# A Comparison of Effective Trading Agents in Double Auction Markets

Shu-Heng Chen and Chung-Ching Tai

AI-ECON Research Center, Department of Economics, National Chengchi University, Taipei 11605, Taiwan

**Abstract.** In this paper, we conducted agent-based simulations in double auction markets with strategies gathered from the literature. The goal is to compare various styles of strategies based on their effectiveness. In shows that adopting heuristics, adaptive strategies, or even more innovative algorithms such as Genetic Programming, can yield different advantages in a profit-variation exchange setup. This result can be the foundation of further comparison between human agents and software agents.

## 1 Introduction

The importance of studying individual strategies in double auction (DA) markets lies in two aspects. First, if observed individual behavior is the realization of Bayesian-Nash equilibrium, it will be necessary to gain better knowledge of the equilibrium trading strategies if they exist (Rust et al. 1994). Second, a series of comparisons between human traders and software trading agents exhibits interesting results, including the superiority of software agents over humans, and the contingent behavior of human traders whether they are notified of the existence of their software counterparts (Das et al. 2001, Taniguchi et al. 2004, and Grossklags & Schmidt 2006). However, before the comparison, we have little knowledge about the characteristics of the software agents. Therefore, a better understanding of trading strategies is also necessary in this vein.

The pioneering work of exploring individual characteristics of effective trading strategies in double auction markets is Rust et al. (1994) 's tournaments held in Santa Fe Institute. Rust et al. raised 30 trading algorithms and categorized them according to whether they were simple or complex, adaptive or nonadaptive, predictive or nonpredictive, stochastic or nonstochastic, and optimizing or nonoptimizing. The result is rather surprising: the winning strategy is simple, nonstochastic, nonpredictive, nonoptimzing, and most importantly nonadaptive. In spite of this, other strategies possessing the same characteristics may performed poorly. As a result, it remains an open question "whether other approaches from the literature on artificial intelligence might be sufficiently powerful to discover effective trading strategies." (Rust et al. 1994, pp. 94–95)

The goal of our study is to examine the same problem from a slightly different point of view. Instead of discovering the essential characteristics of the winning strategy, this study is to discover the benefits of adopting strategies with different characteristics. A different but more general classification of strategies is proposed, and a two-dimensional measurement of profitability and stability is used to evaluate strategies.

This paper is organized as follows: Section 2 depicts the experimental design, including market mechanism, trading strategies, and experiment settings. Classification of strategies is also presented in 2.2. Results, evaluations, and analysis of experiment are presented in section 3; conclusions, in section 4.

## 2 Experimental Design

Experiments in this paper were conducted in AIE-DA (Artificial Intelligence in Economics - Double Auction) platform which is an agent-based discrete double auction simulator with build-in software agents.

### 2.1 Market Mechanism

Although most agent-based simulations are conducted under Continuous Double Auction (CDA) mechanism, we employed discrete double auction markets in our simulations for an important reason described below: The significant advantage of software agents over human traders is commonly found in the studies of agent-human interaction (Das et al. 2001, Taniguchi et al. 2004, and Grossklags & Schmidt 2006), despite quite different software agents were employed in the studies. Although one possible source of this advantage can be attributed to the agents' speed, we cannot distinguish whether it is the speed of calculation or the speed of learning/adaptation. In a discrete time scheme, we can easily identify the attribute, since the advantage of speed of calculation can be extinguished.

AIE-DA is inspired by Santa Fe double auction tournament held in 1990, and in this study we adopted the same token generation process as Rust et al. (1994)'s design. A random seed of 6453 was used to create the demand and supply schedules of the markets. We also adopted AURORA trading rules so that our results can be compared with Rust et al. (1994)'s tournaments. Our experimental markets consisted of four buyers and four sellers. Each of the strategies was randomly assigned as one of these eight traders and stayed in that position for the whole simulation. Considering the vast number of combinations and permutations of traders, we did not try out all possible trader arrangements. Instead, random match-ups were created and the number of match-ups of each series of experiment was set to be 1,000 for shorter simulation(1,000 days) and 300 for longer simulation(7,000 days).

### 2.2 Classification of Strategies

Besides strategies selected from Rust et al. (1994)'s tournaments, we also collected trading strategies from the double auction literature as follows: **Kaplan**, **Ringuette**, and **Skeleton** from Rust et al. (1994)'s tournament; **ZIC** from

Gode and Sunder (1993); **ZIP** from Cliff and Bruten (1997); **Markup** from Zhan and Friedman (2007); **Gjerstad-Dickhaut** (**GD**) from Gjerstad and Dickhaut (1998); **BGAN** from Friedman (1991); **Easley-Ledyard** (**EL**) from Easley and Ledyard (1993); **Empirical** strategy is inspired by Chan et al. (1999), and it works in the same way as Friedman (1991)'s BGAN but develops its belief by constructing histograms from opponents' past shouted prices.[1] We also included **Genetic Programming** (**GP**) as one possible trader candidate, as it can evolve effective trading rules in double auction markets (Chen et al. 2009).

GP agents in this study adopt standard crossover and mutation operations, which means no election, ADFs or other mechanisms are implemented. At the beginning of each trading day, a GP trader randomly picks a strategy from his population of strategies and use it for the whole day. The performance of each selected strategy is recorded, and if a specific strategy is selected more than once, a weighted-average performance will be made to emphasize later experiences. Genetic operations take place and GP traders' strategies are updated every $N$ days, where $N$ is called "select number." To avoid the flaw that a strategy is deserted simply because it was not selected, we set N twice the size of the population so that theoretically each strategy has the chance to be selected twice. Tournament selection is implemented and the size of the tournament is 5 however big the size of the population is.

Although most of the strategies were created for the purpose of studying price formation processes, we still sent them to the "battlefield" because they can represent, to a certain degree, various types of trading strategies we are interested in.

In classifying these strategies, we used a different criterion from Rust et al. (1994)'s method. We categorized strategies based on how they were designed to tackle complex problems. Generally, there are three ways to design strategies to play against others in a double auction game:

- **Heuristic**: Heuristics are rule of thumbs and are fixed throughout the game.
- **Adaptive**: Adaptive strategies are those possessing *rules* of updating strategy parameters or referred statistics.
- **Innovative**: Innovative processes can modify, renovate, or even create brand new structure of the strategies. They are different from adaptive strategies in that adaptive strategies limit their search space in a certain range based on the model they are built upon, while innovative methods have fewer limitations.

Table 1 is the result of the classification of the strategies included in our experiments. Like other taxonomies, this classification is subjective and may lack precise measurement. However, it provides us a way to observe how strategies work and what their limits are.

[1] Named by or after their original designers, these strategies were modified to accommodate our discrete double auction mechanism in various ways. As a result, they may not be 100% the same as they originally were, but they were modified according to their original concepts.

**Table 1.** Classification of Strategies

| Strategy | Classification | Strategy | Classification |
|---|---|---|---|
| Truth Teller | Nonstrategic | Markup | Heuristic |
| Kaplan | Heuristic | GD | Adaptive |
| Ringuette | Heuristic | BGAN | Adaptive |
| Skeleton | Heuristic | EL | Adaptive |
| ZIC | Heuristic | Empirical | Adaptive |
| ZIP | Adaptive | GP | Innovative |

## 3 Results and Discussions

In this study, we conducted two sets of experiments. The first set of experiment excluded GP traders to observe which strategies would win. The second set of experiments added GP traders with different popoluation sizes so as to evaluate how innovative processes perform relatively to existing winning strategies.

### 3.1 Profit Ability and Stability

We evaluated strategies by their profit abilities as well as the variations in their profits. Profit ability is measured in terms of average profits earned and individual efficiencies.[2] In addition to profits, a strategy's profit stability was also considered in this study because in double auction markets, the variation of profits might be considered in human trading strategies, which are determined by human's risk attitudes (Kagel 1995).

Figure 1 is the result of the first set of experiment.[3] As shown in Figure 1, GD is the strongest player in terms of average profit earned. But if we take both profit ability and variation in profits into account, we can see that Markup and Truth Telling also perform well in terms of the exchange relationship between profits and stability. Other things being equal, a strategy with more profits and less variation is preferred. Thus the dashed line in Figure 1 marks the frontier of "effective" strategies that worth consideration.

Figure 2 also demonstrates the profit ability and variation in profits, but the variation here means standard deviations of profits earned by an agent who

[2] In order to evaluate the performance of each strategy, we adapted the notion of *individual efficiency*. Considering the inequality in each agent's endowment due to the random matching and the random market environment, individual efficiency is calculated as the ratio of one's actual profits to his/her theoretical surplus which is defined as the sum of the differences between one's intramarginal reservation prices and the market theoretical equilibrium price.

[3] In each of the 1,000 simulations, strategies were assigned to a specific trader throughout the time of 1,000 trading days. At the beginning of each day, traders' tokens were replenished to keep the market demand and supply intact. Thus the standard deviation of profits measures the variation of profits earned by a strategy encountering a fixed set of opponents in a specific market environment.

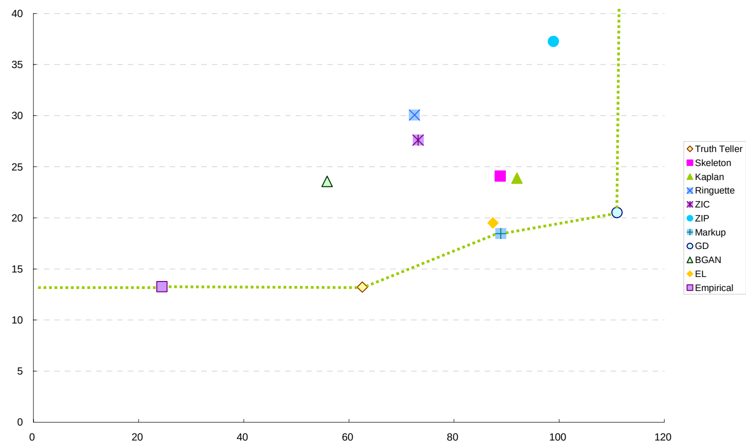


**Fig. 1.** Evaluation of agents, with the horizontal axis standing for their profitibility and the vertical axis standing for the average of standard deviation of their performances within trading days in each simulation.

encounters different opponents in different markets. We can see that strategies lying in the frontier are GD, ZIP, Skeleton, and Truth Telling.

In a brief summary, our study shows that adaptive strategies earn more profits but are also more volatile; heuristics earn median profits and are moderately volatile, while truth telling, which is non-strategic, earns less profits but is highly consistent in its performances.[4] In short, if these agents are presented and are available to humans, it can be expected that GD will be chosen by risk lovers, while Truth Telling will be employed by conservative decision makers.

### 3.2 Learning Agents

In the second set of the experiments, we added GP traders into the tournaments. We are interested in the following questions

– Can GP traders defeat other strategies?
– How many resources are required for GP traders to defeat other strategies?

[4] We want to emphasize that the real power of adopting heuristics or adaptive strategies can be seen if comparing those at the frontier, instead of all strategies. Making bids or asks is intelligent behavior, and the reward depends on both the styles of strategies and the underlying models or reasoning. For example, a winning strategy in double auction markets, either heuristic, adaptive, or innovative, may become one of the worst simply by adding a fixed amount of number in its shout price. We will be comparing too many things if we put all strategies together. We think only strategies that allow characteristics such as heuristic or adaptivity to bring their potential into full play are worth investigating. This may cause us to lose some information, but will help us clarify important features.

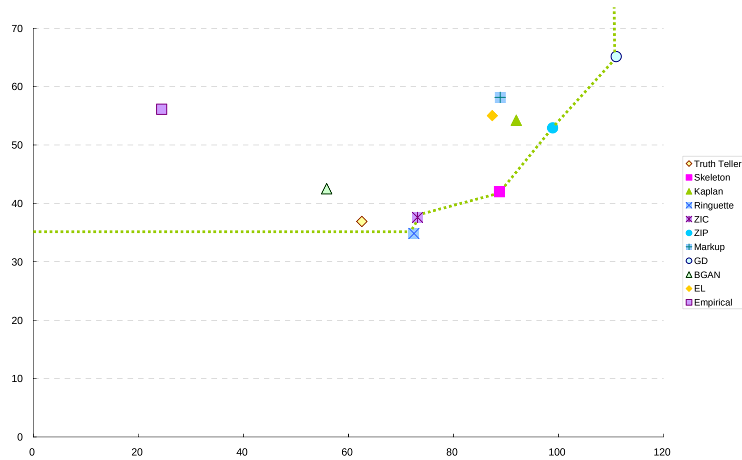


**Fig. 2.** Evaluation of agents, with the horizontal axis standing for their profitibility and the vertical axis standing for the standard deviations of their performances across all simulations.

We set GP's population size at 5, 20, and 50 to test how many resources are required to perform well.[5] Figure 3 illustrates the evolution of strategies' performance in games where GP was present, and we have the following findings:[6]

1. No matter how big the population is, GP can gradually improve and defeat other strategies(even GD) as the winner. (When the population is only 20, GP appoaches GD but does not overtaken it "on average.")
2. GP trader can still improve itself even under the extreme condition of a population of only 5. More interestingly, the fact that the tournament size is also 5 means that strategies in the population might converge very quickly. Figure 4 shows the evolution of average complexity of GP strategies. In the case of pop 5, the average complexity almost equals to 1 at the end of the experiment, meaning that GP could still gain superior advantages by constantly updating its strategy pool composed of very simple heuristics. We also notice that there seems to be two different kinds of evolution taking place. In contrast with the case of population 5, in the case of bigger population, GP develops more complex strategies as time goes by.
3. What is worth noticing is that GP might need a period of time to evolve. The bigger the population, the fewer generations are needed to defeat other

[5] The corresponding select number $N$ were set at 10, 40, and 100 respectively. Briefly speaking, the number of selection is the evaluation cycle for each GP generation. We controlled the entire trading days to be 7,000 days, 7 times longer than those in the first set of experiment. We lengthen the simulation so that we have a better chance to observe what learning agents can achieve.

[6] Figure 3 is unrelated to Figure 1 & 2, for a corresponding treatment to experiments with learning agents, please refer to Figure 5.

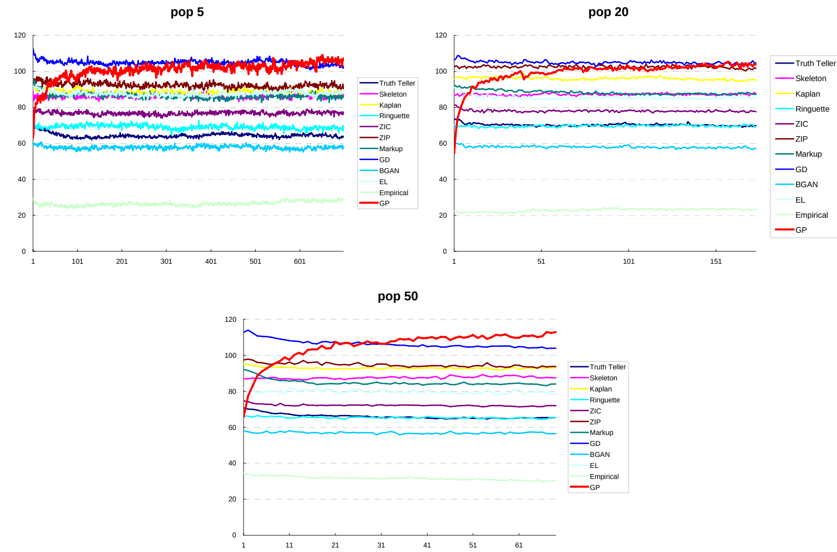


**Fig. 3.** the evolution of agents' performances, evaluated in terms of individual efficiency.

strategies. In any case, it takes hundreds to more than a thousand days to achieve good performances for GP traders. This may be a problem if we are to conduct experiments with both human traders and GP traders.

4. Figure 5 shows the results of this experiment when the population size is 50. We find that GP, replacing GD, has occupied one end of the frontier, which means that the innovative GP traders can outperform other adaptive strategies, even if they may be more sophisticatedly designed.

**Fig. 4.** Evolution of average GP complexity when the population sizes are 5, 20, and 50 respectively (from the left panel to the right panel).

To investigate the issue of learning speed, we notice that psychologists tell us that intelligence of human beings involves the ability to "learn quickly and learn from experiences." (Gottfredson 1997) From an social science point of view, we

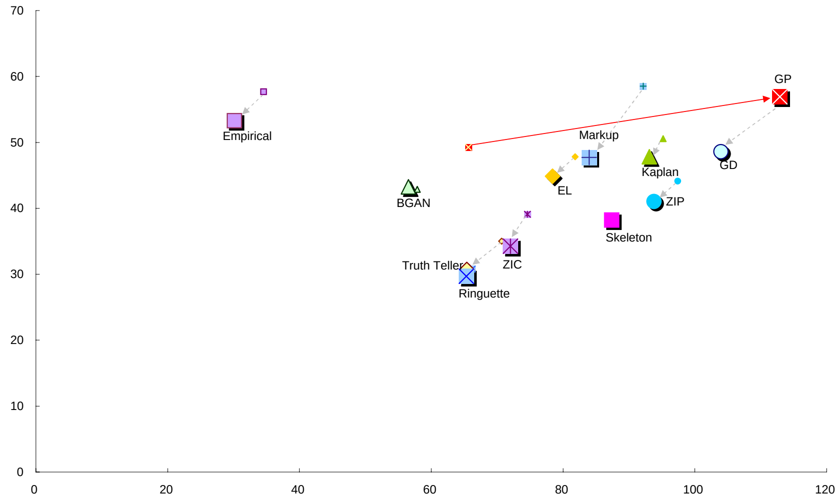


**Fig. 5.** Evolution of Agents' performances when GP's population size is 50. The arrow shows the direction of change from the beginning to the end.

**Table 2.** GP traders' strategy complexity when they are facing truth tellers.

| | pop 5 | pop 20 | pop 50 |
|---|---|---|---|
| Average Complexity | 1 | 2.61 | 2.36 |
| S.D. of Complexity | 0 | 4.15 | 2.15 |
| Performance | 109.81 | 125.67 | 116.52 |
| S.D. of Performance | 39.17 | 55.51 | 29.91 |

can think of a GP's population size as a proxy of his IQ. Then what is the influence of GP agents' IQ on their learning speed?

Figure 6 is the time required to achieve each performance level. We extend our experiments to include population size from 5, 20, 30, ... to 100. We can see clearly from this figure that the bigger the population size, the fewer time GP traders will need to perform well. Hence, with a higher IQ (a bigger population size), GP traders can learn faster, and gain more wealth consequently.

We also observe the fact that GP traders with different population size seem to evolve strategies of different complexity. Does it have anything to do with economic reasoning, or is it simply the phenomenon of "code bloat"(or complexity drift) of Genetic Programming?

We devise another benchmark experiments where GP agents are trading against truth tellers. Table 2 presents the performance as well as strategy complexity of GP agents with different population size. Wilcoxon Rank Sum tests

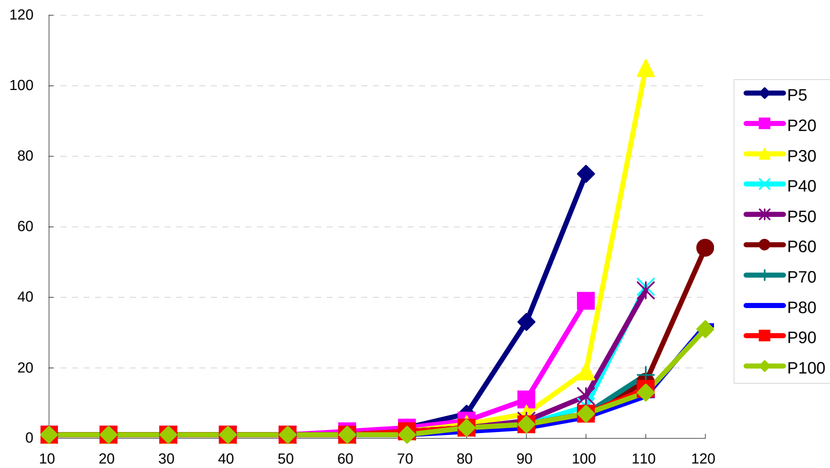


**Fig. 6.** Time required to achieve specific profit levels. The horizontal axis is level of individual efficiency; the vertical axis is the number of generations taken to achieve a specific individual efficiency level.

show that complexity of GP agents of size 20 and 50 are significantly larger than GP agent with population size 5.

Compared with Figure 4, we can find that agent of population 20 use more complex strategies when encountering truth tellers, while GP agent of population 50 employs simpler strategies instead. This is not what we should observe if "code bloat" is a universal and the most striking factor which results in the increase of GP's strategy complexity. We doubt that there are reasons other than GP's "code bloat" properties that cause agents with different population sizes to adopt different style of strategies.

Figure 7 presents another set of benchmark experiments where GP traders with different population sizes are put in an environment comprises solely one kind of opponents. When we examine the relationship between the profit ability of GP agents' opponents and the average complexity of GPs' strategy when facing them, we find that when the population size is 50, GP will choose strategies of different complexities according to his rivals, and the stronger the rival, the more complex the strategy he uses. On the contrary, GP agents with population size 5 and 20 do not exhibit properties like this. This observation brings about the question whether GP evolves more complex strategies not just because of the problem of "code bloat," but also out of the possibility that he can yield better strategy with more complex structure. But unless we have more techniques to trace the complex dynamics of GP's genetic operations, we can never be sure. What we know for sure now is that GP agents do exhibit strong ability to learn, and individual intellectual endowments do play a role in GP agents' learning processes.

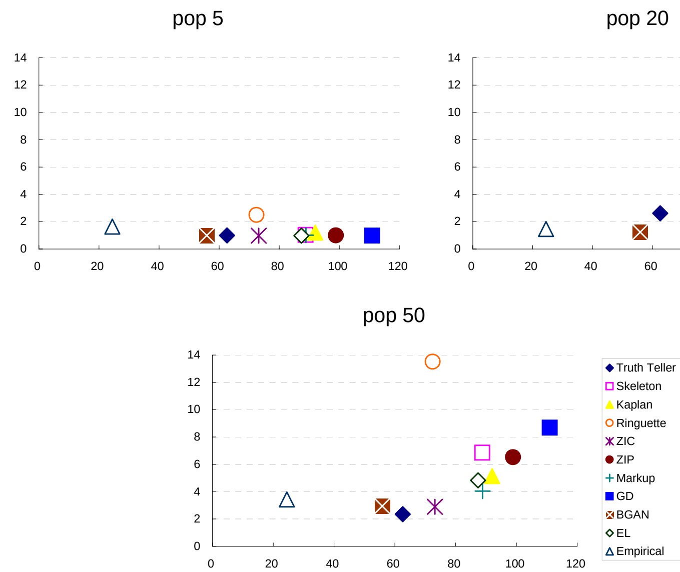


**Fig. 7.** The relationship between opponents' ability and GPs' strategy complexity. The horizontal axis is the performance of GPs' opponents; the vertical axis is the average complexity when GP traders face a particular kind of opponents.

## 4 Conclusion

In this paper, we propose a classification of strategies based on whether they are heuristics, adaptive strategies, or innovative learning processes. We use the term "heuristics" to refer to strategies that are fixed and ready to be used at any time, the adjective "adaptive" to refer to their ability to update parameters or statistics, and the adjective "innovative" to refer their ability to create new structures for themselves.

The results of our experiments show that each style of strategies has its own advantage in a profit-variation space. Along the frontier from high profit and high variation to low profit but consistent performance, we can find the niche for innovative process, adaptive strategies, heuristics, and non-strategic behavior respectively. Interpreting the results in terms of our classification of strategies help us better understand the potential as well as the limit of these kinds of strategies.

The results of this study can also be viewed as the foundation of further experiments involving human subjects. As Rust et al. (1994) pointed out in their discussions, humans' unique characteristics such as "intuitive leap" or "adaptiveness" make them distinct from artificial strategies. Whether humans and software agents differ and how they are different from each other are expected to be answered with the research method proposed in this paper.

## References


Chan, N. T., LeBaron, B., Lo, A. W., Poggio, T.: Agent-based models of financial markets: A comparison with experimental markets. MIT Artificial Markets Project, Paper No. 124, September. Available at http://citeseer.ist.psu.edu/chan99agentbased.html

Chen, S.-H., Tai, C.-C.: On the selection of adaptive algorithms in ABM: A computational-equivalence approach. Computational Economics **28** (2006) 51–69.

Chen, S.-H., Zeng, R.-J., Yu, T.: Co-evolving trading strategies to analyze bounded rationality in double auction markets, to appear in Riolo, R., Soule, T. and Worzel, B. (eds.) Genetic Programming Theory and Practice VI (2009) Springer.

Cliff, D., Bruten, J.: Zero is not enough: On the lower limit of agent intelligence for continuous double auction markets. Tech. Rep. no. HPL-97-141, Hewlett-Packard Laboratories (1997). Available at http://citeseer.ist.psu.edu/cliff97zero.html

Das, R., Hanson, J. E., Kephart, J. O., Tesauro, G.: Agent-human interactions in the continuous double auction. Proceedings of the 17th International Joint Conference on Artificial Intelligence (IJCAI), San Francisco, CA: Morgan-Kaufmann (2001).

Easley, D., Ledyard J.: Theories of price formation and exchange in double oral auction. In: Friedman, D., Rust, J. (eds.) The Double Auction Market-Institutions, Theories, and Evidence. Addison-Wesley (1993).

Friedman, D.: A simple testable model of double auction markets. Journal of Economic Behavior and Organization **15** (1991) 47–70.

Gjerstad, S., Dickhaut, J.: Price formation in double auctions. Games and Economic Behavior **22** (1998) 1–29.

Gode, D., Sunder, S.: Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. Journal of Political Economy **101** (1993) 119–137.

Gottfredson, L.: Mainstream science on intelligence: An editorial with 52 signatories, history, and bibliography. Intelligence **24**(1997) 13–23.

Grossklags, J., Schmidt, C.: Software agents and market (in)efficiency - A human trader experiment. IEEE Transactions on System, Man, and Cybernetics: Part C, Special Issue on Game-theoretic Analysis & Simulation of Negotiation Agents, IEEE SMC **36** (2006) 56–67.

Kagel, J.: Auction: A survey of experimental research. In: Kagel, J., Roth, A. (eds.) The Handbook of Experimental Economics. Princeton University Press (1995).

Rust, J., Miller, J. H., Palmer, R.: Characterizing effective trading strategies: Insights from a computerized double auction tournament. Journal of Economic Dynamics and Control **18** (1994) 61–96.

Taniguchi, K., Nakajima, Y., Hashimoto, F.: A report of U-Mart experiments by human agents. In: R. Shiratori, Arai, K., Kato, F. (eds.) Gaming, Simulations, and Society: Research Scope and Perspective, pp.49–57, Springer (2004).

Zhan, W., Friedman, D.: Markups in double auction markets. Journal of Economic Dynamics and Control **31** (2007) 2984–3005.