

Some preliminary results on competition between markets for automated traders

Jinzhong Niu and Kai Cai

Department of Computer Science
Graduate Center
City University of New York
365, 5th Avenue
New York, NY 10016, USA
{jniu, kcai}@gc.cuny.edu

Simon Parsons and Elizabeth Sklar

Dept of Computer and Information Science
Brooklyn College
City University of New York
2900 Bedford Avenue
Brooklyn, NY 11210, USA
{parsons, sklar}@sci.brooklyn.cuny.edu

Abstract

Real market institutions, stock and commodity exchanges for example, do not occur in isolation. Company stock is frequently listed on several stock exchanges, and futures exchanges make it possible for dealers in a particular commodity to offset their risks by trading options in that commodity. While there has been extensive research into agent-based trading in individual markets, there is little work on agents that trade in multiple market scenarios. Our work seeks to address this imbalance, providing an analysis of the behavior of trading agents that are free to move between a number of parallel markets, each of which has different properties.

Introduction

The market mechanisms known as *auctions*, are widely used to solve real-world resource allocation problems, and in structuring stock or futures exchanges like the New York Stock Exchange (NYSE) and the Chicago Mercantile Exchange (MCE). When well designed (Klemperer 2002), auctions achieve desired economic outcomes like high *allocative efficiency* whilst being easy to implement. Research on auctions originally interested economists and mathematicians. They view auctions as games of incomplete information and have successfully applied traditional analytic methods from game theory to some kinds of auctions (Maskin & Riley 1985; Vickrey 1961). The high complexity of other auction types, especially *double-sided auctions* (Friedman 1993) (DAs)¹, however makes it difficult to go further in this direction (Madhavan 1992; Satterthwaite & Williams 1993).

As a result, researchers turned to experimental approaches. For example, (Smith 1962) showed that in *continuous double auctions* (CDAs)², even a handful of human traders can lead to high overall efficiency, and transaction prices can quickly converge to the theoretical equilibrium.

With real trade increasingly contracted by automated “program traders”, experimental work has followed suit. Gode and Sunder (1993) introduced the *zero intelligence*

trading strategy³ ZI-C — which bids randomly but avoids making a loss — and showed that it generates high efficiency solutions (Gode & Sunder 1993). (Cliff & Bruten 1997) then provided an adaptive trading strategy called *zero intelligence plus* (ZIP), and showed that it outperformed ZI-C, generating high efficiency outcomes and converging to the equilibrium price. This led to the suggestion that ZIP embodies the minimum intelligence required by traders. Subsequent work has led to the development of further trading strategies, including that proposed by (Roth & Erev 1995), and that suggested by (Gjerstad & Dickhaut 1998), commonly referred to as GD.

This work on trading strategies is only one facet of the research on auctions. Gode and Sunder’s results suggest that the structure of the auction mechanisms plays an important role in determining the outcome of an auction, and this is further borne out by the work of (Walsh *et al.* 2002) (which also points out that results hinge on both auction design and the mix of trading strategies used). For example, if an auction is *strategy-proof* traders need not bother to conceal their private values, and in such auctions complex trading agents are not required.

Despite the variety of this work, it has one common theme — it all studies single markets. In contrast, real market institutions, like the stock and commodity exchanges mentioned above, do not occur in isolation. Company stock is frequently listed on several stock exchanges. Indian companies, for example, can be listed on both the National Stock Exchange (NSE) and the Bombay Stock Exchange (BSE) (Shah & Thomas 2000). US companies may be listed on both the NYSE, NASDAQ and, in the case of larger firms, non-US markets like the London Stock Exchange (LSE).

Such multiple markets for the same goods induce complex interactions. The simplest example of this is the work of *arbitrageurs* who exploit price differences between markets to buy low in one and sell high in another, thus evening the prices between markets.⁴ More complex dynamics oc-

Copyright © 2009, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹In DAs, both competing sellers and buyers can make offers, in contrast to the most common auction mechanisms, such as the English auction, where only buyers bid.

²A CDA is a continuous DA in which any trader can accept an offer and make a deal any time during the auction period.

³In a recent paper (Sunder 2004), Sunder reveals that they came up with these simple strategies in the face of demands from students whom they had challenged to create automated strategies, saying that “Our motivation for the ZI-C strategy was part jest: it was sure to lose to the student strategies, but we could still save face with such an obviously simple and silly strategy”.

⁴In addition, futures exchanges make it possible for dealers in

cur when markets compete, as when the NSE opened and proceeded to claim much of the trade volume from the established BSE (Shah & Thomas 2000), or when the newly created Singapore International Monetary Exchange (SIMEX) did the same to Japanese markets for index futures on Nikkei 225 (Shah 1997) in the late 1980s. These changes took place over a long period of time, but inter-market dynamics can have much shorter timescales, as was the case in the flow between the CME and the NYSE during the global stock market crash of 1987 (Miller *et al.* 1988). This kind of interaction between markets has not been widely studied, least of all using automated traders.

The work described in this paper starts to address this imbalance between experimental work and what happens in the real world, providing an analysis of scenarios in which trading agents choose between a number of parallel markets, while the markets simultaneously decide how to profit from the traders. In common with much work in computational economics (Friedman 1998), the strategies used both by traders to choose between markets, and markets to decide how to charge traders, are very simple — the idea is that using more sophisticated strategies might obscure our view of what happening in the complex setting of double auction markets.

Background

To experiment with multiple markets, we used a variant of the Java Auction Simulator API (JASA)⁵. JASA provides the ability to run continuous double auctions populated by traders that use a variety of trading strategies, and has been used for a variety of work in analysing auctions, for example (Niu *et al.* 2006; Phelps *et al.* 2006). Auctions in JASA follow the usual pattern for work on automated trading agents, running for a number of trading *days*, with each day being broken up into a series of *rounds*. A round is an opportunity for agents to make offers to buy or sell⁶, and we distinguish different days because at the end of a day, agents have their inventories replenished. As a result, every buyer can buy goods every day, and every seller can sell every day. Days are not identical because agents are aware of what happened the previous day. Thus it is possible for traders to learn, over the course of several days, the optimal way to trade.

We run a number of JASA markets simultaneously, allowing traders to move between markets at the end of a day. In practice this means that traders need a decision mechanism that picks which market to trade in and we have implemented several — these are discussed below. Using this approach, agents are not only learning how best to make offers, which they will have to do anew for each market, but they are also learning which market is best for them. Of course, which market is best will depend partly on the properties of different markets, but also on which other agents

a particular commodity to offset their risks by trading options — commitments to buy or sell at a future date at a certain price — in that commodity, and provide further opportunities for arbitrage.

⁵<http://sourceforge.net/projects/jasa/>

⁶Offers to buy are also called *bids*, and offers to sell are also called *asks*. Both are called *shouts*.

are in those markets.

We allow markets to levy charges on traders, as real markets do. In doing this, our work has a different focus from the other work on market mechanisms we have mentioned. That work is focused on how the performance of traders helps achieve economic goals like high efficiency (Gode & Sunder 1993) and trading near equilibrium (Cliff & Bruten 1997), or how traders compete amongst themselves to achieve high profits (Tesauro & Das 2001). In contrast, we are interested in competition between *markets*, and what the movement of traders is when they are faced with a variety of markets.

Experimental Setup

The experiments we carried out explore how traders move between markets of different properties and what effect their movement has on the profits of those markets.

Traders

Our traders have two tasks. One is to decide how to make offers. The mechanism they use to do this is their *trading strategy*. The other task is to choose market to make offers in. The mechanism for doing this is their *market selection strategy*. The trading strategies are:

- ZI-C: (Gode & Sunder 1993) which picks offers randomly but ensures the trader doesn't make a loss.
- GD: (Gjerstad & Dickhaut 1998) which estimates the probability of an offer being accepted from the distribution of past offers, and chooses the offer which maximises its expected utility.

The market selection strategies are:

- T_r : the trader randomly picks a market; and
- T_ϵ : the trader treats the choice of market as an n -armed bandit problem which it solves using an ϵ -greedy exploration policy (Sutton & Barto 1998). A T_ϵ trader chooses what it estimates to be the best market, in terms of daily trading profit, with probability $1 - \epsilon$, and randomly chooses one of the remaining markets otherwise. ϵ may remain constant or be variable over time, depending upon the value of the parameter α (Sutton & Barto 1998). If α is 1, ϵ remains constant, while if α takes any value in $(0, 1)$, ϵ will reduce over time.
- T_τ : the trader uses the softmax exploration policy (Sutton & Barto 1998). A T_τ trader does not treat all markets other than the best exactly the same. If it does not choose the best market, it weights the choice of remaining market so that it is more likely to choose better markets. The parameter τ in the softmax strategy controls the relative importance of the weights a trader assigns markets, and similarly to ϵ , it may be fixed or have a variable value that is controlled by α .

Thus all our traders use simple reinforcement learning to decide which market to trade in⁷, basing their choice on the

⁷Though we have results, not presented here, which suggest that more complex forms of reinforcement learning, like the Roth-Erev

expected profit suggested by prior experience, and making no use of any other information that may be available about the markets. As mentioned above, we deliberately chose this simple decision mechanism in order to make the comparison between markets as clear as possible.

Markets

While we can set up markets to charge traders in a variety of ways, we have concentrated on charging traders a proportion of the surplus on a transaction in which they are involved — that is a proportion of the difference between what the buyer bids and the seller asks. We focus on this because it mirrors the case of the competition between the NSE and the BSE (Shah & Thomas 2000) where the BSE, had a much higher charge on transactions than the new market.

We experimented with four basic charging mechanisms, one fixed and three simple adaptive mechanisms:

- Fixed charging rates, typically 20%, 40%, 60% and 80% of the surplus on a transaction.
- Pricecutting (PC): since traders will, all else being equal, prefer markets with lower charges, a pricecutting market will reduce its charge until it is 80% of the charge of the lowest charging market.
- Bait and switch (B&S): the market cuts its charge until it captures 30% of the traders, then slowly increases its charge (adjusting its charge downward again if its market share drops below 30%).
- Zero intelligence (ZIP): a version of the ZIP strategy for markets. The market adjusts its charge to be just lower than that of the market that is the most profitable. If it is the most profitable market, it raises its charges slightly.

Again, our choice of market strategies was driven by the desire to first establish the relative performance of simple charging policies, and thus the basic structure of the problem of competing markets, before trying more complex policies.

Each of the experiments is setup in the following way. The experiment is run for 200 or 400 trading days, with every day being split into 10 rounds, each of which is one second long. The markets are populated by 100 traders, evenly split between buyers and sellers. Each trader is permitted to buy or sell at most one unit of goods per day, and each trader has a private value for that good which is drawn from a uniform distribution between \$50 and \$150.

Results

Results are given in Figures 1(a) to 9(d). These show values averaged over 100 runs of each experiment.

Fixed charge markets

The first set of experiments explore the properties of markets with fixed charges. Figures 1(a) and 1(b) show that traders that pick markets randomly have no discernable pattern of movement between markets, just as we would expect. As a result, the market with the highest charges makes the most

approach (Roth & Erev 1995) do not perform significantly differently.

profit. In contrast, Figure 1(c) and 1(d), when traders pick markets based on their personal profits, they move towards the market with lowest fixed costs. While markets with high charges make initial windfall profits, the trend is for the lower charging market to gain greater cumulative profit as the number of trading days increases.

Figures 2(a)–2(d) show that these results are robust against the ability of traders to make sensible trades since broadly the same results are observed when some or all of the traders make their bidding decisions randomly. Figures 3(a)–3(d) test the sensitivity of the results to the strategy that traders use to learn which market to choose. Decreasing ϵ over time (Figures 3(a) and 3(b)) does not seem to have much effect, but switching to the softmax strategy reduces the attractiveness of the lowest charging market since traders can still make good profits in higher charging markets.

Finally, for the experiments with only fixed charges, Figures 4(a)–4(d) show that the results obtained so far are very sensitive to the length of time agents have to learn about the markets. When some traders start learning afresh every day, simulating traders leaving and entering the markets (4(c) and 4(d)), the lowest charging market might still capture most of the traders, but it captures less of them, and the remaining markets attract enough traders to have the same profit profile as when there is no learning (Figures 1(a) and 1(b)).

Thus, for the fixed charge markets, provided that there is no turnover of traders, it is a winning strategy to undercut the charges of the other markets.

Homogeneous markets

Turning to the adaptive charging strategies, we first tested them against copies of themselves. In these experiments we ran four copies of each kind of market against each other with different initial profit charges (20%, 40%, 60% and 80%). In each experiment we also provided a “null” market which made no charges and executed no trades — the idea of this is to allow traders who cannot trade profitably with a mechanism for not trading — and, for completeness, carried out the same experiment with the fixed price markets. For all of these experiments, and all subsequent experiments, we used traders that made bids using GD, selected markets using an ϵ -greedy policy, and continued learning for all 400 days. The results of these experiments can be seen in Figures 5(a) to 6(d). The fixed price markets, in Figures 5(a) and 5(b), attract less traders in the presence of the null market, but make similar profits (since the traders who tend to the null market do not often trade). The price cutting markets, in Figures 5(c) and 5(d), get into a price war which, unlike the myopic pricebots from (Greenwald & Kephart 1999), they do not have the intelligence to get out of, and the bait-and-switch markets (Figures 6(a) and 6(b)) are similarly unable to generate a significant profit. The ZIP markets, Figures 6(c) and 6(d), adjusting their profit margins to fit what the traders will allow, manage to do better, but generate nowhere near as much profit as the fixed price markets do.

Heterogeneous markets

While the homogenous market experiments give some idea of market performance, it is more interesting to examine

		Many Price Cutting	Many Bait and Switch	Many Zero Intelligence
1-Price Cutting	Profit stdev. relationship		0.8 – 84.1 7.5 – 105.6 <	6502.2 – 6043.6 1527.1 – 2159.7 >
1-Bait and Switch	Profit stdev. relationship	82.0 – 0.7 56.7 – 6.8 >		6545.7 – 5743.8 2325.0 – 1581.8 >
1-Zero Intelligence	Profit stdev. relationship	2289.6 – 0.8 1118.9 – 8.5 >	1773.5 – 166.9 633.0 – 264.8 >	

Table 1: Results of one-to-many experiments. For each experiment, the table gives the cumulative profit of the “one” strategy followed by the cumulative profit of the best of the “many”, and an indication of whether the “one” is greater or less than the “many” at the 90% confidence level (determined by a t-test).

		Many Price Cutting	Many Bait and Switch	Many Zero Intelligence
1-Price Cutting	Profit stdev. relationship		0 – 7.2 0 – 33.5 <	1727.5 – 1475.3 438.8 – 610.6 >
1-Bait and Switch	Profit stdev. relationship	5.9 – 0 40.2 – 0 >		2048.0 – 1397.7 829.3 – 432.1 >
1-Zero Intelligence	Profit stdev. relationship	206.1 – 0 173.4 – 0 >	147.2 – 70.2 54.4 – 227.6 >	

Table 2: Results of one-to-many experiments over the latter days of the run. For each experiment, the table gives the cumulative profit of the “one” strategy over the last 100 days of the experiment followed by the cumulative profit of the best of the “many” and an indication of whether the “one” is greater or less than the “many” at the 90% confidence level (determined by a t-test).

how the adaptive charging strategies work in competition against one another. To explore this, we carried out a series of mixed market experiments along the lines of the trading strategy work of (Tesauro & Das 2001). For each of the four charging strategies, we ran an experiment in which all but one market used that strategy and the remaining market used another strategy, carrying out one such “one-to-many” experiment for each of the other strategies. In other words, we tested every “one-to-many” combination. For all these experiments, we measured the cumulative profit of a market using the charging strategies, and ran the markets alongside the same null market as before.

Table 1 gives the results of “one-to-many” experiments, giving the cumulative profits of the “one” market against the best performing “many” market for each combination of the adaptive markets. The table also indicates which profit is the greater at 90% confidence (as determined by a t-test). “>” means the “one” market is better than the best “many” market at 90% confidence and “<” means the best “many” market is better. The day by day results are also given in Figures 7(a)–9(d). The results show that one price-cutting market is effective against many zero-intelligence markets, since it can capture more traders, as Figure 8(a) shows. In such a case, both types of market generate good profits. However, when there is more than one price-cutter, such markets get into a price war and drive their charges down to zero.

The bait-and-switch strategy was envisaged as a more sophisticated version of PC, one that exploited its market share by increasing charges on traders it had attracted through low charges. The results in Table 1 suggest that B&S achieves

this intention, outperforming PC both when one bait-and-switch takes on multiple price-cutters, and when a single price-cutter competes against multiple bait-and-switch markets. However, as is the case with PC, when there is more than one market using B&S, they may end up cutting charges in a futile attempt to increase market share and hence do not make much profit — this is what happens when there are many bait-and-switch markets running against a single zero-intelligence market.

The zero-intelligence strategy, designed to get out of price wars by increasing charges when it can, performs well against both PC and B&S markets when it is in the minority. When there is only one price-cutter or bait-and-switch against many zero-intelligence markets, the PC and B&S may outperform ZIP. However, even when this is the case, as Figure 8(b) and Figure 9(b) show, ZIP can still make more profit than the other market strategies in the short run (before 200 days have elapsed).

The results in Table 1 are cumulative over the entire 400 days of the experiment. Since the early days of the experiment often contain a lot of noise from the initial exploration of the traders, it is interesting to also look at the profits over the just the later stages of the experiments, when trader movement has settled down. Such results are presented in Table 2. These results suggest that when it is in the majority, the zero intelligence strategy is clearly outperformed by both a single price-cutter and a single bait-and-switch market. Since Figures 8(a) and 9(a) suggest that there is no longer much movement of traders at this point, the results in Table 2 simply reinforce those in Table 1.




Figure 1: Baseline experiments. GD traders, (a) and (b) with random market selection, (c) and (d) with T_ϵ market selection ($\epsilon = 0.1, \alpha = 1$). $M_{0.2}$: dashed line with solid dots; $M_{0.4}$: solid line; $M_{0.6}$: dashed line; $M_{0.8}$: dotted line.

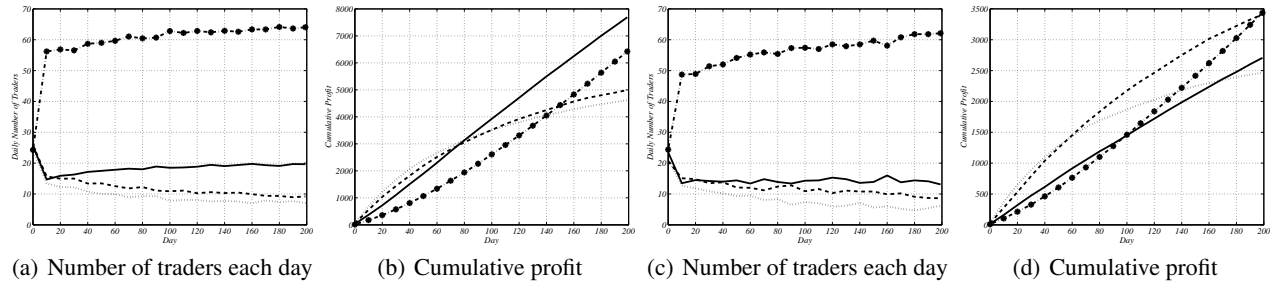


Figure 2: Robustness experiments. (a) and (b) show ZIC traders, and (c) and (d) show a mixture of GD and ZIC traders, all traders use T_ϵ market selection ($\epsilon = 0.1, \alpha = 1$). $M_{0.2}$: dashed line with solid dots; $M_{0.4}$: solid line; $M_{0.6}$: dashed line; $M_{0.8}$: dotted line.

Related work

In our experiments, market performance depends on the mix of market strategies being considered. This suggests that, as is the case for trading strategies (Tesauro & Das 2001), it may be hard to find a dominant strategy for deciding market charges, though such a conclusion must wait until market strategies have been investigated further. This is particularly important since the strategies that we have considered were, quite intentionally, about the simplest we could imagine (starting with simple strategies seemed a good way to understand the problem we are considering).

As mentioned above, there has been little work on the problem of choosing between multiple markets. Our work is similar to (Ladley & Bullock 2005), but differs in that our work assesses the impact of different market charges while (Ladley & Bullock 2005) are concerned with the information available to traders. (Ladley & Bullock 2005) are also concerned with markets that are spatially separated, so that traders' access to trading partners is limited by their location. This is similar to the concern of (Gaston & desJardins 2005). In comparison, our traders are able to find any partner, but the mobility of traders means that they can be separated temporally rather than spatially.

Our work also has similarities to that of (Greenwald & Kephart 1999). In the latter, shoppers choose between different merchants, and the merchants set prices that depend on the prices set by other merchants. While some of the results obtained in (Greenwald & Kephart 1999), especially

the price wars induced by myopic price-setting, look similar to some of ours, the scenario we are considering is considerably more complex. For one thing, the traders in our scenario — the analogs of the buyers in (Greenwald & Kephart 1999) — learn rather than making the same market choice at every trading opportunity. Secondly, and more importantly, the markets in (Greenwald & Kephart 1999) have prices set by the merchants, while in our case the prices are determined by the traders. As a result, when traders pick a market in our scenario, they do not know for sure if they will even be able to trade, much less what prices good will change hands at. From the perspective of the markets, it is possible to attract many traders who, because of their value for the commodity being traded, do not end up trading. We are in the process of investigating the effect of these subtleties.

Conclusions

This paper has described some of our initial work examining the dynamics of trading when agents can choose between different markets. While we are wary of drawing too many conclusions from our results, because we are still at a very preliminary stage in our investigation, we can distinguish some broad trends. These show that, even when they are limited in their ability to make good trades and limited in their learning about markets, traders will gravitate to the lowest charging markets rather quickly, and, as a result, markets with lower charges generate higher profits. However, the advantages of low charges are somewhat

brittle. The advantages evaporate, for example, when not all traders are experienced, and it appears that the best charging strategies are both adaptive and, like the simple “zero intelligence” and “bait-and-switch” strategies that we introduce, quick to increase charges when they can. Clearly there are many other possible charging strategies, and it remains to be seen whether these conclusions hold when other strategies are tested.

Our future, and, indeed, current, work is aimed at further untangling the behavior of competing markets. First, we want to repeat the existing experiments over longer periods, ensuring that the results we have are representative of what happens in the steady state, after all start-up effects are removed. Second, we want to try to optimise the simple adaptive strategies — the behavior of each is determined by some simple parameters (for example the market share that the bait-and-switch market looks to capture), and it seems likely that suitable adjustment of these parameters can improve performance. Third, we aim to investigate additional market strategies with the aim of discovering one that is dominant, moving from the “one-to-many” analysis performed here to the kind of evolutionary game theoretic analysis used in (Walsh *et al.* 2002).

Acknowledgments The authors are grateful for financial support received from the National Science Foundation under grant NSF IIS-0329037.

References

- Cliff, D., and Bruten, J. 1997. Minimal-intelligence agents for bargaining behaviours in market-based environments. Technical report, Hewlett-Packard Research Laboratories.
- Friedman, D. 1993. The double auction institution: A survey. In Friedman, D., and Rust, J., eds., *The Double Auction Market: Institutions, Theories and Evidence*. Cambridge, MA: Perseus Publishing. 3–25.
- Friedman, D. 1998. Evolutionary economics goes mainstream: A review of “The Theory of Learning in Games”. *Journal of Evolutionary Economics* 8(4):423–432.
- Gaston, M. E., and desJardins, M. 2005. Agent-organized networks for dynamic team formation. In *Proceedings of the 4th International Conference on Autonomous Agents and Multiagent Systems*.
- Gjerstad, S., and Dickhaut, J. 1998. Price formation in double auctions. *Games and Economic Behavior* 22:1–29.
- Gode, D. K., and Sunder, S. 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.
- Greenwald, A. R., and Kephart, J. O. 1999. Shopbots and pricebots. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence*.
- Klemperer, P. 2002. How (not) to run auctions: The European 3G telecom auctions. *European Economic Review* 46:829–845.
- Ladley, D., and Bullock, S. 2005. Who to listen to: Exploiting information quality in a ZIP-agent market. In *Proceedings of the Workshop on Trading Agent Design and Analysis*.
- Madhavan, A. 1992. Trading mechanisms in securities markets. *The Journal of Finance* 47(2):607–641.
- Maskin, E. S., and Riley, J. G. 1985. Auction theory with private values. *American Economic Review* 75:150–155.
- Miller, M. H.; Hawke, J. D.; Malkiel, B.; and Scholes, M. 1988. Findings of the Committee of Inquiry Examining the Events Surrounding October 19, 1987. Technical report, Chicago Mercantile Exchange.
- Niu, J.; Cai, K.; Parsons, S.; and Sklar, E. 2006. Reducing price fluctuation in continuous double auctions through pricing policy and shout improvement. In *Proceedings of the 5th International Conference on Autonomous Agents and Multi-Agent Systems*.
- Phelps, S.; Marcinkiewicz, M.; Parsons, S.; and McBurney, P. 2006. A novel method for automated strategy acquisition in n-player non-zero-sum games. In *Proceedings of the 5th International Conference on Autonomous Agents and Multi-Agent Systems*.
- Roth, A. E., and Erev, I. 1995. Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. *Games and Economic Behavior* 8:164–212.
- Satterthwaite, M. A., and Williams, S. R. 1993. The Bayesian theory of the k -double auction. In Friedman, D., and Rust, J., eds., *The Double Auction Market: Institutions, Theories and Evidence*. Cambridge, MA: Perseus Publishing. 99–123.
- Shah, A., and Thomas, S. 2000. David and Goliath: Displacing a primary market. *Global Financial Markets* 1(1):14–21.
- Shah, A. 1997. Competing exchanges: The international dimension. *Business Standard*.
- Smith, V. L. 1962. An experimental study of competitive market behaviour. *Journal of Political Economy* 70(2):111–137.
- Sunder, S. 2004. Markets as artifacts: Aggregate efficiency from zero-intelligence traders. In Augier, M., and March, J. G., eds., *Models of a Man: Essays in memory of Herbert A. Simon*. Cambridge MA: MIT Press.
- Sutton, R. S., and Barto, A. G. 1998. *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
- Tesauro, G., and Das, R. 2001. High-performance bidding agents for the continuous double auction. In *Proceedings of the 3rd ACM Conference on Electronic Commerce*.
- Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* 16:8–27.
- Walsh, W.; Das, R.; Tesauro, G.; and Kephart, J. O. 2002. Analyzing complex strategic interactions in multi-agent systems. In *Proceedings of Workshop on Game-Theoretic and Decision-Theoretic Agents*.

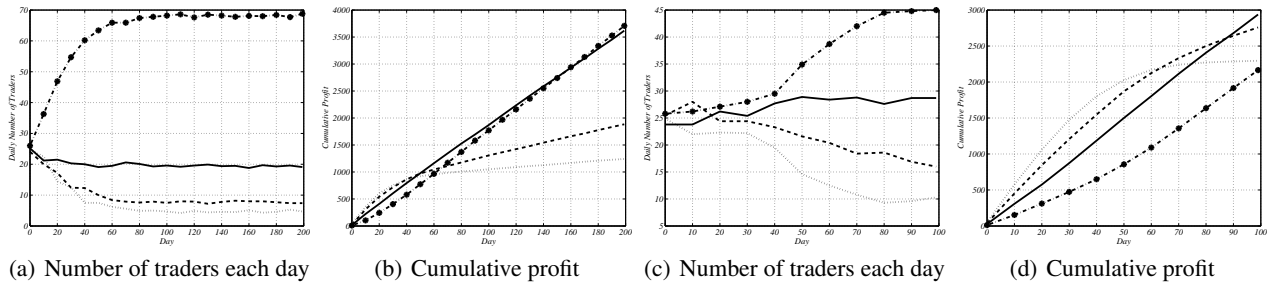


Figure 3: Learning experiments. GD traders, (a) and (b) with T_ϵ traders ($\epsilon = 1, \alpha = 0.95$), (c) and (d) with T_τ traders ($\tau = 1, \alpha = 0.95$). $M_{0.2}$: dashed line with solid dots; $M_{0.4}$: solid line; $M_{0.6}$: dashed line; $M_{0.8}$: dotted line.




Figure 4: Population experiments. GD traders, all traders use T_ϵ market selection ($\epsilon = 0.1, \alpha = 1$). In (a) and (b), all traders learn continuously through the experiment. In (c) and (d), 10% of the traders re-start learning every day. $M_{0.2}$: dashed line with solid dots; $M_{0.4}$: solid line; $M_{0.6}$: dashed line; $M_{0.8}$: dotted line.




Figure 5: Homogeneous markets: fixed price markets and pricecutting markets. GD traders, all traders use T_ϵ market selection ($\epsilon = 0.1, \alpha = 1$). (a) and (b) are 4 homogeneous fixed price markets, (c) and (d) are 4 homogeneous pricecutting markets. The crossed line denotes traders that choose not to enter any market. Market with 20% initial charge: dashed line with solid dots; Market with 40% initial charge: solid line; Market with 60% initial charge: dashed line; Market with 80% initial charge: dotted line.




Figure 6: Homogeneous markets: bait-and-switch markets and zero intelligence markets. GD traders, all traders use T_ϵ market selection ($\epsilon = 0.1, \alpha = 1$). (a) and (b) are 4 homogeneous bait-and-switch markets, (c) and (d) are 4 homogeneous zero intelligence markets. The crossed line denotes traders that choose not to enter any market. Market with 20% initial charge: dashed line with solid dots; Market with 40% initial charge: solid line; Market with 60% initial charge: dashed line; Market with 80% initial charge: dotted line.

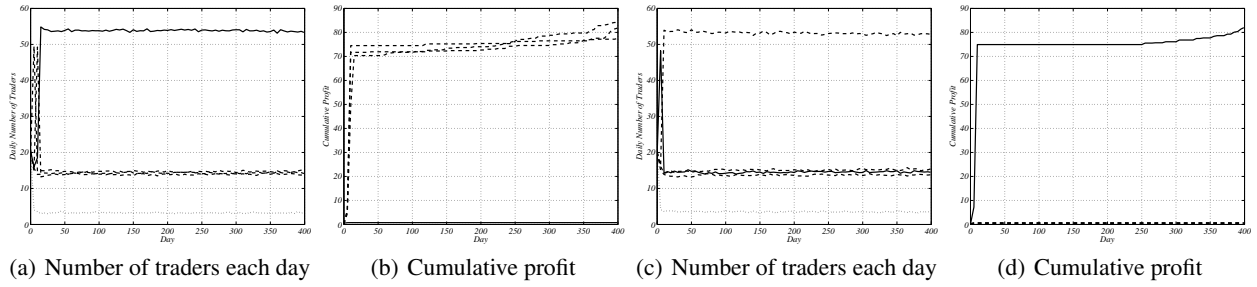


Figure 7: Heterogeneous markets: pricecutting against bait-and-switch. GD traders, all traders use T_ϵ market selection ($\epsilon = 0.1$, $\alpha = 1$). In (a) and (b) one pricecutting market (solid line) competes with 3 bait-and-switch markets, in (c) and (d), one bait-and-switch market (solid line) competes with 3 pricecutting markets. The grey line denotes traders that choose not to enter any market.




Figure 8: Heterogeneous markets: pricecutting against zero intelligence. GD traders, all traders use T_ϵ market selection ($\epsilon = 0.1$, $\alpha = 1$). In (a) and (b) one pricecutting market (solid line) competes with 3 zero intelligence markets, in (c) and (d), one zero intelligence market (solid line) competes with 3 pricecutting markets. The grey line denotes traders that choose not to enter any market.




Figure 9: Heterogeneous markets: bait-and-switch against zero intelligence. GD traders, all traders use T_ϵ market selection ($\epsilon = 0.1$, $\alpha = 1$). In (a) and (b) one bait-and-switch market (solid line) competes with 3 zero intelligence markets, in (c) and (d), one zero intelligence market (solid line) competes with 3 bait-and-switch markets. The grey line denotes traders that choose not to enter any market.