games type, models and representations are not given but emerge endogenously, price dynamics is truly evolutionary in the sense that it is an open-ended dynamics (e.g. Silverberg (1988)).

The market is thus a locus where different processes are combined:
* Interaction between heterogeneous agents who are meeting and exchanging goods and informations following procedures determined by institutional rules and habits and behavioural norms current on this specific market[1].
* Individual learning and adaptation to changes in the environment, and especially to changes in the behaviours of the others: the evolution of the representations of each agent depends upon what he perceives and understands of the representations of the others. Hence, agents do not only interact "materially" through their actions, but also "mentally", in their specific foresight procedures .
* Selection, determining at each moment which action is the most performing one. On a speculative market, this process entails a dynamics, as the performance of a specific behaviour depends upon guesses on the distribution of behaviours and in turn affects the latter. This positive feedback can generate cumulative, path-dependent processes[2].

Indeed, "the reference by which agents are determining their behaviour is not a norm exterior to the process analyzed; it is produced by the process itself, by the average opinion" (Orléan (1988), p. 15). Most often, average opinion is emerging from the multitude of bilateral encounters happening during transactions, and from the network of reciprocal influences linking individual representations. The latter are not stable, but evolve, thus modifying average opinion and the definition of what a "good"behaviour is. In fact, actions and representations co-evolve, the market selecting at each moment a particular correspondence between models of the world and performances, i.e. the "best" decision rule in a particular configuration of the market. We suggest that such a self-referential functioning of the market may either generate a) cumulative, path-dependent dynamics, whereby arbitrary rules impose themselves because everybody believes they are the "best" ones, or b) relatively stable situations, similar to what Keynes named 'conventions', when agents have heterogeneous rules but implicitly agree on the boundaries of the price dynamics, or finally c) different speculation processes amongst

[1]These institutional rules will be detailed in the next section. We will not enter into the discussion of the emergence of these rules, and the link it has with habits and norms, albeit this question is certainly central to the understanding of the coexistence of different forms of market.
[2]In addition to the authors already mentioned above, see, for such models on financial markets, Orléan (1990) and Kirman (1991)).

groups of agents characterized by diverse 'visions of the world' and yielding unpredictable collective dynamics.

III A model of pure speculation with adaptive agents

In this section we present a model of an artificial speculation market where a population of artificial agents, modelled as classifiers systems, trade a given asset. For simplicity we can assume that we deal with an exchange rate market where only two currencies exist, A and B. We will assume currency A as the *numéraire* of the market. Trade takes place at discrete time (t = 0,1,2,....) and at each moment in time each agent has three possible actions at his disposal:

- buy one unit of currency A, paying the corresponding amount of B quoted by the market;
- sell one unit of currency A, receiving the corresponding amount of currency B quoted by the market;
- hold the present position without engaging in any trade.

We assume that trade is centralized: once all agents have posted their intended actions, an auctioneer will compute the new market price at which all possible transactions will take place. Thus agents engage in trade before knowing the price at which transactions will take place, price being in fact the outcome of the decisions taken by the entire population, and in particular of the relation between number of buyers and number of sellers.

In what follows we first present a simplified version of classifiers systems by means of which we model our artificial agents and then we describe the institutional mechanisms which regulate the transactions in our artificial market. In the next section we will present some of the most significant results from simulations of this model of artificial market.

Artificial agents.

We claim that classifiers systems [1] provide a valuable model of learning artificial agents who are primarily engaged in adaptively revising the model of the world through which they formulate expectations on the future evolution of the market.

Classifiers systems are highly general learning systems which process a set of condition-action rules in order to achieve high adaptation to complex and largely unknown environmental conditions. The very low requirement of *a priori* knowledge, the high generality and simplicity of the methodology, combined with the complexity of the

---

[1] A presentation of the Classifier Systems methodology and its main applications can be found in the works by John Holland (see especially Holland (1975) and (1986), Holland *et al.* (1986)); a discussion of some possible applications to economics can be found in Arthur (1991).

patterns of behaviour they can produce, make them very attractive for applications in behavioural and social sciences.

Classifiers systems model an artificial learning agent as a set of condition-action rules which are processed in a typically evolutionary fashion, as they are subject to a process of selection and a process of mutation.

The first element which characterizes our classifier system is the message (signal) agents receive from the environment. Such a message - which contains some information about the recent history of the market - is freely available to every agent, but has to be interpreted and connected to a consequent action according to a model of the world which differs across agents and is always subject to possible revisions. In particular, we suppose that at each time t agents can observe the exchange rate at time t-1 (which we will indicate by p_{t-1}), the moving average of the exchange rate in the last k periods and the ratio between the number of buyers and the number of sellers (as a measure of the degree of "optimism" of the market). These three variables - referred to time t-1 - are encoded as binary strings of given length:

$$m_{11}\, m_{12}.....m_{1n} \mid m_{21}\, m_{22}.....m_{2n} \mid m_{31}\, m_{32}.....m_{3n} \quad \text{with } m_{ij} \in \{0.1\}$$

Each agent is modelled as a set of condition-action rules which are processed in a parallel fashion. Each rule makes a particular action conditional upon the fulfilment of a condition concerning the present state of the world (which in our case is represented by the input message containing the value of the three variables). The condition part is therefore actually made up of three strings (one for each variable) of symbols which encode a subset of the states of nature and is activated when the last detected state of the world falls into such a subset. Thus the condition part is composed by three strings of n symbols (as many as the bits of each component of the environmental message) over the alphabet {0,1,#}:

$$c_{11}c_{12}.....c_{1n} \mid c_{21}c_{22}.....c_{2n} \mid c_{31}c_{32}.....c_{3n} \qquad \text{with } c_{ij} \in \{0.1,\#\}$$

The condition is satisfied when, either $c_{ij} = m_{ij}$ or $c_{ij} = \#$; i.e. the symbol # acts as a "don't care" symbol which does not pose any constraint on the corresponding bit of the environmental message.

It can be easily shown that this way of codifying conditions amounts to defining sets of intervals on the axis of the corresponding variable. Such intervals can be interpreted as categories or information cells which contain all the states of the world which are indistinguishable to the agent. A set of condition defines therefore a model of the world, i.e. a subset of the power set of the set of states of the world; only as a special case may this subset be a partition of the set of states of the world, as required by Bayesian learning models.

To each condition corresponds an action, which is simply a ternary bit which encodes the three possible actions (buy, sell or hold):

$$a_l \qquad \text{with } a_l \in \{B,S,H\}$$

All in all, each agent in our artificial market is represented by a set of such condition-action rules:

$$R = \{R_1, R_2, \ldots\ldots R_q\}$$

where:

$$R_i : c_{11}c_{12}\ldots\ldots c_{1n} \mid c_{21}c_{22}\ldots\ldots c_{2n} \mid c_{31}c_{32}\ldots\ldots c_{3n} \dashrightarrow a_l$$

In addition each rule is assigned a "strength" and a "specificity" (or its reciprocal "generality") measure. <u>Strength</u> basically measures the past usefulness of the rule, that is the payoffs cumulated every time the rule has been applied; <u>specificity</u> measures the strictness of the condition: in our case the highest specificity (or lowest generality) value is given to a rule whose condition does not have any "#" symbol and therefore is satisfied only by one particular value of the input variables, whereas the lowest specificity (or the highest generality) is given to a rule whose condition is entirely formed by "#" symbols and is therefore always satisfied by the occurrence of any state of the world.

At the beginning of each simulation agents are supposed to be absolutely ignorant about the characteristics of the environment, as they are endowed with a set of randomly generated rules. Decision makers are also assumed to have limited computational capabilities, therefore the number of rules which model each of them is kept constant over time and is relatively "small" in comparison to the complexity of the problem which is being tackled.

This set of rules is processed in the following steps throughout the simulation process:

1) <u>Condition matching</u>: a message is received from the environment which informs the system about the last state of the world. Such a message is compared with the condition of all the rules and the rules which are matched, i.e. those which apply to such a state of the world, enter the following step.

2) Competition among matched rules: all the rules whose condition is satisfied compete in order to designate the one which is allowed to execute its action. To enter this competition each rule makes a bid based on its strength and on its specificity. In other words, the bid of each matched rule is proportional to its past usefulness (strength) and its relevance to the present situation (specificity):

$$\text{Bid } (R_{i,t}) = k_1 (k_2 + k_3 \text{ Specificity } (R_i)) \text{ Strength } (R_{i,t})$$

Where k_1, k_2 and k_3 are constant coefficients.

The winning rule is chosen randomly, with probabilities proportional to such bids.
3) Action and strength updating: the winning rule executes the action indicated by its action part and has its own strength reduced by the amount of the bid and increased by the payoff that the action receives, given the occurrence of the "real" state of the world. If the j-th rule is the winner of the competition, we have:

$$\text{Strength } (R_{j,t+1}) = \text{Strength } (R_{j,t}) + \text{Payoff } (t) - \text{Bid } (R_{j,t})$$

4) Generation of new rules: the system must be able not only to select the most successful rules, but also to discover new ones. This is ensured by applying "genetic operators" which, by recombining and mutating elements of the already existing and most successful rules, introduce new ones which could improve the performance of the system. In this way new rules are constantly injected into the system and scope for new search is always made available.

Genetic operators generate new rules which explore other possibilities in the proximity (in a sense which we are going to define precisely) of the presently most successful ones, in order to discover the elements which determine their success and exploit them more thoroughly: the search is not completely random but influenced by the system's past history. New rules so generated substitute the weakest ones, so that the total number of rules is kept constant.
Two genetic operators have been used for the condition part and one for the action part. The latter is simply a mutation of the existing action. It is applied with a given (small) probability and implies that the action included in the newly generated rule is randomly chosen between the two actions different from the one appearing in the parent rule.

The two operators used for the condition part deserve more attention because of their role in modelling the evolution of the state of knowledge embedded into the system. They operate in opposite directions:

a) Specification: a new condition is created which increases the specificity of the parent one: wherever the parent condition presents a "#", this is mutated into a "0" or a "1" (randomly chosen) with a given (small) probability.
b) Generalisation: the new condition decreases the specificity of the parent one: wherever the latter presents a "0" or a "1", this is mutated into a "#" with a given (small) probability.

Specification and generalisation are two possible cognitive attitudes which tend to drive the learning system towards, respectively, specific rules which apply to narrower intervals of values of the variables and more robust rules which instead cover a wider set of states of the world. Different degrees of specification and generalisation can be simulated both by means of different combinations of these two genetic operators and by varying the coefficient k_3 with which specificity enters the bid equation: the higher this coefficient, the more highly specific rules will be likely to be selected against general ones. The simulations discussed in the rest of the paper will use a specificity coefficient to summarize the overall inclination of the system toward the search for specific rules, such coefficient will represent both the value k_3 in the bid equation and the probability of application of the genetic operator "specification" every time the genetic operators routine is called.

The Market

In the simulations which we present in the next section we model a simple artificial market, characterized by the absence of transaction costs and wealth effects and by a one-period-ahead expectation structure. We also suppose that the market is created out of nothing at time t=0, with a fictitious starting price p_0. The market is populated by a relatively large number of agents (N=100) modelled as classifiers systems and who, at each moment in time, take one out of the three possible decisions and communicate it to the auctioneer. The auctioneer can thus summarize the state of the market by means of the numbers N_B, N_S and N_H which indicate respectively the number of buyers, the number of sellers and the number of holders (by definition $N = N_B + N_S + N_H$). The auctioneer can now allow a number N_T of transactions to take place:

$$N_T = \min (N_B, N_S)$$

at a price which is set according to the following rule which makes it vary proportionally to the disequilibrium between the number of buyers and the number of sellers:

$$p_t = p_{t-1} + [(N_B - N_S)/N] \, p_{t-1}$$

If $N_B \neq N_S$ and some rationing is necessary, the individuals who are unable to perform the desired transaction are randomly chosen.

Once all the transactions have been carried out at the new market price, each agent receives a payoff according to his decision and to the price variation. The payoff structure is reported in the following matrix:

Transaction:	A : $\Delta p_t > 0$	B : $\Delta p_t < 0$
Sell	$\pi_{s,A} \cdot \lvert\Delta p_t\rvert$	$- \pi_{s,B} \cdot \lvert\Delta p_t\rvert$
Buy	$-\pi_{b,A} \cdot \lvert\Delta p_t\rvert$	$\pi_{b,B} \cdot \lvert\Delta p_t\rvert$
Hold	$- \lvert\Delta p_t\rvert$ $c > 0$ when $\Delta p_t = 0$	$- \lvert\Delta p_t\rvert$
Rationed	$- B_t \cdot \lvert\Delta p_t\rvert$	$-(1/B_t) \cdot \lvert\Delta p_t\rvert$

where π_s and π_b are constant parameters and $B_t = N_B/N_S$

This payoff structure entails that agents are rewarded when they buy if the price has decreased and sell if it has increased, i.e. our agents behave like Friedmanian agents. The other alternative would have been to reward them when they buy if the price increase and sell otherwise, expecting the increase to continue. Such "positive feedback trading"is believed to be one of the source of market instability (see e.g. De Long et al. (1990), but it requires an expectation structure of at least two periods to be modeled. With a one-period ahead expectation structure, taking into account positive feedback strategies would be like 'forcing' bubbles into the model; hence the Friedmanian agents. It should nevertheless be noted that:

-One period-ahead strategies is sufficient to give an account of the central role of average opinion. Agents who make profits are those who sell when the expect the majority to buy and vice-versa.

-Fitness is endogenous, as the definition of what a 'good' behaviour is depends upon the distribution of behaviours across the population.

-Finally, if speculation is destabilizing when agents are Friedmanian, it is reasonable to suppose that it would be even more so if agents were positive feedback traders.

IV Simulation results

The results of the simulations we present here are very preliminary but some interesting features do nonetheless emerge. First, the market is more or less all the time at a "quasi-temporary equilibrium" in the sense that it almost clears, i.e. the number of rationed agents is very small. Second, the price dynamics generated by this simple model are quite rich, exhibiting periods characterized by price stability and others by bubble-like phenomena. Third, the features of the learning process of the individual agents, and particularly the "exploitation/exploration" dimension, have a significant effect on the price dynamics.

Let us first recall the main parameters of the model and give the parameters' values corresponding to the reference situation.

-Profit parameters:

$\pi_{s,A} = \pi_{b,B} = 10$

$\pi_{s,B} = \pi_{b,A} = -2$

-Learning parameters:

$k_3 = 0.5$

probability of specification: p(s) = 0.5

probability of generaliztion: p(g) = 0.3

Market clearing

The organization of the market with centralized information and centralized transactions allows it to reach a quasi market-clearing price almost all the time, even though quantities are fixed to one unit per transaction. Indeed, fixed quantities introduces a bias as, for the market to clear, the number of buyers must be exactly equal to the number of sellers, and no quantity adjustment is allowed. There is always a small percentage of rationed agents (between 5% and 10%), but surprisingly low for such a model where agents do not have any common equilibrium model on which to coordinate. Figure (1) shows the volume of transaction during two sub-periods of simulation (1) (which is presented below).The upper part of the chart corresponds to the volume of the 200 first iterations, where the price dynamics is quite turbulent. At the beginning a lot of agents are choosing not to transact, they wait; but slowly, as they learn their environment, bids and asks increase and symmetrically the number of agents who are not trading drops[1]. Nonetheless, during the entire sub-period, bids and asks are remarkably close and rationed agents represent only around 5% of the population. The lower part of the chart shows the transaction volume corresponding to the last 200 iterations of the same simulation, where the price is

[1]This feature is consistent with some 'no-trade' results where uncertainty is the main factor explaining why agents don't transact, see for instance Bossaerts (1992).

much more stable. Here, transaction volume is higher (80% of the agents are transacting) and more stable (the mean of the number of transaction is around 40 all the time). Again, rationed agents are few (less than 10%) and bids and asks are close to each other (but for very short-term discrepancies).

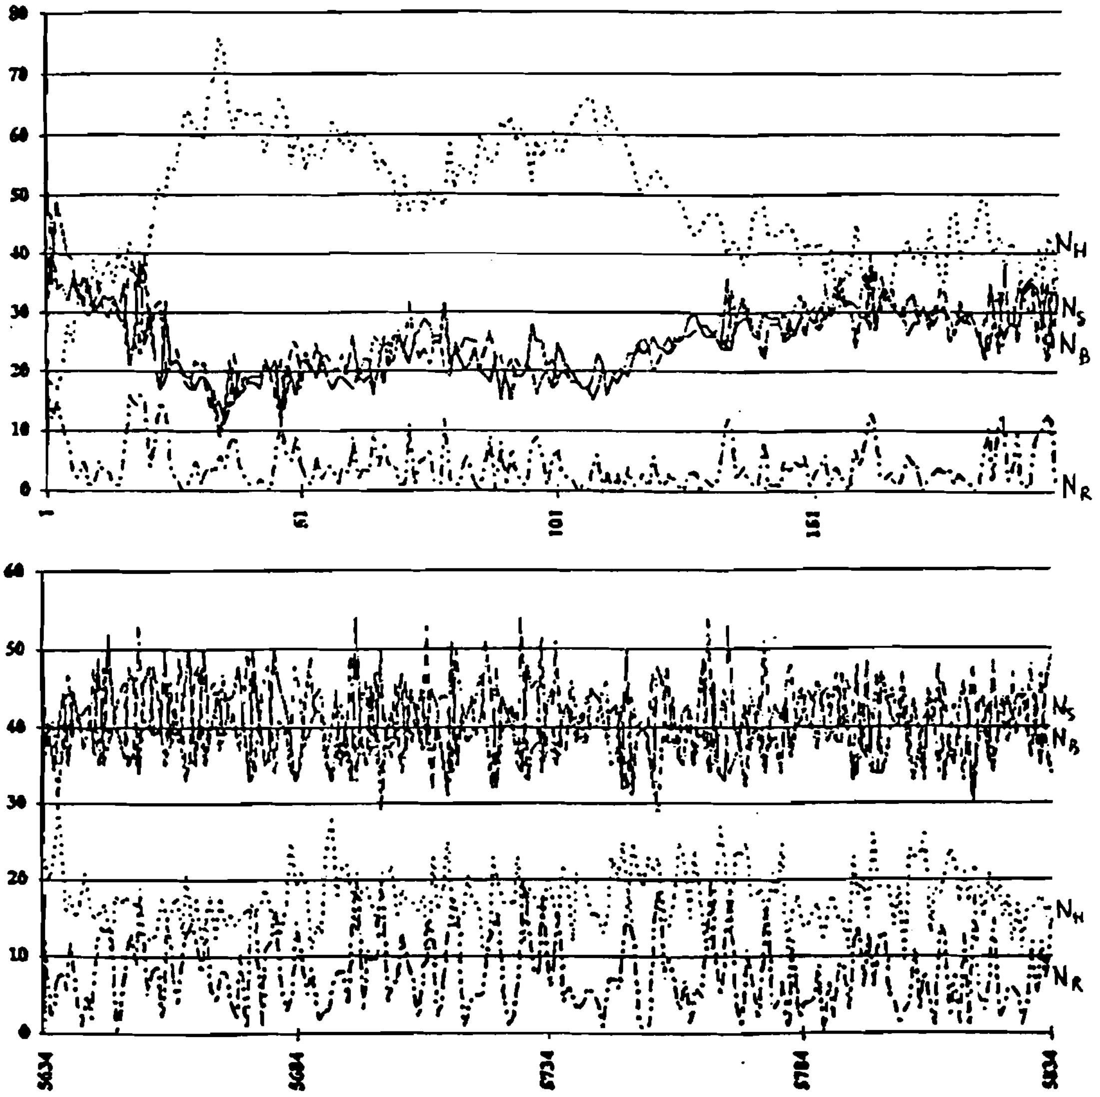

N_B : --- N_S : - - - N_H : ··· $N_R = |N_B - N_S|$:-··-··

Volume of transaction

Figure (1)

The market-clearing properties of this kind of market organization (quite similar to a double auction) are now well documented[1], but the reasons for this efficiency are still not well-understood. Why and how agents coordinate their behaviour in such a way that bids and asks are more or less always equal? It is particularly surprising here where, contrary to experimental studies where agents are endowed with the same equilibrium model, our artificial agents are truly heterogeneous and do not have a common model of the world.

The reasons for this "spontaneous" coordination are twofold. First, the structure of the payoff matrix entails some kind of "Friedmanian" reinforcement: agents who are bullish when the market is bearish get a positive payoff, whereas agents on the majority side of the market are penalized. As we already said, such a payoff structure rules out more complicated strategies like positive feedback trading. Second, our artificial agents seem to be learning about their environment. If they were not, the payoff structure would not condition their actions.

It should be noted that, if this "regressive" payoff structure allows the market to reach a "quasi-temporary equilibrium", it is not enough to warrant price stability, contrary to Friedman's argument. Indeed, all the simulations show that price dynamics is subject to turbulence and bubble-like events, even if the market almost clears. In other words, a regressive payoff structure does not warrant stability if agents are heterogeneous and do not have all the same idea of an equilibrium price.

Different market regimes

Two simulations of the reference model are shown in figure (2) and (3). One feature of these simulations is the sequence of bubble-like events and periods of relative stability, which tend to confirm the conjecture that speculative price dynamics can be characterized by two "polar opposite regimes", one regime of turbulence and one regime of stability, both of them generated by the same market mechanisms. Intuitively, this is consistent with recent results concerning the behaviour of "complex adaptive systems": order and fluctuations are two aspects of the same dynamical process (see for instance the work of Bak & Chen on self-organized criticalities).

These simulations also show that price stability occurs for different periods of time (from 800 to 3000 iterations) and different levels of prices. Hence, there is not any convergence, but only transient stability away from a stationary state. Following Bak & Chen (1991) and Lane (1992), we will qualify this transient stability as *metastability*.

---

[1]See for instance the experimental work of Smith *et al.* (1988), and for a survey, Hey (1991).

Figure (2)

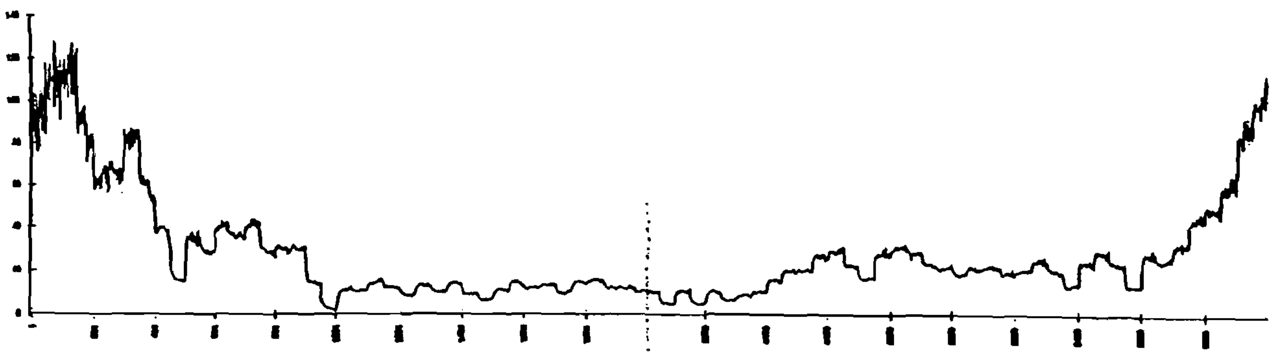

Figure (3)

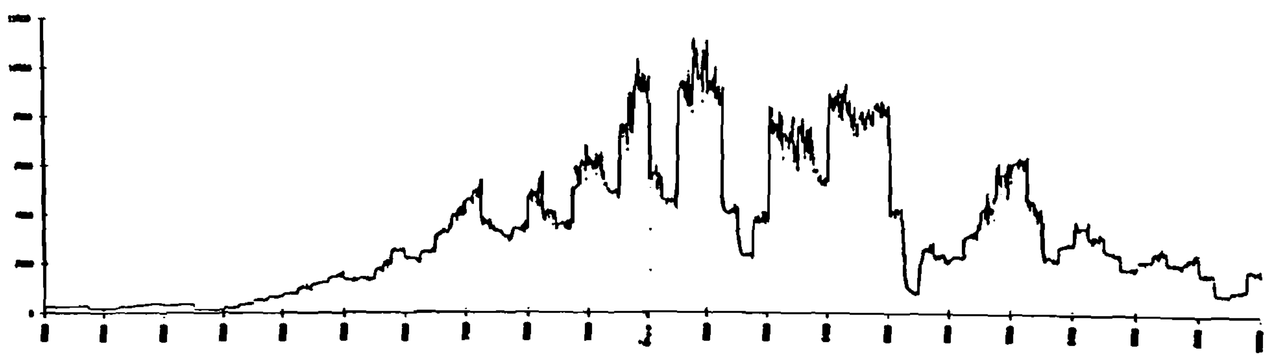

Figure (3) (continued)

In this model, the formation of beliefs and representations are at the core of speculative dynamics: prices are determined by a complicated structure of cross-expectations of the Keynesian 'Beauty Contest' type. As we did not endow agents *ex ante* with any particular model of the world, their representations are emerging and evolving through continuous interaction with their environment. This self-referential process generates a path-dependent price dynamics where both "polar opposite regimes" are self-reinforcing. In the metastable regime, the environment does not change conspicuously, agents adapt to it and modify their beliefs only marginally, hence reinforcing stability. Symmetrically, when agents do not manage to learn about their environment because it is changing too fast, their representations are evolving all the time, increasing instability[1].

These results are consistent with a) what has been obtained in contagion models like the ones by Orléan (1990) and Kirman (1991), and b) econometric evidence on the non-stationarity of the distributions from which price variations are drawn s(e.g. Friedman & Vandersteel (1982) and Bollerslev et al. (1991)). They question the ability of a speculative market to reach an equilibrium when agents have different visions of the world.

Exploitation vs exploration

The dichotomy between exploitation of existing knowledge and exploration of novel hypotheses is central in the modelling of microbehaviours proposed here. There are basically two forces in our learning model which -together- set a particular balance between exploitation and exploration. The first concerns the reward/reinforcement mechanism and is controlled by the reward which the environment assigns to acting rules. High positive rewards to 'good' rules will tend to increase the likelihood that they will be used again in the future; conversely, low negative rewards to 'bad' rules will decrease the likelihood of their future use, loosing in this way exploratory feature they can possibly contain.

Similarily, the frequency and intensity of application of the genetic algorithm clearly acts upon the balance between exploitation of existing rules and generation of new ones.All in all, high relative rewards and low frequency genetic algorithms will tend to produce highly exploitative systems, whereas low relative rewards and high frequency genetic algorithms will tend to produce explorative systems.

---

[1]One of the challenges in such a model is to specify the turning (or critical) points, i.e. what makes the market enter into a phase of turbulence when it was before relatively stable, and vice versa. Unfortunately we are unable to do it now.These points certainly involve thresholds in individual representations and other discontinuities in learning processes, but a more precise investigation will have to be delayed until further research is accomplished on the modelling of learning processes.

Figure (4) shows two simulations; the upper simulation is one from the reference model, and the lower one corresponds to an identical situation but for the values of the profit parameters that have been doubled ($\pi_{s,A} = \pi_{b,B} = 20$ and $\pi_{s,B} = \pi_{b,A} = -4$). The

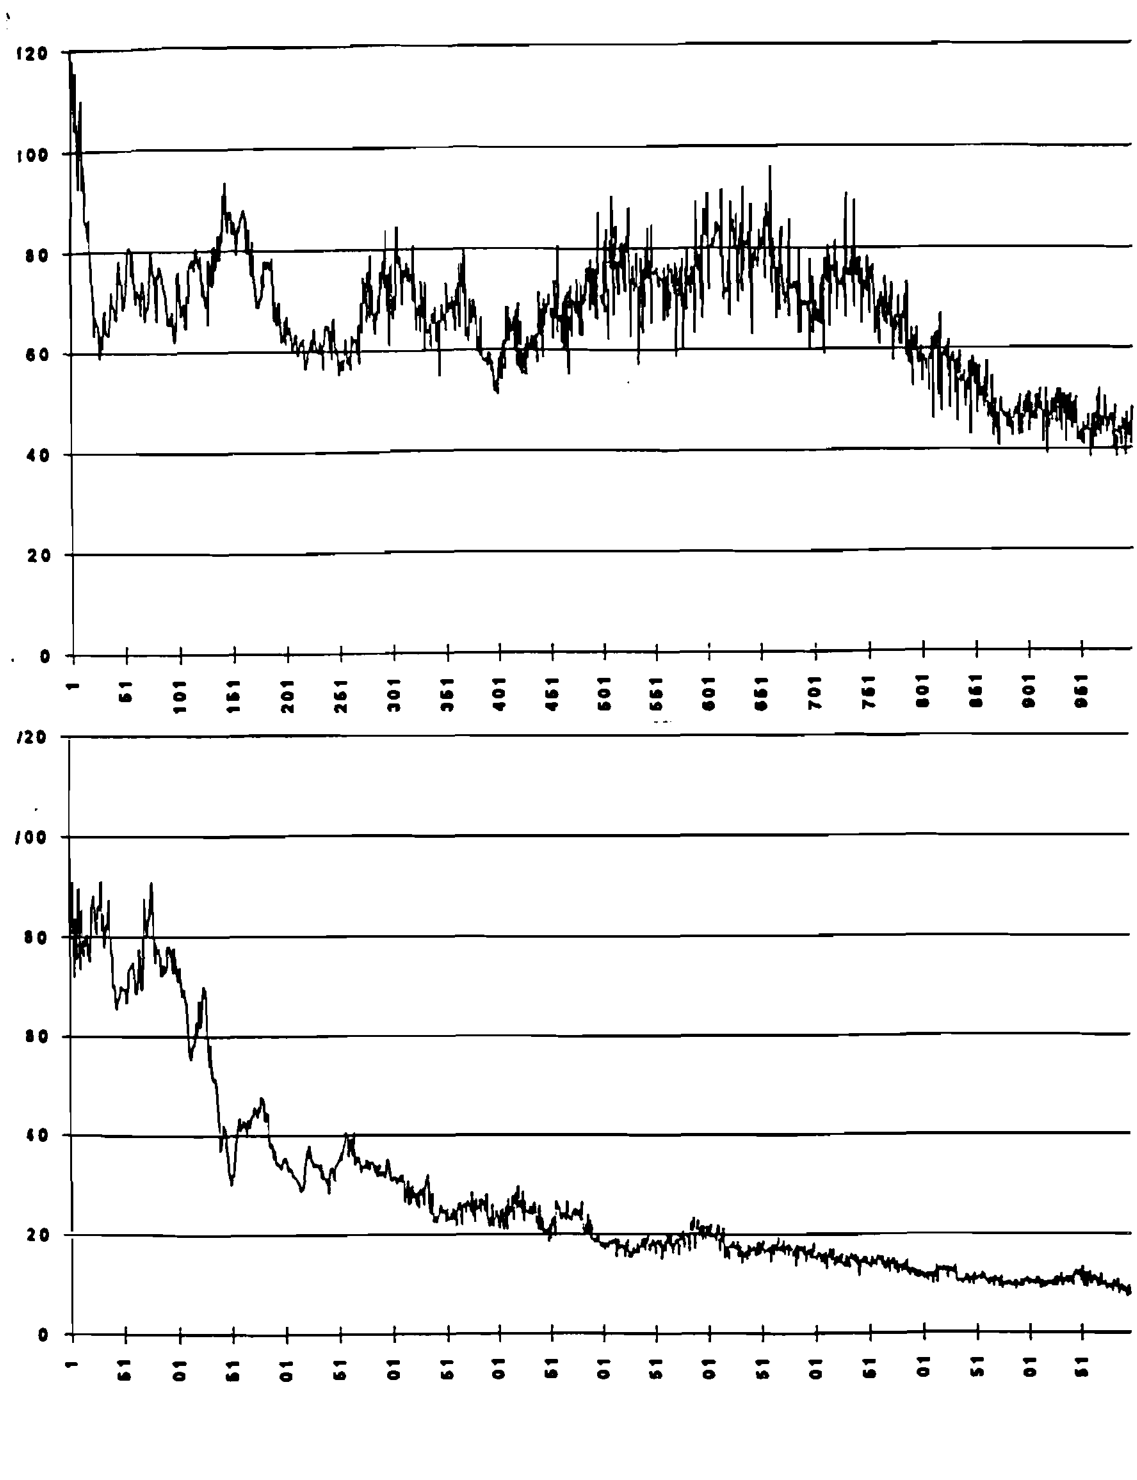

Exploitation vs exploration

Figure (4)

comparison between these two simulations clearly shows that a higher weight given to exploitation relative to exploration tends to stabilize price dynamics. As the 'good' rules are strongly reinforced, agents will stick to them and not change their representations too often. Since in this model, the environment of each agent is made up exclusively by the others, if all agents behave in a relatively inertial manner, the environment will be quite stable, and rules that were 'good' will continue to be 'good'. Such a self-reinforcing mechanism in the formation of representations can lead to situations of 'lock-in' (see e.g. Arthur (1992)); in which the region where the price will stabilize itself depends upon the content of the rules which emerged as 'good' rules at the beginning of the process.

Generality and specifity of rules

By varying k_3 (the specificity parameter in the bucket brigade) and the probabilities given to generalization and specification in the genetic algorithm, we can study the impact of these two 'cognitive attitudes' on the price dynamics. General rules are robust rules, in the sense that they associate the same action to a wide set of states of the world. In providing the same automatic response to a range of signals, they can be thought of a kind of routine[1]. Alternatively, general rules can simply mean that agents are ignorant, and that they cannot decode their environment successfully.

Routinized behaviour on financial markets, as chartism and technical analysis, is believed to have a destabilizing effect on price dynamics. Figure (5) shows two simulations; the upper one is from the reference model and in the lower one, k_3 has been decreased from 0.5 to 0.3. A lower reward to specificity seems here to increase short-term volatility (the same result has been obtained by increasing p(g) relative to p(s) in the genetic algorithm). In other words, routines (in the narrow meaning it is given here) would seem to be destabilizing.

[1]The emphasis on the automaticity of behaviour entailed by a routine has been made by e.g. Nelson & Winter (1982) and March & Simon (1993).

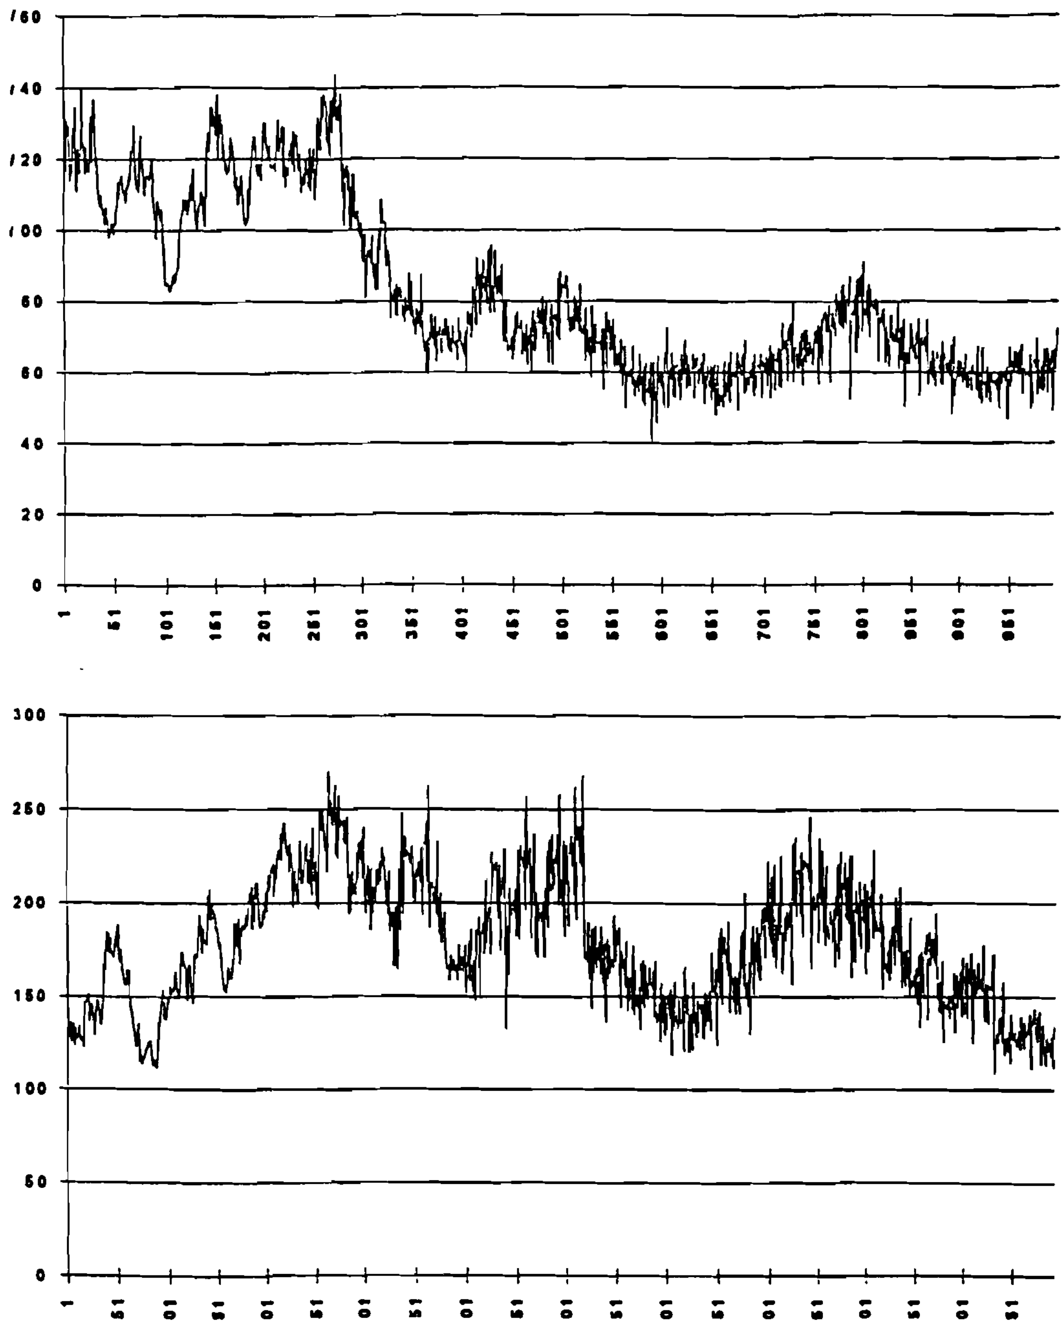

Generality and specificity
Figure (5)

V Concluding remarks

This exercise is intended to show that exchange rate dynamics close to the one observed could be generated by an alternative mechanism where the central role is given to the formation of expectations and, more generally, to learning processes of imperfectly rational agents. The results obtained seem quite promising, but some open questions and unsolved issues remain, which should be pursued in future research.

First, the model is very simple: agents have only one period-ahead expectations and very little information. In turn, one has to introduce a centralizing institution to allow agents to transact. How much more 'intelligence' our agents should be endowed with to be able to coordinate in a decentralized market would be an interesting issue to explore.
Second, classifier systems, if they allow us to model an evolutionary learning process, are essentially black boxes. Building indicators to track more precisely what is going on inside these systems (e.g. in terms of inertia and heterogeneity of the models of the world) should be on the agenda
Finally, there is a whole array of methodological issues to be tackled concerning inference in simulation models. For instance, as Lane (1992) puts it, "Can we argue that the real world aggregate regularity is indeed 'caused by' the entities and interactions we abstracted out of it and built into the artificial world, in which the analog of that regularity was identified as an emergent property?" (p. 9). And how long should a specific feature persist to be considered as an emergent regularity? No rigourous answers are yet available on those topics but it does not seem too farfetched to assume that an identity between observed and emergent properties increases the plausibility that the phenomenon under scrutiny is 'caused' by the modelled mechanisms. Such an identity has not been shown here, and a thorough statistical treatment of the price series generated by this model is the next step of our research.

References

Aronson, E., (1972), *The Social Animal*, San Francisco, W. H. Freeman and Co.

Arrow, K. J., (1986), "Rationality of Self and Others in an Economic System", in R. M. Hogarth & M. W. Reder (Eds.), pp. 201-16.

Arthur, W. B., (1989), "Competing Technologies, Increasing Returns and Lock-in by Historical Events", *The Economic Journal*, 99 (394), pp. 116-31.

Arthur, W. B., (1991), "Designing Economic Agents that Act Like Human Agents: A Behavioral Approach to Bounded Rationality", *American Economic Review*, Papers and Proceedings, 81 (2), pp. 353-59.

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

Arthur, W. B., J. Holland, B. LeBaron, R. Palmer & P. Taylor (1993), Santa Fe Institute, to appear.

Arthur, W. B., Y. M. Ermoliev & Y. M. Kaniovski, (1987), "Path-Dependent Processes and the Emergence of Macro-Structure", *European Journal of Operational Research,* 30, pp. 294-304.

Azariadis, C., (1981), "Self-Fulfilling Prophecies", *Journal of Economic Theory*, 25, pp. 380-96.

Baillie, R. T. & P. C. McMahon, (1989), *The Foreign Exchange Market. Theory and Econometric Evidence*, Cambridge, UK, Cambridge University Press.

Bak, P. & K. Chen, (1991), "Self-Organized Criticality", *Scientific American*, 264 (1), pp.46-53.

Bell, D. E., H. Raiffa & A. Tversky (Eds), (1988), *Decision Making. Descriptive, Normative and Prescriptive Interactions*, Cambridge UK, Cambridge University Press.

Beltratti, A., & S. Margarita, (1992), "Simulating an Artificial Adaptive Stock Market", Universita di Torino, mimeo.

Bilson, J. O., (1981), "The 'Speculative Efficiency' Hypothesis", *Journal of Business*, 54 (3), pp.435-51.

Blanchard, O. J. & M. W. Watson, (1982), "Bubbles, Rational Expectations and Financial Markets", in P. Wachtel (Ed.), *Crisis in the Economic and Financial Structure*, Lexington Books.

Blume, L. E & D. Easley, (1989), "Wealth Dynamics and the Market Selection Hypothesis", Cornell University, mimeo.

Bollerslev, T., R. Y. Chou & K. F. Kroner, (1991), "ARCH modeling in Finance", *Journal of Econometrics*, 52 (1/2), pp. 5-59.

Boothe, P. & D. Glassman, (1987), "The Statistical Distribution of Exchange Rates. Empirical Evidence and Economic Implications", *Journal of International Economics*, 22 (3/4), pp. 297-319.

Bossaerts, P., (1992), "Asset Prices in a Speculative Market", California Institute of Technology, mimeo.

Bray, M, (1982), "Learning, Estimation, and the Stability of Rational Expectations", *Journal of Economic Theory*, 26 (2), pp. 318-39.

Bray, M., (1983), "Convergence to Rational Expectations Equilibrium", in R. Frydman & E. S. Phelps (Eds.), pp.123-32.

Camerer, C. F., (1987), "Do Biases in Probability Judgment Matter in Markets? Experimental Evidence", *American Economic Review*, 77 (5), pp. 981-97.

David, P. A., (1988), "Path-Dependence: Putting the Past into the Future of Economics", *IMSSS Technical Report*, # 533, Stanford University.

De Grauwe, P. & D. Decupere, (1992), "Psychological Barriers in the Foreign Exchange Market", *CEPR Working Paper* #621.

De Long, J. B., A. Shleifer, L. H. Summers & R. J. Waldmann, (1990), "Positive Feedback Investment Strategies and Destabilizing Rational Speculation", *The Journal of Finance*, 45 (2), pp. 379-95.

Dosi, G. & M. Egidi, (1991), "Substantive and Procedural Uncertainty", *Journal of Evolutionary Economics*, 1 (2), pp. 145-68.

Dosi, G., C. Freeman, R. Nelson, G. Silverberg & L. Soete (Eds), 1988, *Technical Change and Economic Theory*, London, Pinter Publishers.

Figlewski, S., (1978), "Market 'Efficiency' in a Market with Heterogeneous Information", *Journal of Political Economy*, 86 (4), pp. 581-97.

Flood, R. P. & P. M. Garber, (1980), "Market Fundamentals vs Price-Level Bubbles: The First Tests", *Journal of Political Economy*, 88 (4), pp. 745-70.

Frankel, J. A. & K. A. Froot, (1987), "Using Survey Data to Test Standard Propositions Regarding Exchange rate Expectations", *American Economic Review*, 77 (1), pp. 133-53.

Friedman, D. & S. Vandersteel, (1982), "Short-Run Fluctuations in Foreign Exchange Rates. Evidence from the Data 1973-79", *Journal of International Economics*, 13 (1/2), pp. 171-86.

Friedman, M., (1953), *Essays in Positive Economics*, Chicago, The University of Chicago Press.

Froot, K. A., & T. Ito, (1989), "On the Consistency of Short-Run and Long-Run Exchange Rate Expectations", *Journal of International Money and Finance*, 8 (4), pp. 487-510.

Frydman, R. & E. S. Phelps, (1983), *Individuel Forecasting and Aggregate Outcomes.'Rational Expectations'Examined*, Cambridge UK, Cambridge University Press.

Gigerenzer, G., Z. Swijtink, T. Porter, L. Daston, G. Beatty & L. Krueger, (1989), *The Empire of Chance*, Cambridge UK, Cambridge University Press.

Granovetter, M., (1978), "Threshold Models of Collective Behaviour", *American Journal of Sociology*, 83 (6), pp. 1420-43.

Granovetter, M., (1985), "Economic Action and Social Structure : The Problem of Embeddedness", *American Journal of Sociology*, 91 (3), pp. 481-510.

Grossman, S. J. & J. E. Stiglitz, (1980), "On the Impossibility of Informationally Efficient Markets", *American Economic Review*, 70 (3), pp. 393-408.

Hayek, F. A., (1937), "Economics and Knowledge", *Economica*, 4, pp. 33-54.

Hey, J. D., (1991), *Experiments in Economics*, Oxford, Basil Blackwell.

Hirshleifer, J., (1975),"Speculation and Equilibrium: Information, Risk, and Markets", *Quarterly Journal of Economics*, 89 (4), pp. 519-42.

Hogarth, R. M. & M. W. Reder (Eds.), (1986), *Rational Choice. The Contrast between Economics and Psychology*, Chicago, The University of Chicago Press.

Holland, J. H. & J. H. Miller, (1991), "Artificial Adaptive Agents in Economic Theory", *American Economic Review*, 81 (2), pp. 365-70.

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

Holland, J. H., (1986), "Escaping Brittleness : The Possibilities of General-Purpose Learning Algorithms Applied to Parallel Rule-Based Systems", in R. S. Michalski, J. G. Carbonell & T. M. Mitchell (Eds), *Machine Learning II* , Los Altos, CA., Morgan Kaufmann, 1986, pp. 593-623.

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

Hsieh, D. A., (1991),"Chaos and Nonlonear Dynamics: Applications to Financial Markets", *The Journal of Finance*, 46 (5), pp. 1839-77.

Ito, T., (1990), "Foreign Exchange Rate Expectations: Micro Survey Data", *American Economic Review*, 80 (3), pp. 434-49.

Johnson-Laird, P. N., (1983), *Mental Models* , Cambridge, MA., Harvard University Press.

Johnson-Laird, P. N., (1991), "A Model Theory of Induction", *Journal of Technology and Social Studies on Society and Technology*, .

Kaldor, N., (1939), "Speculation and Economic Stability", *Review of Economic Studies*, 7 (1) , pp. 1-27.

Keynes, J. M., (1921), *A Treatise on Probability*, in The Collected Writings of J. M. Keynes, Vol. 8, New-York, St Martin's Press for the Royal Economic Society, (1973).

Keynes, J. M., (1936), *The General Theory of Employment, Interest and Money,* New-York, Harcourt Brace & World.

Kindleberger, C. P., (1978), *Manias, Panics and Crashes*, New-York, Basic Books.

Kirman, A. P., (1991), "Ants, Rationality and Recruitment", mimeo, march, European University Institute, Florence.

Kirman, A. P., (1992), "Whom or What Does the Representative Individual Represent?", *Journal of Economic Perspectives*, 6 (2), pp. 117-36.

Lakoff, G., (1987), *Women, Fire, and Dangerous Things. What Categories Reveal about the Mind*, Chicago, TheUniversity of Chicago Press.

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

Langton, C. G., C. Taylor, J. D. Farmer & S. Rasmussen (Eds.), (1991), *Artificial Life II* , Redwood City, CA., Addison-Wesley.

Lesourne, J., (1991), *Economie de l'ordre et du désordre*, Paris, Economica.

Mackay, C. (1841), *Extraordinary Popular Delusions and the Madness of Crowds*, New-York, Harmony Books, (1980).

March, J. G. & H. A. Simon, (1993), "Introduction to the Second Edition of *Organizations*", mimeo, and Oxford, Blackwell Publishers, (1994).

March, J. G., (1991), "Exploration and Exploitation in Organizational Learning", *Organization Science*, 2 (1), pp. 71-87.

Marengo, L., (1991), *Knowledge, Coordination and Learning in an Adaptive Model of t he Firm*, Ph. D. Dissertation, SPRU, University of Sussex, september.

Milgrom, P. & N. Stokey, (1982), "Information,Trade and Common Knowledge",*Journal of Economic Theory*, 26 (1), pp. 17-27.

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

Orléan, A., (1988), "L'auto-référence dans la théorie keynésienne de la spéculation", *Cahiers d'Economie Politique*, 14-15.

Orléan, A., (1990), "Contagion mimétique et bulles spéculatives", in J. Cartelier (Ed.), *La formation des grandeurs économiques*, Paris, Presses Universitaires de France, pp. 285-321.

Polya, G., (1945), *How to Solve It* , Princeton, NJ., Princeton University Press, (second edition, 1957).

Rosser, Jr., J. B., (1991), *From Catastrophe to Chaos: A General Theory of Economic Discontinuities*, Boston, Kluwer Academic Publishers.

Schelling, T. C., (1978), *Micromotives and Macrobehavior*, New-York, W. W. Norton and Co.

Shiller, R. J., (1989), *Market Volatility*, Cambridge MA., MIT Press, 3rd edition (1991).

Silverberg, G., (1988), "Modelling Economic Dynamics and Technical Change: Mathematical Approaches to Self-Organization and Evolution", in G. Dosi et al. (Eds.), pp. 531-59.

Smith, V. L, G. L. Suchanek & A. W. Williams, (1988), "Bubbles, Crashes and Endogenous Expectations in Experimental Spot Asset Markets", *Econometrica*, 56 (5), pp. 1119-51.

Tirole, J., (1982), "On the Possibility of Speculation under Rational Expectations", *Econometrica*, 50 (5), pp. 1163-81.

Tordjman, H., (1994), *Dynamiques spéculatives, hétérogénéité des agents et apprentissage : le cas des taux de change*, Ph. D. Dissertation, CEFI, Aix-Marseille II, january.

Tversky, A. & D. Kahneman, (1986), "Rational Choice and the Framing of Decisions", in R. M. Hogarth & M. W. Reder (Eds.), pp. 67-94.