

Artificial Software Agents on Thin Double Auction Markets – A Human Trader Experiment

Jens Grossklags*, Carsten Schmidt⁺

School of Information Management and Systems, University of California, Berkeley*
jensg@sims.berkeley.edu

Max Planck Institute for Research into Economic Systems, Jena, Germany⁺
Strategic Interaction Group
cschmidt@mpiew-jena.mpg.de

December 17, 2002

Abstract

This paper studies how software agents influence the market behavior of human traders. Programmed traders with a passive arbitrage seeking strategy are introduced in a double auction market experiment with human subjects in the laboratory. As a treatment variable, the influence of information on the existence of software agents is investigated. We found that common knowledge about the presence of software agents triggers more efficient market prices in the presence of the programmed strategy whereas an effect of the information condition on behavioral variables could not be observed. Surprisingly, the introduction of software agents results in lower market efficiency in the no information treatment when compared to the baseline treatment without software agents.

Keywords: *programmed traders, laboratory experiment, software agents, information aggregation*

We thank Werner Güth, John Morgan, Axel Ockenfels, Jonathan Siegel, Shyam Sunder and the participants of seminars at Berlin, Mannheim, the 2002 GEW conference in Wittenberg, the Workshop on “Experimental Economics” in Jena 2002 and the INFORMS 2002 annual meeting in San Jose for helpful suggestions and comments. The financial support of the Max Planck Society is gratefully acknowledged.

*Corresponding author: University of California Berkeley, School of Information Management and Systems, 102 South Hall, Berkeley, CA 94720-4600.

1 Introduction

Obvious examples where human traders and software agents participate alike are to be found in financial markets. In the early '90s neural networks, genetic algorithms, fuzzy logics, chaos theory, and other approaches were applied to automate trading. It seems the hype has disappeared and "black box" traders are managing rather small funds on Wall Street. This can also be attributed to still open research questions about the impact of artificial traders in situations of market instability. Exemplary, Leland and Rubinstein (1988) and Varian (1998) discussed the role of artificial traders that followed 'price insensitive' strategies such as portfolio insurance which might have contributed to the 1987 stock market crash. Currently, the majority of funds are managed by human traders that are supported by software aids filtering and aggregating information.¹

One exemplary application where humans and software agents participate alike is eBay (Ockenfels and Roth 2002). The auction format used by eBay is an open bid second price auction with a fixed ending rule. In this situation bidders have a strategic incentive to delay their bids (Roth and Ockenfels 2002) and to bid at the very end of the auction, the so-called sniping. Most of the bidders place their bid using the graphical user interface of eBay. Lately services such as esnipe and auctionblitz² have started to offer placing a bid in the very last minute on the bidder's behalf. This automated bidding supports human bidders in a routine task which is thus performed more precisely. In this example, the software agent exploits human shortcomings (of bidders not recognizing the strategic incentives of late bidding) and does the job without the error proneness of human interaction with software systems.

In a different context we are posing the following questions that can be stated in the eBay environment as well: Do software agents influence the market behavior of human traders? And does individual knowledge about software agents influence human behavior and the market outcome?

¹ See, for example, articles by Davidson, C. (1999) Securities Industry News, Vol. 11: "The Black Box: For Better or for Worse?" (May 24), "Military Technologists Aim Their Software At the Markets" (March 8) and "Still fuzzy after all these years" (June 14).

² <http://www.esnipe.com>, <http://www.auctionblitz.com>

The intention of this study is to concentrate on a stylized, controlled environment and to understand and disentangle driving economic and psychological variables of human-agent interaction. The experiment presented in this paper relies on a market institution quite common in financial markets - the continuous double auction (CDA). The framework which allows software agents to interact with a market is described in Grossklags et al. (2000), where a simulation exclusively with software agents was conducted. In the present paper, human traders are introduced into this framework. The first pilot experiments combining the participation of human and software traders simultaneously were run in June 2001 to decide which trading strategy to select for human-agent interaction. The complexity of the market environment led us to the conclusion to refrain from using a zero intelligence type of agent. A natural candidate to explore human-agent interaction is the arbitrageur. It is a programmed trader who constantly scans the market in order to exploit riskless profit opportunities, which arise as the result of price differences of a contract in different markets. From a behavioral point of view, the arbitrageur can be described as a passive, rather parasitic strategy sitting in the background and earning profits yielded from the imperfection of other traders. In the following sections, in order to distinguish human traders from software agents, acting as traders as well, the term “*traders*” will be used for human participants and the term “*agents*” for software participants.

The main contribution of the paper is the introduction of the information condition in a human-agent experiment. Two treatments were conducted with experimental parameters held constant except for the information available on software agents: in one treatment the participation of the programmed trader was made common knowledge and in the other treatment subjects were not informed about the existence of software traders. In addition, the data is compared to a baseline treatment without software agents.

To form a hypothesis of what reaction can be expected from human traders when information on software agents is provided, we are drawing an analogy to the lemons market scenario (Akerlof 1970). Similar to quality uncertainty in the lemons market, the human traders’ uncertainty about the agents’ capabilities - the speed to calculate strategies and to process transactions - leads them to crowd out of the market. This is a strong hypothesis that requires human traders not to trade at all when information about the existence of software agents is available. In the context of the double auction market institution, traders

cannot observe if they are trading with a human or a robot. Thus, an alternative hypothesis can be formulated according to which humans compare themselves with other human traders only and neglect the existence of software traders. This hypothesis would predict no difference in human behavior when information is introduced.

Closely related to our experiment Das et al. (2001) conducted an experimental series where human traders interacted with software agents, too. They closely followed the traditional design proposed by Smith (1962) where participants are assigned fixed roles as either buyer (submitting only bids) or seller (submitting only asks) and receive a private valuation (cost) for the traded good as a buyer (seller). In their study the experimental conditions of supply and demand are held constant over several successive trading periods and are then imposed to a random shock that changes market parameters. Experimental sessions involved 6 human traders and 6 agents. In addition, a baseline session with 12 human traders was run. Two types of agents were used that applied either a modified Zero-Intelligence-Plus strategy (Gode and Sunder 1993 a&b, Cliff and Bruten 1997) or a modified Gjerstad-Dickhaut (1998) algorithm.

While in general the human-agent markets in Das et al. (2001) showed convergence to the predicted equilibrium, the price patterns showed strong scalloping behavior in comparison to pure human or agent markets. Markets tended to have a lopsided character in which either buyers consistently exploited sellers, or vice versa. In their experiment, between 30 and 50 percent of the trades are done between agents and human traders. The agent instances reach average profits well above those of human traders. However, the experiments also showed that in human-agent environments artificial traders can be subject to exploitation by humans.

Compared to Das et al. (2001) the present paper introduces software agents to a more complex environment. This includes traders acting both as buyers and sellers, information about the fundamental value of the securities changes in every round, orders allow for multiple units of a specific contract, and the market institution provides no spread improvement rule. A further direct comparison of findings seems difficult since Das et al. (2001) only report selected data about their markets and focus on the differences between the two types of software agents.

The main results of our study are as follows. We find human traders not to crowd out in the treatment with common knowledge on software agents. Instead, common knowledge on the presence of software agents has a significantly positive effect on human traders' ability to converge to equilibrium in the presence of the arbitrageur agent. Intuition would suggest higher efficiency in an environment with software agents when compared to no software agents. Surprisingly, when compared to the baseline treatment the introduction of an arbitrage seeking type of software agent results in lower market efficiency in the no information treatment.

In the next section the experimental setup is presented, including the market and the software agents' strategy. The experimental results are described in Section 3. Related work is discussed in Section 4. Finally, concluding remarks are given and open research questions for further exploration are sketched in Section 5.

2 Experimental setup

2.1 Market institution and information

The market institution was designed by using a continuous double auction, i.e. an auction in which sellers and buyers may submit bids and asks simultaneously and asynchronously. More precisely, sellers and buyers are free to accept bids and asks at any time during the experiment. CDA market designs are very popular among financial markets, both real and virtual, and are described as having the remarkable quality of being fast and efficient (Friedman 1984, 1993). In contrast to markets where the issue of securities is organized by an initial public offering, this is implemented on this particular market via a so-called bundle mechanism and therefore resembles closely the design of the Iowa Electronic Markets (Forsythe et al. 1992, 1999)³. The bundle consisted of a standardized unit-portfolio where the sum of each different contract carries a fixed price. This bundle can be bought from or sold to the bank at any time and any quantity. The market foresaw three valid operations: (1) posting market orders (bids implement buying orders and asks implement selling orders), (2) deleting own market orders, and (3) buying/selling bundles at the bank. Submitted orders remained open until they were traded, or the experiment ended. No restrictions to the posted prices were made (no spread-improvement rule).

The market implements an American futures market, where contracts can be traded on some kind of event. The outcome of the event determines, depending on the market rules, the payoff of the different contracts. For the experiment described, a payoff scheme similar to “vote share” election markets has been used, where each contract pays off a percentage of the total bundle. The experiment was conducted with a market that contained five contracts each representing one firm. Three contracts represented relatively more valuable firms (contracts A, B, and C) and two representing relatively less valuable firms (contracts D and E).

The value of each different contract was characterized by a strength measure given in points. For example, if the firm is doing well it will gain points, while if it performs poorly it will lose points. Furthermore, the strength points provide an indicator of its relative performance compared to the other contracts in the bundle. This implies that an increase in points in one contract results in a proportional decrease of the equilibrium price of the other different contracts on the market as well. The equilibrium price equals the fundamental value of a contract and can be calculated by dividing the points of a contract by the sum of points of all different contracts in the market multiplied by 100.

During the experiment, the participants received the information via the screens of their computers. Initially, all participants were given the same information in the instructions and a trading time of three minutes. Afterwards, the information was sent by the following schedule: reception of private information on the contracts’ points, 4 minutes time of trading, reception of public information, and 2 minutes time of trading. This schedule was repeated 12 times. Altogether the market was open for 75 minutes and each subject received 13 public and 12 private information messages.

A storyboard was designed in order to provide a constant environment for all sessions (Figure 1 presents the equilibrium price over time for each of the 5 contracts). It was determined by the experimenters with a rolling dice. The storyboard describes the change of points and the corresponding fundamental value of the different contracts. In addition, each trader, say for example trader number 2, received the same information throughout all

³ The market software for the experiment uses Web technology and has been used for example in Hansen et al. (forthcoming) and Schmidt and Werwatz (2002).

sessions. For the sake of simplicity, only one contract changed its points in each round, so that if the points of contract A were going to change in the first round the other contracts didn't change in points. Still, this scenario implies a change of the fundamental value of all the different contracts in the market. Each of the six traders in this storyboard received different private information about the change in points. The mean of all private information sent to the traders was equal to the actual change. This true change was made available to the traders with the public information.

2.2 Programmed trader

The market-agent interface (XML) and several implemented programmed traders (in Java) are described in detail by Grossklags et al. (2000). For the purpose of Grossklags et al. the programmed traders were selected under the premise of being pure and simple strategies, which resemble a real world analogy. A simulation tested the market-agent interface and pointed out strategies that are successful when competing with programmed traders. Out of the successful strategies of the simulation only those that are expected to be profitable in a market experiment populated with human traders are of interest for the current study. The *arbitrageur* strategy has been selected that takes advantage of different prices between market and bank. From a behavioral point of view the arbitrageur can be described as passive in that it sits in the background and waits to exploit arbitrage opportunities.

Arbitrageur: Aim of this agent is to profit from arbitrage opportunities that arise because of the difference between the market price and bank price. In detail, the arbitrageur exploits the difference between the market price of a bundle - more correctly, the bids and the asks of the single contract forming the bundle - and the bank price for the same bundle, buying the bundle from whomever sells it at the lowest price (the bank or the market) and reselling it to whoever is ready to buy it at the highest price (the bank or the market). When the agent is able to conduct all transactions, in the experiment this includes 5 market transactions and one bank transaction, a guaranteed profit can be achieved.

The arbitrageur continuously scans the market and applies the following algorithm:

- if $\sum Ask_{Min}[Stock_i] < 100ECU^4$, then buy unit-portfolios from the market and resell them to the bank as a bundle
- if any combination of stocks (one stock alone, two stocks, three stocks, etc.) is requested at a price, which exceeds 100 ECU, then buy a bundle from the bank, split the bundle and sell the stocks separately to the market

The simulation in Grossklags et al. (2000) displayed that in the sample environment with 11 different programmed strategies the arbitrageur was not in every case able to complete the whole set of transactions necessary in order to gain a sure profit and avoid risk. Still in the simulation the arbitrageur agent gained on average positive payoffs.

2.3 Experimental procedure and hypotheses

The experiment is designed to separate the influences of the programmed strategy and information on the participation of software agents. Therefore, the agent treatment has been run with and without information on software agents. In addition, a baseline treatment with human traders only has been run. Each session consisted of a market with 6 human participants and 6 streams of information. The passive trading strategy *arbitrageur* does not use point information on individual contracts in order to apply its strategy (only market prices and quantities). Consequently in the *arbitrageur* treatment the programmed trader was used in addition to the 6 human traders. Apart from the presence or absence of software agents and the provision of information about the existence of programmed traders there was no other difference between the individual markets. Altogether 18 sessions have been run with 108 human participants and 12 programmed traders and thus six independent observations for each treatment were collected.

The human participants were recruited among students of the University of Jena. The laboratory sessions took place in June 2002 in the experimental laboratory of the Max Planck Institute in Jena/Germany. In the experiment software agents face the same budget constraint as human traders do. Each participant - traders as well as agents - was given an initial endowment of 100,000 ECU = 10 Euro. At the beginning of the experiment the instructions were available on the computer screen and read out loud by the experimenter. A

⁴ ECU = Experimental Currency Unit

demonstration of the trading screen including sample transactions for all three valid market operations was provided with a video projector.

Two hypotheses with regards to the influence of software agents on human traders were formulated. First, a crowding out of human traders can be suspected in the treatment with public information on software agents. As already mentioned this hypothesis was derived from an analogy to the lemons market scenario (Akerlof 1970). This hypothesis requests human traders not to trade at all when information about the existence of software agents is available. Second, software agents are expected to improve market efficiency. They follow predefined rules and do not make mistakes with respect to their algorithm. In addition, software agents can process more data in a given time span and interact faster with the software interface than human traders are able to interact with the graphical user interface. For the evaluation of this hypothesis efficiency deviations from the equilibrium price and volatility measures of the different treatments are evaluated.

3 Results

In a first step the payoffs of human and software traders are compared. It seems important that programmed traders do not make losses on average. Otherwise agents just distribute money to human traders and the experimenter loses control. A zero sum market has been used; therefore each different programmed strategy should at least regain the invested capital of 100,000 ECU on average. Out of 12 sessions, programmed traders achieved positive payoffs in 11 cases. One *arbitrageur* agent made a zero profit due to missing arbitrage possibilities. Profits of the programmed traders differed significantly from zero (see Table 1): on average the *arbitrageur* traders made 0.3% profits during the 75 minutes period of time. Human traders did lose this percentage in the corresponding treatments but this effect is not significant from zero profits ($T = -0.963$; $p < 0.33$). Furthermore the variability of the software agents' profits is significantly lower when compared to human traders monetary payoffs ($F = 3.596$; $P < 0.000$).

Next behavioral variables are analyzed. The first aspects are the number of trades and portfolio restructuring activities of human traders and agents. In the following average values of the 6 independent observations for each treatment are compared and if not otherwise noted a permutation test is used in order to test for statistically significant differences. For

each of the agent treatments the number of trades is not significantly different from the baseline treatment. The average number of trades declines in the *arbitrageur* treatment when compared to the baseline treatment but this effect is not significant what can be attributed to a high variability of the individual sessions' averages.

It can be observed that trades between agents and human participants crowd out human-to-human transactions. The number of human-to-human transactions is significantly lower in the treatments involving a software trader when compared to the baseline treatment. The percentage of human-to-agent (h2a) trades in the *arbitrageur* regime is 16% in the no information and 22% in the information treatment.

The information on software agents does not have a significant impact with respect to human portfolio restructuring (Table 1). On the individual level we observed 4 human traders, each of them in a different session, who did not trade at all (even though they had to stay for the complete experiment). This human behavior could be observed two times in both the information and the no information treatment. Thus, a crowding out of human trades in the presence of information on software agents could not be observed and this hypothesis can be rejected.

Result 1: *There is no crowding out of human traders when public information on software agents is available.*

Altogether five different contracts were available for trading. During each round informational changes in points were given to the participants for one contract type only. The market design implies that informational changes in one contract result in changes of the equilibrium price of all other contracts as well. Due to the random character of the storyboard such point information for the contract B was first provided after 69 minutes of the experiment; that is in the last period only. Note that trading a contract without direct informational change requires a further level of reasoning, e.g. if the points and thus also the price of contract A change, this implies that the price of contract B, C, D, and E change as well. Therefore, we regard the time when trades and prices in all five different contracts are first realized as a proxy for trader's rationality. The baseline treatment provides prices for all different contracts after less than 10 minutes of the experiment. In the *arbitrageur* regime

prices for all different contracts are available roughly 25 minutes after the start of the experiment. This difference between the *arbitrageur* and the baseline treatment is significant. Behavioral differences in this context could not be observed with respect to the information conditions.

Result 2: *With respect to behavioral variables a difference in the first time when prices for all 5 different contracts were realized can be observed. We infer that the introduction of the passive agent does not guide human traders to rationalize the connection between the 5 different contracts.*

In the following the focus of this paper is on several efficiency measures. In a first step arbitrage opportunities between the market and the bank are evaluated. The bank promises during the experiment to buy and to sell the unit portfolio for a fixed price of 100 ECU. Therefore, the aggregated price of one unit of each different contract on the market should be 100; lower market prices are an indicator for undervaluation, and higher market prices for overvaluation. The *arbitrageur* agent explicitly scans the market for immediate arbitrage opportunities. It can be suspected that the market price of a bundle should be close to 100 in this regime. Table 1 provides evidence that on average in the arbitrageur regime the unit portfolio is not significantly different from 100.

The information provided in the private information phases allows calculating the equilibrium price of a contract from the six different pieces of information provided to six traders. During the public information phase the points allowed to calculate the prices of the contracts directly. Price deviations from equilibrium price will be considered as inefficiencies. The following two measures will be used. First, the deviation of the market price from equilibrium is calculated. We apply normalization of prices to account for differences in contract prices. Second, a variance measure is used that accounts for price variability and different contract prices as well. The calculation of the Ψ -variance is shown in (1).

$$\Psi_{\text{var}} = \frac{\sum_{\#Trades} \left(1 - \frac{\text{fundamental value}}{\text{market price}} \right)^2}{\#Trades_{\text{Rounds}}} \quad (1)$$

With respect to the variance and the deviation measure it can be observed that the baseline and the *arbitrageur* no information treatments differ significantly from the arbitrageur information treatment. This effect is more pronounced in earlier rounds. We conclude that the information condition has a significant effect on the participants in case of the passive agent: human participants are observed to trade closer to equilibrium in the information condition.

Result 3: *Information condition: The public information on the presence of software agents has a significant positive effect on human traders' ability to converge to equilibrium in the presence of the arbitrageur agent.*

To further explore the determinants of non equilibrium trade we run a GLS regression using the complete data including the baseline treatment. The treatment variables, time, and price are included in the regression to measure deviations from equilibrium trade. Significant negative signs imply trade closer to equilibrium.

The dependent variable deviation takes 0 in equilibrium and is larger when the deviation from the fundamental value is larger. Since deviations in each of the 18 sessions are likely to be correlated, and in a statistical sense are not strictly independent, an error components econometric model with the session as the random component is used. Table 2 reports the results. Confirming our discussion on deviation from equilibrium prices we find that deviations increase in the *arbitrageur* regime without information about the agents' presence. Deviations decline over time and are more pronounced in lower priced contracts. We attribute smaller deviations in later periods to learning and larger deviations for lower priced contracts to rounding of prices in the market.

Result 4: *Baseline treatment: Surprisingly, the introduction of software agents results in lower market efficiency in the no information treatment when compared to the baseline treatment.*

4 Related Work

Experimental asset market literature is vast in size. In the following some experimental asset market literature related to this experiment is reviewed. Common to most market experiments is the incorporation of public as well as private information characteristics, e. g. Copeland and Friedman (1991), Forsythe and Lundholm (1990), and Plott and Sunder (1988). In these experiments, trade is motivated by differences in both private information

and private valuations. Further, Smith et al. (1988) and Peterson (1993) have undertaken pure common value markets. In these cases, traders are endowed with only public and no private information regarding the expected common value. These markets motivate market trading through a combination of different risk attitudes of the traders, and different expectations concerning the other traders' strategies.

Much work has also been done about automated agents in the context of electronic commerce. A good starting point on agent mediated electronic commerce can be found in Guttman et al. (1998). Hereafter, the focus is on related work that is concerned about automated negotiation. In particular market-based approaches are reviewed, which provide a market institution and a set of rules to do the negotiation. In this context agents negotiate in a competitive environment, yet there are other approaches in the AI community, such as collaborative agents, that will not be reviewed here.

In the artificial intelligence community agent tournaments are conducted in an increasingly complex environment, see for example, the Trading Auction Competition (TAC) described by Wellman et al. (2001, 2002). In TAC traders arranged in groups of eight are assigned the role of travel agents charged with the task of arranging and automatically shopping for trips. The challenging part for agents' design is to address the interdependence of the tasks necessary to complete a trip, and the ability to reason about others strategies in a thin market of automated traders and in a continuous time frame.

In experimental economics community work on programmed strategies has also been done by conducting tournaments (Abreu and Rubinstein 1988, Rust et al. 1993, 1994, Selten et al. 1997). Rust et al. report on the Santa Fe Double Auction tournament, where researchers were invited to submit software agents that compete on a CDA market against one another. The focus was to find successful trading strategies out of the submitted set of agents. The most successful strategy in this tournament can be described as rather parasitic: sitting in the background and exploiting the strategies of other traders. In addition, they report about an evolutionary tournament, where the percentage of traders was adjusted in accordance to the success of a strategy over time. Parallel to the tournament there has been a discussion on the lower bound of traders' intelligence to act similarly to human traders in a market institution (Gode and Sunder 1993a&b, Cliff and Bruten 1997).

In the more expanding field of agent-based computational finance multifarious work can be found. Lettau (1997) investigates how closely evolutionary (genetic algorithm) techniques can achieve the optimum in a purchase situation for a risky asset. Other early applications for these techniques are a genetic algorithm environment for learning to construct a test for general equilibrium in a foreign exchange market scenario (Arifovic 1996), and a learning algorithm in Routledge (1994) that addresses investors' optimal choice in a repeated one-shot decision situation for a portfolio when costly information signals are available. Alternatively, approaches using neuronal network based agents as in Beltratti and Margarita (1992) and Beltratti et al. (1996) can find valuable applications in decentralized price-finding institutions. A further starting-point for agents' design can be found in the economic mechanism design literature, see Varian (1995) for an introduction to this discipline. Interested readers might also review references on possible adverse effects on the employment of artificial traders on real financial markets that can be found in Leland and Rubinstein (1988), and Varian (1998). Gennotte and Leland (1990) provide a rational expectations model that draws on these experiences and aim to explain financial instability and discontinuities.

5 Conclusions

This paper reports on an experiment where human subjects and software traders participate in a double auction market institution simultaneously. In this environment traders can buy and sell American futures. The experiment was conducted in a controlled laboratory environment, and six statistically independent observations for each treatment were collected. The experiment was designed to disentangle the effects of the introduction of software agents and the psychological effect of the public announcement. The main result is that human traders do not crowd out when the participation of software agents is made public. Moreover, the public information on the presence of software agents has a significantly positive effect on human traders' ability to converge to equilibrium in the presence of the arbitrageur agent. Surprisingly, the introduction of software agents results in lower market efficiency in the no information treatment when compared to the baseline treatment.

The focus of this research is to shed light on economic and psychological effects imposed on human beings when interacting with software agents in a competitive environment. The

first studies conducted are to serve as a starting point to obtain a deeper insight in how to apply technically well studied software agents in an environment with bounded rational human beings. On a methodological level we are concerned with the rather high variability of individual session averages for efficiency and behavioral variables observed in this and other market experiments. We feel confident that our design and the statistical analysis using the permutation test and random effects GLS regressions provide a good description of the underlying effects. Evidence on CDA markets relying on a single independent observation for each treatment should be treated carefully and may require more repetitions.

It can be observed that behavioral and economic effects can be attributed to different experimental conditions. Behavioral differences can be observed between agent sessions and the baseline treatment. More specific, subjects start trading (all different contracts) at a later point in time in the presence of the arbitrageur agent when compared to the baseline treatment. The change in human behavior due to the introduction of a programmed trader might be agent specific. An environment using more active zero intelligence traders might support human traders to start trading at an earlier point in time. With respect to the information condition human traders are observed to act more efficiently in a market environment when information on software agents is available. This might be the most surprising result of the study since standard economic theory would predict no treatment effects. To generalize the results, the introduction of different types of agents in the current framework might be interesting. Further, it seems that auction experiments with human and artificial traders might be a promising area as well.

6 References

- Abreu, D. and Rubinstein, A. (1988) "The Structure of Nash Equilibrium in Repeated Games with Finite Automata", *Econometrica*, 56(6), 259-1281.
- Akerlof, G.A. (1970) "The Market for "Lemons": Quality Uncertainty and the Market Mechanism", *Quarterly Journal of Economics*, 84(3), 488-500.
- Arifovic, J. (1996) "The behavior of the exchange rate in the genetic algorithm and experimental Economies", *The Journal of Political Economy*, 104(3), 510-541.
- Beltratti, A. and Margarita, S. (1992) "Evolution of trading strategies among heterogeneous artificial economic agents", in Meyer, J.A., Roitblat, H.L. and Wilson, S.W. (Eds.) "From Animals to Animats 2", MIT Press, Cambridge, MA, 494-501.
- Beltratti, A., Margarita, S. and Terna, P. (1996) "Neural Networks for Economic and Financial Modeling", International Thomson Computer Press, London, UK.
- Cliff, D. and Bruten, J. (1997) "Minimal-intelligence agents for bargaining behaviors in market-based environments", HP Labs Technical Reports HPL 97-91, Hewlett-Packard Laboratories, Bristol, England.
- Copeland, T.E. and Friedman, D. (1991) "Partial Revelation of Information in Experimental Asset Markets", *Journal of Finance*, 46(1), 265-295.
- Das, R., Hanson, J.E., Kephart, J.O. and Tesauro, G. (2001) "Agent-Human Interactions in the Continuous Double Auction", in Proceedings of the International Joint Conferences on Artificial Intelligence (IJCAI), Seattle, USA (August, 2001).
- Forsythe, R. and Lundholm, R. (1990) "Information Aggregation in an Experimental Market", *Econometrica*, 58(2), 309-347.
- Forsythe, R., Nelson, F., Neumann, G.R. and Wright, J. (1992) "Anatomy of an Experimental Political Stock Market", *American Economic Review* 82(5): 1142-1161.
- Forsythe, R., Rietz, T.A. and Ross, T.W. (1999) "Wishes, expectations and actions: a survey on price formation in election stock markets", *Journal of Economic Behavior & Organization*, 39(1), 83-110.
- Friedman, D. (1984) "On the efficiency of Experimental Double Auction Markets", *American Economic Review* 74(1): 60-72.
- Friedman, D. (1993) "The Double Auction Market Institution: A Survey", in Friedman, D. and Rust, J. (Eds.) "The Double Auction Market – Institutions, Theories, and Evidence", Santa Fe Institute studies in sciences of complexity 14, Addison-Wesley, 3-25.
- Gennotte, G. and Leland, H. (1990) "Market Liquidity, Hedging, and Crashes", *American Economic Review*, 80(5), 999-1021.
- Gode, D.K. and Sunder, S. (1993a) "Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality", *Journal of Political Economy*, 101(1), 119-137.
- Gode, D.K. and Sunder, S. (1993b) "Lower Bounds for Efficiency of Surplus Extraction in Double Auctions", in Friedman, D. and Rust, J. (Eds.) "The Double Auction Market – Institutions, Theories, and Evidence", Santa Fe Institute studies in sciences of complexity 14, Addison-Wesley, 199-219.
- Grossklags, J., Schmidt, C. and Siegel, J. (2000) "Dumb software agents on experimental electronic markets" in Gesellschaft für Informatik e.V. (eds.), *Informatiktage 2000*, Konradin Verlag, Leinfelden-Echterdingen, Germany.
- Guttman, R.H., Moukas, A.G. and Maes, P. (1998) "Agent-mediated Electronic Commerce: A Survey", *The Knowledge Engineering Review*, 13(2), 147-159.

- Hansen, J., Schmidt, C. and Strobel, M. (forthcoming) "Manipulation in Political Stock Markets - Preconditions and Evidence", *Applied Economics Letters*.
- Leland, H. and Rubinstein, M. (1988) "Comments on the Market Crash: Six Months After", *Journal of Economic Perspectives*, 2(3), (Summer 1988), 45-50.
- Lettau, M. (1997) "Explaining the facts with adaptive agents: The case of mutual fund flows", *Journal of Economic Dynamics and Control*, 21(7), 1117-1148.
- Ockenfels, A. and Roth, A.E. (2002) "The Timing of Bids in Internet Auctions: Market Design, Bidder Behavior, and Artificial Agents", *Artificial Intelligence Magazine*, Fall 2002, 79-87.
- Peterson, S. (1993) "Forecasting Dynamics and Convergence to Market Fundamentals. Evidence from Experimental Asset Markets", *Journal of Economic Behavior and Organization*, 22(3), 269-284.
- Plott, C.R., and Sunder, S. (1988) "Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets", *Econometrica*, 56(5), 1085-1118.
- Routledge, B.R. (1994) "Artificial selection: genetic algorithms and learning in a rational expectations model", Technical Report 1997-47, Graduate School of Industrial Administration, Carnegie Mellon University, Pittsburgh.
- Roth, A.E. and Ockenfels, A. (2002) "Last Minute Bidding and the Rules for Ending Second Price Auctions: Evidence from eBay and Amazon Auctions on the Internet", *American Economic Review*, 92(4), 1093-1103.
- Rust, J., Miller, J. and Palmer, R. (1993) "Behavior of Trading Automata in a Computerized Double Auction Market", in Friedman, D. and Rust, J. (Eds.) "The Double Auction Market – Institutions, Theories, and Evidence", Santa Fe Institute studies in sciences of complexity 14, Addison-Wesley, 155-198.
- Rust, J., Miller, J.H. and Palmer, R. (1994) "Characterizing effective trading strategies. Insights from a computerized double auction tournament", *Journal of Economic Dynamics and Control* 18(1), 61-96.
- Schmidt, C. and Werwatz, A. (2002) "How well do markets predict the outcome of an event? The Euro 2000 soccer championships experiment", Discussion Papers on Strategic Interaction 2002-9, Max Planck Institute for Research into Economic Systems, Jena, Germany.
- Selten, R., Mitzkewitz, M. and Uhlich, G.R. (1997) "Duopoly Strategies Programmed by Experienced Players", *Econometrica*, 65(3), 517-555.
- Smith, V.L. (1962) "An Experimental Study of Competitive Market Behavior", *The Journal of Political Economy*, 70(2), 111-137.
- Smith, V.L., Suchanek, G. and Williams, A. (1988) "Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets", *Econometrica*, 56(5), 1119-1151.
- Varian, H.R. (1995) "Economic mechanism design for computerized agents", in Proceedings of USENIX Workshop on Electronic Commerce, New York.
- Varian, H.R. (1998) "Effect of the Internet on Financial Markets", School of Information Management and Systems, University of California, Berkeley.
- Wellman, M.P., Wurman, P.R., O'Malley, K., Bangera, R., Lin, S., Reeves, D. and Walsh, W. (2001) "Designing the market game for a trading agent competition", *IEEE Internet Computing*, 5(2), 43-51.
- Wellman, M.P., Greenwald, A., Stone, P. and Wurman, P.R. (2002) "The 2001 Trading Agent Competition", in Proceedings of Fourteenth Innovative Applications of Artificial Intelligence Conference (IAAI-02), Edmonton, Canada, 935-941.




Figure 1: Model of the fundamental value of each firm's contract

Table 1: Summary of descriptive results, average of 6 cohorts (standard deviation)

		No Information	Information
	No Agent	Arbitrageur	Arbitrageur
Average Payoffs (in ECU)			
Software Agents	-	*100323 (248.9)	*100282 (285.7)
Human Traders	100000 (7447.5)	99946 (5626.8)	99953 (7469.9)
Number of Market Trades			
Total	141.7 (34.8)	108.2 (71.9)	119.2 (73.0)
Human to Human (h2h)	141.7 (34.7)	91.2 (64.2)	92.5 (34.5)
Human to Agent (h2a)	- -	17.0 (13.9)	26.7 (40.3)
Unit Portfolio at Market Prices			
Overall	97.82 (5.23)	102.11 (5.39)	99.61 (3.52)
Rounds 1-6	106.70 (5.87)	108.52 (11.95)	104.17 (2.43)
Rounds 7-12	90.73 (8.34)	98.63 (2.86)	97.96 (3.32)
Time until Market Prices for all Contracts were Available			
Time in sec.	616.17 (130.59)	1556.50 (954.08)	1587.33 (1031.80)
Time (% Total Experiment)	12.71%	34.59%	35.27%
Deviation from Fundamental Value			
Overall	0.21 (0.05)	0.31 (0.22)	0.17 (0.02)
Round 1-6	0.17 (0.08)	0.27 (0.28)	0.13 (0.02)
Round 7-12	0.24 (0.03)	0.33 (0.22)	0.21 (0.04)
Theta-Variance			
Overall	0.09 (0.06)	1.33 (2.50)	0.05 (0.01)
Round 1-6	0.09 (0.12)	2.30 (5.50)	0.03 (0.01)
Round 7-12	0.09 (0.03)	0.89 (1.29)	0.06 (0.02)

*** p<0.001; ** p < 0.01; * p < 0.1 t-test payoff<>100000

Table 2: Random-effects GLS regression

Overall Model Fit		
Number of Observations	2214	
Number of Groups	18	
Observations per Group Avg.	123.0	
R2 (within)	0.025	
R2 (between)	0.187	
R2 (overall)	0.037	
Parameter Estimates		
Independent Variables	Dependent Variable: Deviation	P > z
Constant	0.523 (0.058)	***
Arbitrageur, no information	0.141 (0.067)	*
Price	-0.017 (0.002)	***
(standard error)		
*** p<0.001, ** p < 0.01, * p < 0.05		

Appendix A: Tick prices and fundamental value (6 cohorts)




Figure 2: Tick prices and fundamental value contract A




Figure 3: Tick prices and fundamental value contract B




Figure 4: Tick prices and fundamental value contract C




Figure 5: Tick prices and fundamental value contract D

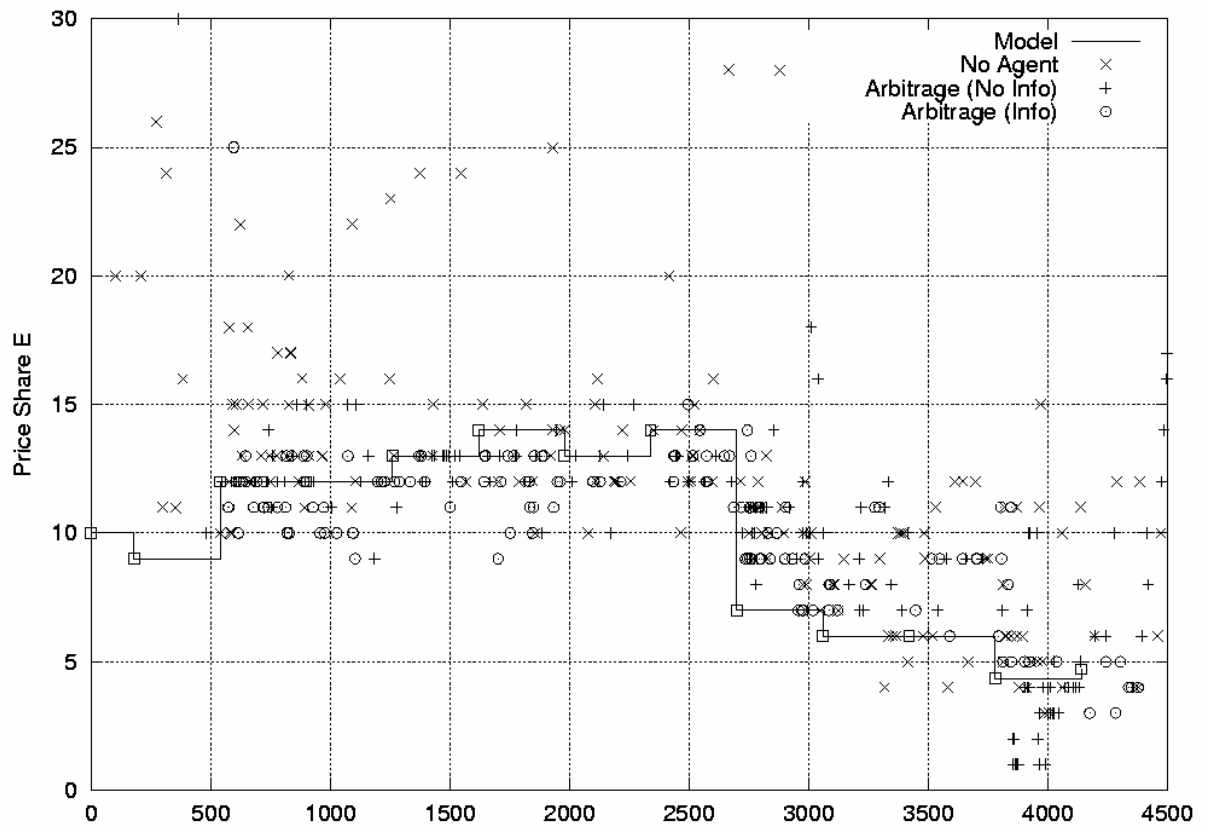


Figure 6: Tick prices and fundamental value contract E

Appendix B: Translation of the Instructions

Dear participants,

Thanks for taking the time to participate to the experiment. In first place: please turn off your cellular phone.

WHAT IS THIS EXPERIMENT ABOUT?

During the following experiment you buy and sell contracts on an electronic market. In order to trade you will be provided with an initial endowment of 10 Euro = 100.000 ECU (Experimental Currency Unit). This amount is the same for each participant.

WHAT IS THE STRUCTURE OF THE EXPERIMENT?

Each trader gets information about the initial value of each stock (in points). During the experiment each trader receives information about the stock. The experiment is divided in 10 rounds. Each round is structured as follows:

- step 1. reception of a private information
- step 2. four minutes for trading
- step 3. reception of a public information
- step 4. two minutes for trading

HOW DO YOU TRADE?

In this experiment you are able to trade 5 different stocks. These stocks represent 5 firms. From those 3 of the firms have a relatively high (initial) value (stock A, B and C). The other two companies are compared to that of a lower value (stock D and E). Stocks can be bought and sold as follows:

- Separately on the market, this means trade between the traders.
- As a bundle from and to the Bank. One bundle equals one stock of each firm (unit-portfolio), in this experiment one piece of stock A, B, C, D, and E. A bundle can be bought from and sold to the bank at every point in time during the experiment at 100 ECU.

VALUE OF A STOCK

Each stock is characterized by:

- Points (P), which show the importance of a firm relative to all other firms in the market
- Monetary value on the market (in ECU)

Initial value of the Stocks:

Stock A	26 Points	=	26 ECU
Stock B	26 Points	=	26 ECU
Stock C	26 Points	=	26 ECU
Stock D	12 Points	=	12 ECU

Stock E	10 Points	=	10 ECU
Sum	100 Points		100 ECU

Attention: After the first information has been broadcast, 1 Point is not equal 1 ECU anymore!

THE INFORMATION

There are two type of information available that alternate: private information (each trader receives his own personal information) and public information. After receiving private information you have until the broadcast of the public information a so-called “insider information”. Your private information won’t be 100% accurate, yet as a general rule it will reflect the trend the points of a firm will change. The public information might be regarded as answer since it presents the real change in information. The public information will be broadcasted simultaneously and with the same content to each trader.

During the experiment you will receive information (in points) which provide information about the value of a stock. An increase in points means an increase in the value of the stock. Beware that an increase in 1 point does not mean regularly a proportional increase of the firm value of 1 ECU. Rather a change in information of one firm results into a new weighting of the stock in the bundle. This new weighting is responsible for the current value of a stock.

Example: Stock X has initially 26P. You receive the private information: “PRIVATE so far: A – 26, B –26, C – 26, D – 12, C – 10: stock A: change of + 10”. This means that according to your information stock A has now 36 points. To calculate the value stock A in ECU apply the following rule: Add up the points of all stocks (according to your information 110 points). This means that 110 points have do be distributed to 100 ECU. Stock A is therefore $36P \cdot 100 / 110 = 32.7$ ECU worth. After a few minutes all traders receive the public information: “PUBLIC information: stock A changes by +5P”). Every trader has now the information that stock A is 31 points worth! You may find the value of stock A applying the above method ($31P \cdot 100 / 105 = 29.5$ ECU).

PARTICIPATION OF A SOFTWARE TRADER ⁵

A software trader is participating in this experiment. The “programmed trader” has the same information like an average human participant.

WHAT KIND OF HELP DOES THE SOFTWARE PROVIDES ?

Please refer to the link “Help” on the screen (top right)

PAYOFF RULES

At the end of the experiment your cash balance and your stock portfolio will be paid out. Stocks will be liquidated according to the rules given in the example. This means that the points of a contract will be multiplied by 100 and divided by the sum of the points of all different firms. To calculate the Euro value of your assets divide by 10,000.

⁵ This paragraph was only included in the treatment with information on software agents.