Working Paper

The Formation of Beliefs on Financial Markets: Representativeness and Prototypes

Hélène Tordjman
CREA, Ecole Polytechnique, and CREI, University of Paris XIII

WP-96-87
July 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

The Formation of Beliefs on Financial Markets: Representativeness and Prototypes

Hélène Tordjman
CREA, Ecole Polytechnique, and CREI, University of Paris XIII

WP-96-87
July 1995

To appear in Lazaric and Lorenz (Eds.), *The Economics of Trust and Learning*, E. Elgar Publishers, 1996.

Working Papers are interim reports on work of the International Institute for Applied Systems Analysis and have received only limited review. Views or opinions expressed herein do not necessarily represent those of the Institute, its National Member Organizations, or other organizations supporting the work.

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

Acknowledgement

I would like to thank Edward Lorenz and André Orléan for their valuable comments and suggestions on an earlier version.

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

The Formation of Beliefs on Financial Markets: Representativeness and Prototypes

I. Introduction

As rational expectations, and more recently rational learning, have become the usual way to embody the formation of individual beliefs and their evolution in economic models, empirical evidence gathered from market surveys and experiments is accumulating which shows systematic departure from the rationality hypothesis. These departures from rationality are especially salient on financial markets, where numerous studies (e.g. Frankel & Froot, 1987; Shiller, 1989; De Bondt & Thaler, 1990) have highlighted systematic biases in the formation of agents' beliefs[1].

This kind of biases have been studied in different settings by cognitive psychologists since a long time (see Kahneman, Slovic & Tversky (Eds.), 1982, from now on KST, 1982). Their pervasiveness and their similarity across domains, i.e. their robustness, suggest that agents don't behave in a rational fashion while constructing their judgements and expectations, neither on financial markets nor elsewhere. More specifically, the way they revise their beliefs in front of new information is not Bayesian: they tend to be insensitive to prior knowledge and to sample size (Camerer, 1987; KST, 1982), and to be overconfident in their judgments relatively to the available evidence (Shiller, 1989; KST, 1982; Griffin & Tversky, 1992). Hence, from a descriptive point of view, the evolution of individual beliefs does not seem to be adequately modelled by a rational learning process which postulates -among others- Bayesian revision and convergence of beliefs toward a unique "true" model of the world[2].

As put forward by Kahneman & Tversky, the fact that these judgements are systematically biased in the same directions shows that we use other procedures than probabilistic updating to form our judgements. These procedures are generally simpler, such as rules of thumb and heuristics. They have identified few of these heuristics which have since been extensively studied (see KST, 1982), e.g. representativeness, availability, anchoring...

This paper is about the representativeness heuristic, i.e. judging by closeness and/or familiarity with a typical case (a more precise definition will be given below). The discussion will be restrained to representativeness and its consequences on the formation of beliefs, but further investigation should show that other heuristics, like availability, anchoring and simulation play a role in markets too.

The purpose of this paper is centered around the following question: as asked by Arrow, 1982, is the representativeness heuristic at work on financial markets, and if yes, could it help to explain some of the so-called irrationality of these markets? In other words, is it possible that phenomena like overreaction to new information, cumulative deviations and bubbles be generated by the fact that agents on these markets use representativeness instead of Bayes' rule in constructing their judgements? The findings of this study tend to favour positive answers to these questions: "heuristics matter".

The theoretical implications of these "appreciative" results[3] are not trivial. Actually, if economic agents do use a heuristic such as representativeness to form and revise their beliefs, the Bayesian learning approach does no longer offer an adequate framework for representing individual knowledge and its evolution. Thus, one needs another representation of learning processes which would, among other, at least be compatible with the empirical evidence mentioned above, and account for the fact that individuals mainly -but not exclusively- use heuristics instead of probability theory in constructing their beliefs. Some recent theoretical developments in cognitive psychology and Artificial Intelligence could shed light on the cognitive mechanisms generating behaviours like judgement by representativeness. These approaches are not unified, but they share some basic concepts such as prototypical categories and mental models (see for instance Johnson-Laird, 1983, Holland, Holyoak, Nisbett & Thagard, 1986, and Lakoff, 1987), and can give some theoretical foundations as to why individuals use heuristics like representativeness in forming and revising their beliefs.

The paper is organized as follows. Section II presents some empirical evidence about the representativeness heuristic. The consequences of the use of this heuristic on financial behaviours and asset price dynamics are discussed in section III. Section IV introduces theoretical issues concerning representativeness and processes of categorization. Section V concludes and highlights perspectives for future research.

II. Empirical evidence about representativeness

Definition

Before presenting some empirical results about the use of the representativeness heuristic in belief formation, it is necessary to define precisely in what sense this term will be used here. Tversky & Kahneman (in KST, 1982, ch. 6) distinguish two different meanings of the concept. In the first acceptation, representativeness refers to a sample and its quality: "how representative of the population under study is the sample on which I base my judgement?". It has been shown that judgements *of* the representativeness of a sample are generally wrong, giving rise to biases such as the gambler fallacy, and to what Tversky & Kahneman labelled "the law of small numbers" (KST, 1982, ch. 2).

Failures in judgements of representativeness can certainly be spotted in economic behaviours, but it is not the purpose of this paper. Instead, I wish to explore the consequences of representativeness in its second acceptation, which is less statistical and more cognitive. In this second sense, representativeness is a mean (a heuristic) by which individuals construct subjective probability judgements, and we will talk about judgements *by* representativeness. More precisely, a probability of an event constructed in this way is evaluated "by the degree to which [this event] is: (i) similar in essential properties to its parent population; and (ii) reflects the salient features of the process by which it is generated. [...] In many situations, an event A is judged more probable than an event B whenever A appears more representative than B" (Kahneman & Tversky in KST, 1982, ch. 3, p. 33).

For instance, if one has to evaluate the probability that a given animal is a dog, the more the encountered animal resembles a typical dog, the higher the probability that it is a dog. Thus, similarity is essential in the way individuals form beliefs about the belonging of an object A to a class X, i.e. in categorization. Representativeness is also present in attributing causality, in evaluating the probability that an event A is caused by a process X. Let us take a description of a manager and try to predict if his firm makes high profits. If the manager is described in terms such as "young, ambitious and well-educated", the probability that his firm makes high profits will be judged higher than if the manager is "old, conservative and has not been to college", because the stereotype of a successful firm usually includes, in our times, a representation of a young and dynamic manager. It should be noted that these stereotypes often have a strong social dimension. The stereotype of a successful manager is not the same in France and in Japan, and these representations are different now from what they were in the 19th century. In other words, cognition -at least some part of it- is *embedded* (Granovetter, 1985), or*embodied* (Lakoff, 1987) in more global social structures such as institutions, norms, and collective representations (Durkheim, 1896). How are such stereotypical categories constructed and how do they evolve is an issue of great importance involving different theoretical questions, some of which will be evoked in section IV.

Empirical evidence

Experiments about the formation of beliefs by representativeness are numerous and many systematic biases frequently observed seem to be generated by the use of this heuristic (see Tversky & Kahneman, KST, 1982, ch. 1 for a brief introduction). Their results show that the Bayesian model of belief revision (in its Rational Learning acceptation) is not descriptively correct. In particular, it is well-established that (i) people tend to ignore prior information, and (ii) underestimate the importance of the *quality* of the evidence, its precision (they misjudge the representativeness of the sample, are prone to small sample bias etc...). At the same time, they often give too much weight to the *strength* of the evidence, i.e. to the semantic content of an information, which can lead to an overestimation of probability judgements and to behaviours of overconfidence. Both of these cases will successively be reviewed in sections (i) and (ii), and a final point will draw some implications of the existence of such biases in "real world" settings.

Before presenting some experiments, let us briefly recall Bayes' rule.
Let H_A and H_B be two independent hypotheses such that $H_A \cap H_B = \varnothing$ and $H_A \cup H_B = \Omega$. Let C be an event one can observe, i.e. an information. $P(H_A)$ is the prior probability of having hypothesis A true; $P(H_A|C)$ is the posterior probability of having hypothesis A verified when the event C is observed; $P(C|H_A)$ is the conditional probability of observing the event C when hypothesis H_A is true, and P(C) is the unconditional probability of the event C. If we know $P(H_A)$ -it can be whatever subjective assessment-, P(C) and $P(C|H_A)$, then, Bayes' rule allows us to calculate the posterior probability of hypothesis A knowing C (the same holds of course for H_B):

$$P(H_A|C) = \frac{P(C|H_A)\,P(H_A)}{P(C)}$$

where $P(C) = \sum_{i=1}^{n} P(C|H_i)\,P(H_i)$ and $H_i = 1...n$, i.e. n mutually exclusive and collectively

exhaustive hypothseses (here, n=2, H_A and H_B). Let us just remark that, if the calculus is very simple, it is nonetheless difficult to perform because the prerequisites in terms of knowledge are quite high. In order to calculate $P(H_i|C)$, the individual has to know P(C) and $P(C|H_i)$, which may not be obvious.

(i) Neglect of base rates information

A well-known experiment done by Kahneman & Tversky (KST, 1982, ch. 4) highlights the tendency to be insensitive to prior information, contrary to the prediction of Bayes' rule. Two groups of subjects[4] were given some identical descriptions of persons (age, family status, tastes and hobbies...). They were told that people could be of two "types", engineers and lawyers. In addition, they were given the prior probabilities attached to the two groups, lawyers and engineers, but the first group was told that there was 30% of lawyers and 70% of engineers and the second group was told the reverse, i.e. that there was 70% of lawyers and 30% of engineers. They were then asked to evaluate the probability that the person described was belonging to the group of engineers. Altogether, subjects were presented with five descriptions of persons, and the results were identical in all cases: prior information concerning the relative frequency of

engineers in the population was not deemed relevant. One of such description was the following:
"Jack is a 45-year-old man. He is married and has four children. He is generally conservative, careful, and ambitious. He shows no interest in political and social issues and spends most of his free time on his many hobbies which include home carpentry, sailing, and mathematical puzzles".

Given the configuration of the priors in the two groups of subjects, the prediction of Bayesian theory would be that the first group would give a higher probability to the fact that Jack is an engineer, when the second group would estimate it to be less probable. The results obtained by Kahneman & Tversky do not exhibit such pattern. Indeed, the two groups gave more or less the same estimation of the probability for Jack to be an engineer, revealing that they were not taking their prior information into account[5].

Different criticisms have been addressed to this kind of experiment. For instance, it is often argued that such verbal descriptions are not randomly drawn from the population, and this could explain the observed biases. Indeed, when the draws are verbal descriptions, it is difficult to establish if these descriptions reflect the characteristics of the population at large. Moreover, there is room for interpretation, and framing effects or the like could explain deviations from the Bayesian rule. Another criticism, often made by economists, concerns the monetary incentives individuals have in making the right choice. Economists usually claim that on real markets, people are experienced and must make the correct choice if they want to survive, otherwise they are selected out of the market[6], a constraint which is not overwhelming in experiments[7].

Others experiments about representativeness and the neglect of priors have been designed to avoid these problems of "false" randomness and of monetary incentives, e.g. Grether, 1980, and Camerer, 1987. In order to compare the relative descriptive adequacy of Bayes' rule vs. representativeness in evaluating subjective probabilities, Grether made the following experiment. He used three "bingo cages", i.e. urns, containing each six balls. The first cage, X, contained six balls numbered one through six. The second cage, A, contained four balls marked with an "N" and two balls marked with a "G", whereas in the third cage, B, there were three "N" balls and three "G" balls. The first urn, X, was used as a prior: a ball was drawn from X, and if it was numbered one to four, the urn A would be chosen; when the ball was numbered five or six, the urn B would be used. Then, six balls were drawn, with replacement, from the chosen urn. The subjects knew how the whole procedure was functioning, and the content of the three cages, but they were just observing the final draw of six balls ("Ns" and "Gs") without knowing from which urn, A or B, it had been drawn. For each sample of six balls, they were asked to estimate the probability that the sample had been drawn from A. As they also knew the content of X and the rule determining the choice between A and B, they knew the prior probability that the sample was coming from A[8]. In order to evaluate the incidence of monetary incentives, the subjects were divided into two groups: the first group was told that they would be paid a sure gain at the end of the experiment; in the second group, people would be paid conditionally to the correctness of their answers.

The results obtained by Grether are not so clear-cut as the ones from the lawyers / engineers experiment. First, in most experiments, prior information is taken into account by the subjects, although its importance is generally underweighted; but even if they are not pure Bayesians, agents do use base rates information in estimating probabilities. However, there is a systematic bias in the formation of beliefs which is generated by something akin to the representativeness heuristic. For instance, whenever the final outcome (a draw of six balls) was exactly representative of urn A (four "Ns" and two "Gs") or urn B (three "Ns" and three "Gs'), the priors were almost ignored and the probability that the sample was drawn from urn A (respectively urn B) was very high, whatever the priors[9]. Second, a rather puzzling result of Grether's experiment, at least for economists, is that monetary incentives do not seem to be of much importance: the general patterns were the same in the two groups, whether there was an incentive to make the correct choice or not.

(ii) Overconfidence

Another frequent behaviour which is at odds with Bayes' rule is overconfidence, i.e. the tendency to be "more confident in [our] judgements than is warranted by the facts" (Griffin & Tversky, 1992, p. 411). This bias is often provoked by an underestimation of the weight of the evidence (its quality or precision) and, simultaneously, an overestimation of its strength (its semantic content).

An example of the combination of weight and strength of the evidence in a simple and well-defined setting is given by the heads and tails game, where weight is sample size and strength is sample proportion. Even in such a simple setting, Griffin & Tversky, 1992, showed that people tended to be overconfident when strength was high and weight was low. People were presented with series of heads and tails resulting from spinning a coin. As the subjects were told, "unlike tossing, which (on average) yields an equal number of heads and tails, spinning a coin leads to a bias favouring one side or the other because of slight imperfections on the rim of the coin" (p. 414). People were given the value of the bias, the coin fell on one side three times out of five, but they did not know if the favoured side was heads or tails. They were then presented with twelve different samples of heads and tails, with varying sizes (from three to thirty three) and numbers of heads (from two to nineteen). All samples favoured heads. The subjects were asked to estimate the probability for each sample that the bias was in favour of heads. The results of this experiment highlights the way judgement by representativeness can induce overconfidence: each time the sample was small (from three to five) and the relative proportion of heads high, people's estimates were above their Bayesian values. Thus, overconfidence can be associated with an insufficient attention given to sample size (weight of the evidence) and too much attention given to sample proportion (strength of the evidence). But the symmetric is also true: when the sample was larger and the proportion of heads relatively smaller, people exhibited underconfidence, i.e. their estimates always lied under the Bayesian values.

The same pattern was obtained in others experiments (Griffin & Tversky, 1992; KST, 1982) and, as for the neglect of priors, seems relatively robust. Some of these other experiments were done in more complicated settings, real-life situations involving representations about social, institutional and cultural events. The same biases were found in social situations than in pure chance settings. These findings, according to Griffin & Tversky, 1992, thus apply also when people have to form "judgements about uncertain events such as who will win an upcoming election, or whether a given book will make the best-seller list. When people assess the probability of such events they evaluate [...] their impression of the candidate or the book" (pp. 422-23). However, as we shall see now, the interpretation of results obtained in "real world" situations is quite more complicated.

(iii) "Real world" settings

The comparison between the use of the representativeness heuristic in "real life" situations and in pure chance settings highlights two facts. First, the more complicated the environment, the more people use this heuristic. Second, representativeness is more difficult to define precisely, and to spot, when the environment is complicated and uncertain.

Let us go back to the results about the neglect of prior information. The results of Kahneman & Tversky in the "lawyer/ engineer" study are more radical towards a rejection of the Bayesian hypothesis than the ones obtained by Grether, 1980, and Camerer, 1987, in pure chance settings. But these results are less contradictory to what it seems from a first look. Indeed, when the problems are set up in terms of urns, as in Grether and in Camerer, it is easier to build a "good" representation of the problem because all the data is given and well-defined, and the necessary knowledge to solve the problem is contained in the data. The evidence at hand does not really ask for interpretation: the randomness is only deriving from the fact that the draws can be made from two different urns. In the engineer experiment, the knowledge mobilized and processed to solve the problem is much more complex, and has to do with the model one can have about how an engineer looks like, which in turn comes from our particular past experience -encounters with specific engineers- and the collective images entailed by social structures and "culture" at large.

It should be noted that in the engineer case, people know that there are two types of agents, engineers and lawyers, and the only uncertainty stemms from the ambiguity concerning the "boundaries" of these two categories. Hence, one could argue that there is no true uncertainty, and that, if the environment is more complex than in urns problems, it is not radically uncertain. This may be true, depending on the acceptation of the notion of uncertainty. I do not wish to enter into these debates here; for the purpose of this paper, it suffices to consider that there is true uncertainty when it becomes impossible to define standard probabilities over the set of possible states of the world, which is the case in the engineer problem[10]. To summarize, it may be possible to be almost Bayesian in well-defined, "close-worlds" problems, but the more the environment becomes uncertain, the richer the necessary knowledge, and the higher the probability to use simplifying heuristics instead of sophisticated mathematical constructions to form our beliefs.

Representativeness in a complicated and evolving environment is much more difficult to define precisely than in pure games of chance, where the similarity criterion is unique, well-defined and it is thus possible to talk about "exact representativeness" (e.g. Camerer). In the real world, "the factors that make particular task and problem characteristics the salient ones with respect to which representativeness is judged" (Olson quoted by Bar-Hillel in KST, 1982, ch. 5, p. 73) cannot be defined a priori. In experiments where there were two possible criteria of similarity, people would order them differently, albeit the problem was well-defined, and arrive at different conclusions although they were following the same heuristic. Some studies have been done to make people "reveal" their representativeness criteria (e.g. Bar-Hillel in KST, 1982, ch. 5) by estimating similarities instead of probabilities. But this type of approach is, it seems, bound to fail because of the infinity of possible criteria. Typical categories (stereotypes) are context dependent and history dependent, and, if there are agreements about some basic features of some basic categories -communication would otherwise be impossible- the extent to which such agreements exist is very difficult to evaluate. I will come back to these issues in section IV. Another way to approach the problem could be to try to identify, in a given domain, some representations that seem to be commonly shared, i.e. some stereotypes, and to analyse the formation of beliefs of individuals that are not randomly chosen, but share similar concerns (as it has been done for instance by Kahneman & Lovallo on small businesses). The next section is devoted to such a (preliminary) tentative about representativeness on financial markets.

III. Representativeness and financial instability

There are basically two views of the functioning of financial and speculative markets. In the classical approach, agents are supposed to be fully rational; they trade financial assets only when their prices are going away from their equilibrium values, and by doing so, prices are pushed back to their fundamental levels. In this view, well exemplified by

Friedman, 1953, markets are rational and speculation is stabilizing. However, such a view of stable financial markets is at odds with the empirical evidence, which shows that price series exhibit high volatility and bubble-like events[11]. Another approach to speculation, which goes back to Keynes, 1936, considers speculation to be destabilizing, mainly because individuals are not fully rational. They have cognitive limitations and "drawbacks" but are generally more or less reasonable: their behaviours are intentional and they have a certain idea of adequacy between means and ends. They may nonetheless be prone to euphoria and panic, and more generally to different kind of influences which may either be "environmental" (influences of collegues, friends, so-called experts, newspapers) or "historical" (their own past experience, collective representations of historical events...). Indeed, in this view, agents are continuously evolving, adapting to a changing and uncertain environment, an environment which in turn is modified by their actions. As put forward by Keynes in his "Beauty Contest" metaphor, these feedback relations between agents and their environment are at the core of financial dynamics. In such a perspective, the cognitive mechanisms generating beliefs and orienting actions can be a source of endogenous instability, without having to postulate complete irrationality. Actually, studies in cognitive psychology such as the ones presented before can help to understand some features of financial markets.

More precisely, as it has been claimed by Arrow, 1982, Camerer, 1987, and others, heuristics like representativeness play a role in market functioning. For reasons evoked earlier, representativeness in real-life situations is quite difficult to spot: categories are complex, similarity criteria are many and different across individuals. It is nonetheless possible to exploit the numerous empirical literature about speculative episodes and tentatively identify the representations financial operators seem to have and the specific biases they display.

This empirical literature mainly takes two forms: the first one consists in historical studies of "great bubbles" (Mackay, 1841; Kindleberger, 1978; Galbraith, 1990); the second one analyses the results of markets surveys and questionnaires elaborated to understand how financial operators are forming their expectations (Frankel & Froot, 1987; Shiller, 1989; De Bondt & Thaler, 1990). They both display biases in the formation of beliefs on financial markets and other "anomalies". Two of these anomalies will successively be studied, namely excessive volatility and speculative bubbles, and I will try to highlight the way these features of financial dynamics are linked to the representativeness heuristic. First, the tendency to judge by similarity without taking into account prior information produces more extreme judgements than what they should be if individuals were perfect Bayesians. Thus, agents very often overreact to new information, which produces high volatility. Second, the overconfidence generated by an overweighing of the strength of the evidence relatively to its quality can explain at the same time the unrealistic optimism which is characteristic of the growing of a bubble (the euphoric phase) and the crash (the explosion of a bubble).

Overreaction and volatility

One of the explanation often proposed to "rationalize" the excessive volatility of asset prices is in terms of "news". In this view, price volatility is provoked by the process of information aggregation taking place on the market: each new information is rationaly interpreted and produces a modification of the previous expectations; price dynamics should thus follow a process similar to the process of arrival of new information. As there is no reason to think that the arrival of new information is not random, asset prices should follow a random walk. Unfortunately, this is not exactly the case (e.g. LeRoy, 1989). An alternative explanation of the high volatility experienced by financial prices is that individuals overreact to this new information: the way they revise their beliefs to take this information into account is too drastic. In other terms, in according too much importance to new information, people do not behave in a Bayesian fashion.

Since a number of years, some public and private institutions have been recording periodically the predictions of some groups of traders and/or analysts concerning the evolution of stock prices, exchange rates and other financial prices. We thus start to have at our disposal some data bases of individual forecasts that can be exploited to test alternative hypotheses about expectations formation on financial markets.

For instance, De Bondt & Thaler, 1990, studied one of these data base concerning one-year and two-years forecasted changes in earning-per-share made by security analysts. As data for actual returns is available, it is possible to test the accuracy of these predictions. One of the reasons why they chose security analysts' forecasts was that, as they put it (p. 52), "the precision of analysts expectations represents a natural upper bound to the quality of the earnings forecasts of less sophisticated agents". If the predictions of so-called experts are biased, it is reasonable to think that it will be even more so for "standard" operators.

Their results can be summarized as follows. First, they found evidence of a systematical bias in forecasts, compared to the rational expectations predictions. Regressions of actual changes of earnings-per-share on forecasted changes shed light on the regularity of expectational errors. Actually, De Bondt & Thaler found that the forecasts errors are predictable from the forecasted changes in earnings and that forecasts have the right sign but are too extreme relatively to the changes in earnings that effectively occured. In other words, security analysts seem to overreact to information and for instance judge the probability of "big events" higher than it appears to be ex post. Second, De Bondt & Thaler wanted to know if this phenomenon of overreaction was correlated with uncertainty about the future, and they compared the results obtained with one-year and two-years ahead forecasts, the idea being that two-years forecasts are more uncertain than one-year's. They found that the two-years forecasts were more extreme than the one-year ones, i.e. overreaction seems to be more pronounced when uncertainty is higher. As noted by Arrow, 1982, (p. 5), representativeness entails a "tendency to underestimate uncertainties" which is at the base of the general phenomenon of excessive reaction to information.

This feature of overreaction has been found in other data bases. For instance, Frankel & Froot, 1987, and Froot & Frankel, 1989, have extensively analysed predictions (one month, three months, six months and one year) of traders on foreign exchange markets. Frankel & Froot, 1987, did not look for overreaction, but found that expectations about future exchange rates were systematically biased, although they could not find a simple alternative model of expectations (adaptive, extrapolative, regressive, and so on...) that fitted the data. They concluded on the complexity of the process that generates the predictions... Froot & Frankel, 1989, using the same data base, studied the relation between the spot exchange rate and the forward rate, i.e. the well-known "forward bias" problem. They first decomposed the bias into a time-varying risk premium and an expectational error, and showed that the bias was essentially stemming from a failure of rational expectations, and was not correlated with any risk premia. Then, they looked for overreaction (they labelled it "excessive speculation") and regressed the expectational error on the expected depreciation of the exchange rate. Their results are similar to the ones obtained by De Bondt & Thaler with earnings predictions: the expectational error is positively correlated to the expected depreciation, and forecasted changes are again too extreme. As for the security analysts, there is evidence that the Forex traders are prone to overreaction too.

The link between this behaviour of overreaction and the representativeness heuristic is quite straightforward. To overreact to a new information means to give more weight to this information in the construction of a probability judgement than would be justified by Bayes' rule. Its corollary is insensitivity to prior information: the posterior judgement reflects mainly the last information, as if no summary data about the phenomenon was available. As we saw in section II, such a neglect of prior information may, and is often, a result of constructing judgements by representativeness. Indeed, when a new information arrives which fits with the general representation of the market currently held (e.g. "the market is bullish", or "web-related software firms have exponentially rising earnings"), such an information will have a disproportionate weight in the formation of judgements, and the knowledge contained in the priors will tend to be ignored. This is because this new information will be taken as "typical" of a state of the market; hence, the probability attached to this state will be overestimated. For instance, the security analysts from the De Bondt & Thaler's study display overreaction in letting their predictions of future earnings be too much influenced by the most recent information, whereas earnings dynamics is also determined by structural factors which could be known through a fundamental analysis. Even in the case of total ignorance about these fundamentals, a professional analyst would -at least this is postulated by rational expectations economics- know something about the distribution of earnings across the population of firms, and use these known frequencies in the construction of his or her belief. Instead, they seem to take each new information at face value, excessively focusing on the semantics of this information and seeing it as a clue that their current representation is true (hence the extreme values of their probability judgements). This seems to be quite similar to the engineer / lawyer story: if the received signal is part of a more general picture of a successful firm (a young and dynamic manager, a new high-tech product, a restructuration plan etc...), the predicted future earnings will be high, and, as shown by De Bondt & Thaler, higher than the actual ones. Thus, overreaction and high volatility can be generated by the fact that individuals acting in the market use a heuristic akin to representativeness.

Overconfidence and positive feedback trading

As it expresses itself in extreme judgements, overreaction often entails "unrealistic optimism" -or pessimism (De Bondt & Thaler, 1990). But there is more to unrealistic optimism, and to overconfidence in general, than a neglect of prior information. Studies of overconfidence such as the ones exposed before show that overconfidence is produced simultaneously by an insufficient attention to the quality of the evidence and an overestimation of its strength. This pervasive bias is part of the mechanisms by which a speculative bubble both appears and bursts.

Let us briefly recall what a bubble is. For most economic historians and practioneers, a speculative bubble is a cumulative deviation of an asset price from its believed "normal" or equilibrium value; it is characterized by a slow but accelerating rise of the price followed by a crash. In this view, psychological factors such as "euphoria", crowd behaviour and contagion are emphasized. The initiation of a bubble usually comes from a "displacement", i.e. "some outside event that changes horizons, expectations, profit opportunities" and induces confidence (Kindleberger, 1978, p. 41). Then, "optimism sets in. Confident expectations of a steady stream of prosperity, and of gross profits, make portfolio plunging more appealing. Financial institutions accept liability structures that decrease liquidity, and that in a more sober climate they would have rejected. The rise is under way, and may feed on itself until it constitutes a mania" (*ibid.*, pp. 29-30). A speculative "mania" or "euphoria" has a positive feedback dynamics: the more people have confidence in the value of an asset, the more they buy it, the more its price increases and the more they believe their optimism is justified. At this point, whatever the information, agents interpret it in a way which is consistent with their expectations.

Examples of such a cognitive dissonance are numerous (see the changing interpretations of exchange rates movements witnessed by financial newspapers). But cognitve dissonance is not the whole story; I will try to show how representativeness can also be invoked to give an account of speculative euphoria.

Speculative mania are inevitably followed by a brutal decrease in prices: this is the bursting of the bubble, or the crash. Suddenly, the previous confidence cracks down, for reasons that are not really well-understood. An information or a rumor casts some doubts about the reasons to be optimistic; once agents are starting to be diffident and suspicious, everything is set up for a panic. Some agents will feel uncomfortable with their excessive exposure and start to sell, the price will go down, other agents will follow, and the panic will "feed on itself" the same way the mania did. This pattern is similar to bank runs and other self-fulfilling prophecies described by Merton, 1936.

Both phases of a speculative bubble, the euphoria and the panic, are characterized by positive feedback dynamics on one hand, and overconfidence on the other (both phenomena are in fact related)[12]. Indeed, as Keynes, 1936, ch. 12, forcefully argued, these dynamics are inherent to the functioning of a speculative market. The main reason for this type of dynamics has to be searched in the way individual and collective beliefs are formed. As shown by historical accounts and expectations surveys, and unlike what is postulated by most economic theorists, traders and other market's practitioneers are well aware that price dynamics is mainly generated by the collective opinion they, as a group, have on the future evolution of the price. The main activity of a speculator is to guess what the average opinion will be, i.e. what are the dominant representations of others and how they will use them in interpreting new information. Hence, the specific representations and models of the world agents have are of primary importance in financial price dynamics, and the way these models are formed and evolve strongly determines this dynamics.

Market surveys, especially the one made by Shiller, 1989, just after the October 1987 crash, teach us quite a lot about the kind of representations agents seem to use and the way they use them. Actually, if as claimed here, people mainly construct their judgements by representativeness, the specificity of these models can give some insights about the typical categories and similarity criteria shared by financial operators which condition their beliefs[13].

During the year 1986, Shiller sent questionnaires to private and institutional investors each time there was a sharp move in the Dow Jones. Answers were sometimes difficult to interpret, but slowly, Shiller modified the formulation of the questions, and in October 1987, he was ready with a questionnaire he sent to about 2000 investors -private and institutional- on Tuesday the 20th, just after the Dow Jones had dropped of 508 points, i.e. 22, 6%, on Monday the 19th. Among the many questions that were asked, some were pertaining to the beliefs, the intuitions, and the "visions of the world" those investors had on this day and the week before the crash, and what they expected for the next future. The answers to this survey may, for our purpose, be separated into two groups: the first one concerns the general representations agents have about the functioning of the market and the mechanisms underlying speculative crisis (i); the second one has to do with the way individuals judge the evidence at hand and how they are prone to overconfidence (ii).

(i) General representations concerning market functioning

The difference between practitioneers' and theorists' view of financial markets functioning is quite striking. The broad picture which appears from Shiller's survey makes clear that, for traders and other operators, market psychology is central in the determination of price dynamics. One of the question asked on the 20th of October was: "Which of the following better describes your theory about the declines: a theory about investor psychology or a theory about fundamentals such as profits or interest rates?" (Shiller, 1989, p. 394). A large majority of private and institutional investors (around 65%) chose the "market psychology" explanation. One of the main features of this "psychology" is the tendency to be affected by the behaviours of the others; however, this phenomenon of mutual influences and imitation is not completly irrational: as people know that their collective behaviour is at the origin of price movements, what the others are thinking and doing is of first importance. During the October 1987 crisis, investors reported having talked to others about what was going on "very many times" or "continuously", i.e. interpersonal communication was high (*ibid.*, note 5). "Moreover, 23% of the individual investors and 40.2% of the institutional investors reported experiencing a contagion of fear from other investors" (p. 388). Hence, people seem to be well aware of the positive feedback process characterizing price dynamics.

Another quite striking feature of many answers is the frequent reference to the 1929 crisis, showing that the memory of this crash is still vivid half a century later although of course today's investors have not experienced it personally. Both crashes happened a Monday, "after a preceding week of great market turmoil" and many people reported having symptoms of anxiety that were building up through the week-end (Shiller, 1989, p. 388 and p. 393). Thus, on Monday morning, the possibility of a 1929 type crash was judged high and, in such a context of defiance, whatever small event or unimportant news could produce a panic. As Shiller puts it, "Investors had expectations before the 1987 crash that something like a 1929 crash was a possibility, and *comparisons with 1929 were an integral part of the phenomenon.* It would be wrong to think that the crash could be understood without reference to the expectations engendered by this historical comparison. In a sense, many people were playing out an event again that they knew well" (p. 399, my emphasis). The 1929 crash appears to be, at least in troubled times, the typical representation of a financial crisis, the model of reference by which agents construct their judgements. As noted by Shiller, in 1987, agents compared the situation they were going through with the 1929 crisis, and the more they felt it was similar, the higher the probability judgement they formed about the eventuality of a crash, whatever the low priors of such a dramatic event and the weight of the evidence (often being rumours, second or third-hand information etc...). These pessimistic beliefs were self-fulfilling because many people started to sell on such judgements, and the crash happened.

Of course, it is impossible to be sure that this story exactly describes what happened in October 1987, we do not have enough data about the general representations individuals have concerning market functioning and speculative phenomena. Moreover, if the notion of representativeness is precisely defined from a theoretical point of view, its empirical manifestations in real life situations are much more complicated to trace down. Actually, we do not know if similarity criteria are identical, and it may well be that different individuals, having different cultural and educational backgrounds, and different sources of information, also have different models of what a financial crisis is and what 1929 was about[14]. From the little data we have, it seems nonetheless that some basic features of the representation of market functioning, namely, the role of market psychology and of positive feedback trading, are shared by the majority of financial operators (historical accounts of speculation tell us the same kind of story, see for instance Mackay, 1841, Kindleberger, 1978, and Galbraith, 1990). Because agents use heuristics like representativeness to construct their beliefs, the pecularities of these representations and models could shape financial dynamics.

(ii) The role of overconfidence in speculative euphoria and crashes

After having exposed some general background representations, let us now turn to the mechanisms by which judgement by representativeness may produce euphoria and crashes. Note that these mechanisms are very similar: bubbles and crashes have in common the same feedback processes, and the later are possible because people share the view that asset price dynamics is mainly determined by the so-called market psychology.

In section II, we saw that overconfidence may be an outcome of judgement by representativeness. By focusing mostly on the magnitude of events, with little attention to the quality of evidence, people generally give a lot of weight to some salient traits of events whereas they do not take into account "statistical" data. This appears very clearly in the answers to Shiller's questionnaire. As long as the price is rising, such a rise is interpreted as a sign that markets are optimistic or bullish, even though some agents may be thinking that the asset is overvalued. Either people believe that such an increase is justified by fundamentals (they seem to be more numerous at the begining of the euphoria than just before the crash), or they know that the asset is overvalued but they follow the trend because they understand the role of market psychology. In both cases, investors display some kind of overconfidence: in the first one because they interpret information at face value, without taking into account prior information; in the second one because they are sure they can predict the turning point and "beat the market"[15]. To the question "Did you have a sense just before the crash (around October, 12, 1987) that the market was overpriced relative to fundamental value? (Try hard to remember what you thought *then*)" (Shiller, p. 391), 71.7% of the private investors and 84.3% of the institutional ones answered yes. At the same time, between September, 12 and October, 12, among the net buyers, 68.1% of the private investors and 93.1% of the institutional investors recognized that the market was overpriced (p. 392). Thus, agents were following the trend, which implies that their were confident enough in their ability to "time the crash", i.e. to go out of the market before the others. The discrepancy between private and institutional investors evaluations could be interpreted in the same way: institutional investors were more overconfident than the private ones, because they felt that they, as "experts", could not be wrong (on expertise and overconfidence, see Griffin & Tversky, 1992).

This feature is also patent in the behaviour of foreign exchange traders analysed by Frankel & Froot, 1987, during the bubble on the dollar in the mid 80's: they had short-term "extrapolative" expectations, they followed the trend, but judged the dollar overpriced in a long-term, fundamentalist view. As for October 1987, it seems that people knew the market was overvalued, but the rise of the dollar continued nonetheless for more than two years (until February 1985), even though all fundamental information would, if it had been taken into account, have pushed the dollar down.

The kind of overconfident behaviour just described is impossible to disentangle from positive feedback trading. It is because each individual believes he is more intelligent than his collegues that he follows the trend, being sure he will be able to "jump off the train" before the crash. When generalized, this behaviour engender a self-referential functioning of the market, in which whatever judgement can be validated ex post if enough people share it. Contagion is thus another important part of the story, but this goes beyond the scope of this paper (see Orléan's chapter 6 in this volume).

As just shown, positive feedback trading may produce euphoria, but it also plays an important part in crashes and in their timing. Indeed, one of the standing puzzles in financial economics concerns the factors at the origin of market's turning points. In other words, what causes the price, after a long and steady rise, to collapse brutally? As usual in finance and in the absence of a more compelling explanation, "news" are called in. When the overvaluation is patent and anyone is more or less aware of it, the confidence in the future vanishes, everybody now believe that it has been "too long too far" for the price not to go back to its equilibrium level. In such a situation, whatever news that can be interpreted as a sign that "markets become normal again" may produce the crash.

The answers to Shiller's questionnaire cast some doubt about this news theory. He explicitely asked to investors to rate different news by order of the importance they had for them on October, 19th (on this precise day only, not after the crash). Among these news, there was for instance "trade deficit figures announced Wenesday, October, 14, 1987", "producer price index figures announced Friday, October 16, 1987" and "the 200 points drop in the Dow the morning of Monday, October 19" (Shiller, 1989, p. 384). What is striking in the answers is that news about exogenous events are not deemed very important, whereas the most important news are price movements themselves. The 200 points drop in the Dow Jones comes in first position, followed by the price drops of the week before. The major news that provoked the decline of stock prices is the fall of prices itself. Shiller also refers to a survey done in 1946, just after a 6% drop of stock prices, and where the first reason invoked by investors was, there also, the decline in prices the same morning. Hence, it seems that the most important information for prediction is contained in price changes. Note that this is different from efficiency market theory in that here, the change in price is interpreted as a signal of further change, whereas in efficiency market theory, price is unpredictable, and whatever could be interpreted is already contained in price changes.

The fact that agents are giving a great importance to the magnitude of some events (here, the percentage of price decline), to the detriment of the quality of the evidence at hand (e.g. the priors: "how many times such a decline provoked a crash?", or the source of the information and its credibility: "is it a rumour or was it published by the Wall Street Journal?") is inherent to positive feedback trading. The more the price changes, the more it is expected to change (because of the overestimation due to the strenght of the evidence), and the more this judgement is made with confidence (when the strenght of the argument is high, people judge by similarity and forget the weight of their information, its precision). Overconfidence and more generally representativeness thus seems to matter on speculative markets.

IV Representativeness and categorization

The preceding sections have highlighted the role of the representativeness heuristic on financial, speculative markets. The so-called anomalies of these markets are not, in the perspective developed here, anomalies, but clues that individual decision making does not proceed as postulated by the rational / Bayesian approach. Indeed, representativeness is not a pathology, but a learning heuristic which helps individuals to form beliefs about the occurrence of uncertain events in a non-transparent and always changing environment. When problems are ill-defined, it is easier to assess a "probability" by comparing the closeness of the possible event to a typical case than to use Bayes' rule to revise a probability judgement. Of course, beliefs obtained by judging by closeness and similarity have no reason to follow the axioms of probability theory. From the point of view of the rational / Bayesian approach, they are systematically biased. In real-life situations however, they seem to be quite intuitive and robust procedures to form beliefs. Actually, representativeness suggests a stable procedure of how to construct judgements in uncertain and complex environments: the "probability" of the occurrence of an event depends on the similarity of this event with a typical event which is part of the knowledge base of the agent. The study of representativeness thus leads to a study of what kind of events and situations are considered "typical" by individuals. This cannot be done without defining precisely what is a typical situation and what are the mechanisms by which agents build such typical categories. More generally, pursuing in this line entails at least a gross understanding of how agents form and evolve representations of their ever-changing environment.

Such a study about the role of heuristics on financial markets is thus part of a more general research concerning the microfoundations of financial dynamics. It is our belief that making progress in modelling financial markets first implies to radically reconsider the way individual decision making is taken care of in financial models. More precisely, one needs a representation of microbehaviours which would be able to give an account, at least, as to why people use heuristics. In order to do that, one needs to go back to "first principles" about the mechanisms of categorization and induction. What follows is thus more a research agenda where I will try to underline some of the basic issues that should be tackled, than an achieved theory of representativeness.

There are basically two polar views of categories and concepts (e.g. Lakoff, 1987, Kleiber, 1990). In the first approach, the "objectivist" one, categories are defined by a list of their qualities, i.e. a set of necessary and sufficient conditions: "a dog is hairy, has four legs, sweet eyes, is moving its tail each time it is happy, etc..." An encounter with an agressive dog or a Mexican hairless dog contradicts the whole categorie of a dog, and the "object" is impossible to identify. One of the main problems with this objectivist view is that it cannot cope with the variety of the real world: either the object exactly corresponds to the known category, or it is discarded as unknown. In this perspective, similarity between two objects or events is a notion that does not exist: either it is a dog, with all the well-known properties of dogs, or it is something else. There is no possible measure of resemblance, of some fuzziness of the basic category which would allow to account for diversity; it is either all or nothing.

Following the work of Rosch in the 70's (e.g. Rosch, 1975), another approach to categories has been developed. In this view, categories are formed by some prototypes accepting variations or defaults. For instance, every individual forms a prototype of what a dog looks like, (it is generally hairy and has sweet eyes etc..), but it exists hairless dogs, agressive dogs, and the like. A category is thus defined by a prototype assorted of a set of defaults, i.e. exeptions which do not contradict the category of a dog. The procedure of categorization here is the same as in the lawyer / engineer experiment described in section II, where each individual seems to have models of lawyers and engineers which are representative instances of these classes of "objects", and the belonging to a class, namely categorization, is judged by closeness and similarity of the encountered object with the representative instance of the class. As lawyers are different one from another, and more generally as there is variation inside a class, categorization must allow for some diversity while preserving basic or essential features which make the specificity of the class. This flexibility is obtained by the defaults that are judged acceptable, namely the characteristics of the object or the event which can vary without questionning the standing classification. These defaults are usually organized in a pyramidal or hierarchical way. For instance, the category "dog" is first defined by some basic features (hairy mammal, sweet eyes...) and progressively refined because of encounters with hairless, agressive, big or small, yellow or brown... dogs. These operations of specification consist in additioning defaults to the prototype in a hierarchical way: first, dogs are hairy mammals (but it exists some hairless dogs); then, if they are hairy, they can be either yellow or brown, and so on. Theoretically, there is no limit to this kind of specification procedure. This organization of defaults is labelled a "default hierarchy", and allow individuals to form and test hypotheses about the environment[16]

Such a method of categorization is more flexible and more global than the one built upon necessary and sufficient conditions; individuals can form categories of varied, changing, and possibly ill-defined objects and events. Experiments on the acquisition of concepts by children tend to confirm that they use something like prototypes to classify objects, which could explain their ability to recognize objects very early, eventhough they do not know all their necessary and sufficient conditions (see Johnson-Laird, 1983, pp.121-30).

In order to precise the relations between the representativeness heuristic and prototypical categorization, as well as some theoretical implications of these relations, let us briefly make some remarks about the nature of prototypes (i), and some of their properties (ii).

(i) The nature of prototypes: typical features and family resemblance

The prototypical approach to categories emphasizes their fuzziness, the fact that "they have no clear-cut boundaries...[but] 'blurred edges'" (Johnson-Laird, 1983, p. 189), which is why they can accomodate diversity. Hence, prototypes are not defined by their essence, but by "schemata of their most characteristic members" (*ibid.*, p. 190). In Johnson-Laird's terminology, a schema is "a model that underlies the ability to form an image"; thus, a category can be represented by a "schema that specifies, *not* a set of necessary and sufficient conditions, but the typical or default characteristics of the items it subsumes" (*ibid.*, p 190). In other words, a prototype, as defined here, is a model of an object or an event that embodies their more typical features. For instance, saying that 1929 is a prototype of a financial crisis refers to a model embodying typical features like a collapse in asset prices after a long period of euphoria, individuals and institutions ruined in two days, banks and firms going bankrupt, and popular images such as a panicking crowd at the New-York Stock Exchange, queues of people at the doors of banks, and a bit later, a pervasive misery.

The difficulty that is instantaneously coming to mind concerns the way to define "the more typical features" of an object and "the most characteristic members" of a class. As we saw in section II, this is still quite feasible in pure chance settings as the heads and tails game, where the typical features of, say, a fair draw, are determined by probability theory. In these cases, exact representativeness has a precise meaning. However, the task becomes much more complicated in real-life situations where the simplests objects we have to deal with are dogs, houses and cities. Here, representativeness cannot be exact, i.e. there is not one single way to define what is typical. Indeed, typical features are by definition different from necessary and sufficient conditions: they are sufficient but not necessary. For instance, most of the birds have feathers and are flying, hence feathers and the ability to fly can be considered as typical attributes of birds eventhough some birds do not fly and some non-birds do. In the standard approach to prototypical categorization, typical attributes will be chosen by what is labelled their *cue validity*, i.e. their degree of predictability. An attribute like feathers has a high cue validity because it is very frequent in birds and almost absent in all other categories. On the contrary, the feature "having legs" does not help much in categorizing an animal because a lot of different animals have legs: such an attribute has a low cue validity, it is not a typical feature[17]. A prototype (or a stereotype, those two terms being often used as synonyms) is thus a "best example" of a given category, it contains more typical attributes than any other object of the class (e.g. a sparrow is often considered as a prototype of a bird, whereas this never happens to the penguin). In this view, judgement by representativeness is a judgement of the *degree* of prototypicality of the object considered.

The definition of similarity as closeness to a prototype goes back to Wittgenstein and his notion of family resemblance, i.e "the idea that members of a category may be related to one another without all members having any properties in common that define the category" (Lakoff, 1987, p. 12). Trying to identify what kind of commonality may ground the definition of a concept, Wittgenstein notes: "we see a complicated network of similarities overlapping and criss-crossing: sometimes overall similarities, sometimes similarities of detail" (quoted by Johnson-Laird, 1983, p. 189). Hence, as hinted to by the results of the experiments exposed in section II, similarity occurs at different levels, has multiple meanings, and it seems impossible to define one single criterion of similarity that would allow to apprehend representativeness, at a theroretical level as well as in its empirical manifestations.

(ii) Properties of prototypical categories: embeddedness and path-dependence

There are many debates about the possible origin of the criteria on which concepts are built. The nature of concepts and of meaning is related to the origin of these criteria, and the many questions underlying this issue are reccurent in the history of philosophy. They certainly lie outside the scope of this paper, but let us, for our purpose here, briefly recall the general terms of the debate. A long tradition in philosophy, from Plato to Frege, claim that concepts exist outside human mind, and that, to put it grossly, they can be defined by a set of objective, necessary and sufficient conditions. The criteria by which to judge a dog are well-defined and reflect the intrinsic properties of dogs, something like their essence. Of course, humans can only have a subjective idea, a partial knowledge about dogs, but the concept of a dog exists nonetheless outside individuals. This realist approach (to which corresponds the objectivist view of categories evoked earlier) is in opposition with a psychologist view of concepts and meaning emphasizing the fact that criteria are social constructs, conventions in the terminology of Wittgenstein, and depend among others on the structure of the mind. These conventions are wide societal agreement allowing individuals to communicate, to be sure that when they say "dog", they mean the same hairy animal who has the ability to bark and moves its tail when happy[18]. Certainly, as these fundamental similarity criteria do not completly describe a dog, many defaults and exceptions will slowly be added to the prototype by inductive mechanisms, and the concept will become richer with every encounter with a dog, i.e. with experience (for a more detailed account of this debate, see Johnson-Laird, 1983, ch. 9).

In this perspective, criteria are images and models that are embedded in collective, social representations, and as such, contingent to the specific society and epoch in which they emerge. Let us go back once more to the lawyer / engineer experiment. Similarity criteria in this case are impossible to define at an abstract level: a lawyer cannot be define as a set of necessary and sufficient conditions, but one has to pin down more or less "intuitively" what are the basic characteristics of lawyers. To perform this operation, we usually mobilize explicit and implicit knowledge pertaining to the usual task of lawyers in our societies; lawyers as a social group; images of known lawyers (friends or movie stars); the last "big legal affair" which was everyday in the newspapers and of which Mr X., lawyer, was the outstanding public figure; and so on. The articulation of this knowledge forms a mental model of a typical lawyer. The same kind of process may be spotted behind Shiller's traders' answers when they evoked 1929 as a reference they were using in making their predictions about what was next. Again, even if at a theoretical level it is possible to define a financial crisis as a set of necessary and sufficient conditions (indeed, it is our task as scientists), "intuitive" predictions do not seem to be built on such conditions. Instead, the traders and analysts interrogated by Shiller displayed they were using a typical case of a financial crisis, namely 1929, to construct their judgements; they were judging the actual situation by its similarity to what they knew about this typical case. In turn, their state of knowledge about financial crisis was mainly determined by what they learned at college, the books they read about the subject, stories friends and friends of friends could report concerning 1929 and other dramatic speculative episodes, and so on. In both the lawyer / engineer and the speculative stories, similarity criteria are shaped by the specific social history people, as a collectivity, have gone through, and by "culture" at large -with some marginal variations stemming from particular individual experiences. Because prototypes, at the heart of the processes of categorization, are mainly social constructs, it seems that it is possible to talk about an *embodied cognition* (e.g. Lakoff, 1987). This notion of an embodied cognition is quite close to the one of *embeddedness*, which has been developed by Granovetter, 1985, concerning the embeddedness of economic behaviours in social structures and culture.

A second property of prototypical categorization, related quite closely to the one of embeddedness, is path-dependence: categories at instant t depend on what they were before, i.e. on their history. In fact, a distinction has to ba made here. As just seen, fundamental criteria on which categories are built are mainly determined by something akin to social conventions, and as such, they depend on history because they are social constructions: the typical characteristics of an engineer today are different from what they were during Da Vinci's time. This dependence on history comes from the embodiement of cognition in social structures and in what is usually denominated as "culture". The process of prototypical categorization is path-dependent in a different way: it depends on history because its mechanisms are mainly inductive ones, and induction is a path-dependent process (e.g. Holland et al., 1986, and Johnson-Laird, 1991, and for a discussion in economics, Arthur, 1992). Remember that a prototypical category is composed by similarity criteria plus a default hierarchy, i.e. a complicated but nonetheless organized network of exceptions which do not contradict the category. Similarity criteria are tought by parents, learned at school or through television at an early age, and, as just said, they depend on history because they are embedded. But, at the level of one's individual experience, they are more or less given. Default hierarchies are path-dependent in a more direct way, because, as they are slowly built through experience, the precise manner in which they will evolve, their semantic content, is an outcome of the specific history of the individual, something like a summary of his or her individual history. My default hierarchy may include a default for Mexican hairless dogs because I encountered few of them, but other people may have totally different defaults hierarchies, depending on the kind of dogs they encountered and the experience they had with them.

Holland *et al.*, 1986, highlighted the fact that the mechanisms by which default hierarchies are constructed are mainly inductive ones. Different approaches to induction have been developed; the so-called syntaxic approaches are trying to define the formal conditions under which an induction is a valid inference. Bayesian revision of beliefs is one of such syntaxic apparatus; when one uses Baye's rule to take into account a new information, one can be sure that the resulting inference is valid. Unfortunately, even if one admits that it is possible to define a corpus of formal rules controling the validity of inductive inferences (the numerous debates in probability theory cast some doubt about this statement), it is not because an inference is valid "that it is of the slightest interest" (Holland *et al.*, 1986, p. 5). The many theoretical paradoxes of a purely syntaxic theory of induction can only be solved by taking into account the meaning of inferences, i.e. their semantics, and the goal individuals are pursuing in making these inferences (see Holland *et al.*, 1986; Johnson-Laird, 1983 and 1991; and Lakoff, 1987). In this perspective, relevant inferences depend on the aim, the task, the role, or whatever activity in which the individual is engaged, and the feasible inferences are determined *by what is already known* (you cannot build a category for "financial crisis" if you do not have some for "market' and "crowd behaviour").

The aim of this section was to highlight the basic cognitive mechanisms which are behind the use of the representativeness heuristic. To summarize, one can say that similarity and family resemblance are at the heart of such mechanisms of categorization; hence judgements by representativeness are easily understandable and do not imply that individuals are irrational. Moreover, one could draw general properties from the process of categorization, and, as representativeness stems from prototypical categorization, conjecture that judgement by representativeness has the same properties of embeddedness and path-dependence.

V Conclusion

This paper has tried to highlight the role of the representativeness heuristic on financial markets. As this alternative way of constructing judgements in complicated and evolving environments is much less demanding in terms of computational and cognitive abilities than Bayes' rule, it is not surprising that it should be pervasively used by agents, especially in environments such as financial markets. Let us summarize the main findings of this preliminary study:

* If individuals majoritarily use representativeness in forming their expectations on speculative markets, overreaction to new information and bubbles have not anymore to be

considered as "anomalies". These phenomena are the aggregate results of the interaction of agents following the "rules" of prototypical categorization.

* At a more micro level, representativeness is a clue that learning does not proceed as stated by Bayes' rule; instead, it hints at a radically different vision of learning and microbehaviours. Indeed, if, as it seems, representativeness is a learning heuristic -"a rule of how to learn new rules"-, prototypical categorization becomes a central learning mechanism, and its properties could be considered as more general learning properties. We saw that embeddedness and path-dependence of cognition typifies the inductive process of building categories and models. It follows that learning dynamics depends on the specific social structures the agent is embedded in, and the particular history he or she will have gone through. Such properties of learning processes may greatly affect price dynamics on financial markets, as the latter are basically "markets for representations" (like for instance the market for art) [19].

* Finally, many theoretical questions about categorization and inductive learning processes are still unsolved. For instance, how are similarity criteria formed? They are conventions, and as such, social constructs, but how do these conventions emerge? They may evolve through learning and collective experience, and maybe through individual experience too, but what are the precise mechanisms of their evolution? What are the links between individual and collective learning dynamics? Another related set of questions concerns the characterization of the dynamics of inductive learning. Categories are composed by a prototype and a hierarchy defaults, and organized to form mental or cognitive models. As we have seen, the evolution of these models is embedded and path-dependent, i.e. the specific trajectories individual cognition will follow depend on the environment of the agent and of his/her history. How can one define learning in such a context? Is it only an evolution of representations and models, whatever the "direction"? The notion of learning is usually impossible to dissociate from the one of amelioration, but when one abandons the reference to a "true" model of the world, the very definition of amelioration becomes problematic.

Notes

[1]The focus is here on individual beliefs; for studies of the formation of collective beliefs from individual interactions, see for instance Orléan, 1990, and ch. 6 in this volume. These two issues are nonetheless linked, and individual biases which are observed at the agent's level certainly have some impact at the collective level, although it is not much studied yet.

[2]The literature on rational learning is immense; for a good overview, see Frydman & Phelps (Eds.),1983.

[3]The terminology follows Nelson's "appreciative theorizing", see Nelson, 1995 (?).

[4]The term "subject" is, as Hey remarks, rather unfortunate, but "it is now part of the experimental folklore", (Hey, 1991, p. 4).

[5]As emphasized by Grether, 1980, commenting this experiment, for another proposed description, the median estimate of the probability that the person was an engineer was close to 0.05, and still for another one, it was roughly 0.95, and this, for both groups of subjects. The fact that the two groups made nearly identical estimates clearly shows that prior information was, if not ignored, at least grossly underweighted. A closely related argument is that the extremeness of these estimates (0.05 for one description and 0.95 for another) is in itself favourable to the representativeness hypothesis: these descriptions were stereotypes of lawyers and engineers, and the estimates were mainly determined by the closeness of the description with the stereotype, i.e. the strength of the evidence, whatever the base rates.

[6] This is the well-known "as if" argument Friedman, 1953, uses to show that speculation is stabilizing.

[7] For a discussion of the methodological issues attached to experimental economics, see for instance Hey, 1991.

[8]In this type of setting, there is no room for interpretation, and the only randomness individuls have to face is intrinsic to the game: the criticism that biases are due to interpretation of verbal descriptions cannot hold here.

[9]The same results have been obtained by Camerer, 1987, in a similar setting; priors are underweighted but not ignored, and there is evidence of what Camerer calls "exact representativeness", i.e. a neglect of priors when there is a perfect adequacy between the outcome of the draw and the content of the urn.

[10]Indeed, in this experiment, the difficulty comes from the inference one can draw from informations like "John likes mathematical puzzles". Such an information could concern an engineer as well as a lawyer: thus, the two categories are collectively exhaustive but not mutually exclusive.

[11]Standard theory has recently developed the notion of a rational bubble, which is a transitory but cumulative deviation compatible with rational expectations. This quite peculiar phenomenon is obtained by exploiting the property of multiplicity of equilibria in rational expectations models, but its microfoundations (a rational representative agent) are difficult to reconciliate with empirical data. See for instance Blanchard & Watson, 1982, and, for a survey, Camerer, 1989.

[12]A similar relation between positive feedback and representativeness has been highlighted and tested on real estate markets by Clapp & Tirtiroglu, 1994.

[13]As just said, speculative dynamics are positive feedback dynamics where the dominant models or collective representations (in the sense of Durkheim) coordinate and "anchor" individual beliefs. The identification of such collective models is thus central to the comprehension of the particular processes of judgement formation which characterize financial markets.

[14]This issue of a possible heterogeneity of beliefs, which contradicts rational expectations and its "true" model of the world, is at the core of an important debate in financial economics, namely, the "no trade" paradox. The question is about why so much speculative trade occurs on financial markets if people are rational and not enough different for contradictory beliefs to exist in equilibrium. Seminal contributions are Milgrom & Stokey, 1982, and Tirole, 1982. Note that if beliefs are mainly formed by following heuristics such as representativeness, as beliefs come from categories and mental models, and in turn the

formation and evolution of these models are engendered by contextual and path-dependent processes, it is easy to give an account of heterogeneity (see section IV below).

[15]Galbraith, 1990, notes that investors, especially at the beginning of the mania, are encline to speculative euphoria because, first, they become rich or richer, and second, they believe their increased wealth is the sign of a "superior intelligence", a common belief on financial markets and elsewhere. This bias in attributing causality is labelled by psychologists the "fundamental attribution error", which consists in underestimating "the impact of situational factors and [in overestimating] the role of dispositional factors in controlling behaviour".(Ross & Anderson in KST, ch. 9, p. 135).

[16]For more details on default hierarchies, see Johnson-Laird, 1983, and Holland *et al.*, 1986. The latter also present tentative ways of modelling the emergence of such default hierarchies, using connexionnist networks. Following the methodology of "Artificial Life", simulation models have been built to study the emergence of aggregate phenomena from the interaction of agents modelled as connexionnist networks, i.e. agents endowed with such a "prototypical" process of categorization. For first applications of this type of modelling in economics, and particularly on financial markets, see Palmer et al. ,1994, and Marengo & Tordjman, 1996.

[17]There are many debates about this issue in linguistics and cognitive psychology, and it is not clear that cue validity is the best way to define typica traits (see for instance Lakoff, 1987). The "standard approach" just mentioned refers to the work of Rosch and her followers in the 70's.

[18]The study of the mechanisms by which such norms and conventions appear and perpetuate is outside the scope of this paper. See e.g. Lewis, 1969, and Orléan (Ed.), 1994.

[19]The consequences of these general properties of learning processes on financial dynamics have been studied elsewhere (see Marengo & Tordjman, 1996)

References

Arhur, W. B., 1992,"On Learning and Adaptation in the Economy", *Santa Fe Institute WP* # 92-07-038.

Arrow, K. J., 1982, "Risk Perception in Psychology and Economics", *Economic Inquiry*, 20 (1), pp. 1-9.

Bar-Hillel, M., 1982, "Studies of representativeness", in D. Kahneman, P. Slovic & A. Tversky, (Eds.), 1982, Ch. 5, pp. 68-83.

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

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

Camerer, C., 1989, "Bubbles and Fads in Asset Prices", *Journal of Economic Surveys*, 3 (1), pp.3-41.

Clapp, J. M. & D. Tirtiroglu, 1994, "Positive feedback trading and diffusion of asset price changes: Evidence from housing transaction", *Journal of Economic Behavior and Organization*, 24 (3), pp. 337-55.

De Bondt, W. F. M. & R. H. Thaler, 1990, "Do Security Analysts Overreact?", *American Economic Review*, 80 (2), pp. 52-57.

Durkheim, E, 1896, *Formes élémentaires de la vie religieuse*, Paris, Presses Universitaires de France.

Einhorn, H. J., 1980, "Learning from experience and suboptimal rules in decision making", in D. Kahneman, P. Slovic & A. Tversky, (Eds.), 1982, Ch. 19, pp. 268-83.

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

Froot, K. A & J. A. Frankel, 1989, "Forward Discount Bias: Is It an Exchange Risk Premium?", *Quarterly Journal of Economics*, 104 (1), pp. 139-62.

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

Galbraith, J. K., 1990, *A Short History of Financial Euphoria*, Paris, Editions du Seuil,1992, for the french translation.

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

Grether, D. M. , 1980, "Bayes rule as a descriptive model: the representativeness heuristic", *Quarterly Journal of Economics*, 95 (4), pp. 537-57.

Griffin, D. & A. Tversky, 1992, "The Weighing of Evidence and the Determinants of Confidence", *Cognitive Psychology*, 24 (3), pp. 411-35.

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

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.

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*

Kahneman D. & A. Tversky, 1972, "Subjective probability: A judgement by representativeness", in D. Kahneman, P. Slovic & A. Tversky, (Eds.), 1982, Ch. 3, pp. 32-47.

Kahneman D. & A. Tversky, 1973, "On the psychology of prediction", in D. Kahneman, P. Slovic & A. Tversky, (Eds.), 1982, Ch. 4, pp. 48-68.

Kahneman, D, P. Slovic & A. Tversky, (Eds.), 1982, *Judgement under uncertainty: Heuristics and biases*, Cambridge UK, Cambridge University Press.

Kahneman, D. & D. Lovallo, 1993, "Timid Choice and Bold Forecast : A Cognitive Perspective on Risk Taking, *Management Science*, 39 (1).

Keynes, J. M., 1936, *A General Theory of Employement, Interest, and Money*, Paris, Payot, 1969, for the second french edition.

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

Kleiber, G., 1990, *La sémantique du prototype*, Paris, Presses Universitaires de France;

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

Lane, D. A., 1994, "Artificial Worlds and Economics", *Journal of Evolutionary Economics*, ?

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

LeRoy, S. F., (1989), "Efficient Capital Markets and Martingales", *Journal of Economic Literature*", 27 (4), pp. 1583-1621.

Lewis, D., (1969), *Convention. A Philosophical Study*, Cambridge MA., Harvard University Press.

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

Marengo, L. & H. Tordjman, 1995, "Speculation, heterogeneity and learning: A model of exchange rates dynamics", *IIASA WP* # 95-17, forthcoming in *Kyklos*, 1996.

Merton, R. K., (1936), "The Self-Fulfilling Prophecy", in*Social Theory and Social Structure*, New-York, The Free Press, Second Edition (1968), pp. 475-90.

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

Nelson, R. R., (1995), "Recent Evolutionary Theorizing about Economic Change", *Journal of Economic Literature*, 33(1), pp. 48-90.

Orléan, A., (Ed.), 1994, *Analyse économique des conventions*, Paris, Presses Universitaires de France.

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.

Orléan, A., (1996), "The ambivalent role of mimetism in collective decentralized learning", Ch. 6 in this volume.

Palmer, R. G, W. B. Arthur, J. H. Holland, B. Le Baron and P. Tayler, 1994, "Artificial Economic Life: A Simple Model of a Stockmarket", *Physica D*, 75, pp.264-74.

Rosch, E., 1975, "Cognitive Representation of Semantic Categories", *Journal of Experimental Psychology*, 104, pp.192-233.

Ross, L. & C. A. Anderson, 1977, "Shortcomings in the attribution process: On the origins and maintenance of erroneous social assessments", in D. Kahneman, P. Slovic & A. Tversky, (Eds.), 1982, Ch. 9, pp. 129-52.

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

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

Tversky, A. & D. Kahneman, 1971, "Belief in the law of small numbers", in D. Kahneman, P. Slovic & A. Tversky, (Eds.), 1982, Ch. 2, pp. 23-31.

Tversky, A. & D. Kahneman, 1974, "Judgement under uncertainty: Heuristics and biases", in D. Kahneman, P. Slovic & A. Tversky, (Eds.), 1982, Ch. 1, pp. 3-20.

Tversky, A. & D. Kahneman, 1982, "Judgements of and by representativeness", in D. Kahneman, P. Slovic & A. Tversky, (Eds.), 1982, Ch. 6, pp. 84-98.