On the effects of competition between agent-based double auction markets*

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

Simon Parsons
Department of Computer and Information Science
Brooklyn College, City University of New York
2900 Bedford Avenue, Brooklyn, NY 11210, USA
parsons@sci.brooklyn.cuny.edu

November 13, 2008

Abstract

Real market institutions, stock and commodity exchanges for example, do not occur in isolation. The same stocks and commodities may be listed on multiple exchanges, and traders who want to deal in those goods have a choice of markets in which to trade. While there has been extensive research into agent-based trading in individual markets, there is little work on this kind of multiple market scenario. Our work seeks to address this imbalance, This paper examines how standard economic measures, like allocative efficiency, are affected by the presence of multiple markets for the same good. We find that while dividing traders between several small markets typically leads to lower efficiency and convergence than grouping them into one large market, the movement of traders between markets, and price incentives for changing markets, can reduce these losses.

1 Introduction

An *auction*, according to [7], is a market mechanism in which messages from traders include some price information — this information may be an offer to buy at a given price, in the case of a *bid*, or an offer to sell at a given price, in the case of an *ask* — and which gives priority to higher bids and lower asks. The rules of an auction determine, on the basis of the offers that have been made, the allocation of goods and money between traders. When well designed [11], auctions achieve desired economic outcomes like high *allocative efficiency* whilst being easy to implement. Auctions have been widely used in solving real-world resource allocation problems [14], and in structuring stock or futures exchanges [7].

There are many different kinds of auction. One of the most widely used kinds is the *double auction* (DA), in which both buyers and sellers are allowed to exchange offers simultaneously. Since double auctions allow dynamic pricing on both the supply side and the demand side of the marketplace, their study is of great importance, both to theoretical economists, and those seeking to implement real-world market places. The *continuous double auction* (CDA) is a DA in which traders make deals continuously throughout the auction. The CDA is one of the most common exchange institutions, and is in fact the primary institution for trading of equities, commodities and derivatives in markets such as the New York Stock Exchange (NYSE) and Chicago

*A revised and expanded version of a paper that appeared in the Proceedings of the 10th Workshop on Agent-Mediated Electronic Commerce

Mercantile Exchange[1]. Another common kind of double auction market is the *clearing-house* (CH) in which the market clears at a pre-specified time, allowing all traders to place offers before any matches are found. The CH is used, for example, to set stock prices at the beginning of trading on some exchange markets.

Our focus in this paper is on the behavior multiple auctions for the same good. This interest is motivated by the fact that such situations occur in the real world. Company stock is frequently listed on several stock exchanges. 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). Indian companies, for example, can be listed on both the National Stock Exchange (NSE) and the Bombay Stock Exchange (BSE) [30]. Until their merger in 2008, many commodities could be traded on both the CME and the New York Mercantile Exchange. Such multiple markets for the same goods have a complex dynamics. The simplest example of this is, of course, 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[2]. More complex, and less predictable dynamics occur in situations like that when the newly created Singapore International Monetary Exchange (SIMEX) claimed much of the trade in index futures on Nikkei 225 from Japanese markets in the late 1980s or when the NSE opened and proceeded to claim much of the trade volume from the established BSE [30], Changes like this take place over a long period of time, and stem from considerations such as the liquidity provided by the markets, or the (lack of) regulation that the market is subject to. Inter-market dynamics can also 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 [15].

This kind of interaction between markets has not been widely studied as yet, but seems to be increasingly relevant. Indeed, the number of online markets offering the same goods and services, suggests that studying single markets is becoming irrelevant — the default situation for electronic markets is that they are in competition with one another for participating traders and this is precisely the scenario that we study here. In addition, we focus on the markets populated by automated traders, since we believe that such traders will become increasingly widely used in electronic markets.

2 Background

Double auctions have been extensively studied using both human traders and computerized agents. Starting in 1955, Smith carried out numerous experiments investigating the behavior of such markets, documented in papers such as [31, 32, 33, 34, 35]. The experiments in [31], for example, involved human traders and showed that even with limited information available, and only a few participants, the CDA can achieve very high efficiency, comes close to the theoretical equilibrium, and responds rapidly to changing market conditions. This result was in contrast to classical theory, which suggested that high efficiency would require a very large number of traders, and led some to suggest that the form of the market itself was sufficient to ensure efficiency. In other words, Smith's results led to the suggestion that double auction markets are bound to lead to efficiency irrespective of the way that traders behave. Gode and Sunder [9] tested this hypothesis, introducing two automated trading strategies which they dubbed "zero-intelligence". The two strategies Gode and Sunder studied were *zero intelligence without constraint* (ZI-U) and *zero intelligence with constraint* (ZI-C). ZI-U traders make offers at random, while ZI-C traders make offers at random, but are constrained so as to ensure that traders do not make a loss (it is easy to see that ZI-U traders can make a loss, and so can easily lead to low efficiency markets). In the experiments reported in [9], the ZI-C traders gained high efficiency and came close enough to the performance of human traders that Gode and Sunder claimed that trader intelligence is not necessary for the market to achieve high efficiency and that only the constraint on not making a loss is important [3].

This position was attacked by Cliff [4], who showed that if supply and demand are asymmetric, the

[1] Historically this was not the case in markets, like the NYSE, that were largely controlled by specialists [13], but this is increasingly true as less and less trade is carried out on the trading floor, and more and more is carried out electronically in a way that is very similar to that permitted by the experimental PLATO system, as described in [35]

[2] 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 these provide further opportunities for arbitrage.

[3] In fact, for the markets tested in [9], even the ZI-U traders achieved pretty high efficiency, they were just outperformed by ZI-C traders in this regard.

average transaction prices of ZI-C traders can vary significantly from the theoretical equilibrium [4]. Cliff then introduced the *zero intelligence plus* (ZIP) trader, which uses a simple machine learning technique to decide what offers to make based on previous offers and the trades that have taken place. ZIP traders outperform ZI-C traders, achieving both higher efficiency and approaching equilibrium more closely across a wider range of market conditions (though [4][page 60] suggests conditions under which ZIP will fail to attain equilibrium), prompting Cliff to suggest that ZIP traders embodied the minimal intelligence required [5]. A range of other trading algorithms have been proposed — including those that took part in the Santa Fe double auction tournament [28], the *Roth-Erev* approach (RE) [27] which is form of reinforcement learning *Gjerstad-Dickhaut* approach (GD) [8] which seeks to maximise expected profit — and the performance of these algorithms evaluated under various market conditions.

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 [24, 40] 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. [22, 37, 40] also make it clear that the results hinge on both auction design and the mix of trading strategies used). This complex interaction between mechanisms and traders helps to make the task of designing mechanisms a complex one, and several authors have suggested computational approaches to this task [6, 22, 23, 39].

Despite all of this work, there has been no systematic study of the use of automated traders in multiple connected markets [38], though a number of different scenarios have been investigated. [2] uses agent-based methods to examine the effects of linked markets on financial crises, while [41] looks at the effect of different trade routes on price convergence. [16, 17] study the bull-whip effect [12][6] in supply chains. In addition, some initial results on multiple auctions that compete for traders were presented in [20] and the design of such auctions is the focus of the TAC Market Design competition analyzed in [21]. The work we report here further extends the use of agent-based computational economics to study groups of connected markets.

3 Experimental Setup

The experiments we carried out explore the economic effects of having a number of parallel markets, and the consequences of allowing traders to move between these markets.

3.1 Software

To experiment with multiple markets, we used the Java-based server platform JCAT [10, 18]. JCAT provides the ability to run multiple double auction markets populated by traders that use a variety of trading strategies, and was used to support the 2007 TAC Market Design competition [3]. Auctions in JCAT 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 (shouts) to buy or sell, and we distinguish different days because at the beginning 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 on 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 JCAT 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. 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.

[4] The experiments in [9], while reflecting typical market conditions, might be considered to represent easy conditions from which to attain equilibrium. In contrast, the experiments in [31] show convergence to equilibrium from a much wider range of initial conditions.

[5] A point maybe undermined by the ZIP variant proposed by Preist and van Tol [26], which is arguably simpler and hence more minimal.

[6] Where small fluctuations in supply in one market can have an effect that magnifies through the network.

3.2 Traders

Traders in our experiments 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.* We studied markets in which all the traders used the same trading strategy, and considered three such strategies:

- Gode and Sunder's zero intelligence with constraint (ZI-C) strategy [9];
- Cliff's zero intelligence plus (ZIP) strategy [4]; and
- Roth and Erev's reinforcement learning strategy (RE) [27].

The reason for picking the first of these is that given by [19, 39], that since ZI-C is not making bids with any intelligence, any effects we see have to be a result of market structure, rather than a consequence of the trading strategy, and hence will be robust across markets inhabited by different kinds of trader. The reason for picking ZIP and RE is that given by [24]. The first of these strategies is typical of the behavior of automated traders, while the second is a good model of human bidding behavior. Using both will give us results indicative of markets with both human and software traders.

The market selection strategy is based on a simple model for reinforcement learning. Traders treat the choice of market as an n-armed bandit problem that they solve using an ϵ-greedy exploration policy [36]. Using this approach the behavior of the agents is controlled by two parameters ϵ and α. A 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 [36]. If α is 1, ϵ remains constant, while if α takes any value in $(0, 1)$, ϵ will reduce over time. For these experiments, we set α to 1, and ϵ to 0.1. The results from or previous work on the interactions between multiple markets [20] suggest that market selection behavior is rather insensitive to the parameters we choose here.

JCAT is typically set up to use the market selection strategy to decide which market each trader should participate in at the start of each day. Since this facility can be disabled, however, we could experiment with two different kinds of trader movement:

- Mobile: traders choose a market at the start of each day (this may be the same market in which the traders participated the previous day).
- Stationary: traders always remain in the same market.

Each trader is permitted to buy or sell at most five units of goods per day, and each trader has a private value for these goods, a value which is drawn from a uniform distribution between $50 and $150. A given trader is assumed to have the same private value for all goods that it trades throughout the entire experiment.

3.3 Markets

While JCAT allows us to charge traders in a variety of ways, we used just four kinds of charge in the work reported here:

- Shout fees, charges made by the market for each shout made by a trader.
- Information fees, charges made by the market for information about shouts made by other traders in the market.
- Transaction fees, charges made by the market for each transaction executed by a trader.
- Profit fees, charges made by the market on the bid/ask spread of any transactions they execute.[7].

[7]The name arose since the bid/ask spread is the transaction surplus, and with the $k = 0.5$ rule we usually use for allocating the surplus, the surplus is thus directly related to the profit realised by both agents.

We set shout, information and transaction fees to constant, low, figures — \$0.1, \$2 and \$0.1 respectively. These are values typical of those adopted by entrants in the 2007 TAC Market Design Competition, and, as [21] discusses, are sufficient to provide a small negative reinforcement that encourages traders to leave markets in which they are not managing to make trades.

We used three different mechanisms for setting the profit fees:

- Free: no profit fees are charged.
- Fixed: a constant proportion, typically 10%, 20%, 30%, 40% and 50% of the surplus on a transaction, is taken as a fee.
- Lure-or-learn fast (LL): a version of the ZIP strategy for traders [5] adapted for markets and introduced by [20] under the name "zero intelligence"[8]. A LL market adjusts its charges 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.

In all of our experiments the markets are populated by 100 traders, evenly split between buyers and sellers. We run five markets — named M0, M1, M2, M3 and M4. When the markets have fixed profit charges, M0 has the lowest charge, and the chrages increase from M1 to M2 to M3 to M4. When the markets use LL, the initial charges used by each market follow the same pattern.

3.4 Experiments

Our main aim in this work was to answer the questions "what is the economic effect of running a number of parallel markets?", and "what is the effect of different charging regimes?", so our basic comparisons are between the situation in which all traders transact in a single market, and the situation in which traders are split across a number of markets for different charging mechanisms. We were also interested in the effect of traders moving between markets — the results we published in [20] tell us that traders move between markets due to the charges imposed by markets, but it does not say anything about the effect of that movement on the overall performance of the markets in economic terms.

These considerations led us to compare the performance of the single market, and the multiple markets in different scenarios. We considered six different scenarios — one scenario for each combination of charging mechanism (free, fixed and LL) and traders that are either mobile or stationary. For a given trading strategy, we considered all six of these scenarios for both the CH and the CDA.

Thus we ran a total of 36 experiments, six scenarios for the two different kinds of market and the three different trading strategies. For each experiment we obtained results for both trades split across five markets and all the traders concentrated in one market. Each of these 36 experiments was run for 400 trading days, with each day being split into 50 0.5-second-long rounds. We repeated each experiment 50 times.

3.5 Measurements

The effectiveness of a market can be measured in a number of different ways. We compare markets in our experiments in terms of their allocative efficiency and coefficient of convergence. *Allocative efficiency*, E_a, is used to measure how good a market is at generating global profits [25]. The *actual overall profit*, P_a, of an auction is:

$$P_a = \sum_i |v_i - p_i| \tag{1}$$

for all agents who trade, where p_i is the price of a trade made by agent i and v_i is the private value of agent i. The *equilibrium profit*, P_e, is:

$$P_e = \sum_i |v_i - p_0| \tag{2}$$

for all buyers whose private value is no less than the *equilibrium price*, p_0, and all sellers whose private value is no greater than p_0. The equilibrium price is the price at which the number of goods sold equals

[8]The name is inspired by Bowling's "win or learn fast" [1] and prevents us from confusing zero-intelligence traders with zero-intelligence markets

the number of goods bought and can be computed from the private values of the traders assuming that no trader makes a loss. E_a, is then P_a/P_e expressed as a percentage.

$$E_a = \frac{P_a}{P_e} \times 100 \tag{3}$$

This tells us how close a market is to theoretical equilibrium in terms of profits made — the more efficient the market, the closer E_a is to 100. However, it says nothing about how close a market is to trading at the equilibrium price. For the latter we use the *coefficient of convergence* α [31]. α measures the deviation of transaction prices from the equilibrium price:

$$\alpha = \frac{\sqrt{\frac{1}{n}\sum_i (p_i - p_0)^2}}{p_0} \times 100 \tag{4}$$

where n is the number of traders in the market. The closer a market is to trading at equilibrium, the closer α is to 0.

For the multiple market experiments, we measure the efficiencies and convergence of each individual market. We also compute what we call the *global* values which assess the measurements across all the parallel markets. The calculation of these global values is similar to the calculation of the measurements for a single market. Global efficiency E_a^g is computed as:

$$E_a^g = \frac{\sum_j \sum_i |v_i^j - p_i^j|}{\sum_j \sum_i |v_i^j - p_0^g|} \tag{5}$$

where v_i^j is the private value of agent i in market j, p_i^j is the price paid by agent i in market j, and p_0^g is the equilibrium price that would hold were all the traders in a single market. The global value of α is computed similarly:

$$\alpha = \frac{\sqrt{\frac{1}{n^g}\sum_j \sum_i (p_i^j - p_0^g)^2}}{p_0^g} \times 100 \tag{6}$$

where n^g is the total number of traders in all markets. These global measurements give us a way of examining the properties of a set of markets — we can compare the global measurement across the set of markets with the mesurement for a market that contains the same traders as are spread across the set. They are appropriate measures for our purposes because we are interested in examining how far multiple markets diverge from the performance of an individual market. However, it should be noted that they are rather harsh measures since the markets can, for example, be individually efficient but globally somewhat inefficient because of the distribution of traders.

4 Results

Our analysis of the results of the experiments starts by examining the changes in markets over time, before considering the differences in behavior between single and multiple markets. This leads us to the conclusion that charging in markets leads to a segmentation of traders that improves market efficiency and convergence. As a result, we examine the change in efficiency in some detail, especially the relationship between efficiency and the number of traders in each market.

4.1 Change in markets over time

Figure 1 gives one view of the the experiments with mobile ZIP traders. It shows the number of traders leaving each of the five markets every day for both CDA and CH markets, for markets that don't charge, for markets that use fixed profit charges, and for markets that adjust their profit charges using the LL mechanism. The lines plotting these numbers for each of the markets are superimposed over each other since the performances of the markets in this regard are indistinguishable. The reason for plotting these values is that they summarise the amount of "churn" (to use the marketing term) in the markets. When the situation

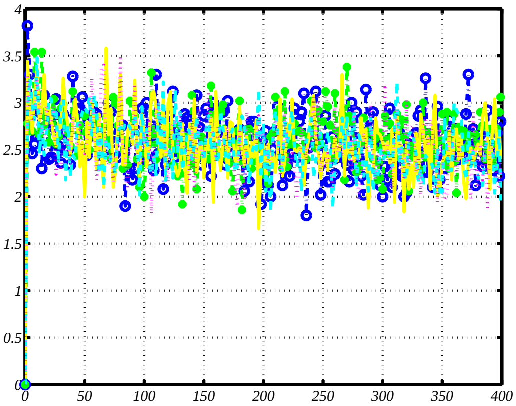

(a) CDAs with no charges

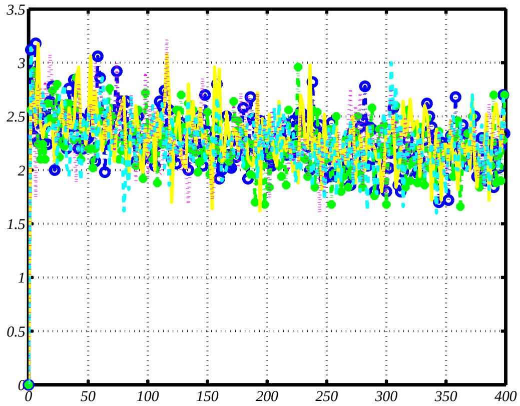

(b) CHs with no charges

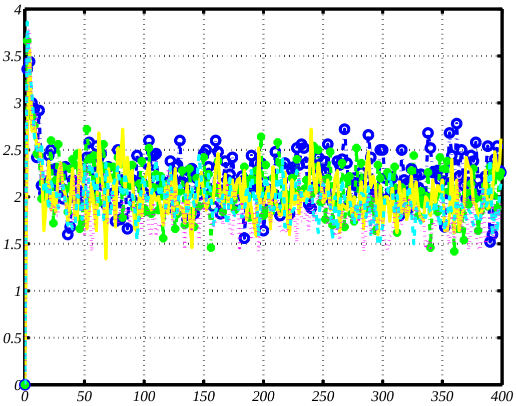

(c) CDAs with fixed charges

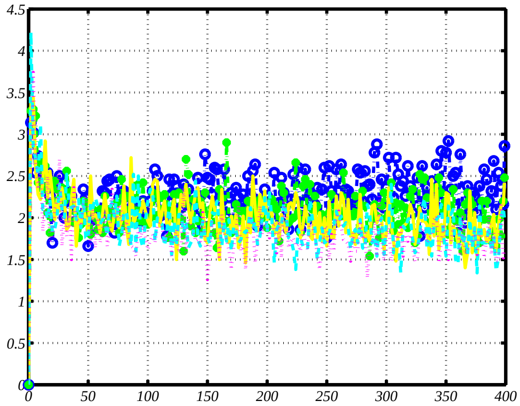

(d) CHs with fixed charges

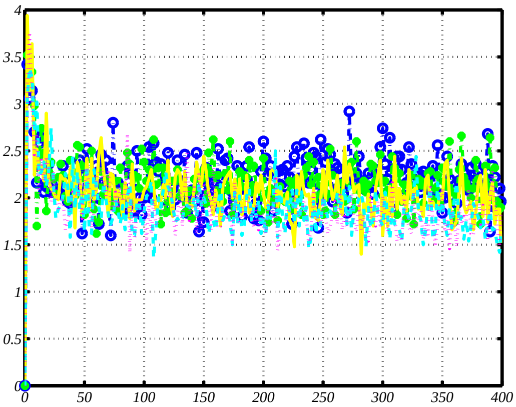

(e) CDAs with LL charges

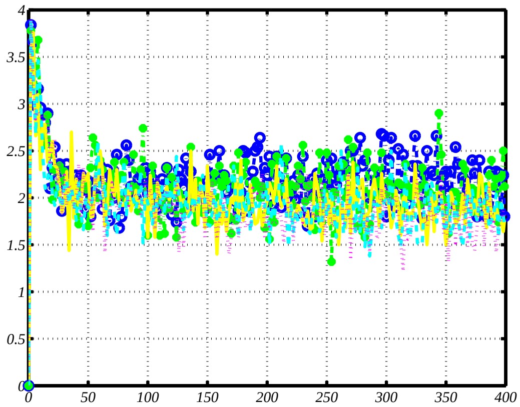

(f) CHs with LL charges

Figure 1: How individual markets change over time — the number of traders changing market. The x-axis gives the trading day, and the y-axis gives the number of traders that leaving a given market on that day. All markets features mobile ZIP traders. The markets in (a), (c) and (e) are CDAs with different charging patterns, while the markets in (b), (d) and (f) are CHs with different charging patterns.

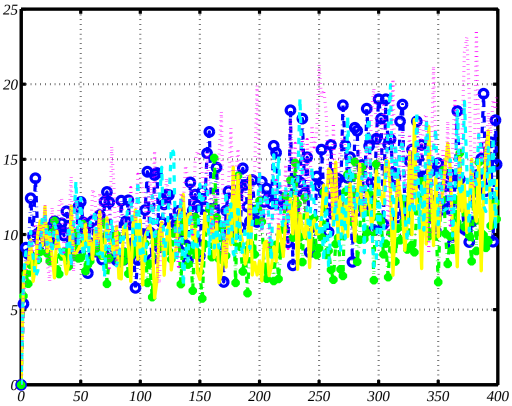

(a) CDAs with no charges

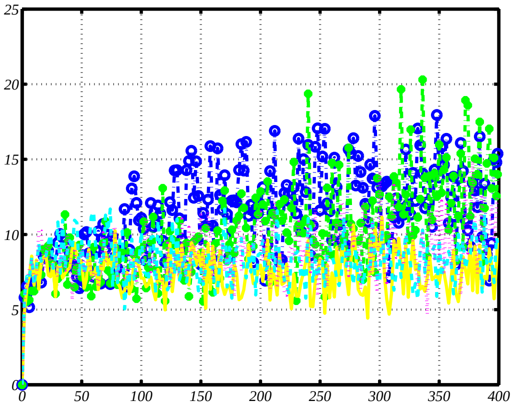

(b) CHs with no charges

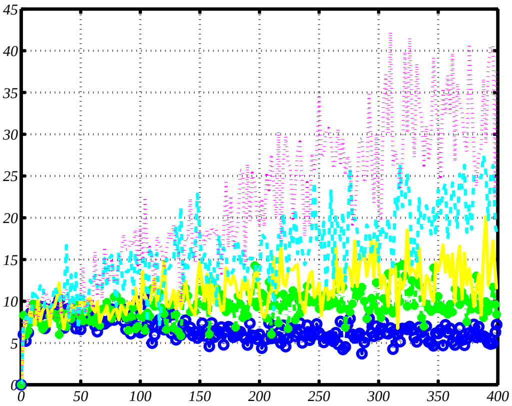

(c) CDAs with fixed charges

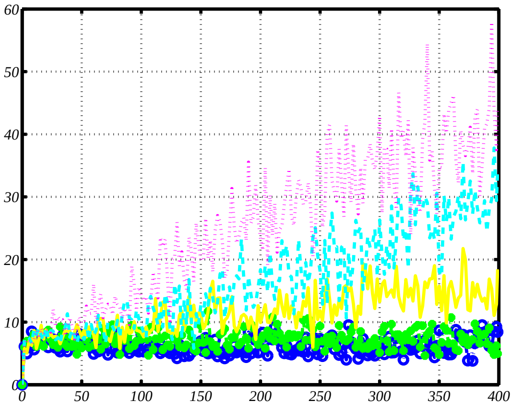

(d) CHs with fixed charges

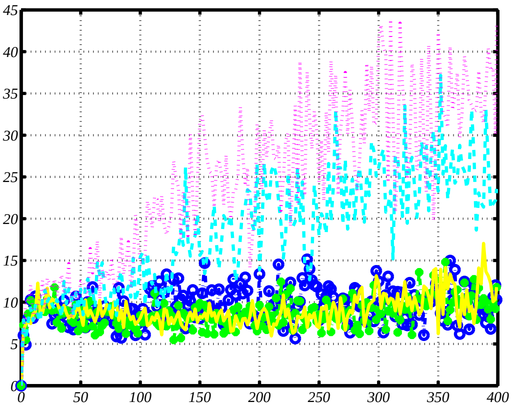

(e) CDAs with LL charges

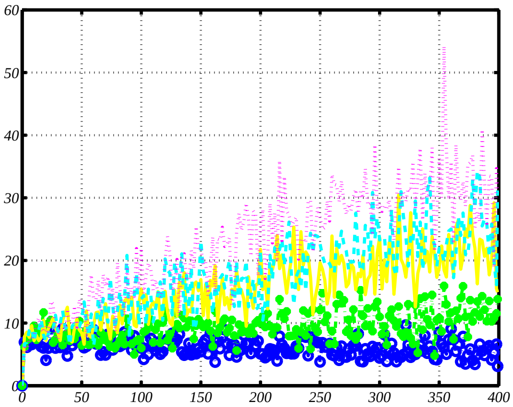

(f) CHs with LL charges

Figure 2: How individual markets change over time — the change in equilibrium price. The x-axis gives the trading day, and the y-axis gives the change in equilibrium price between that day and the previous day. All markets features mobile ZIP traders. The markets in (a), (c) and (e) are CDAs with different charging patterns, while the markets in (b), (d) and (f) are CHs with different charging patterns.

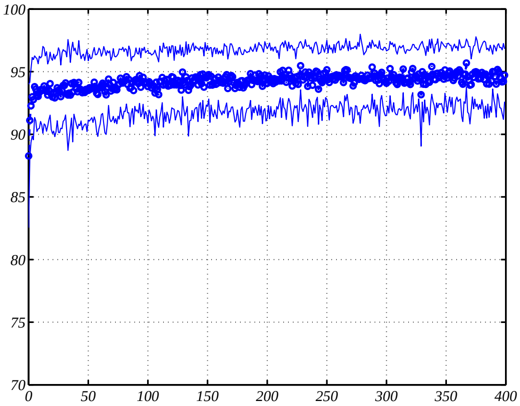

(a) CDAs with no charges

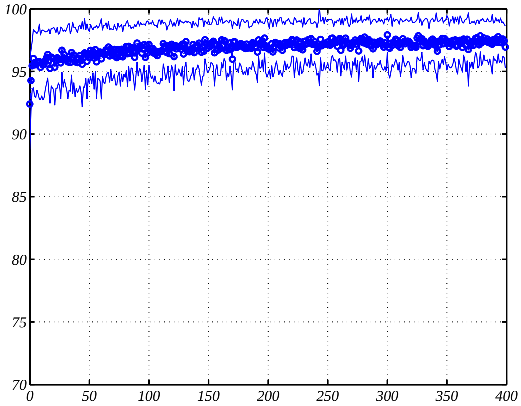

(b) CHs with no charges

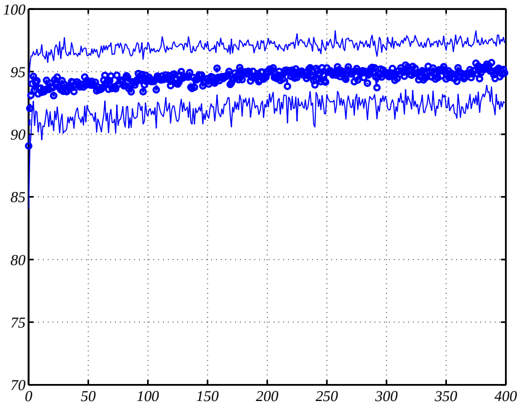

(c) CDAs with fixed charges

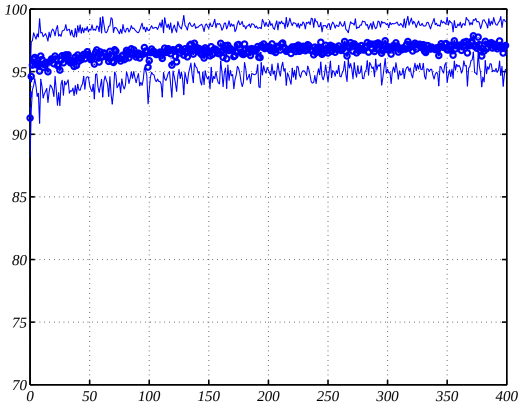

(d) CHs with fixed charges

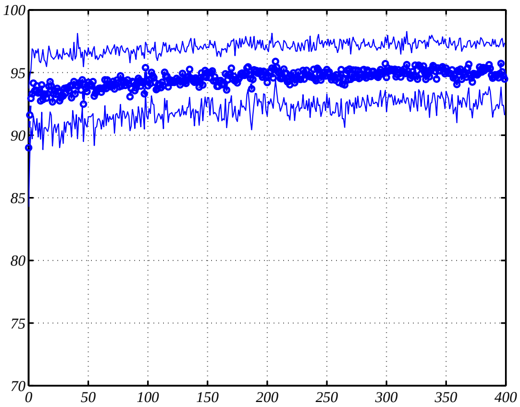

(e) CDAs with LL charges

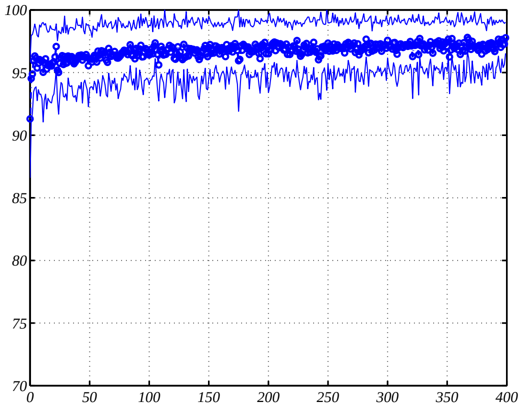

(f) CHs with LL charges

Figure 3: How markets as a group change over time — global efficiency. The x-axis gives the trading day, and the y-axis gives the global efficiency on that day. All markets features mobile ZIP traders. The markets in (a), (c) and (e) are CDAs with different charging patterns, while the markets in (b), (d) and (f) are CHs with different charging patterns. Each figure shows the average value and one standard deviation.

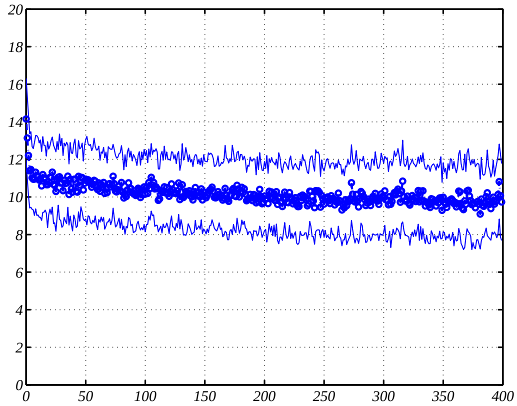

(a) CDAs with no charges

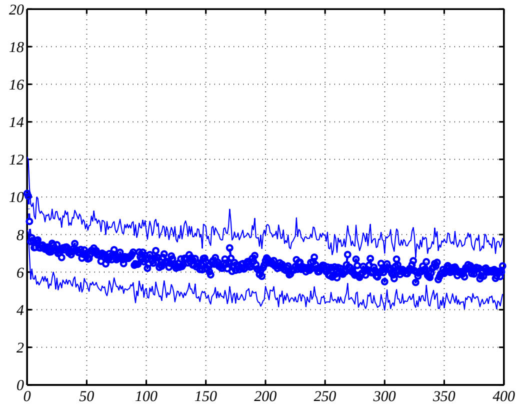

(b) CHs with no charges

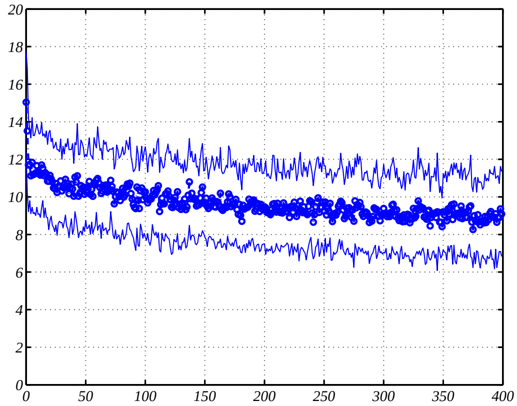

(c) CDAs with fixed charges

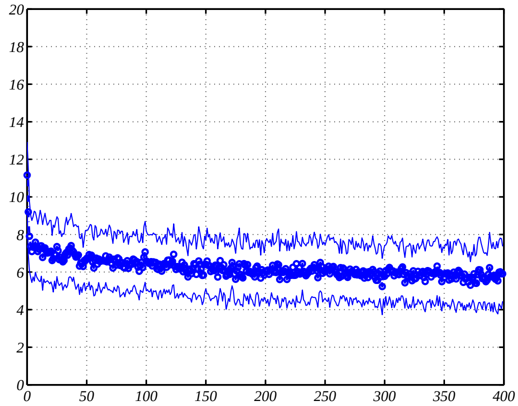

(d) CHs with fixed charges

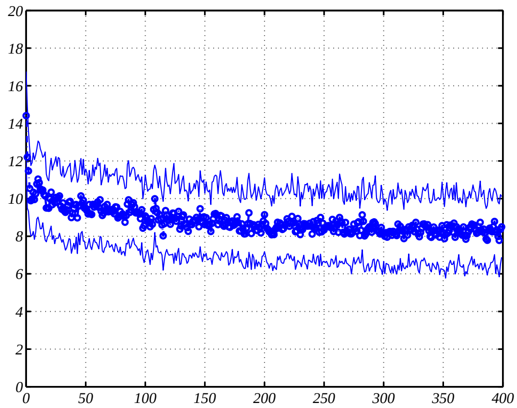

(e) CDAs with LL charges

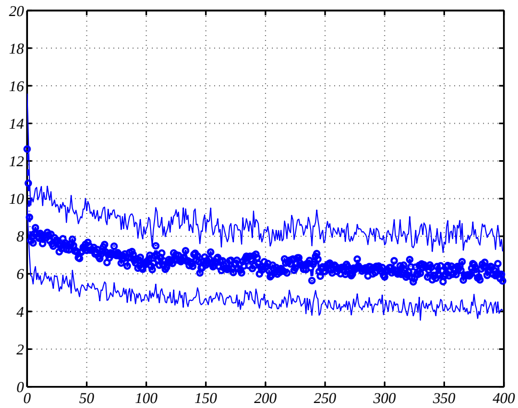

(f) CHs with LL charges

Figure 4: How markets as a group change over time — global coefficient of convergence. The x-axis gives the trading day, and the y-axis gives the global coefficient of convergence on that day. All markets features mobile ZIP traders. The markets in (a), (c) and (e) are CDAs with different charging patterns, while the markets in (b), (d) and (f) are CHs with different charging patterns. Each figure shows the average value and one standard deviation.

is settled, the number of traders moving will be low, when the situation is unsettled, the number of traders moving will be higher. All six experiments show that movement decreases over time. It decreases most in the LL markets, and least in the markets that don't charge and it decreases much more quickly in the markets that do charge than those that don't. It never drops much below 2, which is consistent with having ϵ in the market selection mechanism fixed to 0.1 so that traders continue to explore throughout each experiment [9].

This movement of traders necessarily has a effect on the trading that takes place in each of the markets. Whereas we would expect a single market to rapidly approach equilibrium after just a few days at most — whether traders are software agents [8, 26] or humans [31] — in the multiple market case, this does not happen. Figure 2, which plots the daily *change* in equilibrium price in each market in Figure 1, is testimony to the way that that the markets don't have a settled equilibrium. Every market has a non-zero daily change, even at the end of the 400 day period. However, we do see a certain level of stability emerge in the markets that charge — by 300 days or so, while there are still changes from day to day, the trend is for the average change in equilibrium price to settle towards a limit. This limit varies from market to market. For example in the experiment on clearing houses that use LL (Figure 2 (f)) the limit ranges from around \$5 in M0 to around \$30 in M4, while in for the experiment on CDAs that use fixed charges, the range is from around \$5 in M0 to around \$50 in M4.

In case the changes in equilibrium price smack of anarchy in individual markets, consider Figures 3 and Figures 4. These plot the global values of efficiency and the coefficient of convergence for the same experiments as in Figure 1 and Figure 2. As described above, global efficiency is computed by summing actual trader profits and then dividing by the theoretical profit that would be made *if all the traders were in the same market*. It thus gives us a picture of our set of markets taken as a whole, and shows that, despite the churn in individual markets, the overall picture has settled down after around 200 days. The global coefficient similarly settles down, though it takes closer to 300 days to stop changing significantly. The global coefficient of convergence measures the average distance of each transaction from where theory says it should take place, once again, if all the traders were in the same market. The fact that it settles down to a reasonably constant value suggests that most transactions are taking place at a similar value from day to day, and the fact that the global coefficient of convergence would not be unreasonably high for a single market suggests that — despite the fluctuation in transaction price in individual markets — most transactions that take place do so close to the theoretical equilibrium.

Though we have only shown the results for ZIP trades here, the other experiments have very similar results. A full set of results is available from the authors on request. All these results parallel those we reported in [20], and the together these results suggest that the effects we see hold for a wide range of charging regimes and market selection strategies.

4.2 Comparing single and multiple markets

Having sketched the overall behavior of the markets in our experiments, the main results of this paper are given in Tables 1–6. These give, for each of the experiments outlined above, both the efficiencies and coefficients of convergence for markets M0 to M4, a single market containing all the traders, and the global values computed as in Section 3.5. The single market value differs from the global measure in that the actual trader profits and equilibrium prices are obtained in the single market rather than in the individual markets (while the theoretical profit and global equilibrium price is the same in both cases). The values of the efficiency and coefficient of convergence given is averaged over the last 100 days of each experiment as well as across the 50 runs of each experiment.

The first point to make is that, just as one would expect from usual theoretical analysis, say [29], the efficiency of the single market of 100 traders is generally greater than the global efficiency (though there is an exception). Not only is this in agreement with the theory, but it is not surprising. The theoretical profit is the same in both cases, so for the global efficiency to be higher, the individual markets would have to do a better job of matching traders than the single market. Clearly the churn will make any optimal matching hard to sustain even if it occurs in the first place. Similarly, the coefficient of convergence of the single market is typically lower than that of the global market (the exception here is the performance of RE in the CDA markets, but this is against the backdrop of a generally poor performance in this market). Again

[9] With 100 traders that have learnt the best market to trade in, we'd still expect 10 (100×0.1) to change market each trading day on average, which is, of course, approximately two for each market.

this is not surprising. The variation measured by the coefficient of convergence is against the equilibrium price given all trader private values, and this will only be the equilibrium price of each individual market if traders are even distributed by private value. Spreading traders across markets we'd expect that even if the equilibrium price was the same on average, its variance would increase as the number of markets increased[10].

Some other interesting points emerge. First, looking just at the global efficiency values, we see that across all three trading strategies, markets with mobile traders are more efficient than markets with stationary traders. It therefore seems to be the case that trader mobility leads to higher efficiency. Traders that move to maximize their own expected profit, which is the effect of the market selection strategy we use, end up improving the performance of the markets as a whole. Second, again across all three trading strategies, the best performing (in terms of efficiency) individual markets, with mobile traders, that make charges on profits outperform any of the corresponding individual markets that do not charge[11]. Thus, not only does it seem that mobility leads to higher efficiency, but it also seems that charging does.

Third, the effect of charging is strong enough that with ZIP and RE traders (the ones that might be considered more rational because they pick offers that aim to maximize their profits) these best performing individual markets do so well that they lift the global performance of the charging markets with mobile traders above that of the markets that don't charge. (This despite the fact that the higher charging individual markets have considerably lower efficiencies than the markets that do not charge). Thus, not only do individual markets benefit from the charges, but it seems that *overall* the markets benefit — they certainly manage to extract more total profits that way.

The results for the coefficient of convergence show some of the same features, but they are less marked. In all cases except that of ZIP traders in CH markets, mobile traders generate lower coefficients of convergence than stationary traders, and the better performing charging markets with mobile traders outperform (in the sense of having a lower coefficient of converegnce) the markets that do not charge.

4.3 The effects of charging on efficiency

The results so far suggest that when markets charge, it has some effect on their efficiency. However, it is not clear whether this is a direct effect or an indirect effect. In particular, it is clear that when markets charge at different rates, the number of traders in the markets will change — for example [20] shows that, in exactly the experimental setup we have here, traders tend tomove from higher charging markets to lower charging markets (as one would expect). Thus it might be the case that the changes in efficiency that we see are only due to the changing numbers of traders rather than some more complex sorting of traders. To investigate this, we ran a further set of experiments.

In these latter experiments, we only considered the two charging mechanisms — fixed and LL — that impose fees, and focused once again on the profit fees. We did, however, consider additional levels of charging, running experiments with the following charging regimes:

- M0 charges 5%, M1 charges 10%, M2 charges 15%, M3 charges 20% and M4 charges 25%;
- M0 charges 10%, M1 charges 20%, M2 charges 30%, M3 charges 40% and M4 charges 50%;
- M0 charges 15%, M1 charges 30%, M2 charges 45%, M3 charges 60% and M4 charges 75%.

As before, these were the values used throughout the experiments by the fixed markets, and for the initial trading day by the LL markets. Carrying out these experiments for both CDA and CH markets, with all three kinds of trader, and using fixed and LL markets with three different charging levels gives us 36 separate experiments. Each experiment was run over 400 trading days, and repeated 100 times.

The results of these experiments for ZIP traders and fixed charging markets are shown in Figures 6–5 and Tables 7–8. As before we only included the results from some of the experiments in the interests of space — these are representative, and the full set of results can be obtained from the authors.

Table 7 and Table 8 give, respectively, the relationship between the average number of traders in the lowest charging market and the average global allocative efficiency, and the relationship between the average number of traders in the lowest charging market and the average allocative efficiency in that market. Both

[10] You can see this effect in the fact that individual markets with ZIP traders often have a considerably lower coefficient of convergence than the global market.

[11] In other words, M0 under fixed and LL charging has greater efficiency than any of the markets which are free.

Table 1: Allocative efficiency for markets with ZIC traders in single-market and multiple-market scenarios.

			multiple markets						single market
			M0	M1	M2	M3	M4	global	
CDA	Mobile	Fixed	87.14	80.67	71.47	65.90	64.99	85.45	$\star$***88.86***
			11.96	*20.07*	*27.04*	*29.85*	*30.99*	*3.49*	*2.05*
		LL	92.34	92.17	***94.69***	90.43	92.19	$\star$**93.16**	87.58
			13.54	*14.22*	*9.48*	*17.20*	*15.27*	*2.62*	*2.35*
		Free	78.80	76.37	78.27	79.36	78.24	85.66	$\star$***88.92***
			22.46	*25.48*	*22.41*	*22.22*	*23.31*	*3.03*	*2.01*
	Stationary	Fixed	83.10	82.71	83.59	82.91	83.86	77.02	$\star$***88.86***
			11.11	*9.19*	*9.19*	*8.25*	*8.40*	*5.80*	*2.05*
		LL	82.38	84.70	81.65	80.51	81.51	77.18	$\star$***87.58***
			10.63	*10.42*	*12.08*	*13.49*	*14.91*	*6.05*	*2.35*
		Free	81.20	81.83	81.65	80.58	81.20	77.25	$\star$***88.92***
			11.05	*10.86*	*12.48*	*11.29*	*12.55*	*5.37*	*2.01*
CH	Mobile	Fixed	***84.99***	75.05	69.12	57.41	55.83	81.16	$\star$**81.99**
			20.01	*24.85*	*30.87*	*30.87*	*31.30*	*3.20*	*2.99*
		LL	***95.70***	94.57	95.45	94.20	94.84	$\star$**94.21**	81.30
			4.80	*8.25*	*6.18*	*9.88*	*7.34*	*2.01*	*2.63*
		Free	74.58	76.38	71.83	72.94	77.90	83.72	$\star$***83.89***
			24.73	*22.10*	*24.96*	*25.31*	*21.37*	*3.14*	*2.76*
	Stationary	Fixed	86.40	86.26	85.56	86.74	***87.67***	77.80	$\star$**81.99**
			8.47	*8.85*	*7.63*	*8.72*	*8.72*	*5.11*	*2.99*
		LL	79.78	81.08	***81.72***	78.62	77.69	76.09	$\star$**81.30**
			9.50	*9.95*	*7.99*	*12.24*	*13.73*	*5.13*	*2.63*
		Free	79.35	80.77	82.46	80.32	81.29	76.86	$\star$***83.89***
			11.82	*10.48*	*9.11*	*10.12*	*11.86*	*4.66*	*2.76*

Italic numbers are standard deviations, **bold** numbers indicate the better of the global and single market values, ***bold italic*** identifies the largest value on each line, and $\star$ denotes that where these comparisons are significant at the 95% level. The charges on profit rise linearly from M0 (10%) to M4 (50%). In the case of the LL markets, these are the figures from which charges start.

Note that in a single market it makes no sense for traders to move since there is no market to move to or from. As a result, figures for mobile and stationary traders are the same.

Table 2: Allocative efficiency for markets with ZIP traders in single-market and multiple-market scenarios.

			multiple markets						single market
			M0	M1	M2	M3	M4	global	
CDA	Mobile	Fixed	97.06	96.24	96.11	94.60	93.24	94.53	$\star$***97.93***
			5.59	*7.54*	*7.76*	*11.84*	*14.38*	*2.59*	*1.27*
		LL	96.14	95.66	95.18	94.92	95.38	94.48	$\star$***98.55***
			7.60	*9.13*	*10.71*	*10.30*	*9.82*	*2.96*	*1.04*
		Free	96.04	96.39	96.17	95.88	95.63	94.22	$\star$***99.49***
			8.19	*6.80*	*7.34*	*8.26*	*9.18*	*2.63*	*0.47*
	Stationary	Fixed	97.47	97.86	97.46	***98.05***	96.98	91.14	$\star$**97.93**
			3.03	*3.23*	*3.34*	*4.96*	*4.16*	*4.16*	*1.27*
		LL	97.66	97.85	97.80	97.97	97.87	90.37	$\star$***98.55***
			2.96	*2.72*	*2.76*	*3.25*	*2.65*	*4.15*	*1.04*
		Free	97.27	97.59	97.60	97.55	97.54	89.62	$\star$***99.49***
			4.11	*3.75*	*3.49*	*4.57*	*4.16*	*5.10*	*0.47*
CH	Mobile	Fixed	98.85	98.53	97.52	96.38	95.09	96.62	$\star$***99.74***
			4.74	*8.25*	*11.25*	*13.75*	*13.75*	*2.10*	*0.52*
		LL	98.32	97.73	98.30	97.13	97.73	96.78	$\star$***99.68***
			4.68	*8.17*	*5.13*	*9.10*	*7.14*	*2.20*	*0.49*
		Free	97.96	97.79	98.41	98.24	98.17	96.91	$\star$***99.75***
			6.77	*7.62*	*4.60*	*5.02*	*5.98*	*2.06*	*0.49*
	Stationary	Fixed	99.04	99.01	99.36	99.22	99.01	90.54	$\star$***99.74***
			3.45	*2.06*	*3.40*	*3.96*	*4.98*	*4.98*	*0.52*
		LL	99.35	99.16	99.21	99.32	99.03	92.50	$\star$***99.68***
			1.79	*2.82*	*2.67*	*2.04*	*4.11*	*4.19*	*0.49*
		Free	99.29	98.56	99.06	99.06	99.19	91.34	$\star$***99.75***
			2.55	*5.66*	*3.35*	*2.97*	*2.91*	*4.76*	*0.49*

Italic numbers are standard deviations, **bold** numbers indicate the better of the global and single market values, ***bold italic*** identifies the largest value on each line, and $\star$ denotes that where these comparisons are significant at the 95% level. The charges on profit rise linearly from M0 (10%) to M4 (50%). In the case of the LL markets, these are the figures from which charges start.

Note that in a single market it makes no sense for traders to move since there is no market to move to or from. As a result, figures for mobile and stationary traders are the same.

Table 3: Allocative efficiency for markets with RE traders in single-market and multiple-market scenarios.

			multiple markets						single market
			M0	M1	M2	M3	M4	global	
CDA	Mobile	Fixed	***89.89***	88.62	79.54	68.81	68.57	85.79	⋆**89.14**
			9.29	*29.06*	*39.19*	*40.06*	*3.07*	*3.07*	*1.68*
		LL	***87.88***	87.35	86.94	87.45	87.11	86.42	⋆**87.39**
			11.79	*14.49*	*14.21*	*13.91*	*15.33*	*3.12*	*2.46*
		Free	86.97	87.29	85.85	85.37	84.93	85.59	⋆***89.37***
			14.74	*12.08*	*17.89*	*18.11*	*18.58*	*3.00*	*1.69*
	Stationary	Fixed	88.47	***89.79***	88.17	88.26	89.40	82.07	⋆**89.14**
			4.85	*4.80*	*5.33*	*4.70*	*4.92*	*4.92*	*1.68*
		LL	87.75	87.62	87.12	86.97	***88.09***	81.42	⋆**87.39**
			5.53	*7.25*	*6.74*	*5.66*	*5.49*	*5.49*	*2.46*
		Free	88.64	***89.53***	87.93	88.74	87.72	81.15	⋆**89.37**
			5.94	*5.18*	*5.65*	*4.98*	*5.59*	*5.26*	*1.69*
CH	Mobile	Fixed	99.01	97.73	94.52	89.83	87.90	95.90	⋆***99.33***
			5.30	*15.90*	*24.81*	*27.67*	*27.67*	*2.94*	*0.86*
		LL	97.56	97.24	97.63	97.46	97.15	94.66	⋆***99.42***
			4.93	*6.62*	*4.27*	*6.35*	*6.89*	*3.32*	*0.78*
		Free	97.18	97.87	97.41	97.23	97.27	95.51	⋆***99.20***
			6.28	*8.34*	*8.84*	*8.54*	*8.54*	*2.90*	*0.92*
	Stationary	Fixed	98.46	98.51	98.50	98.56	98.89	91.99	⋆***99.33***
			2.79	*2.73*	*2.62*	*2.41*	*4.60*	*4.60*	*0.86*
		LL	98.65	98.66	98.58	98.81	98.84	88.13	⋆***99.42***
			2.49	*2.36*	*2.57*	*2.48*	*2.13*	*6.42*	*0.78*
		Free	98.44	98.66	98.73	98.65	98.59	89.48	⋆***99.20***
			2.58	*2.30*	*2.52*	*2.86*	*5.59*	*5.59*	*0.92*

Italic numbers are standard deviations, **bold** numbers indicate the better of the global and single market values, ***bold italic*** identifies the largest value on each line, and ⋆ denotes that where these comparisons are significant at the 95% level. The charges on profit rise linearly from M0 (10%) to M4 (50%). In the case of the LL markets, these are the figures from which charges start.

Note that in a single market it makes no sense for traders to move since there is no market to move to or from. As a result, figures for mobile and stationary traders are the same.

Table 4: Coefficient of convergence for markets with ZIC traders in single-market and multiple-market scenarios.

			multiple markets						single market
			M0	M1	M2	M3	M4	global	
CDA	Mobile	Fixed	15.33	14.98	14.35	***13.71***	14.01	15.56	⋆**13.94**
			2.92	*4.46*	*6.05*	*6.44*	*7.36*	*0.95*	*0.71*
		ZIP	16.51	16.44	16.52	15.80	16.54	17.25	⋆***14.79***
			5.80	*5.74*	*4.60*	*5.96*	*5.57*	*1.11*	*0.83*
		Free	14.65	14.50	***14.44***	14.66	14.46	15.66	⋆**14.69**
			4.43	*4.68*	*4.41*	*4.36*	*4.66*	*0.96*	*0.78*
	Stationary	Fixed	16.23	16.63	16.92	15.67	17.66	17.93	⋆***13.94***
			4.97	*4.75*	*4.78*	*5.72*	*5.54*	*1.43*	*0.71*
		ZIP	16.36	17.18	16.02	15.41	17.39	16.68	⋆***14.79***
			5.80	*6.34*	*5.41*	*4.67*	*8.29*	*1.36*	*0.83*
		Free	16.45	15.51	16.80	16.37	16.18	16.85	⋆***14.69***
			5.60	*4.00*	*5.76*	*5.91*	*5.14*	*1.37*	*0.78*
CH	Mobile	Fixed	11.27	11.67	11.49	11.35	11.33	11.14	⋆***7.89***
			3.06	*3.85*	*4.41*	*5.64*	*5.92*	*1.42*	*0.89*
		ZIP	13.36	13.41	13.15	13.59	13.15	13.95	⋆***7.16***
			3.73	*4.03*	*3.49*	*4.07*	*3.91*	*1.13*	*0.79*
		Free	12.51	12.75	12.71	12.83	12.70	12.10	⋆***7.87***
			4.61	*4.73*	*4.93*	*5.19*	*4.61*	*1.47*	*0.88*
	Stationary	Fixed	16.32	16.86	18.34	16.96	17.75	14.65	⋆***7.89***
			4.98	*7.45*	*9.87*	*5.41*	*7.21*	*1.46*	*0.89*
		ZIP	12.80	13.84	13.02	13.25	12.35	13.24	⋆***7.16***
			3.11	*4.50*	*4.04*	*4.85*	*3.22*	*1.29*	*0.79*
		Free	12.59	12.93	13.12	12.18	13.57	12.77	⋆***7.87***
			4.25	*3.94*	*3.92*	*3.97*	*5.47*	*1.39*	*0.88*

Italic numbers are standard deviations, **bold** numbers indicate the better of the global and single market values, ***bold italic*** identifies the best value on each line, and ⋆ denotes that where these comparisons are significant at the 95% level. The charges on profit rise linearly from M0 (10%) to M4 (50%). In the case of the LL markets, these are the figures from which charges start.

Note that in a single market it makes no sense for traders to move since there is no market to move to or from. As a result, figures for mobile and stationary traders are the same.

Table 5: Coefficient of convergence for markets with ZIP traders in single-market and multiple-market scenarios.

			multiple markets						single market
			M0	M1	M2	M3	M4	global	
CDA	Mobile	Fixed	8.58	8.67	9.10	8.95	8.93	9.67	⋆***4.54***
			4.26	*4.63*	*4.89*	*5.18*	*5.68*	*2.26*	*2.01*
		ZIP	8.04	8.02	8.11	7.92	7.92	9.32	⋆***4.00***
			4.14	*4.36*	*4.52*	*4.60*	*4.43*	*2.17*	*2.02*
		Free	8.41	8.40	8.28	8.21	8.17	10.16	⋆***3.96***
			3.99	*3.83*	*3.83*	*3.88*	*3.89*	*1.97*	*2.24*
	Stationary	Fixed	7.09	6.69	6.86	6.84	7.87	14.47	⋆***4.54***
			4.21	*3.79*	*3.95*	*3.80*	*6.49*	*3.39*	*2.01*
		ZIP	6.14	6.07	6.99	6.24	6.18	13.58	⋆***4.00***
			3.58	*3.45*	*6.00*	*4.20*	*3.57*	*2.92*	*2.02*
		Free	7.60	7.06	7.97	7.89	7.20	16.53	⋆***3.96***
			4.22	*4.09*	*5.24*	*4.28*	*4.03*	*4.59*	*2.24*
CH	Mobile	Fixed	5.09	5.28	5.31	5.52	5.56	6.22	⋆ ***2.87***
			3.12	*3.37*	*3.69*	*4.05*	*4.38*	*1.62*	*1.65*
		ZIP	5.02	5.05	5.19	5.33	5.39	6.22	⋆***2.78***
			3.69	*3.71*	*3.78*	*3.91*	*4.09*	*1.85*	*1.55*
		Free	5.50	5.47	5.54	5.49	5.58	6.46	⋆***2.81***
			3.64	*3.61*	*3.64*	*3.61*	*3.64*	*1.71*	*1.57*
	Stationary	Fixed	3.43	3.94	3.42	3.85	3.91	15.45	⋆***2.87***
			2.32	*3.28*	*2.40*	*2.95*	*3.36*	*4.16*	*1.65*
		ZIP	2.74	3.09	3.11	2.80	***2.68***	9.66	⋆**2.78**
			1.87	*2.16*	*2.29*	*1.96*	*2.02*	*3.01*	*1.55*
		Free	3.77	4.23	4.06	4.05	3.71	13.44	⋆***2.81***
			2.56	*3.47*	*3.20*	*2.77*	*2.69*	*3.76*	*1.57*

Italic numbers are standard deviations, **bold** numbers indicate the better of the global and single market values, ***bold italic*** identifies the best value on each line, and ⋆ denotes that where these comparisons are significant at the 95% level. The charges on profit rise linearly from M0 (10%) to M4 (50%). In the case of the LL markets, these are the figures from which charges start.

Note that in a single market it makes no sense for traders to move since there is no market to move to or from. As a result, figures for mobile and stationary traders are the same.

Table 6: Coefficient of convergence for markets with RE traders in single-market and multiple-market scenarios.

			multiple markets						single market
			M0	M1	M2	M3	M4	global	
CDA	Mobile	Fixed	15.23	15.68	14.74	12.78	***12.67***	$\star$**14.59**	17.04
			2.67	*5.39*	*8.88*	*9.56*	*9.62*	*1.29*	*1.44*
		ZIP	13.92	13.63	14.14	13.59	13.82	14.27	$\star$***13.05***
			4.35	*4.96*	*5.14*	*4.96*	*4.97*	*1.30*	*1.39*
		Free	15.19	15.70	15.45	15.42	15.46	$\star$***14.93***	15.96
			5.05	*4.96*	*5.78*	*5.36*	*5.76*	*1.30*	*1.29*
	Stationary	Fixed	14.96	14.72	16.71	14.44	***14.37***	$\star$**14.77**	17.04
			5.46	*5.80*	*7.10*	*5.35*	*5.31*	*1.21*	*1.44*
		ZIP	15.37	16.28	14.30	15.90	14.38	16.85	$\star$***13.05***
			7.78	*7.16*	*6.16*	*6.40*	*5.09*	*1.90*	*1.39*
		Free	16.38	***15.38***	15.45	17.27	17.24	$\star$**15.85**	15.96
			4.86	*5.36*	*5.10*	*7.90*	*6.83*	*1.60*	*1.29*
CH	Mobile	Fixed	6.53	6.97	7.74	7.54	7.32	5.21	$\star$***2.34***
			3.83	*5.05*	*6.66*	*6.57*	*6.14*	*1.54*	*0.76*
		ZIP	6.56	6.95	7.11	6.97	6.65	4.92	$\star$***2.50***
			4.72	*5.17*	*5.06*	*5.48*	*5.44*	*1.25*	*0.85*
		Free	7.14	7.51	7.47	7.50	7.29	6.49	$\star$***2.29***
			4.58	*4.70*	*4.73*	*4.60*	*4.65*	*1.46*	*0.64*
	Stationary	Fixed	7.94	7.81	6.61	8.34	6.19	9.28	$\star$***2.34***
			4.97	*4.31*	*4.02*	*5.73*	*4.73*	*2.90*	*0.76*
		ZIP	10.76	9.59	9.34	10.36	9.23	12.35	$\star$***2.50***
			7.23	*6.78*	*6.28*	*7.18*	*5.19*	*3.33*	*0.85*
		Free	6.91	7.48	8.69	8.96	8.11	9.70	$\star$***2.29***
			4.75	*4.25*	*5.52*	*6.33*	*4.69*	*2.43*	*0.64*

Italic numbers are standard deviations, **bold** numbers indicate the better of the global and single market values, ***bold italic*** identifies the largest value on each line, and $\star$ denotes that where these comparisons are significant at the 95% level. The charges on profit rise linearly from M0 (10%) to M4 (50%). In the case of the LL markets, these are the figures from which charges start.

Note that in a single market it makes no sense for traders to move since there is no market to move to or from. As a result, figures for mobile and stationary traders are the same.

	Profit Fees: 5%, 10%, 15%, 20%, 25%		Profit Fees: 10%, 20%, 30%, 40%, 50%		Profit Fees: 15%, 30%, 45%, 60%, 75%	
Day	Number of Traders	Efficiency	Number of Traders	Efficiency	Number of Traders	Efficiency
1 - 50	20.93	93.52	20.62	93.33	22.39	95.01
51 - 100	21.93	94.20	22.43	93.92	26.20	95.31
101 - 150	22.86	94.32	23.79	94.19	28.83	95.40
151 - 200	23.03	94.66	24.89	94.45	30.86	95.69
201 - 250	23.74	94.82	25.77	94.63	32.40	95.68
251 - 300	24.08	94.86	26.55	94.66	33.55	95.68
301 - 351	24.27	94.93	27.10	94.57	34.52	95.91
351 - 400	24.22	94.95	27.69	94.78	35.27	95.92
correlation:	0.97		0.96		0.98	

(a) CDA

	Profit Fees: 5%, 10%, 15%, 20%, 25%		Profit Fees: 10%, 20%, 30%, 40%, 50%		Profit Fees: 15%, 30%, 45%, 60%, 75%	
Day	Number of Traders	Efficiency	Number of Traders	Efficiency	Number of Traders	Efficiency
1 - 50	20.86	95.73	22.03	95.44	22.82	95.89
51 - 100	22.26	96.39	25.48	96.13	27.06	96.43
101 - 150	22.98	96.63	27.80	96.49	29.96	96.77
151 - 200	23.56	96.80	29.59	96.64	31.99	96.80
201 - 250	24.38	96.82	30.72	96.72	33.49	96.92
251 - 300	24.70	96.92	31.66	96.77	34.66	96.96
301 - 351	24.90	96.94	32.48	96.82	35.50	97.09
351 - 400	25.55	96.98	33.25	96.83	36.03	97.11
correlation:	0.95		0.97		0.98	

(b) CH

Table 7: The correlation between global efficiency and the number of traders in $M0$ for ZIP traders in markets with fixed charges.

	Profit Fees: 5%, 10%, 15%, 20%, 25%		Profit Fees: 10%, 20%, 30%, 40%, 50%		Profit Fees: 15%, 30%, 45%, 60%, 75%	
day	Number of Traders	Efficiency	Number of Traders	Efficiency	Number of Traders	Efficiency
1 - 50	20.93	96.46	20.62	96.79	22.39	97.05
51 - 100	21.93	96.26	22.43	96.86	26.20	97.41
101 - 150	22.86	96.27	23.79	96.96	28.83	97.53
151 - 200	23.03	96.43	24.89	97.28	30.86	97.62
201 - 250	23.74	96.47	25.77	97.28	32.40	97.62
251 - 300	24.08	96.47	26.55	97.33	33.55	97.70
301 - 351	24.27	96.61	27.10	97.33	34.52	97.83
351 - 400	24.22	96.47	27.69	97.44	35.27	97.90
correlation	0.51		0.96		0.98	

(a) CDA

	Profit Fees: 5%, 10%, 15%, 20%, 25%		Profit Fees: 10%, 20%, 30%, 40%, 50%		Profit Fees: 15%, 30%, 45%, 60%, 75%	
day	Number of Traders	Efficiency	Number of Traders	Efficiency	Number of Traders	Efficiency
1 - 50	20.86	97.89	22.03	98.18	22.82	98.56
51 - 100	22.26	98.02	25.48	98.54	27.06	98.86
101 - 150	22.98	97.95	27.80	98.73	29.96	99.05
151 - 200	23.56	97.57	29.59	98.82	31.99	99.13
201 - 250	24.38	97.76	30.72	98.85	33.49	99.16
251 - 300	24.70	97.26	31.66	98.88	34.66	99.25
301 - 351	24.90	97.13	32.48	98.98	35.50	99.25
351 - 400	25.55	97.33	33.25	98.97	36.03	99.27
correlation	-0.80		0.99		0.99	

(b) CH

Table 8: The correlation between the efficiency of M0 and the number of traders in M0 for ZIP traders in markets with fixed charges.

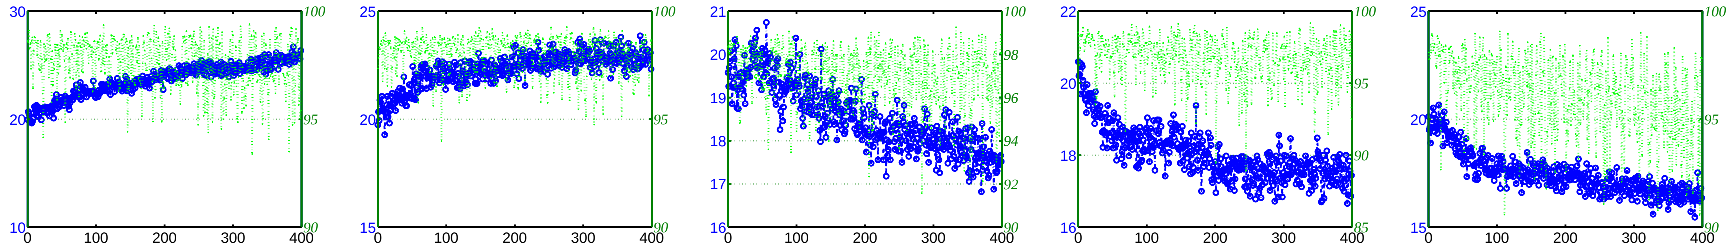

(a) Markets charging, left to right, 5%, 10%, 15%, 20% and 25% profit fees

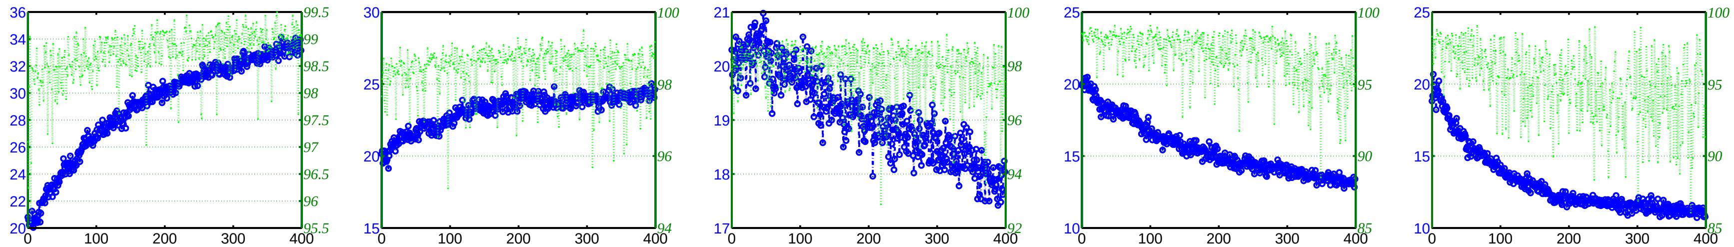

(b) Markets charging, left to right, 10%, 20%, 30%, 40% and 50% profit fees

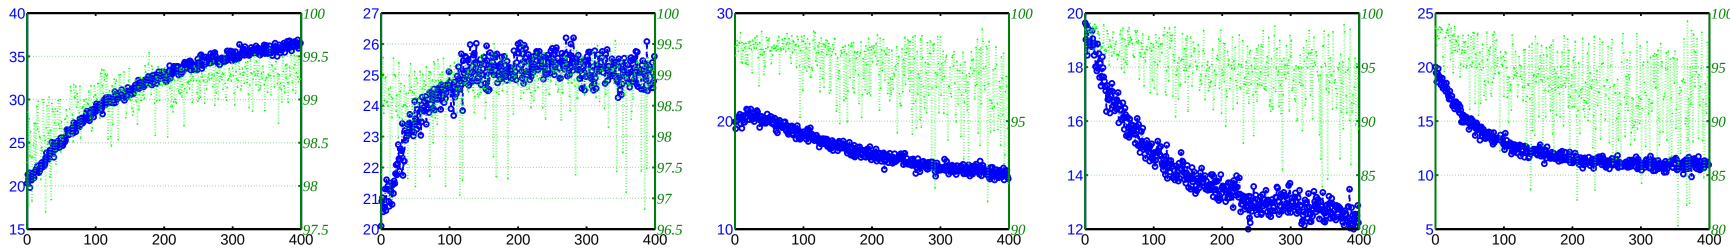

(c) Markets charging, left to right, 15%, 30%, 45%, 60% and 75% profit fees

Figure 5: Number of traders and efficiency for ZIP traders in CH markets with fixed charges on profit

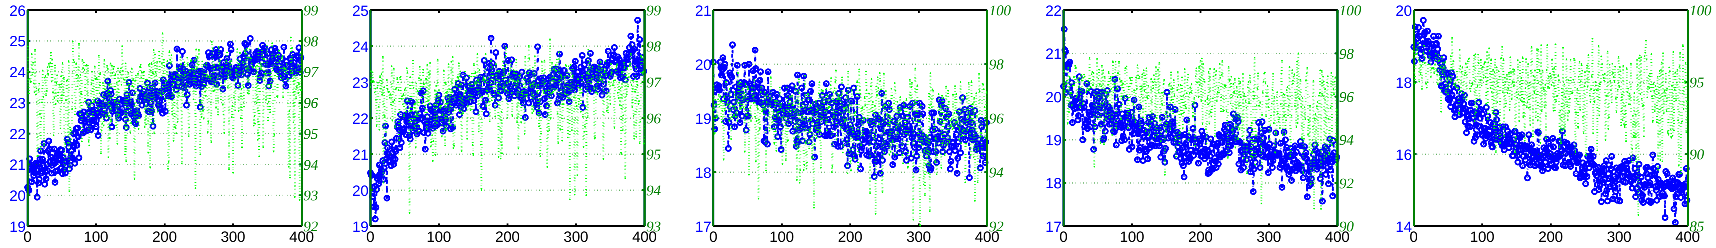

(a) Markets charging, left to right, 5%, 10%, 15%, 20% and 25% profit fees

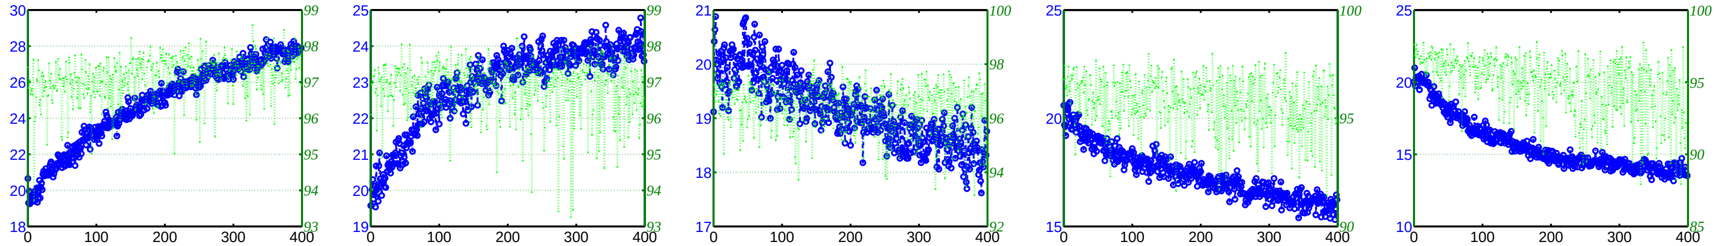

(b) Markets charging, left to right, 10%, 20%, 30%, 40% and 50% profit fees

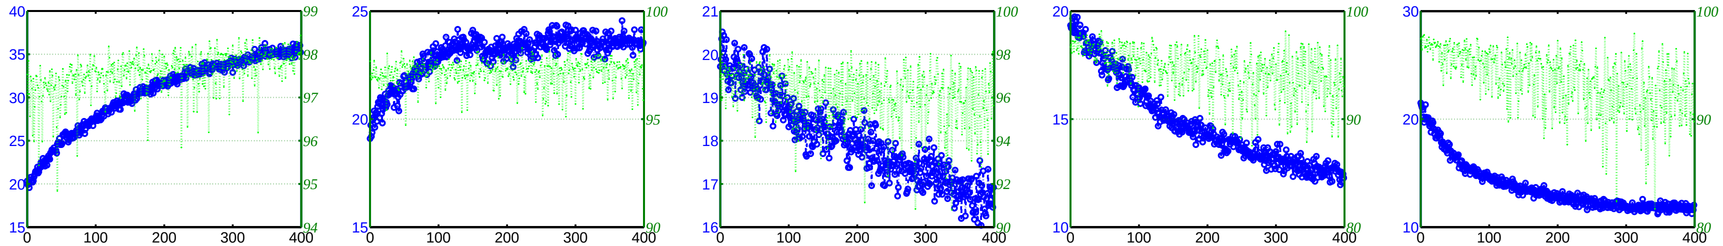

(c) Markets charging, left to right, 15%, 30%, 45%, 60% and 75% profit fees

Figure 6: Number of traders and efficiency for ZIP traders in CDA markets with fixed charges on profit

tables provide results for the three different charging schemes. The results in these tables are provided in terms of the average not only across the 100 repetitions, but across sequences of trading days — the 400 days are broken into 50-day sequences to allow us to see how values change over time. Figure 5 and Figures 6 provide another view of the same data. These figures plot the average daily values of efficiency (dark color) and the number of traders in market M0 (light color).

In all cases, we see that there is indeed a strong relationship between the efficiency of both M0 and the global market and the number of traders in M0 — in most cases the correlations between efficiency and trader numbers are very strong. However, what we are observing is not simply all the traders heading over to M0 since, as Tables 7 and 8 plainly show, even at the highest fee levels, most of the traders are *not* in M0 at the end of the experiment. In addition, note that for the lowest charging combination, there is a much reduced correlation between the efficiency of M0 and the number of traders in the market. Indeed, for the CH, there is actually a negative correlation. Despite this, the efficiency of the global market is positively correlated with the number of traders, so what is happening here is that more traders are being attracted to M0 causing the efficiency to *decrease*, but the overall efficiency of the global market to increase. Again, this suggests that what is happening here is more complex than just all the traders heading to M0 and turning it into an approximation to the global market.

5 Discussion

An explanation for the effects that we see is provided by Figure 9. This compares one typical set of supply and demand curves for the final trading day of five parallel CDA markets, all of which charge. The difference between the two sets is that in one the traders are allowed to move, while in the other they are stationary. Whereas in the markets with stationary traders the numbers of intra-marginal traders (to the left of the intersection between supply and demand curves) and extra-marginal traders (to the right of the intersection) are fairly well balanced, as one would expect of a random allocation of private values, this is not the case in the markets with the mobile traders. In these latter markets the traders have sorted themselves so that market M0 has no extra-marginal buyers, market M2 has no extra-marginal traders at all, M4 has no intra-marginal traders, and M3 has virtually no intra-marginal traders. Since, as [42] points out, the reason that CDA markets lose efficiency is because of extra-marginal traders "stealing" transactions from intra-marginal traders (who for a given transaction will, by definition, generate a larger profit), the segregation that we observe will lead to increased efficiency. In addition, as we observed in [21], charges have the effect of prodding traders that aren't making profits — and so are not adding to the efficiency of a given market — to try different markets, allowing markets to rid themselves of unproductive traders.

In CH markets, of course, extra-marginal traders cannot "steal" trades away from intra-marginal traders (at least not if they make rational offers). However, the movement of traders can still increase profits by allowing a trader that is extra-marginal in one market to become intra-marginal in another. Again, this behavior is encouraged by the combination of the market selection strategy and the charges imposed by the markets.

Finally, we should note that the efficiencies of the individual markets and the global market are rather low, and the coefficients of convergece are rather high, compared with those often reported for the trading strategies we use (in contrast the single market values are much the same as one would expect). We attribute this, at least in part, to churn. When a trader moves from one market to another, any learning it underwent in the old market about how to make offers is no use any more, and may even be detrimental. Similarly, the influx of new traders into a market can invalidate the learning previously undertaken by traders that have not moved. This means that offers are made away from equilibrium, pushing the coefficient of convergence up, and in turn this means that extra-marginal traders can trade, pushing efficiency down.

6 Conclusions

The main conclusion of this paper is that while dividing traders into multiple markets leads to a loss of efficiency and an increase in the coefficient of convergence, these changes are reduced when traders are allowed to move between markets in search of greater profits. In addition, this movement is encouraged by the imposition of fees on the traders, meaning that markets that charge show smaller loses of efficiency and lack of convergence than markets that do not charge.

This result holds because the movement of traders between markets serves to segment those markets. Since the movement is profit-driven, traders migrate towards markets that allow them to make good trades, and overall this increases the total profits of the set of markets, increasing the global efficiency. This effect is sharpened by the application of fees since these tend to reduce profits and so further discourage agents from remaining in markets that are unprofitable for them. A similar effect reduces the coefficient of convergence when traders move. Since traders move away from unprofitable markets, it tends to be the extra-marginal traders that move, and these tend to be the traders that make offers further from the theoretical equilibrium. Removing the possibility that these traders make offers that are accepted reduces the overall distance of accepted offers from the theoretical equilibrium and reduces the coefficient of convergence.

Our current work extends the investigation reported here. We are examining: the robustness of our results against traders who use different algorithms to do market selection[12]; the effect of different levels of charging on the changes in efficiency that we observe; and the influence of network effects, such as restrictions on the mobility of traders, on the effects that we observe here.

Acknowledgments:

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

References

[1] M. Bowling and M. Veloso. Multiagent learning using a variable learning rate. *Artificial Intelligence*, 136:215–250, 2002.

[2] A. Cassar and N. Duffy. Contagion of financial crises under local and global markets. In F. Luna and A. Perrone, editors, *Agent-based Methods in Economics and Finance: Simulations in Swarm*. Kluwer, New York, NY, 2001.

[3] `http://www.marketbasedcontrol.com/`.

[4] D. Cliff. Minimal-intelligence agents for bargaining behaviours in market-based environments. Technical Report HPL-97-91, Hewlett-Packard Research Laboratories, 1997.

[5] D. Cliff and J. Bruten. Less than human: Simple adaptive trading agents for CDA markets. Technical Report HP-97-155, Hewlett-Packard Research Laboratories, 1997.

[6] V. Conitzer and T. Sandholm. An algorithm for automatically designing deterministic mechanisms without payments. In *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 128–135, New York, July 2004.

[7] D. Friedman. The double auction institution: A survey. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories and Evidence*, pages 3–25. Perseus Publishing, Cambridge, MA, 1993.

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

[9] D. K. Gode and S. Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1):119–137, 1993.

[10] `http://jcat.sourceforge.net/`.

[12] Especially since the n-armed bandit algorithm we use is not particularly effective when payoffs change, as they do here

[11] P. Klemperer. How (not) to run auctions: The European 3G telecom auctions. *European Economic Review*, 46:829–845, 2002.

[12] H. L. Lee, V. Padmanabhan, and S. Whang. Information distortion in a supply chain: The bullwhip effect. *Management Science*, 43(4):546–558, 1997.

[13] A. Madhavan and M. Cheng. Price discovery in auction markets: A look inside the black box. *The Review of Financial Studies*, 13(3):627–658, Autumn 2000.

[14] J. McMillan. *Reinventing the Bazaar: A Natural History of Markets*. W. W. Norton & Company, 2003.

[15] M. H. Miller, J. D. Hawke, B. Malkiel, and M. Scholes. Findings of the Committee of Inquiry Examining the Events Surrounding October 19, 1987. Technical report, Chicago Mercantile Exchange, Spring 1988.

[16] T. Moyaux and P. McBurney. Modeling a supply chain as a network of markets. In *Proceedings of the IEEE International Conference on Service Systems and Service Managament*, Troyes, France, 2006.

[17] T. Moyaux and P. McBurney. Reduction of the bullwhip effect in supply chains through speculation. In *Proceedings of the Symposium on Artificial Economics*, Aalborg, Denmark, 2006.

[18] J. Niu, K. Cai, E. Gerding, P. McBurney, T. Moyaux, S. Phelps, D. Shield, and S. Parsons. JCAT: A platform for the TAC Market Design Competition. In Padgham, Parkes, Müller, and Parsons, editors, *Proceedings of the 7th International Conference on Autonomous Agents and Multiagent Systems*, Estoril, Portugal, May 2008. Demo Paper.

[19] J. Niu, K. Cai, S. Parsons, and E. Sklar. 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*, Hakodate, Japan, 2006.

[20] J. Niu, K. Cai, S. Parsons, and E. Sklar. Some preliminary results on competition between markets for automated traders. In *Proceedings of the Workshop on Trading Agent Design and Analysis*, Vancouver, BC, 2007.

[21] J. Niu, P. McBurney, E. Gerding, and S. Parsons. Characterizing effective auction mechanisms: Insights from the 2007 TAC market design competition. In *Proceedings of the 7th International Conference on Autonomous Agents and Multi-Agent Systems*, Estoril, Portugal, 2008.

[22] S. Phelps, M. Marcinkiewicz, S. Parsons, and P. McBurney. 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*, Hakodate, Japan, 2006.

[23] S. Phelps, P. McBurney, S. Parsons, and E. Sklar. Co-evolutionary mechanism design: A preliminary report. In J. Padget, O. Shehory, D. Parkes, N. Sadeh, and W. E. Walsh, editors, *Agent Mediated Electronic Commerce IV: Designing Mechanisms and Systems*, number 2531 in Lecture Notes in Artificial Intelligence, pages 123–142. Springer Verlag, 2002.

[24] S. Phelps, S. Parsons, and P. McBurney. An evolutionary game-theoretic comparison of two double auction markets. In P. Faratin and J. A. Rodríguez-Aguilar, editors, *Agent Mediated Electronic Commerce VI: Theories for and Engineering of Distributed Mechanisms and Systems*, number 3435 in Lecture Notes in Artificial Intelligence, pages 101–114. Springer Verlag, 2004.

[25] C. R. Plott and V. L. Smith. An experimental comparison of two exchange institutions. *The Review of Economic Studies*, 45(1):133–153, February 1978.

[26] C. Preist and M. van Tol. Adaptive agents in a persistent shout double auction. In *Proceedings of the 1st International Conference on Information and computation economies*, pages 11–18. ACM Press, 1998.

[27] A. E. Roth and I. Erev. Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. *Games and Economic Behavior*, 8:164–212, 1995.

[28] J. Rust, J. H. Miller, and R. Palmer. Characterizing effective trading strategies. *Journal of Economic Dynamics and Control*, 18:61–96, 1994.

[29] A. Rustichini, M. A. Satterthwaite, and S. R. Williams. Convergence to efficiency in a simple market with incomplete information. *Econometrica*, 62(5):1041–1063, September 1994.

[30] A. Shah and S. Thomas. David and Goliath: Displacing a primary market. *Global Financial Markets*, 1(1):14–21, 18th June 2000.

[31] V. L. Smith. An experimental study of competitive market behaviour. *Journal of Political Economy*, 70(2):111–137, April 1962.

[32] V. L. Smith. Effect of market organisation on competitive equilibrium. *The Quarterly Journal of Economics*, 78(2):182–201, May 1964.

[33] V. L. Smith. Experimental auction markets and the Walrasian hypothesis. *The Journal of Political Economy*, 73(4):387–393, August 1965.

[34] V. L. Smith and A. W. Williams. An experimental comparison of alternative rules for competitive market exchange. In *Auctions, Bidding and Contracting: Uses and Theory*. New York University Press, New York, 1983.

[35] V. L. Smith, A. W. Williams, W. K. Bratton, and M. G. Vannoni. Competitive market institutions: Double auctions vs. sealed bid auctions. *The American Economic Review*, 72(1):58–77, March 1982.

[36] R. S. Sutton and A. G. Barto. *Reinforcement learning: An introduction*. MIT Press, Cambridge, MA, 1998.

[37] G. Tesauro and R. Das. High-performance bidding agents for the continuous double auction. In *Proceedings of the 3rd ACM Conference on Electronic Commerce*, 2001.

[38] L. Tesfatsion. Agent-based computational economics: Growing economies from the bottom up. *Artificial Life*, 8(1):55–82, 2002.

[39] V. Walia, A. Byde, and D. Cliff. Evolving market design in zero-intelligence trader markets. Technical Report HPL-2002-290, Hewlett-Packard Research Laboratories, Bristol, England, 2003.

[40] W. Walsh, R. Das, G. Tesauro, and J. O. Kephart. Analyzing complex strategic interactions in multi-agent systems. In *Proceedings of Workshop on Game-Theoretic and Decision-Theoretic Agents*, 2002.

[41] A. Wilhite. Bilateral trade and 'small-world' networks. *Computational Economics*, 18:49–64, 2001.

[42] W. Zhan and D. Friedman. Markups in double auction markets. Technical report, LEEPS, Department of Economics, University of Santa Cruz, 2005.

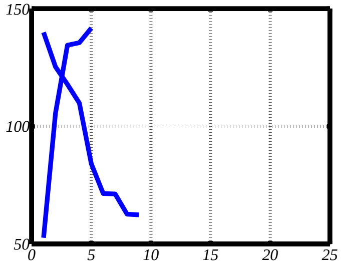

(a) M0

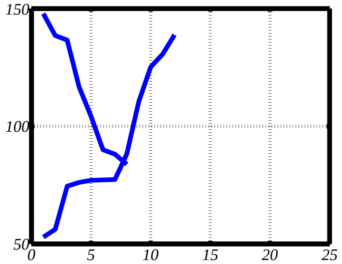

(b) M1

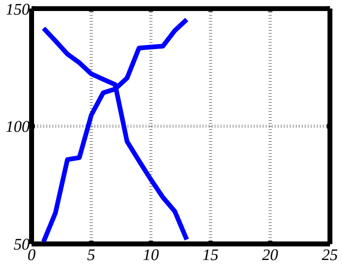

(c) M2

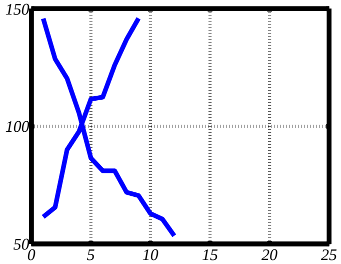

(d) M3

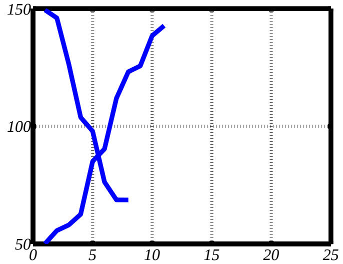

(e) M4

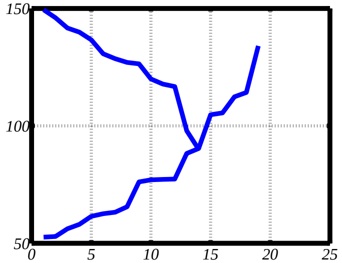

(f) M0

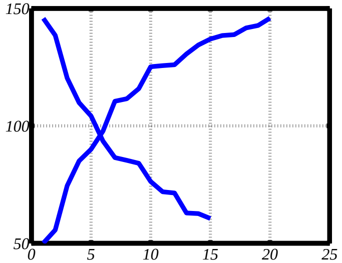

(g) M1

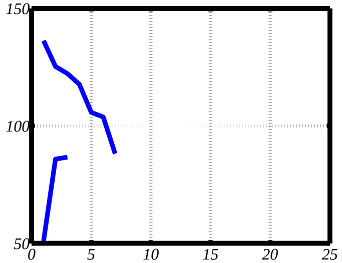

(h) M2

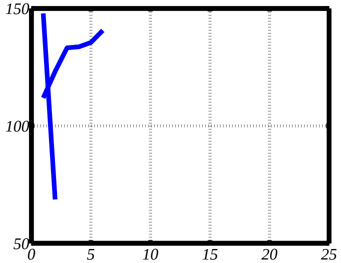

(i) M3

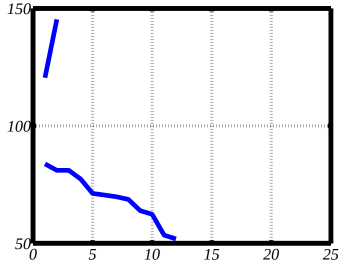

(j) M4

Table 9: Example final day supply and demand curves for the fixed charging CDA markets (a)–(e) with stationary traders and (f)–(j) with mobile traders.