

Interaction of Human and Artificial Agents on Double Auction Markets

Simulations and Laboratory Experiments

Jens Grossklags
University of California at Berkeley
School of Information Mgt. & Systems
102 South Hall
94720 Berkeley, CA, USA
jensg@sims.berkeley.edu

Carsten Schmidt
Max-Planck Institute for Research
into Economic Systems
Kahlaische Str. 10
07745 Jena, Germany
cschmidt@mpiew-jena.mpg.de

Abstract

This paper provides an overview on the simulations and experiments we have done in order to better understand human-agent interaction in a market environment. We find that the introduction of software agents does not necessarily induce a more efficient market. More surprisingly, information on the existence of software agents in the market environment results in more efficient behavior of human traders.

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.

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 auctionblitz 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 the following we are interested in the influence of human-agent interaction on human market behavior. When going back to the eBay example, we found in many seminar presentations a large share of the audience not aware of

software agents acting in this environment. The experimental method gives the opportunity to control whether subjects know about software agents or not. So we are able to observe whether individual knowledge on software agents influences human behavior and the market outcome.

The intention of this paper is to provide an overview on the simulations and experiments we have done in order to better understand human-agent interaction in a market environment. In the next section the experiment infrastructure is sketched. In Section 3 a simple simulation with selected software agents is presented. The main experimental results from human-software agent interaction are presented in Section 4. Related work is briefly discussed in Section 5. Finally, concluding remarks are given and open research questions for further exploration are sketched in Section 6.

2. Market Institution and Information

The artificial stock market uses the continuous double auction (CDA) with queues. The implemented mechanism allows submitting bids and asks simultaneously and asynchronously without restrictions made by a spread improvement rule. The emission of contracts on this market is implemented via a bundle mechanism and the market therefore resembles closely the design of the Iowa Electronic Markets [Forsythe *et al.*, 1992 & 1999]. The bundle consists of a standardized unit-portfolio of one piece of each contract at a fixed price: this can be bought from or sold to the bank at any time and any quantity.

The agents are allowed three valid operations: (1) posting market orders (bids implement buy orders and asks sell orders), (2) deleting own market orders, and (3) buying/selling bundles at the bank.

The marketplace 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.

In order to apply a strategy human and software agents require information about the market. One of the design goals of the market-agent interface was to give the same information to human and to artificial agents. For this purpose we analyzed the information available to human traders and categorized the information into the following macro-categories:

- public rule information (payoff rules of the Futures, stop time of market)
- public market information (Prices, best 5 bids and asks, server time)
- private information (portfolio, liquid money)
- transaction status information (validity and success of transaction)

Out of the above categories we designed the framework for the market-agent interface. EXtensible Markup Language (XML) was chosen as high level data exchange format which allows for a continuous and asynchronous trading of human and programmed trader.

3. Simulations using basic programmed strategies

We conducted a simple simulation experiment [Grossklags et al., 2000]. During a 60 minute period programmed traders could trade 16 different futures each with a 1/16 a priori probability of paying off. Beside the initial probability of a future software agents only relied on market price information in order to trade. One randomly determined contract (which was then held constant) paid off the price of the bundle, all other 15 contracts 0.

The main concern with regards to the simulation is which strategies are incorporated. We deferred from running a tournament, where other researchers could send in their strategies, and implemented strategies that we have observed by human traders in former political stock markets. In either case, there can be always the argument made that a specific strategy is not included.

More specific, the following strategies have been implemented: (a) a stochastic strategy, which is aware of the market rules, (b) a “fundamental strategy” in which a trader buys a contract according to his initial information and waits for the final event to happen, (c) A “follower strategy” is one in which a trader echoes the actions of the other subjects of the market based on a pre-defined signal. In our case, the agent reacts to the price of a contract (the signal), buying when the price of the contract is high or low, (d) an “arbitrageur strategy” tries to make instant arbitrage by observing price differences of the unit-portfolio between market and bank, and (e) the speculator, who implements a risky arbitrage strategy in postponing going into cash and take some risk.

A lesson learned from this very simple simulation is that speed matters. To our surprise the arbitrageur only yielded a not significant profit on average and even lost money on occasions. It turned out that the agent could often not complete the whole set of trades to employ its arbitrage strategy. This is due to the market mechanism, which does not allow for atomic transaction for a set of trades in order to gain arbitrage. In the case of the arbitrageur the result was an unbalanced portfolio, which is equal to engaging in a risk.

The most successful strategy (given the participating agents) is a simple price following strategy that buys a well diversified portfolio of relatively cheap contracts. Given the set of implemented agents and the WTA market rule, the most successful strategy is the follower low, who attempts

to buy a portfolio of low priced contracts. This agent starts his first market transaction late, waiting for the other agents to generate prices. At this point the strategy is very effective in conducting market transactions that implement a successful portfolio (see Table 2). Similar to the most successful strategy described by Rust et al. [1994] the agent can be described as “sit in the background”.

4. Human-software interaction

For the experiments with human-agent interaction the passive trading strategy “arbitrageur” out of the strategies used in the simulation was selected to compete with human traders in the laboratory. One argument for that choice is that speed does not matter on a thin laboratory market. In addition, we were concerned that the other implemented strategies might loose money an just distributing “mana from heaven” to the human subjects in the experiment. Most important, Das et al. [2001] already tested the influence of active trading strategies in a similar environment. The influence of a passive trading strategy on human behavior was to our knowledge not tested in human-agent environments, so far.

The experiment was conducted using 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). At the end, each contract pays off a percentage of the total bundle.

A storyboard was designed in order to provide a constant environment for all sessions and treatments that were determined by the experimenters with a rolling dice. The storyboard describes the fundamental value of the five different contracts, and the private information points distributed to the six human traders’ roles. Thus, each human trader role, say for example trader number 2, was assigned the same information throughout all 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 subjects in the public information phase.

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. For this treatment no information about software agents was given.⁴ The experiment was conducted in a controlled laboratory environment, Each session consisted of a market with 6

⁴ Note that in contrast to experiments in psychology it is an enforced general standard in economic experiments not to lie to participants. Therefore, we have not conducted a treatment without agents while still providing information that software agents are present.

human participants and 6 streams of information. In the arbitrageur treatment the software agent 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 software agents there was no other difference between the individual markets. Altogether 18 sessions have been run with 108 different human participants and 12 programmed traders and thus six independent observations for each treatment were collected (for instructions and details see Grossklags and Schmidt [2003]).

The main result is that human traders do not crowd out when the participation of software agents is made public. However, the results show that there is a significant decrease of human-to-human trades in the treatments including an *arbitrageur* when compared to the baseline treatment. 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.

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 specifically, 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 trading agent might be agent-strategy specific. An environment using more active agents might support human traders in starting 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.

5. Related work

Closely related to our experiments 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, Cliff and Bruten, 1997] or a modified Gjerstad-Dickhaut [Gjerstad and 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. 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.

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., 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. 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, 1993, Cliff and Bruten, 1997].

6. Conclusion

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.

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 specifically, 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 trading agent might be agent-strategy specific. An environment using more active agents might support human traders in starting 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 commodity auction experiments with human traders and artificial agents might be a promising area of research as well.

7. References

- [Abreu and Rubinstein, 1988] Dilip Abreu and Ariel Rubinstein (1988) "The Structure of Nash Equilibrium in Repeated Games with Finite Automata", *Econometrica*, 56(6): 259-1281.
- [Cliff and Bruten, 1997] Dave Cliff and Janet Bruten (1997) "Minimal-intelligence agents for bargaining behaviors in market-based environments", HP Labs Technical Reports HPL 97-91, Hewlett-Packard Laboratories, Bristol, England.
- [Das *et al.*, 2001] Rajarshi Das, James E. Hanson, Jeffrey O. Kephart and Gerald Tesauro (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 *et al.*, 1992] Robert Forsythe, Forrest Nelson, George R. Neumann and Jack Wright (1992) "Anatomy of an Experimental Political Stock Market", *American Economic Review* 82(5): 1142-1161.
- [Forsythe *et al.*, 1999] Robert Forsythe, Thomas A. Rietz and Thomas W. Ross (1999) "Wishes, expectations and actions: a survey on price formation in election stock markets", *Journal of Economic Behavior & Organization*, 39(1): 83-110.
- [Gjerstad and Dickhaut, 1998] S. Gjerstad and J. Dickhaut (1998) "Price Formation in Double Auctions", *Games and Economic Behaviour*, 22:1-29, 1998.
- [Gode and Sunder, 1993] Dhananjay K. Gode and Shyam Sunder (1993) "Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality", *Journal of Political Economy*, 101(1): 119-137.
- [Grossklags *et al.*, 2000] Jens Grossklags, Carsten Schmidt and Jonathan Siegel (2000) "Dumb software agents on experimental electronic markets" in Gesellschaft für Informatik e.V. (eds.), *Informatiktage 2000*, Konradin Verlag, Leinfelden-Echterdingen, Germany.
- [Grossklags and Schmidt, 2003] Jens Grossklags and Carsten Schmidt (2003) *Artificial Software Agents on Thin Double Auction Markets - A Human Trader Experiment*, Proceedings of the IEEE/WIC International Conference on Intelligent Agent Technology (IAT-2003), Halifax, Canada.
- [Leland and Rubinstein, 1988] Hayne Leland and Mark Rubinstein (1988) "Comments on the Market Crash: Six Months After", *Journal of Economic Perspectives*, 2(3), (Summer 1988): 45-50.
- [Ockenfels and Roth, 2002] Axel Ockenfels and Alvin E. Roth (2002) "The Timing of Bids in Internet Auctions: Market Design, Bidder Behavior, and Artificial Agents", *Artificial Intelligence Magazine*, Fall 2002, 79-87.
- [Roth and Ockenfels, 2002] Alvin E. Roth and Axel Ockenfels (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 *et al.*, 1994] John Rust, John H. Miller and Richard Palmer (1994) "Characterizing effective trading strategies. Insights from a computerized double auction tournament", *Journal of Economic Dynamics and Control*, 18(1): 61-96.
- [Selten *et al.*, 1997] Reinhard Selten, Michael Mitzkewitz and Gerald R. Uhlich (1997) "Duopoly Strategies Programmed by Experienced Players", *Econometrica*, 65(3): 517-555.
- [Smith, 1962] Vernon L. Smith (1962) "An Experimental Study of Competitive Market Behavior", *The Journal of Political Economy*, 70(2): 111-137.
- [Varian, 1998] Hal R. Varian (1998) "Effect of the Internet on Financial Markets", School of Information Management and Systems, University of California, Berkeley.
- [Wellman *et al.*, 2001] Michael P. Wellman, Peter R. Wurman, Kevin O'Malley, Roshan Banger, Shou-de Lin, Daniel Reeves and William E. Walsh (2001) "Designing the market game for a trading agent competition", *IEEE Internet Computing*, 5(2): 43-51.
- [Wellman *et al.*, 2002] Michael P. Wellman, Amy Greenwald, Peter Stone and Peter R. Wurman (2002) "The 2001 Trading Agent Competition", in Proceedings of Fourteenth Innovative Applications of Artificial Intelligence Conference (IAAI-02), Edmonton, Canada, 935-941.