On Intelligent-Agent Based Analysis of Financial Markets

W. M. van den Bergh K. Boer A. de Bruin U. Kaymak J. Spronk
Erasmus University Rotterdam, Faculty of Economics
P.O. Box 1738, 3000 DR, Rotterdam, the Netherlands
{kboer, kaymak}@few.eur.nl

ABSTRACT: Agent-based computational economics acknowledges the distributed nature of trading in financial markets by modeling the markets as evolving systems of autonomous, interacting agents that correspond to the trading parties. Conventionally, the behavior of traders has been described mathematically, and the market system is analyzed at equilibrium conditions. The dynamics of price formation, however, is influenced by the large diversity in the cognitive structures of the traders (e.g. differences in decision making methods, interpretation of available information and learning capacity), their specific circumstances (e.g. attitude to risk, time horizon) and the organization of the specific market in which the traders operate (e.g. market microstructure). Therefore, we propose to study financial markets by using intelligent agents that have rich cognitive structures borrowed from artificial intelligence research for modeling their decision making behavior. This representation allows us to model the decision making behavior of agents in terms of algorithms, that can represent a more diverse set of behaviors than mathematical formulae only. We discuss the role of intelligent-agents in the analysis of financial markets and speculate on the type of agents that can be expected to be suitable for the analysis and simulation of financial markets. We elucidate our thoughts by exposing the outline of a research project that has started recently at our university. As a first step of our research project, we discuss a classification of adaptation that we proposed recently for agents in agent-based computational economics.

KEYWORDS: Intelligent agents, financial markets, agent-based economics, multi-agent systems.

1 INTRODUCTION

One of the pillars of current financial theory is the *efficient markets hypothesis (EMH)*, which assumes that the asset prices in the financial markets are based on rational expectations of traders who are trying to maximize their expected utility. As a consequence, past information can not help predicting future prices, and the markets are assumed to be free of internal dynamics of their own. In such a setting, mathematical models of price movements in the markets are based on the assumptions of homogeneity of traders (e.g. traders are rational), equilibrium analysis and the possession of all relevant information by the traders. Any deviations from these idealized conditions are considered to be exogenous effects or uncertainty (e.g. noise).

Despite the EMH view, it has been suggested that empirical evidence regarding the price patterns observed in financial markets seems to indicate that markets do have internal dynamics of their own. It has been proposed that various empirical patterns like volatility clustering, fat tails and speculative bubbles can be explained in terms of the trading activities in heterogeneous communities of boundedly rational agents [10]. In recent years, there has been much discussion about whether heterogeneity and behavioral aspects of the traders play an important role in financial markets. It has been observed that markets do not process new information instantaneously [4], and that the markets can overreact as a result of trader optimism or trader pessimism [7]. Furthermore, it has been shown that simple technical analysts may survive competition with rational traders under certain conditions [6]. These results agree with the general consensus amongst practitioners in the financial markets that the market returns are predictable, although the predictability varies over time [16].

One promising way of analyzing the internal dynamics of financial markets is the use of an agent-based approach. In recent years, the agent-based approach to economical and financial analysis has grown into an important research field for developing an understanding of complex patterns and phenomena that are observed in economic systems. The agent-based approach to the analysis of financial systems models financial markets as evolutionary systems of competing boundedly rational agents. The agents adapt, learn and evolve in order to remain successful in their competition with other agents.

Analysis of these agent-based systems has shown that adaptation and evolution can be used to explain the workings of financial markets [11].

Often, the behavior of traders is described, in agent-based computational economics, by using mathematical formulae. This representation makes it possible to investigate the mathematical properties of the agent-based system. When there are a few types of traders (and hence a few types of agents), equilibrium conditions, population dynamics and steady-state solutions can be investigated. However, the traders in financial markets show a large degree of variation in terms of their perception, interpretation of available information, belief structure, attitude to risk, time horizon and decision rules. The dynamics of price formation in the markets is influenced by this variation of traders as well as by the organization of the specific market in which the traders operate. The difference in decision rules, perception and belief of traders can better be represented in an algorithmic framework. In this representation, the agents, which correspond to traders, manipulate available information according to a combination of decision logic, mathematical formulae and context-dependent algorithms that describe the trader's behavior. Both normative and descriptive approaches to traders' decision making can be studied, depending on the type of traders studied and the goals of modeling. The algorithmic representation allows one to study a more diverse set of behaviors than by mathematical description only. Hence, it becomes possible to study price dynamics in a more diverse environment that can be closer to reality, at the cost of more computations.

Intelligent agents possess rich cognitive structures borrowed from artificial intelligence research in order to model and implement a more diverse set of behaviors. Intelligent agents are by definition capable of adapting, learning and evolving, which makes them especially suitable for analysis of financial markets, since adaptation and learning in heterogeneous agent communities are known to be important for describing the behavior of financial markets [11]. We propose to study financial markets by using intelligent agents. In this paper, we sketch the outline of a recent research project that has started at our university regarding intelligent-agent based analysis of financial markets. We consider the markets as information processing entities and discuss the information diffusion process at the markets, which leads to price developments. We highlight the role of adaptation in computational economics and discuss the properties of intelligent agents for computational economics. As a first step of our research project, we propose a classification of adaptation for intelligent agents in agent-based computational economics. We expect this classification to help us understand and model the wide range of adaptive behavior that can be observed in financial markets.

The outline of the paper is as follows. In Section 2, we consider information processing at financial markets and discuss how information diffuses, to eventually lead to the establishment of the actual price. Key components of the information diffusion process are identified. In Section 3, we discuss the role of adaptation in computational economics. The properties of intelligent agents for computational economics are discussed in Section 4. A generic architecture for studying intelligent agents in agent-based economics is given. We expose in Section 5 a taxonomy of agent adaptation that we have proposed recently for studying adaptation in agent-based computational economics [2]. Various mechanisms through which adaptation may occur are considered in Section 6. Finally, we give our conclusions in Section 7.

2 INFORMATION DIFFUSION IN MARKETS

Price formation in financial markets is a result of the interaction between the market mechanism, the trading decisions of multiple parties and the external developments influencing the markets. The trading decisions of the traders depend on the available information. The resulting trade actions determine the price developments. The price developments enter the decision making system again as new data. Figure 1 depicts the general framework for the information diffusion process at financial markets. Actual developments in firms, markets, (macro) economic conditions and other investment opportunities provide new information that enters a pool of available information and data in the market. Usually, only part of all the information is available, or it may be available to only a group of economic agents that select it in different ways. This is effectively a filtering step. The traders in the markets (i.e. different agents) perceive and interpret the filtered information in different ways. Depending on their individual goals, preferences and their interpretation of the available information, the agents take different trading actions, which lead to the price developments through the interaction at the market. The market mechanism and its internal structure determines the nature of this interaction. The price developments are in turn available as further information in the market, upon which the traders may modify their trading decisions.

The participants in the financial markets have different attitudes to risk, different time horizons and different motivations leading to different reactions to unexpected (or even expected) news. Intelligent agents can be used to model these differences, as well as the learning and the adaptation behavior of the traders. It becomes then possible to study explicitly the influence of dynamics at a microscopic level, and to study how these dynamics eventually lead to the macroscopic behavior observed at the markets. In general, one needs to study the information diffusion at the markets and the expectation formation by the agents. The overall process can be studied by varying four main factors.

1. **External environment.** Different types of actual developments can be considered. Smooth changes (e.g. firms with different growth rates) can be studied as well as abrupt changes, shocks and crises. The uncertainty in the

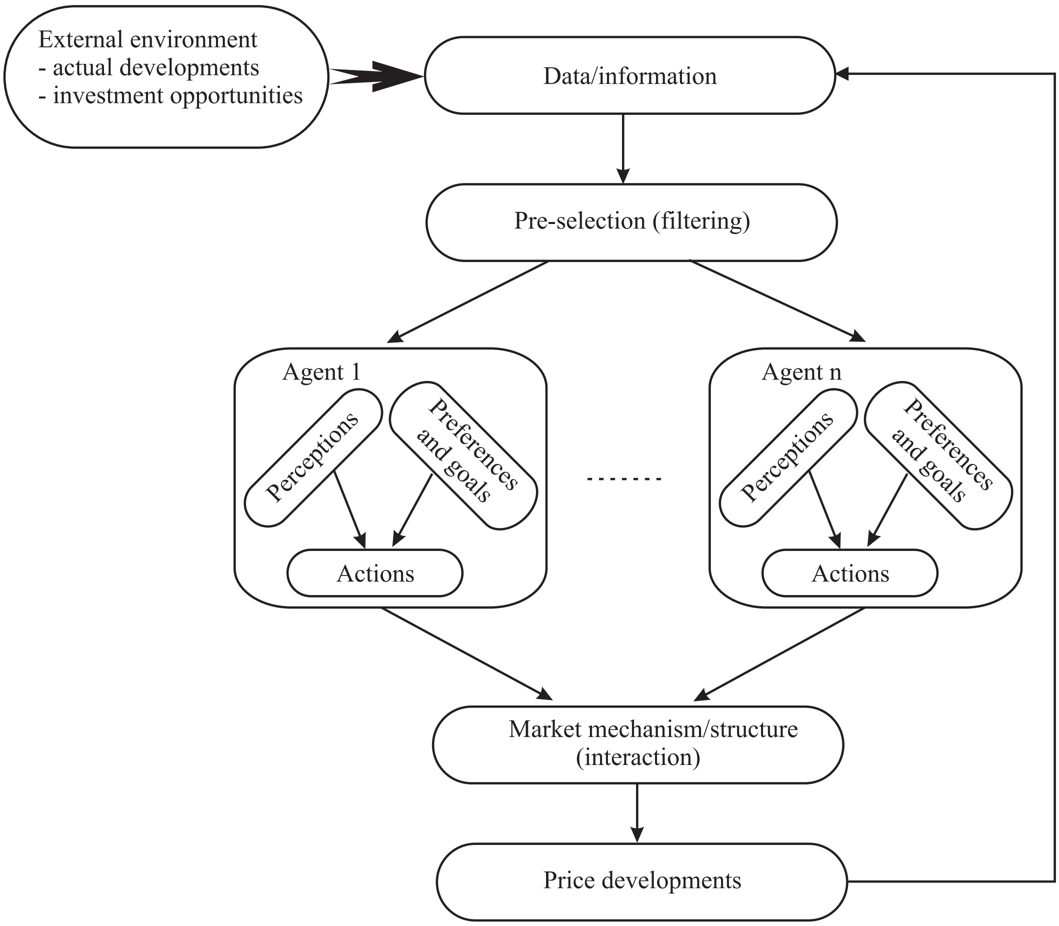

Figure 1: Framework for information diffusion process at financial markets.

external environment can also be considered in a stochastic or possibilistic setting.

2. **Data selection.** The pre-selection of data can lead to different information transitions. One can consider the case where all the agents are the same, i.e. they perform the same pre-selection, or the case where there are multiple groups of agents accessing different data.

3. **Properties of agents.** Agents with different perceptive and reasoning properties can be used, leading to different actions in the marketplace. The price developments will then be influenced by the resultant of the individual actions.

4. **Market mechanisms.** Different types of market mechanisms, market micro-structure and organization can be studied.

In the proposed framework, one can observe the system in controlled experiments, where only one of the factors is modified, and hence one can assess the validity of various hypotheses. For instance, the EMH implies that the feedback loop from the price development to the data pool in Fig. 1 does not contain any information that can be used in an advantageous manner. This hypothesis can be tested by using information theoretic measures, or by observing whether successful agents use that information to their advantage. Similarly, the influence of the agents' adaptation strategies can be studied. For example, what happens if agents mimic the actions of the most successful agents? Such an experiment could help to determine whether evolutionary selection pressure (selection by external criteria) is a necessary component of financial markets as often assumed in agent-based computational economics (as opposed to voluntary learning based on agent's internal criteria). By carefully analyzing the effects of different choices for the four components above, a large number of hypotheses regarding the financial markets can be studied. The market patterns that arise as a result of these choices can be tested against various conditions such as the EMH or the no-arbitrage condition. Some of the questions that could be addressed are the following.

- Can self-organization at financial markets emerge through the actions of boundedly rational, utility maximizing agents?
- Can empirical patterns such as volatility clustering be explained in terms of self-organization and reinforcement?
- How do the properties of agents influence the selection of successful trading strategies in the face of co-adaptation and co-evolution?

As this brief discussion shows, a key part of our research project is trying to understand how agent learning and adaptation influences the dynamics of price formation at the markets. In the next section, we review briefly the literature on the role of adaptation in agent-based computational economics.

3 ADAPTATION IN COMPUTATIONAL ECONOMICS

Economical systems such as financial markets can be considered as multiple auction environments, where the market mechanisms and the interaction amongst various market parties help compute a price for the traded asset. Recent research has focused on the computational properties of the markets themselves, by considering what computations various types of markets make, and how the markets can be designed to perform specified computations such as the minimum asset price [20]. Hence, a market can be viewed as an elaborate machinery that performs computation (e.g. price determination) in a distributed manner. Agent-based computational economics studies such economical systems by modeling them as evolving systems of autonomous, interacting agents [23]. There is a wide recent literature focusing on this area [12, 13, 19, 15]. A number of computational simulation laboraties have also been developed for agent-based economics [1, 14].
Adaptation, learning and evolution are important in the working of economical systems such as financial markets. Financial theory, for example, predicts that the markets are efficient [9], but anomalies are known to exist [5]. There is discussion on whether these anomalies are due to endogenous market processes, due to imperfections in the market or due to an institutional setting. Various empirical patterns in financial markets, such as volatility clustering, fat tails and speculative bubbles have been related to the interaction of heterogeneous, adaptive agents [11]. Agent-based computational economics has focused mainly on evolutionary adaptation of agents in large communities (see, e.g. [13, 12]). In many studies regarding agent-based computational economics, the agents learn winning strategies through evolution [12]. For example, the agents in a simulated market may possess different trading rules and ways of updating these rules. It is assumed that unsuccessful agents will loose money, and hence prove to be 'less fit to survive' and so will be replaced by agents that use more successful trading rules. This analysis, however, fails to acknowledge the internal success criteria for the agents, as well as the differences in their attitude to risk, time horizon, etc. Furthermore, the effects of different

adaptation strategies and mechanisms that the agents can employ are also ignored. Therefore, a more general consideration of adaptation in agent-based systems is useful to obtain a full understanding of the possible adaptive properties of the agents and their influence on the overall behavior of financial markets. Individual learning and social behavior of agents are expected to have a large impact on the emergence of various effects like self-organization, emergent patterns and the formation of belief systems. In order to study these more complicated forms of interaction, the agents themselves must possess a rich cognitive structure, be more sophisticated and have a large degree of autonomy and "intelligence." It is thus worth to study the interaction of agents in societies of intelligent agents, leading to the investigation of smart adaptive agents.

4 INTELLIGENT AGENTS IN COMPUTATIONAL ECONOMICS

An agent-based system involves one or more agents operating to meet their design objectives. Russell and Norvig define an agent very broadly as anything that can be viewed as perceiving its environment (through sensors) and acting upon that environment (through effectors) [17]. Assume that the agent's environment can be characterized as a set of *environment states* $S = \{s_1, s_2, \ldots\}$ that the agent can influence only partially. The influence of the agent is effected through a set $A = \{a_1, a_2, \ldots\}$ of actions that the agent can perform. The agent can then be viewed as a function $action : S \to A$ that maps environment states to actions [25]. The (possibly non-deterministic) behavior of the environment can also be modeled as a function $env : S \times A \to \mathcal{P}(S)$, which maps the current environment state and the action of the agent into a set of environment states. The range of the env function is always a singleton in case the environment is deterministic. Note that this definition of an agent completely parallels the definition of a decision maker in a decision environment. Hence, one could argue that this is a decision-theoretic approach to agents.
Typically, the agents observe the environment states only partially. Therefore, the agent's actions will depend on a set P of percepts, which consists of a subset of the environment states and quantities that can be derived from the environment states. In this representation, the agent can not possess all the information about its environment. Although the agents may act to obtain missing information through communication with other agents (e.g. buying some piece of information), and they may possess some information about the market mechanism (e.g. the amount of transactions costs to be made), no single party can possess all the information at the time needed. An implication of this property is that the agents must necessarily be boundedly rational. Their actions can only be rational within the boundaries of available information and the time available for taking the decisions. This is not an unrealistic assumption. Traders in a financial market do not possess all information either, and they must react within limited time in order to take advantage of the developments in the markets. Although one could compare the response of rational agents to the response of boundedly rational agents, we propose to use boundedly rational agents in general, since they correspond more closely to the actual traders in the markets.
It is part of the agent's design to determine which percepts it can map from the available signals. In the formalism of this section, this mapping can be represented as a function $see : S \longrightarrow P$. The agent's decision making mechanism now maps (sequences of) percepts to the actions of the agent. Let f denote this mapping. Then we have $f(P, \Theta) : P \longrightarrow A$, where Θ denotes a set of parameters with which the mapping f can be parameterized. An agent's function can now be specified by defining its set P of percepts, set A of actions and the mapping $f(P, \Theta)$ from the percepts to the actions as shown in Fig. 2. It is assumed for simplicity that the specification of P also implies the specification of the function see. Hence, the definition of this function is not considered explicitly in this paper. However, it could be included if a more detailed analysis is required.

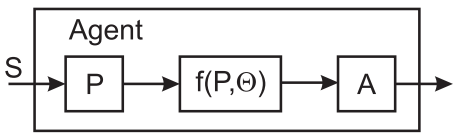

Figure 2: An agent maps its percepts to its actions.

The traders in financial markets pursue a set of goals and their actions are supposed to take them closer to their goals. Hence, it makes sense to consider only goal-directed agents for computational finance. The overall market need not have a goal, or it need not have the same goals as the agents, but the interaction amongst agents occur on basis of agents pursuing their own goals. This also fits the framework of boundedly rational agents, since rational agents must have own goals that determine their own actions. That agents show goal-directed behavior has implications for the specification of the mapping f. There are two important points for the specification of such a mapping f. First, there is the set I of intended outcomes or intentions of the agent. Second, there is the set G of the goals that the agent is pursuing. The goals are often derived from the intentions of the agent. For instance, if the intention of an agent is to trade in low-risk assets, then one of the goals might be "to minimize risk." An agent must take both its intentions and its goals into account. The sets I and G can be defined explicitly as in the belief-desire-intention (BDI) architecture for intelligent agents [25]. In other architectures, I and G may be defined only implicitly through specific choices for the mapping f, such as in the case of agents with the subsumption architecture [3]. An agent's adaptive properties can now be described in terms of modifications to any of the elements described in this section, as shown in Section 5.

Intelligent agents adapt and learn from their interactions and past information to develop rules and strategies that will allow them to achieve their goals and their intentions. The intelligence of agents is attributed to their adaptation, learning and strategy forming capability. Additionally, intelligent agents can also be characterized by their internal social norms, behavioral rules and knowledge acquired from experience. Consequently, intelligent agents have a richer internal cognitive structure and more autonomy than the conventional agents used in economical analysis. Various techniques from artificial and computational intelligence can be used to capture the intelligent behavior of the agents. Neural networks, evolutionary algorithms, fuzzy systems or their combination in hybrid systems are the most popular. The common denominator of all these methods is that they can be used to implement a nonlinear mapping from their inputs to their outputs. Consequently, the specification of the mapping f from the percepts to the actions can be nonlinear, thereby corresponding to a more diverse set of behaviors.

5 A TAXONOMY OF AGENT ADAPTATION

Adaptation in agent-based systems is usually considered in terms of learning or evolution of agents. In general, adaptation denotes all changes to a system (agent and its environment) so that it becomes suitable for a given situation or purpose [24]. Up till now, agent-based computational economics has focused mainly on the evolutionary adaptation of the agents [23]. In many studies regarding agent-based computational economics, the agents learn winning strategies through evolution [12]. Yet, a more general consideration of adaptation in agent-based systems is useful to obtain a full understanding of the possible adaptive properties of the agents and their influence on the overall behavior of the system. As shown below, such a consideration leads to a hierarchical framework for analyzing adaptation in agent-based systems.

We have mentioned that agents in financial markets are typically goal-directed. Depending on the level of complexity for expressing and organizing the goal-directed behavior, the adaptation in agent-based systems can also be classified in a hierarchical manner. This section proposes a hierarchical taxonomy of adaptation in agent-based systems, where the adaptive properties of the agents become stronger and more complex at each level of the hierarchy.

Recall that the generic agent model from Section 4 has five components, namely the set of percepts P, the set of actions A, the mapping $f(P, \Theta)$ between the percepts and the actions, the set of goals G and the set of intentions I. The last two elements, G and I can be explicit in the agent's design, or they can be implicit. The classification scheme proposed below gives the mapping from each level of classification to the elements of the agent that need to be modified.

We distinguish four levels of adaptation in agent-based systems.

- **Weak adaptation.** In this type of adaptation, the agent determines its output from its percepts according to a static mapping $f(P)$. This mapping (along with P and A) is determined at design time, and it remains fixed during the agent's lifetime. The agent can modify its environment in different ways, depending on its percepts. Therefore, the agent itself is not adaptive since it is not modified. A technical trader with fixed trading rules can be regarded as an agent operating at this level (fixed actions influence price formation in the market — the environment).

- **Semi-weak adaptation.** At this level, the mapping from the percepts to the actions can be modified. Either the function is parameterized to $f(P, \Theta)$ and the parameters are adapted, or the class of functions that f belongs to is changed (e.g. from a linear mapping to a quadratic mapping). Note that if the class of functions that f belongs to is parameterized, then changing f can also be formulated as an adaptation of parameters at a meta-level. The sets P and A are assumed to be fixed in semi-weak adaptation. An agent whose decision mechanism is a neural network operates at this level during training or on-line learning. For example, a trader may be learning to adapt its portfolio decisions based on new market conditions.

- **Semi-strong adaptation.** At this level, an agent can modify its goals G. A goal can be represented by a goal function that assigns a utility to various states of the environment and the agent. The goals can be modified by modifying these goal functions (either through their parameters or through their function classes). An interesting modification of the goals is achieved by changing the set P of percepts. Since the agent observes the states of the system through its percepts, changing the percepts will automatically lead to a functional change in the goals of the system (the agent can be assumed to have a different representation of the world). A financial agent that changes its attitude to risk and imposes a limit on the volatility (risk) of the asset prices displays semi-strong adaptation.

- **Strong adaptation.** The agents operating at this level can modify their intentions and manage the strategies for achieving their design goals. Modifications to the priority of the goals imply changed intentions. At this level, the agent may also change its set A of actions. This changes the functionality of the agent completely and corresponds to the strongest form of adaptive systems, namely to *functionally adaptive systems* of Sagasti [18]. Functional (strong) adaptation is often seen in complex systems. For example, when a trader starts trading derivatives in addition to ordinary stocks, we observe functional adaptation.

Note that the classification proposed is incremental, i.e. an agent that displays semi-weak adaptation can also display weak adaptation, an agent that displays strong adaptation can also display all other types of adaptation, etc. Furthermore, the classification is done purely from the point of view of the adaptive properties of the agents themselves (as opposed to the adaptive properties of the agent-environment pair). It does not imply that complex behavior can only be achieved at stronger forms of adaptation. In fact, agents with a subsumption architecture are weakly adaptive, but they can still solve complex problems cooperatively [22]. However, semi-weak, semi-strong and strong forms of adaptation are more relevant for adaptive agents.

6 ADAPTATION MECHANISMS

We have seen in Section 5 that agents can display adaptive behavior by modifying any of their components. An important question is, how this modification can be performed on the different levels of adaptations we introduced. Without going into the details of criteria for evaluating the effects of change, we give in this section a brief list of various mechanisms with which agent-based systems can adapt. Imitation, reaction, reactive learning, generative learning and evolution are the adaptation strategies which can be considered by agents in order to achieve adaptive behavior [8, 21].

- **Imitation.** One way to achieve adaptation is to simply copy observed data, actions or solutions. The mapping from the percepts to the actions can be obtained by copying the action of an observed agent. Agents operating on different levels of adaptation can copy percepts, actions, goals or intentions of other agents. Of course it must be possible to observe other agent's components and actions. In financial markets, traders may mimic the actions of other traders for example, who are known to be successful.

- **Reaction.** Reactions are direct responses to particular events or changes. They can be typically expressed in the form of if-then rules or in the form of mathematical formulae. Agents providing semi-weak adaptation can use the mechanism of reaction for modifying their parameterized function. A trader, for example, that sells an asset when its price falls below a certain value demonstrates adaptation by reaction.

- **Reactive learning.** Past experience plays an important role in the process of reactive learning. Agents that provide reactive learning coherently use received feedbacks when making future decisions. Reactive learning is a technique that can be used for modifying (based on experience) the parameters or the functions at the level of semi-weak adaptation. Goals in semi-strong adaptation or strategies and actions at the strong-adaptation level can also be modified through reactive learning. Technical traders that update their trading models based on past trading data learn reactively.

- **Generative learning.** Generative learning refers to a kind of innovative learning, to an anticipatory modification of the components. This type of learning is more goal-oriented. Anticipating new goals, intentions or actions at the two highest levels of agent-based adaptation implies generative learning. A financial institute that develops a new financial product based on new market conditions such as a change in fiscal system can be said to show generative learning.

- **Evolution.** Evolutionary adaptation modifies various components of the agent gradually during successive generations. It can be used for semi-weak, semi-strong and strong adaptation, provided there is a mechanism to inherit properties across generations for the adaptation of species (phylogenesis). In a financial market, for example, trading rules that lead to a loss of money consistently are selected away by evolutionary mechanisms (e.g. the trader that uses them goes bankrupt or leaves the market forever).

7 CONCLUSIONS

The dynamics of price formation at financial markets is influenced by the large diversity in the cognitive structures of the traders (e.g. their decision making methods and learning capability), the traders' specific circumstances (e.g. attitude to risk and time horizon) and the organization of the specific market in which the traders operate. We propose to model this diversity of traders by using intelligent agents that are capable of complex decision behavior, learning and adaptation. This representation allows us to analyze financial markets in a bottom-up manner, starting from the behavioral characteristics of the agents and leading to the aggregate price patterns observed at the markets.

Intelligent agents have rich cognitive structures borrowed from artificial intelligence research, and they exhibit a varied set of behaviors to model the decision making characteristics of different traders. Furthermore, they can have adaptive, learning and evolution capability that are known to have an important role in the internal workings of the financial markets. We have considered, in this paper, goal-directed agents in a generic beliefs-desires-intentions architecture. In this representation, the adaptive behavior of agents can be obtained by modifying any of their generic components. We have shown that the modification of the generic components of an agent can be organized in a hierarchical manner, leading to a hierarchical classification of adaptation in intelligent-agent based systems. A hierarchical taxonomy of possible adaptation types in intelligent-agent based systems has been introduced for this purpose. In the future, we expect our taxonomy to be useful for analyzing financial markets by considering different possibilities for adaptation in a systematic way. As a result, the increased understanding of the dynamics of financial markets may eventually lead to making policy in order to regulate the markets to achieve global goals such as the maximization of total wealth.

REFERENCES

[1] W. B. Arthur, J. Holland, B. LeBaron, R. Palmer, and P. Tayler. Asset pricing under endogeneous expectations in an artificial stock market. In W. B. Arthur, S. Durlauf, and D. Lane, editors, *The Economy as an Evolving Complex System II*, pages 15–44. Addison-Wesley, Reading, MA, 1997.

[2] K. Boer and U. Kaymak. Adaptation in agent-based systems. In *Proceedings of the* 13^{th} *Belgian-Dutch Conference on Artificial Intelligence*, pages 347–354, Oct. 2001.

[3] R. A. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, 2(1):14–23, 1986.

[4] L. K. C. Chan, N. Jegadeesh, and J. Lakonishok. Momentum strategies. *Journal of Finance*, 51(5):1681–1714, 1996.

[5] W. De Bondt and R. Thaler. Further evidence on investor overreaction and stock market seasonality. *Journal of Finance*, 42(3):557–581, 1987.

[6] J. B. De Long, A. Shleifer, L. H. Summers, and R. J. Waldmann. Noise trader risk in financial markets. *Journal of Political Economy*, 98:703–738, 1990.

[7] G. Dissanaike. Do stock market investors overreact? *Journal of Business Finance and Accounting*, 24(1):27–50, 1997.

[8] E. Droste. *Adaptive Behavior in Economic and Social Environments*. Ph.D. thesis, Center for Economic Research, Tilburg University, Tilburg, the Netherlands, Dec. 1999.

[9] E. Fama. Efficient capital markets: a review of theory and empirical work. *Journal of Finance*, 25(2):383–423, 1970.

[10] A. Gaunersdorfer, C. Hommes, and F. Wagener. Bifurcation routes to volatility clustering. Working paper 00–04, CeNDEF, University of Amsterdam, Sept. 2000.

[11] C. H. Hommes. Financial markets as nonlinear adaptive evolutionary systems. *Quantitative Finance*, 1:149–167, 2001.

[12] B. LeBaron, W. B. Arthur, and R. Palmer. Time series properties of an artificial stock market. *Journal of Economic Dynamics and Control*, 23:1487–1516, 1999.

[13] M. Lettau. Explaining the facts with adaptive agents: the case of mutual fund flows. *Journal of Economic Dynamics and Control*, 21:1117–1147, 1997.

[14] O. Loistl and O. Vetter. Computer-modelled stock exchanges. Technical report, University of Economics and Business Administration, Vienna, Austria, 1997.

[15] P. Noriega and C. Sierra. Auctions and multi-agent systems. In M. Klusch, editor, *Intelligent Information Agents*, pages 153–175. Springer, Berlin, 1999.

[16] M. H. Pesaran and A. Timmermann. Predictability of stock returns: robustness and economic significance. *Journal of Finance*, 4:1201–1228, 1995.

[17] S. Russell and P. Norvig. *Artificial Intelligence*. Prentice Hall, Upper Saddle River, New Jersey, 1995.

[18] F. Sagasti. A conceptual and taxonomic framework for the analysis of adaptive behavior. *General Systems*, XV:151–160, 1970.

[19] Y. Shoham and M. Tennenholtz. On the emergence of social conventions: modeling, analysis, and simulations. *Artificial Intelligence*, 94(1–2):139–166, July 1997.

[20] Y. Shoham and M. Tennenholtz. On rational computability and communication complexity. *Games and Economic Behavior*, 35(1–2):197–211, Apr-May 2001.

[21] A. L. D. Smet. Adaptive models of organization for substance abuse treatment. Technical report, Columbia University, Aug. 1998. http://165.112.78.61/HSR/da-tre/DeSmetAdaptiveModels.htm.

[22] L. Steels. Cooperation between distributed agents through self organization. In Y. Demazeau and J. P. Müller, editors, *Decentralized AI - Proceedings of the First European Workshop on Modeling Autonomous Agents in a Multi-Agent World*, pages 175–196. Elsevier Science Publishers, Amsterdam, 1990.

[23] L. Tesfatsion. Introduction to the special issue on agent-based computational economics. *Journal of Economic Dynamics and Control*, 25:281–293, 2001.

[24] G. Weiß. Adaptation and learning in multi-agent systems: some remarks and a bibliography. In G. Weiß and S. Sen, editors, *Adaptation and Learning in Multi-agent Systems: Proceedings of IJCAI'95 Workshop*, Lecture Notes in Artificial Intelligence, pages 1–21, Montreal, Canada, Aug. 1995. Springer-Verlag.

[25] M. Wooldridge. Intelligent agents. In G. Weiß, editor, *Multiagent Systems : a modern approach to distributed artificial intelligence*, pages 27–77. MIT Press, Cambridge, Massachusetts, 1999.