On the behavior of competing markets populated by automated traders

Jinzhong Niu[1], Kai Cai[1], and Simon Parsons[1 2] and Elizabeth Sklar[1 2]

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

[2] Department 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, allowing traders to potentially trade such stock in different markets. While there has been extensive research into agent-based trading in individual markets, there is little work on agents that trade in such multiple market scenarios. Our work seeks to address this imbalance. Here we provide an initial analysis of the behavior of trading agents that are free to move between a number of parallel markets, where markets are able to charge traders in a variety of ways. We show the movement of traders between markets, sketch some adaptive strategies that markets may use to adjust charges, evaluate the effectiveness of these strategies, and give some results which show the effect of trader movement on properties of the markets.

Key words: Continuous double auction, multiple markets.

1 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 (CME). When well designed [11], auctions achieve desirable 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 [14, 28]. The high complexity of other auction types, especially *double-sided auctions* [4], however makes it difficult to go further in this direction [13, 21] except in special cases such as the *buyer's bid* double auction [10]. In 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 can make offers, and this greatly expands the space of possible trader strategies. To deal with this complexity, researchers turned to experimental approaches to analyse the most common varieties of the double auction, the *continuous double auction* (CDA) — in which any trader can accept an offer and make a deal any time during the auction period — and the *clearing house* (CH) auction — where deals may only be made at the end of the auction period though offers may be continuously exchanged. For example, [24] showed that for CDAs, 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. [8] introduced the *zero intelligence* trading strategy [1] ZI-C — which bids randomly but avoids making a loss — and showed that it generates high efficiency solutions [8]. [3] 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 [20], and that suggested by [7], the latter commonly being referred to as GD after its creators.

This work on trading strategies is only one facet of the research on auctions. The results in [8] suggest that the structure of the auction mechanisms plays an important role in determining the outcome of an auction, and this is further bourne out by the work of [29] (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) [23]. 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. In addition, futures exchanges make it possible for dealers in 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. More complex dynamics occur when markets compete, as when the NSE opened and proceeded to claim much of the trade volume from the established BSE [23], or when the newly created Singapore International Monetary Exchange (SIMEX) did the same to Japanese markets for index futures on Nikkei 225 [22] in the late 1980s. These changes took place over a long period of time, but inter-market dynamics can have much

[1] In a recent paper [25], one of the authors of [8] reveals that they came up with this simple strategy 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".

shorter timescales, as was the case in the flow between the CME and the NYSE during the global stock market crash of 1987 [15]. 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 simultaeously decide how to profit from the traders. In common with much work in computational economics [5], 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 is happening in the complex setting of double auction markets.

2 Background

To experiment with multiple markets, we used a variant of the Java Auction Simulator API (JASA)[2]. 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 [17, 19]. 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 sell[3], 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 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 [8] and trading near equilibrium [3], or how traders compete amongst themselves to achieve high profits [27]. 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.

[2] `http://sourceforge.net/projects/jasa/`

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

3 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.

3.1 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: [8] which picks offers randomly but ensures the trader doesn't make a loss.
– GD: [7] 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 [26]. 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 [26]. 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 [26]. 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 in[4], basing their choice on the 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.

3.2 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 refer to this as a "profit charge"). We focus on this because it mirrors the case of the competition between the NSE and the BSE [23] where the BSE, had a much higher charge on transactions than the new market.

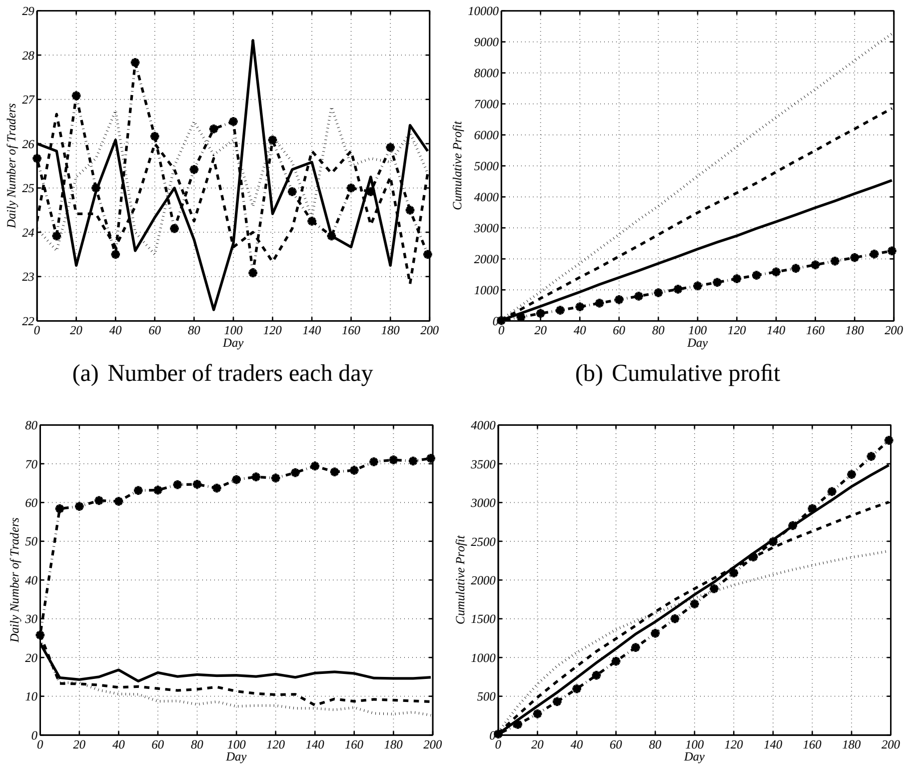

(a) Number of traders each day (b) Cumulative profit

(c) Number of traders each day (d) Cumulative profit

Fig. 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.

We experimented with four basic charging mechanisms, one that imposed fixed charges, and three simple mechanisms for adapting charges:

– 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%).

[4] Though we have results, not presented here, which suggest that more complex forms of reinforcement learning, like the Roth-Erev approach [20] do not perform significantly differently.

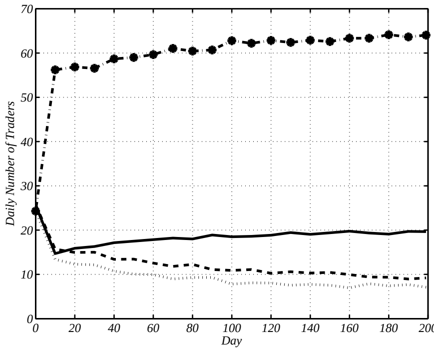

(a) Number of traders each day

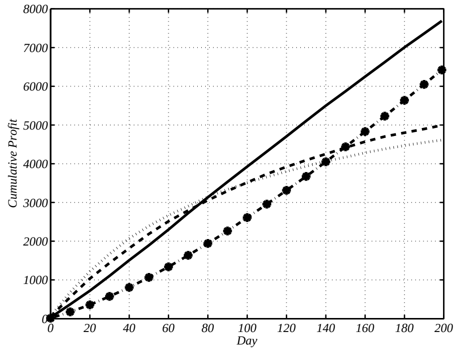

(b) Cumulative profit

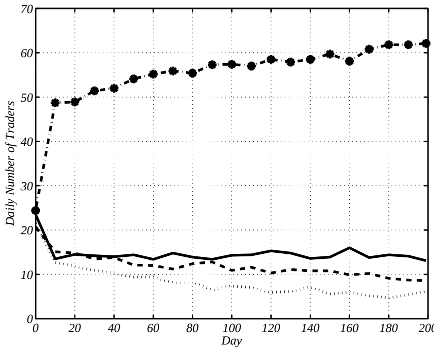

(c) Number of traders each day

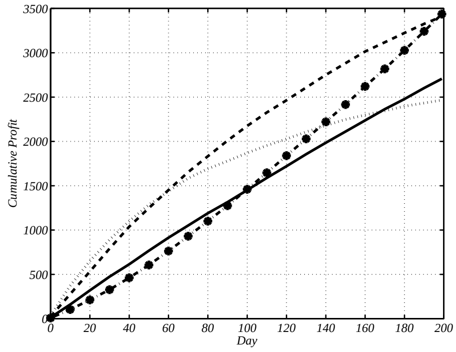

(d) Cumulative profit

Fig. 2. Robustness experiments. (a) and (b) show ZI-C traders, and (c) and (d) show a mixture of GD and ZI-C 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.

– Lure or learn fast[5] (LL): a version of the ZIP strategy, suitably adapted 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 charge slightly.

As with our choice of the mechanism used by traders to choose which market to trade in, 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 beween 100 and 400 trading days, with every day being split into 10 rounds, each of which is 1 second long. The markets are populated by 100 traders, evenly split between buyers and sellers, and initially evenly split between markets. 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

[5] The name is intended as a play on Bowling's "win or learn fast" [1]. We initially called this strategy "zero intelligence" but found we confused it with zero intelligence traders.

good which is drawn from a uniform distribution between $50 and $150. Private values are constant across all the trading days. For most of the experiments, all markets were continuous double auctions, though the very last experiments we discuss in the paper also used clearing house auctions (we will identify these experiments clearly when we discuss them).

4 Results

The results of our experiments are given in Figures 1 to 6 and Tables 1 and 2. These all show values averaged over 100 runs of each experiment.

4.1 Fixed charge markets

The first set of experiments explore the properties of markets with fixed charges. These are the results in Figures 1 to 4.

Figures 1 provides some baseline results. 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 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.

Figure 2 and Figure 3 show the insensitivity of the results we obtain to the choice of bidding strategy and the choice of market selection strategy. Figures 2(a)–2(d) show that 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 using ZI-C rather than using the sophisticated GD strategy. Figures 3(a)–3(d) test the sensitivity of the results to the kind of learning used in the market selection. Decreasing ϵ over time (Figures 3(a) and 3(b)) does not seem to have much effect, but switching to the softmax strategy (Figures 3(c) and 3(d)) reduces the attractiveness of the lowest charging market since some traders can still make reasonable profits in higher charging markets, and so will pick them relatively often.

Finally, Figure 4 tests the effect of allowing populations of traders to learn for different lengths of time. As we can see from Figures 4(a)–4(d), the results obtained so far are very sensitive to the length of time agents have to learn about the markets. When as few as 10% of traders start learning afresh every day, simulating traders leaving and entering the system of 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 — which is all we used in these initial experiments — provided that there is no turnover of traders, it is a winning strategy to undercut the charges of the other markets.

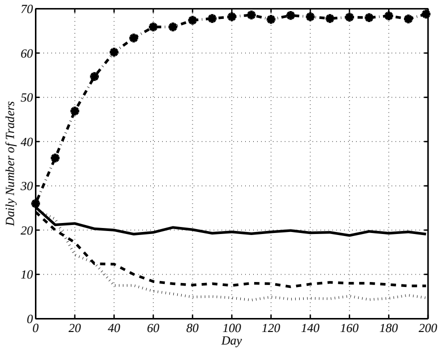

(a) Number of traders each day

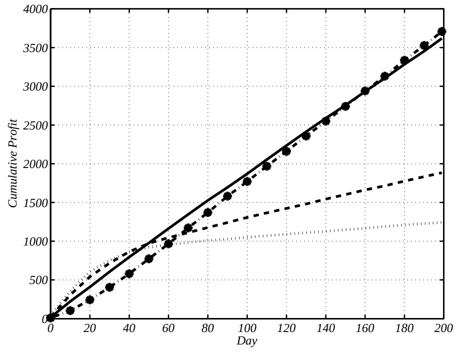

(b) Cumulative profit

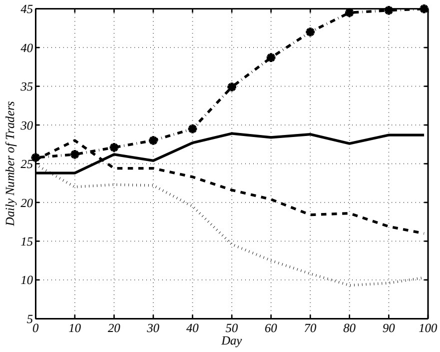

(c) Number of traders each day

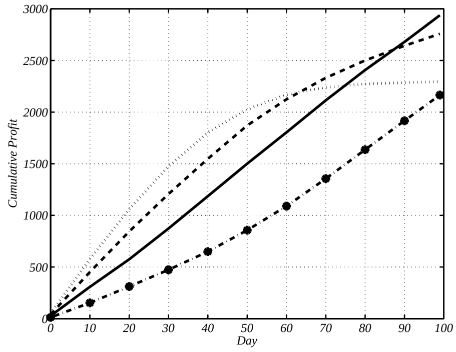

(d) Cumulative profit

Fig. 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.

4.2 Other approaches to market charging

While the experiments described above give us some idea of the interplay between trader movement, market charging and overall market performance, it is perhaps more interesting to examine how the different 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 [27]. For each of the three adaptive charging strategies — pricecutting (PC), bait-and-switch (B&S) and lure-or-learn (LL) — 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.

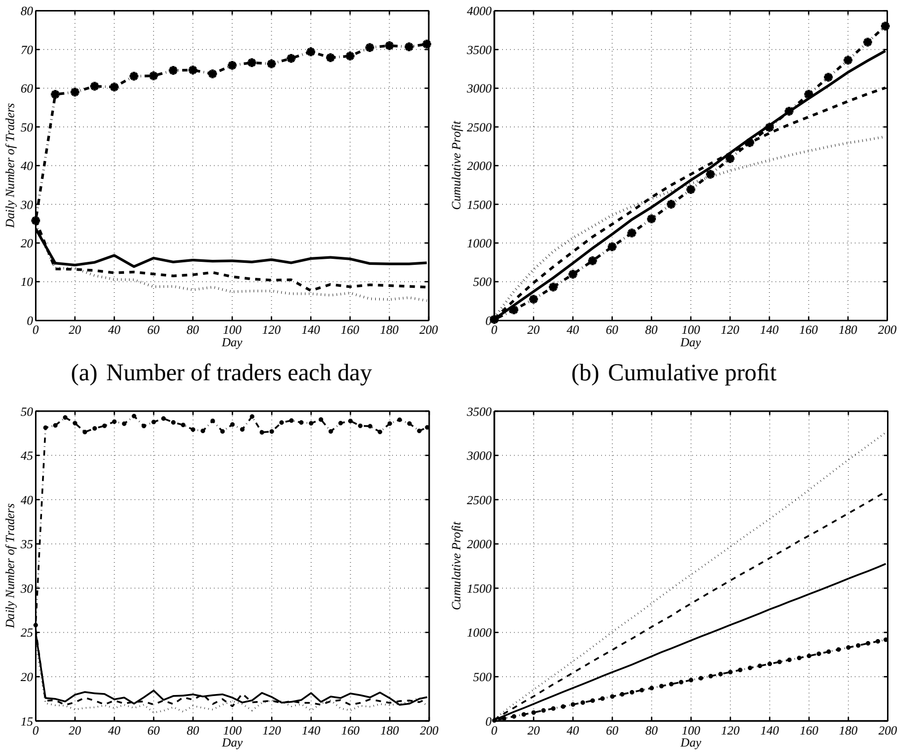

(a) Number of traders each day (b) Cumulative profit

(c) Number of traders each day (d) Cumulative profit

Fig. 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.

For all of these experiments we ran four charging markets. Each only made charges on profits, and each had a different initial charge (each experiment had one market each with a 20%, 40%, 60% or 80% charge). After the first day, each market then changed its charge using whichever of the three adaptive strategies it had been assigned. In each experiment we also provided a fifth "null" market which made no charges and executed no trades — the idea of this was to allow traders that were unable to trade profitably with a mechanism for not trading. For all of these experiments, we used traders that made bids using GD, selected markets using an ϵ-greedy policy ($\epsilon = 0.1$ and $\alpha = 1$), and continued learning for all 400 days.

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

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 PC	Many B&S	Many LL
1-PC	Profit		0.8 – 84.1	6502.2 – 6043.6
	stdev.		7.5 – 105.6	1527.1 – 2159.7
	relationship		$<$	$>$
1-B&S	Profit	82.0 – 0.7		6545.7 – 5743.8
	stdev.	56.7 – 6.8		2325.0 – 1581.8
	relationship	$>$		$>$
1-LL	Profit	2289.6 – 0.8	1773.5 – 166.9	
	stdev.	1118.9 – 8.5	633.0 – 264.8	
	relationship	$>$	$>$	

day by day results for these experiments are not included here for want of space, but may be found in [18]. Table 1 indicates that one price-cutting market is effective against many lure-or-learn markets (the daily results show that it does this by capturing more traders). In such a case, both types of market generate good profits. However, when all the markets are price-cutters, they 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 all the markets use B&S, they may end up cutting charges in a futile attempt to increase market share and hence do not make much profit. Something similar happens when there are many bait-and-switch markets running against a single lure-or-learn market.

The lure-or-learn 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 LL markets, the PC and B&S markets may outperform the LL markets. However, even when this is the case, the daily results reveal that LL 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 lure-or-learn strategy is clearly outperformed by both a single price-cutter and a single bait-and-switch market. This result just reinforces what we could already see in Table 1, and overall none of the relationships change between Table 1 and Table 2 — our results are robust against the initial noise as traders settle down.

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 pthan the "many" at the 90% confidence level (determined by a t-test).

		Many PC	Many B&S	Many LL
1-PC	Profit		0 – 7.2	1727.5 – 1475.3
	stdev.		0 – 33.5	438.8 – 610.6
	relationship		$<$	$>$
1-B&S	Profit	5.9 – 0		2048.0 – 1397.7
	stdev.	40.2 – 0		829.3 – 432.1
	relationship	$>$		$>$
1-LL	Profit	206.1 – 0	147.2 – 70.2	
	stdev.	173.4 – 0	54.4 – 227.6	
	relationship	$>$	$>$	

4.3 Equilibrium in multiple markets

The experiments that we have described so far were intended to assess how traders move between multiple markets that compete to attract traders, and how competition between markets unfolds in terms of the profits made by each market. However, there is another aspect that is of interest — the effect of this competition between markets on the usual economic measures by which we assess markets, measures such as allocative efficiency and proximity to theoretical equilibrium. We therefore examined these measures using the same experimental setup as in the "one-to-many" experiments, although we only used homogeneous mechanisms for choosing market charges (all were LL) and carried out the experiments for both CDA and CH markets. The results are given in Figures 5 and 6, the former being the results when all markets are CDAs, and the latter being the results when all the markets are CHs.

The results show that trader movement (Figures 5(a) and Figure 6(a)) between the different markets has settled down to some extent by around the 100th day of trading. At this point traders are still moving — this is a result of the market selection mechanism which still chooses one of the non-optimal markets 10% of the time (since ϵ is 0.1) — but the average number of traders that move in each market on each day has reduced to an approximately constant level.

Equilibrium price is also still changing by day 100, and continues to change throughout the experiments. Figures 5(b) and Figure 6(b), which plot the *change* in equilibrium price each day, make this clear — there is a non-zero change every day in every market. However, some pattern does emerge. The change in equilibrium price has flattened off for each market by around day 300. Not only has the change stopped changing on average (though it still fluctuates from day to day), but the size of the change has stratified by market — a couple of the markets have an equilibrium price that is changing very little, while others have an equilibrium price that is changing a lot. This seems to be because in the latter kind of market, traders are sparse, and the suppply and demand curves scarcely overlap [2]. In such a market, the movement of traders can have a big effect on the profitable trades, and hence the equilibrium price.

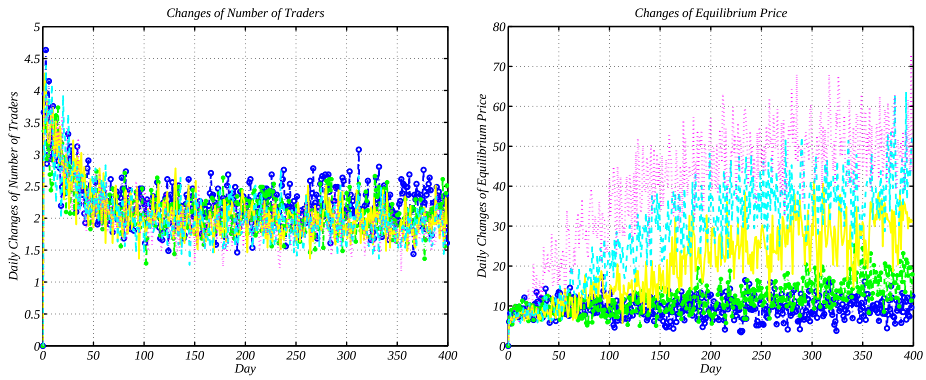

(a) Number of traders changing market each day (b) Change in equilibrium price each day

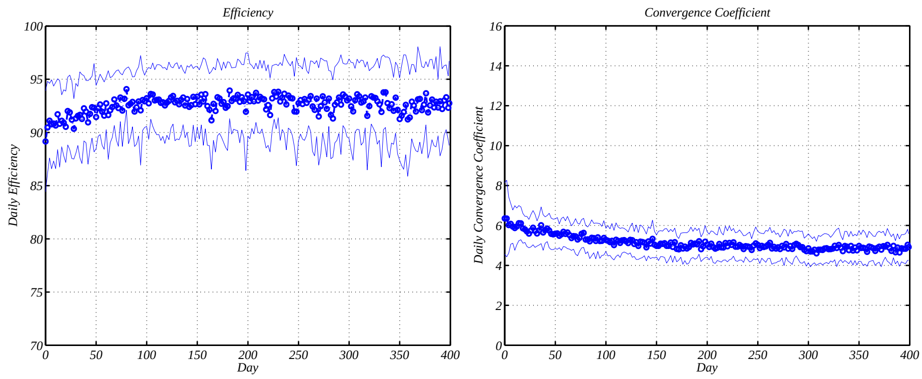

(c) Global efficiency (d) Global coefficient of convergence

Fig. 5. The performance of markets over time. GD traders in multiple CDA markets that use LL to set charges on trader profits. All plots show average value, and plots (c) and (d) show standard deviation also.

In addition to figures for individual markets, we measured the figures for all the markets combined — these provide the "global" results in Figures 5(c), Figure 5(d), Figures 6(c) and Figure 6(d). Global efficiency measures the ratio of the actual profit achieved by the traders as against the profit that theory says would be achieved were all the traders operating in a single market. The global cofficient of convergence is the RMS difference between the equilibrium prices in the individual markets and the equilibrium price that theory says would hold were all the trader in a single market. As the figures show, these values settle down to approximately constant values after between 100 and 200 days (despite the continuing change in equilibrium prices).

Both the global efficiency and the coefficient of convergence improve over time — efficiency rises while the cooefficient of convergence drops, indicating that trading is closer to the theoretical equilibrium. The results in [16] suggest that the rise in efficiency occurs because the charges imposed by markets displace extra-marginal traders,

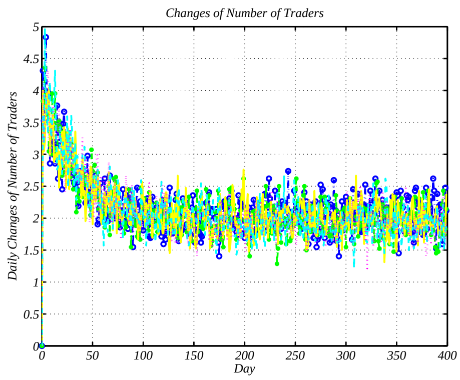

(a) Number of traders changing market each day

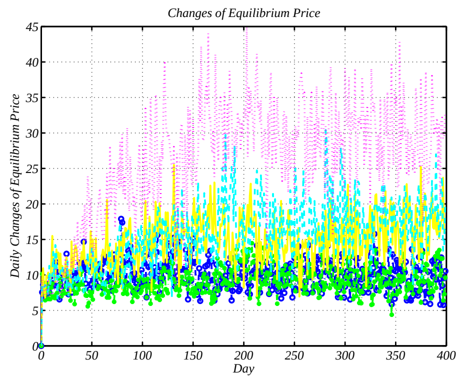

(b) Change in equilibrium price each day

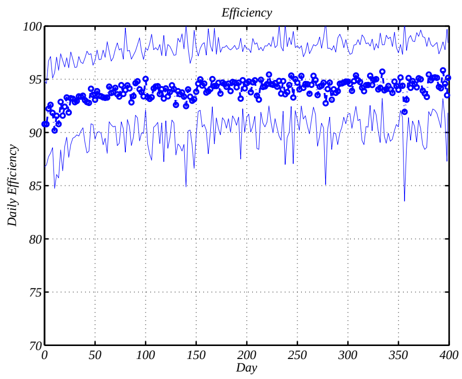

(c) Global efficiency

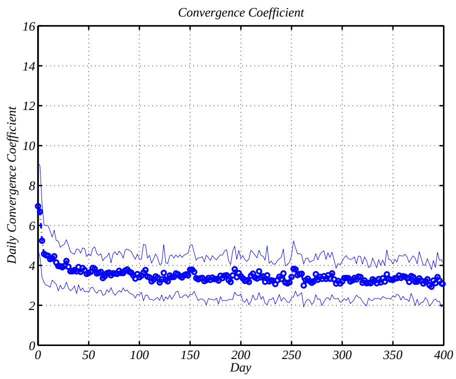

(d) Global coefficient of convergence

Fig. 6. The performance of markets over time. GD traders in multiple CH markets that use LL to set charges on trader profits. All plots show average value, and plots (c) and (d) show standard deviation also.

and thus remove the possibility of inefficient trades. The fact that trades involving extra-marginal traders will tend to occur further from the equilibrium price than trades involving intra-marginal traders explains the fall in the coefficient of convergence.

While the final values for the coefficient of convergence compare well with the values obtained in work on single markets, efficiency looks low compared with the values obtained in single markets. We believe that this is explained by trader movement (which, as we recall, continues throughout the experiment). Trader movement means that the distribution of trader private values in a given market changes every trading day. As a result, traders have to keep relearning optimal offers [6], and this continual learning re-

[6] All traders are affected by this. Traders that move are clearly in a different environment, while traders that do not move will typically have to cope with a market from which some traders have removed themselves, or a market that has new traders to cope with.

duces efficiency when compared with measurements obtained on single markets where the distribution of private values is held constant from day to day.

5 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 [27], 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 [12], but differs in that our work assesses the impact of different market charges while [12] is concerned with the information available to traders. [12] is 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 [6]. 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 [9]. 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 [9], 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 [9] — learn rather than making the same market choice at every trading opportunity. Secondly, and more importantly, the markets in [9] 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.

6 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 an early 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 "lure or learn" 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. (See [16] for more on this topic.) Our results also suggest that while splitting traders across multiple markets tends to mean that the overall system takes much longer to reach equilibrium, even if the equilibrium state that is eventually reached is not so different from that which would be reached by a single large market.

Our future, and, indeed, current, work is aimed at further untangling the behavior of competing markets. First, we want to ensure the robustness of the results we present here, and so are repeating the experiments (a) over longer periods, to be sure that what we have is indicative of performance in the steady state, when all start-up effects are removed, (b) with different market rules, for example with CH as well as CDA, and (c) with traders with different mixtures of trading and market selection strategies. 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 [29]. Finally, we want to examine the effect of different market topologies on our results, rather like [6].

Acknowledgments This work was supported by the National Science Foundation under grant NSF IIS-0329037 *Tools and Techniques for Automated Mechanism Design*, and by the EPSRC under grant GR/T10657/01 *Market Based Control of Complex Computational Systems*.

References

1. Bowling M., Veloso, M.: Multiagent learning using a variable learning rate. Artif. Intell. 136(2), 215–250 (2002)
2. Cai, K., Niu, J., Parsons, S.: On the economic effects of competition between double auction markets. In 10th Int. Workshop on Agent Mediated Electronic Commerce (2008)
3. Cliff, D., Bruten, J.: Minimal-intelligence agents for bargaining behaviours in market-based environments. Technical Report, Hewlett-Packard Research Laboratories (1997)
4. Friedman, D.: The double auction institution: A survey. In Friedman, D., Rust, J., (eds.) The Double Auction Market: Institutions, Theories and Evidence, pp 3–25. Perseus Publishing, Cambridge, MA (1993)
5. Friedman, D.: Evolutionary economics goes mainstream: A review of "The Theory of Learning in Games". J. Evol. Econ. 8(4), 423–432 (1998)
6. Gaston, M.E., DesJardins, M.: Agent-organized networks for dynamic team formation. In 4th Int. Conf. on Autonomous Agents and Multiagent Systems, ACM Press, New York (2005)
7. Gjerstad, S., Dickhaut, J.: Price formation in double auctions. Games & Econ. Behav. 22, 1–29 (1998)
8. Gode, D.K., Sunder, S.: Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. J. Polit. Econ. 101(1), 119–137 (1993)

9. Greenwald, A.R., Kephart, J.O.: Shopbots and pricebots. In 16th Int. Joint Conf. on Artif. Intell. (1999).
10. Kagel, J.H, Vogt, W.: Buyer's bid double auctions: Preliminary experimental results. In Friedman, D., Rust, J., (eds.) The Double Auction Market: Institutions, Theories and Evidence, pp 285–305. Perseus Publishing, Cambridge, MA (1993)
11. Klemperer, P.: How (not) to run auctions: The European 3G telecom auctions. Euro. Econ. Rev. 46, 829–845 (2002)
12. Ladley, D., Bullock, S.: Who to listen to: Exploiting information quality in a ZIP-agent market. In Workshop on Trading Agent Design and Analysis (2005)
13. Madhavan, A.: Trading mechanisms in securities markets. J. Finance. 47(2), 607–641 (1992)
14. Maskin, E.S. and Riley, J.G. Auction theory with private values. Am. Econ. Rev. 75, 150–155 (1985)
15. Miller, M.H., Hawke, J.D., Malkiel, B., Scholes, M.: Findings of the Committee of Inquiry Examining the Events Surrounding October 19, 1987. Technical Report, Chicago Mercantile Exchange, (1988)
16. Niu, J., Cai, K., Gerding, E., McBurney, P., Parsons, S.: Characterizing effective auction mechanisms: Insights from the 2007 TAC market design competition. In 7th Int. Conf. on Autonomous Agents and Multi-Agent Systems, IFAAMAS (2008)
17. Niu, J., Cai, K., Parsons, S., Sklar, E.: Reducing price fluctuation in continuous double auctions through pricing policy and shout improvement. In 5th Int. Conf. on Autonomous Agents and Multi-Agent Systems, ACM Press, New York (2006)
18. Niu, J., Cai, K., Parsons, S., Sklar, E.: Some preliminary results on competition between markets for automated traders. In Workshop on Trading Agent Design and Analysis (2007)
19. Phelps, S., Marcinkiewicz, M., Parsons, S., McBurney, P.: A novel method for automated strategy acquisition in n-player non-zero-sum games. In 5th Int. Conf. on Autonomous Agents and Multi-Agent Systems, ACM Press, New York (2006)
20. Roth, A.E., Erev, I.: Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. Games & Econ. Behav. 8, 164–212 (1995)
21. Satterthwaite, M.A., Williams, S.R. The Bayesian theory of the *k*-double auction. In Friedman, D., Rust, J., (eds.) The Double Auction Market: Institutions, Theories and Evidence, pp 99–123. Perseus Publishing, Cambridge, MA (1993)
22. Shah, A.: Competing exchanges: The international dimension. Bus. Standard (1997)
23. Shah, A., Thomas, S.: David and Goliath: Displacing a primary market. Global Financial Markets. 1(1), 14–21 (2000)
24. Smith, V.L.: An experimental study of competitive market behaviour. J. Polit. Econ. 70(2), 111–137 (1962)
25. Sunder, S. Markets as artifacts: Aggregate efficiency from zero-intelligence traders. In Augier, M., March, J.G. (eds) Models of a Man: Essays in memory of Herbert A. Simon. MIT Press, Cambridge (2004)
26. Sutton, R.S. and Barto, A.G. Reinforcement learning: An introduction. MIT Press, Cambridge (1998)
27. Tesauro, G., Das, R.: High-performance bidding agents for the continuous double auction. In 3rd ACM Conf. on Electronic Commerce, ACM Press, New York (2001)
28. Vickrey, W.: Counterspeculation, auctions, and competitive sealed tenders. J. Finance. 16, 8–27 (1961)
29. Walsh, W., Das, R., Tesauro, G., Kephart, J.O.: Analyzing complex strategic interactions in multi-agent systems. In Workshop on Game-Theoretic and Decision-Theoretic Agents (2002)