From: AAAI-96 Proceedings. Copyright © 1996, AAAI (www.aaai.org). All rights reserved.

The Use of Artificially Intelligent Agents with Bounded Rationality in the Study of Economic Markets

Vijay Rajan and James R. Slagle
Department of Computer Science
University of Minnesota
Minneapolis, MN 55455
(jayrajan|slagle)@cs.umn.edu

Abstract

The concepts of 'knowledge' and 'rationality' are of central importance to fields of science that are interested in human behavior and learning, such as artificial intelligence, economics, and psychology. The similarity between artificial intelligence and economics - both are concerned with intelligent thought, rational behavior, and the use and acquisition of knowledge - has led to the use of economic models as a paradigm for solving problems in distributed artificial intelligence (DAI) and multi agent systems (MAS). What we propose is the opposite; the use of artificial intelligence in the study of economic markets. Over the centuries various theories of market behavior have been advanced. The prevailing theory holds that an asset's current price converges to the risk adjusted value of the rationally expected dividend stream. While this rational expectations model holds in equilibrium or near-equilibrium conditions, it does not sufficiently explain conditions of market disequilibrium. An example of market disequilibrium is the phenomenon of a speculative bubble. We present an example of using artificially intelligent agents with bounded rationality in the study of speculative bubbles.

Introduction

Economics is concerned with agents making choices in situations where information is decentralized and agents have imperfect knowledge and finite resources. As pointed out by Hayek, 'the problem is to show how a solution is produced by the interactions of agents each of whom has partial knowledge' (Hayek 1945). Thus economic models based on market price systems can be used as a paradigm for solving problems in distributed artificial intelligence and multi-agent systems (Malone *et al.* 1988; Waldspurger *et al.* 1992; Wellman 1993; Rajan & Slagle 1995; Clearwater 1996). Here we are interested in the opposite, using artificial agents to understand the working of markets. The formation of prices in a market is influenced by the specific rules that govern trade in the market. These rules are referred to as the 'market institution'. Market institutions have evolved in an ad-hoc manner over the years and economists are still trying to understand them. The 1950s witnessed a big intellectual advancement in economics with the theory of general equilibrium. However, general equilibrium theory is institution free. There is currently no general theory of trading under various market institutions.

Over the centuries economists have advanced various theories of market behavior. The prevailing hypothesis holds that the market is driven to a competitive rational expectations equilibrium that reflects the aggregate of all the information dispersed among the market participants. At this equilibrium there should be no opportunity for arbitrage. This hypothesis has proven to be robust and its predictions supported to some extent by experimental laboratory markets as well as tests based on field data. However, there are certain market conditions that theory fails to explain. One such example is the speculative bubble. Under the condition that all market participants are rational, speculative bubbles would be impossible. Hence, to study such phenomena, the condition of rationality needs to be relaxed. However, incorporating the notion of bounded or limited rationality in economic theory has been extremely difficult. Theoretical models of bubbles have an infinity of solutions satisfying the equilibrium conditions. Hence theory does not provide any information about the process by which bubbles form and collapse.

In this paper we explore the use of artificially intelligent agents with bounded rationality to provide an alternative framework to study the divergence of market prices from the rational expectations equilibrium. We believe that computer simulations using such agents offers a promising alternative to study market conditions that are not easily explained by theory. Empirical research based on simulations is beginning to play an important role in the understanding of many complex systems, both natural and social, that defy precise mathematical characterization. For example, in physics, computer simulations have played an important role in the understanding of complex phenomena such as spin glasses, Ising magnets, and quantum chromodynamics. However, there is an important distinction between complexity in natural systems and social systems. The complexity in social systems arises from the nature of human thought and behavior. Hence the simulation of complex social systems requires an additional component to model the thought, behavior, and learning processes of the participants in the social system. Several areas of artificial intelligence, especially those related to reasoning about knowledge (Fagin, Halpern, & Vardi 1984; Fagin *et al.* 1995), agent theory (Shoham 1993), and goal driven learning (Ram & Leake 1995) can play an important role in the study of social systems including economic markets.

Bubbles

A phenomenon that has plagued markets, since the inception of organized trading, is the speculative bubble. A bubble is defined loosely as a sharp rise in the price of an asset, with the initial rise generating expectations of further rises and attracting new buyers. The rise is usually followed by a reversal of expectations causing a sharp decline in the assets price. Bubbles are of wide spread interest since their consequences are serious often resulting in a financial crisis. Such a crisis can affect millions of people and have a significant impact on society. Little progress has been made so far in understanding this phenomenon. Some examples of bubbles are the Tulip mania that occurred in Holland around 1636-7, the Mississippi bubble that occurred in Paris around 1719-20, the South Sea bubble that occurred in London also around 1720, the Railway mania that occurred in England around 1846, and the stock market crashes that occurred in New York in 1929 and 1987.

As of today, economic theory has failed to explain the phenomena of bubbles. Theorists are still divided about even the existence of bubbles. Recently, experimental laboratory markets conducted with human traders have shown the possibility of the occurrence of speculative bubbles. Smith, Suchanek, and Williams conducted a set of experiments to study the possibility of speculative bubbles in asset markets (Smith, Suchanek, & Williams 1988). These experiments were organized such that the price of the asset was based on the expected value of the dividend stream. The dividend structure and the number of trading periods were publicly announced to all the traders making it common knowledge. They found price bubbles in about half of the experiments they conducted.

Cutler, Poterba, and Summers (Cutler, Poterba, & Summers 1990) suggest that models based solely on existing theory cannot account for the movement of asset prices (such as the crash of 1987) and the high trading volumes observed in modern securities markets. An approach based on 'feedback' or 'noise' traders is used by Cutler, Poterba, and Summers (Cutler, Poterba, & Summers 1990) and by DeLong, Shleifer, Summers, and Waldman (DeLong *et al.* 1989; 1990). The models in (DeLong *et al.* 1989) and (DeLong *et al.* 1990) were developed based on an overlapping generations environment where investors live for two periods. (Cutler, Poterba, & Summers 1990) develops a model of price dynamics for investors using heterogeneous trading strategies. The model emphasizes the existence of 'feedback traders' who behave based on the history of past returns rather than future fundamentals. These models based on noise trading provide some insight into gradual asset price swings. However, they do not account for the sharp price swings observed during the crash of a price bubble in the laboratory.

The failure of theory to explain the phenomena of bubbles has lead to empirical efforts to understand them. Empirical work attempting to identify bubbles using field data (Flood & Garber 1980; Flood, Garber, & Scott 1981) has turned up mixed results and has been inconclusive about the relevance of bubbles. However, studies based on laboratory markets have established the existence of bubbles. As previously mentioned, Smith, Suchanek, and Williams found frequent price bubbles in an experimental asset market that paid a dividend (from a known probably distribution) at the end of each period. Camerer and Weigelt (Camerer & Weigelt 1993) also observed bubbles, especially with inexperienced traders. King, Smith, Williams, and Boening (King *et al.* 1993) conducted further experiments building on the results of (Smith, Suchanek, & Williams 1988). Their goal was to determine whether additional policies such as short selling, buying on margin, and rules that limit price change could prevent price bubbles. None of these factors seemed to affect the possibility of observing bubbles.

The experimental markets conducted with human traders have shown the existence of bubbles but have not been able to provide insight into the process of creation and the collapse of such bubbles. We hope to accomplish this by using artificial agents rather than human subjects. By using artificial agents programmed with specific behaviors we can observe, verify, control, and replicate (Gode & Sunder 1992) the decision rules used by the agents enabling us to better understand the phenomenon of speculative bubbles.

The Study of Speculative Bubbles

A market can be considered as a complex social system that consists of the individuals that participate in the market. The participants in a market act based on their expectations and beliefs about future market outcomes. The market outcomes, however, are themselves affected by the actions of the participants. Hence there is a mapping from beliefs to market outcomes back to beliefs, leading to a form of self-reference. By the efficient market theory the market reaches the rational expectations equilibrium. This equilibrium is a fixed point in the mapping from beliefs to beliefs. At the equilibrium, the actions of the market participants based on current beliefs generate market outcomes that confirm the beliefs. Hence the rational expectations model implies homogeneous beliefs among all the market participants. Our approach to studying phenomena such as the speculative bubble is based on representing a market as a multi- agent system organized as a society of artificially intelligent software agents. Agents can be programmed with different models of belief formation allowing us to study the effects of trader heterogeneity on market prices.

The Market Institution

As mentioned earlier, the market institution plays a significant role in the price formation process. Hence, the notion of a bubble can make no sense in the absence of a precise model detailing the markets operation. The specific model we use is a computerized version of the 'Double Auction'. The double auction is a market institution in which traders make offers to buy and sell. There are several forms of the double auction. The market we use is a restricted form of the Continuous Double Auction (CDA). The lowest ask submitted so far is called the current ask, and the highest bid the current bid. An incoming ask has to be lower than the current ask, and an incoming bid higher than the current bid. If the incoming ask is equal to or less than the current bid a transaction occurs (at the bid price), otherwise it becomes the current ask. Similarly if the incoming bid is equal to or greater than the current ask a transaction occurs (at the ask price), else it becomes the current bid. After a transaction, the current bid and ask are removed, and the first bid (ask) received after the transaction will become the current bid (ask). Most financial and commodities market around the world, such as the New York Stock Exchange, are based on variants of the double auction.

Experimental Design

We use the experimental design of (Smith, Suchanek, & Williams 1988). However, instead of human subject,

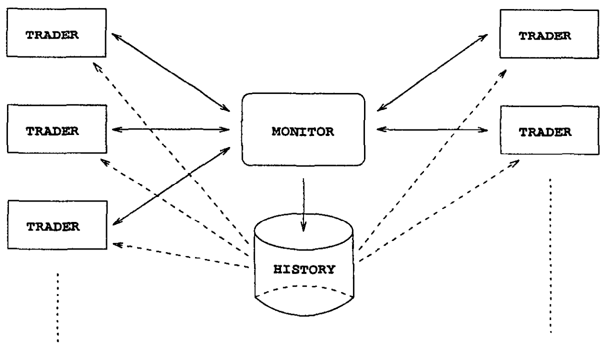

Figure 1: The double auction as a multi-agent system

we use 12 artificial agents as the traders in the market. At the beginning of the experiment, each agent receives an initial endowment of cash and shares. Trading occurs over a sequence of 15 market periods, each period lasting 5 minutes. Agents can make bids to buy or offers to sell shares at any time during a period. The cash and shares owned by an agent at the end of a period is carried over to the next period. In addition, at the end of each trading period, each share earns a dividend. The dividend is drawn from a probability distribution centered around a fixed value. The structure of the probability distribution is common knowledge to all the agents. At the end of the experiment the cash held by an agent is the profit earned for participating in the experiment. All shares are worthless at the end of the experiment.

The Market as a Multi-Agent System

The computerized market used for this study is organized as a multi-agent system. Each agent participating in the market is a UNIX process. In addition there is one more agent, the market monitor, whose role is similar to a specialist in the trading pit of a stock exchange. All market participants send their messages to the market monitor. The market monitor, implements the rules of the double auction; decides what messages are valid, maintains the current bid and ask, and decides when a transaction occurs. Agents can at any time request the current ask, current bid, and the market price of the last transaction. Figure 1 shows the double auction organized as a multi-agent system.

The Agent Architecture

An agent at any time has a choice among five possible actions; enter a bid, enter an ask, accept the current bid, accept the current ask, or wait. In order to choose among the five possible actions, the agent needs to form expectations of the asset's price. Figure 2 shows the architecture for agents participating in the market.

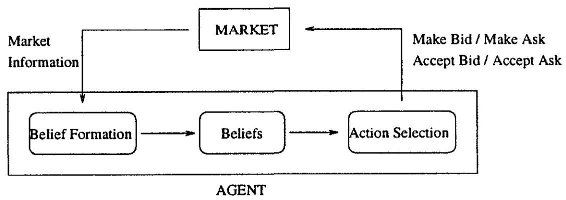

Figure 2: Agent architecture

Rationality requires each agent to build a correct model of beliefs and expectations based on the available market information. The expectation is correct if it is fulfilled by the future behavior of the market. However, an agent due to bounded or limited rationality, may not be able to completely interpret all the available information correctly. Our model is built on agents using only certain subsets of the information available to them. This leads to agents with differing beliefs and expectations. This notion of bounded rationality differs from Simon (Simon 1958). The agents are not limited by computational resources, but rather ignore a certain subset of the information available to them.

Agents

For this study we used three types of agents based on the subset of the information used by them. The behavior of these agents were arrived at by studying the data from (Smith, Suchanek, & Williams 1988). The following are brief descriptions of the three types of agents: fundamental traders, speculative traders, and strategic traders.

Fundamental traders forms their expectations based solely on information about the dividends. Their valuation of an asset's price is the same as the theoretical value. The theoretical value for each share is the product of the expected value of the dividend at the end of a period and the number of trading periods that remain. The fundamental trader ignores all other market information. The valuations are updated once at the beginning of each trading period. A fundamental trader buys shares when the price is below the theoretical value and sells shares when the price is above the theoretical value.

Speculative traders ignore all information about the dividend structure. They are interested purely in short term capital gains – buy low and sell high. In this study, speculative traders make use of information about the transactions in which they participate. If a speculative trader buys a share at a given price, it immediately changes its valuation to that price. Similarly when it sells a share it changes its valuation to the transaction price. Information about transactions conducted by other agents and information about all bids and asks are not used. Based on the number of shares in its possession a speculative trader frequently changes its role between being a buyer or a seller. Each time a speculative trader changes from a buyer to a seller, it increases its valuation of the price so that it can sell the shares at a higher price. If it is unable to sell for an extended period of time, it begins to drop its valuation. When a speculative trader changes from a seller to a buyer, it decreases its valuation of the price so that it can buy assets at a lower price. If it is unable to buy for an extended period of time, it begins to increase its valuation.

Strategic traders build their expectations of price based on a combination of dividend information and information about the current price at which shares are being traded. Strategic traders are similar to fundamental traders, in that they know the theoretical value of the price based on the dividend information. However, they do not begin to sell as soon as the market price exceeds the theoretical value. Instead they build expectations on whether the price is likely to stay at the current level or go higher. If this is likely they can hold on to their shares, collecting dividend at the end of the period, and sell later. In this way they can exploit the behavior of the short term speculators.

Results

We conducted two sets of multi-agent simulations. We performed 100 runs of each set of simulations. However, only illustrative results are shown here. The first set of simulations attempted to study the interaction between fundamental and speculative traders. Figure 3 shows the results from two runs of these experiments. In the first run we used three fundamental traders and nine speculative traders. The three fundamental traders were able to hold the market price close to the theoretical value. They did so by buying assets when the price was below the theoretical value and selling assets when the price was above the theoretical value. The second run used one fundamental trader and eleven speculative traders. In this case too, any divergence from the theoretical value was quickly corrected by the single fundamental trader (in some runs of this case, the single fundamental trader ran out of cash or assets and could no longer influence the market price).

Since fundamental traders were able to hold the market prices close to the theoretical value, the second set of simulations studied the interaction between strategic and speculative traders. Figure 4 shows the prices observed in one of the laboratory experiments conducted by Smith, Suchanek, and Williams (with human traders) versus the prices observed in one of our multi-

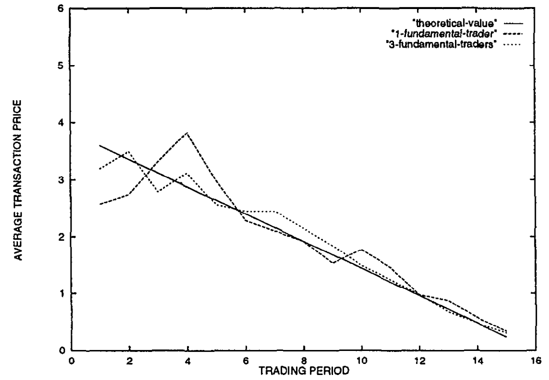

Figure 3: Fundamental traders keep the price close to the theoretical value

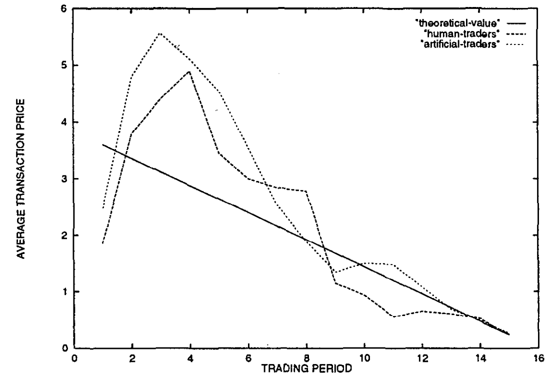

Figure 4: Speculative bubbles

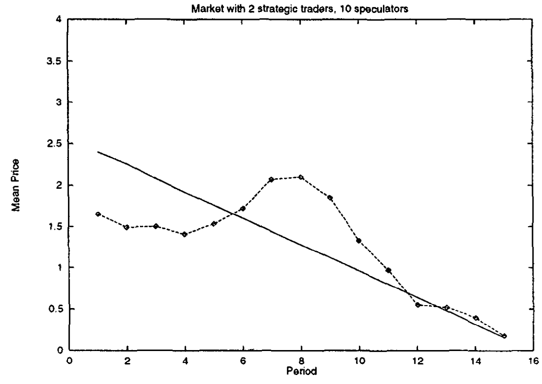

Figure 5: Average prices in a market with 2 strategic traders and 10 speculative traders

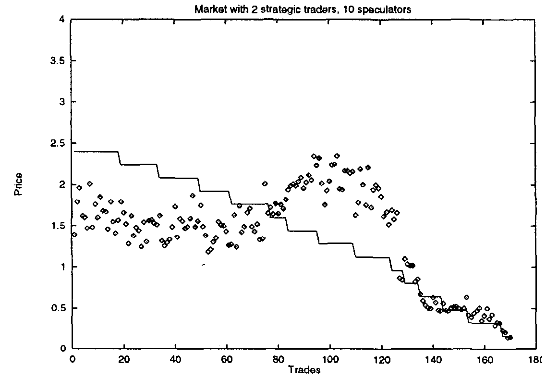

Figure 6: Transaction prices in a market with 2 strategic traders and 10 speculative traders

agent simulations using two strategic traders and ten speculative traders. In this case the simulation produces a speculative bubble similar to the one observed in the laboratory market. The speculative bubble is identified by a rise in the asset's price to a value significantly over the fundamentals followed by a sharp reversal. The strategic traders allowed the price to rise to a high value before beginning to sell, thereby causing a reversal. Figures 5 and 6 show the prices from another typical run of the simulation with two strategic traders and ten speculative traders.

Discussion

The speculative bubble is a good example of market disequilibrium. Since theory does not provide a reasonable explanation of the process of speculative bubbles, we need to rely on the empirical study of this phenomenon. There are three possible approaches to take when trying to empirically analyze the behavior of markets. The first, based on field data, involves studying actual market data. The second, involves the use of human traders in experimental laboratory markets. The third, is based on using artificial agents in computational markets. While field data provides the best source of actual market behavior, due to the large number of variables involved, it is difficult if not impossible to isolate the effect of individual factors on the markets outcome. Over the last two decades, laboratory experiments have proved invaluable in the study of market institutions. However, to specifically study the effects of a single factor such as the agent's behavior on the market outcome it is necessary to observe the participant's decision rules in order to isolate the conditions that impact the market. It is difficult to observe human decision rules both in the field and in experimental laboratory markets. Hence the use of multi-agent simulations based on artificial agents provides a promising alternative where it is possible to control factors such as the agents rationality and decision rules.

In this paper we have used multi-agent simulations, based on artificially intelligent agents, to show one possible explanation for how speculative bubbles can occur within the framework of the double auction market institution. We find that the interaction of two classes of agents, each with simple behavior, often results in a bubble. Further statistical analysis is required to validate this approach.

Acknowledgments

We would like to thank John Dickhaut, Arijit Mukherji, and the three anonymous reviewers for their suggestions. We also acknowledge the help of Shawn LaMaster of the Economic Science Laboratory at the University of Arizona in providing us with the experimental data from (Smith, Suchanek, & Williams 1988).

References

Camerer, C., and Weigelt, K. 1993. Convergence in experimental double auctions for stochastically lived assets. In Friedman, D., and Rust, J., eds., *The Double Auction Market.* Addison Wesley. 333–354.

Clearwater, S. H., ed. 1996. *Market Based Control.* World Scientific.

Cutler, D. M.; Poterba, J. M.; and Summers, L. H. 1990. Speculative dynamics and the role of feedback traders. *American Economic Review, Papers and Proceedings* 80:63–68.

DeLong, J. B.; Shleifer, A.; Summers, L. H.; and Waldman, R. J. 1989. The size and incidence of the losses from noise trading. *Journal of Finance* 44:681–696.

DeLong, J. B.; Shleifer, A.; Summers, L. H.; and Waldman, R. J. 1990. Noise trader risk in financial markets. *Journal of Political Economy* 98:703–738.

Fagin, R.; Halpern, J. Y.; Moses, Y.; and Vardi, M. Y. 1995. *Reasoning about Knowledge.* MIT Press, Cambridge, Massachusetts.

Fagin, R.; Halpern, J. Y.; and Vardi, M. 1984. A model-theoretic analysis of knowledge: preliminary report. In *Proc. of the 25th IEEE Symposium on Foundationsof Computer Science.*

Flood, R., and Garber, P. 1980. Market fundamentals versus price level bubbles: The first tests. *Journal of Political Economy* 745–770.

Flood, R.; Garber, P.; and Scott, L. 1981. Further evidence on price level bubbles. *Working Paper, University of Rochester.*

Gode, D. K., and Sunder, S. 1992. Some issues in electronic modeling of continuous double auctions with computer traders. *Working Paper No. 1992-25, Graduate School of Industrial Administration, Carnegie Mellon University.*

Hayek, F. A. 1945. The use of knowledge in society. *American Econ. Rev.* 34(4):519–530.

King, R. R.; Smith, V. L.; Williams, A. W.; and Boening, M. V. 1993. The robustness of bubbles and crashes in experimental stock markets. In Day, R. H., and Chen, P., eds., *Nonlinear Dynamics and Evolutionary Economics.* Oxford University Press. 183–200.

Malone, T. W.; Fikes, R. E.; Grant, K. R.; and Howard, M. T. 1988. Enterprise: A market like task scheduler for distributed computing environments. In Huberman, B. A., ed., *The Ecology of Computation.* North Holland. 177–205.

Rajan, V., and Slagle, J. 1995. The use of màrkets in decentralized decision making. In *Proceedings of the Eighth International Symposium on Artificial Intelligence.*

Ram, A., and Leake, D. B., eds. 1995. *Goal Driven Learning.* A Bradford book, MIT Press, Cambridge, Massachusetts.

Shoham, Y. 1993. Agent-oriented programming. *Artificial Intelligence* 60(1):51–92.

Simon, H. A. 1958. *Administrative behavior : a study of decision-making processes in administrative organization.* Macmillan, 2nd edition.

Smith, V. L.; Suchanek, G. L.; and Williams, A. W. 1988. Bubbles, crashes, and endogenous expectations in experimental spot asset markets. *Econometrica Vol. 56. No. 5* 1119–1151.

Waldspurger, C. A.; Hogg, T.; Huberman, B. A.; Kephart, J. O.; and Stornetta, S. 1992. Spawn: A distributed computational economy. *IEEE Transactions on Software Engineering* 18:103–117.

Wellman, M. P. 1993. A market-oriented programming environment and its applications to distributed multicommodity flow problems. *Artificial Intelligence Research* 1:1–23.