On the Economic Effects of Competition between Double Auction Markets

Kai Cai[1], Jinzhong Niu[1], and Simon Parsons[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@sci.brooklyn.cuny.edu

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, examining 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 than grouping them into one large market, the movement of traders between markets, and price incentives for changing markets, can reduce this loss of efficiency.

1 Introduction

An *auction*, according to [4], 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 [8], 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 [9], and in structuring stock or futures exchanges [4].

There are many different kinds of auction. One of the most widely used auction 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 Mercantile Exchange. 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 of 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. Indian companies, for example, can be listed on both the National Stock Exchange (NSE) and the Bombay Stock Exchange (BSE) [18]. 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). The interactions between such exchanges can be complex, as when the NSE opened and proceeded to claim much of the trade volume from the established BSE [18], or when unfulfilled orders on the CME overflowed onto the NYSE during the global stock market crash of 1987 [10]. This kind of interaction between markets has not been widely studied, least of all using automated traders.

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 [19, 20]. The experiments in [19], 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 [6] 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 [6], 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.

This position was attacked by Cliff [2], who showed that if supply and demand are asymmetric, the average transaction prices of ZI-C traders can very significantly from the theoretical equilibrium. They 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 [2][page 60] suggests conditions under which ZIP will fail to attain equilibrium), prompting Cliff to suggest that ZIP traders embodied the minimal intelligence required. A range of other trading algorithms have been proposed — including those that took part in the Santa Fe double auction tournament [16], the reinforcement learning *Roth-Erev* approach (RE) [15] and the expected-profit maximizing *Gjerstad-Dickhaut* approach (GD) [5] — 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 the work of [23] (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.

As mentioned above, there has been little work on multiple market scenarios. We have presented some initial results on the dynamics of auctions that compete for traders [12] and the design of such auctions was the focus of the TAC Market Design competition [13]. This paper is a further contribution in the same direction, considering the impact of multiple markets on the efficiency of trading.

3 Experimental Setup

3.1 Software

To experiment with multiple markets, we used the Java-based server platform (JCAT) [7]. 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 [1]. 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.

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 [6];
– Cliff's zero intelligence plus (ZIP) strategy [2]; and
– Roth and Erev's reinforcement learning strategy (RE) [15].

The reason for picking the first of these is that given by [11, 22], 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 [14]. 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 [21]. 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 [21]. 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 [12] 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 profit made by traders on any transactions that they execute.

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 [13] 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:

– Fixed: a constant proportion, typically 10%, 20%, 30%, 40% and 50% of the surplus on a transaction, is taken as a fee.
– Zero intelligence (ZI): a version of the ZIP strategy for traders [3] adapted for markets and introduced by [12]. A ZI 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.
– Free: no profit fees are charged.

In all of our experiments the markets are populated by 100 traders, evenly split between buyers and sellers.

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 published by Niu *et al.* [12] 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 (fixed, ZI and free) 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. *Allocative efficiency*, E_a, is used to measure how good a market is at generating global profits. 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 the number of good bought and can be computed from the private values of the traders assuming that no trader makes a loss. E_a, is then:

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

E_a tells us how close a market is to theoretical equilibrium in terms of profits made. 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, introduced by Smith [19]. α actually measures the deviation of transaction prices from the equilibrium price:

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

For the multiple market experiments, we measure the efficiencies and convergence of each individual market, but also what we call the *global* values which assess the measurements across all the parallel markets. Global efficiency E_a^g is computed as:

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

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 is the equilibrium price of the global market. The global value of α is computed similarly.

4 Results

Figure 1, which summarizes the results of the experiment that places mobile ZIP traders in CH markets that adjust their profit charges using the ZI mechanism, show the typical

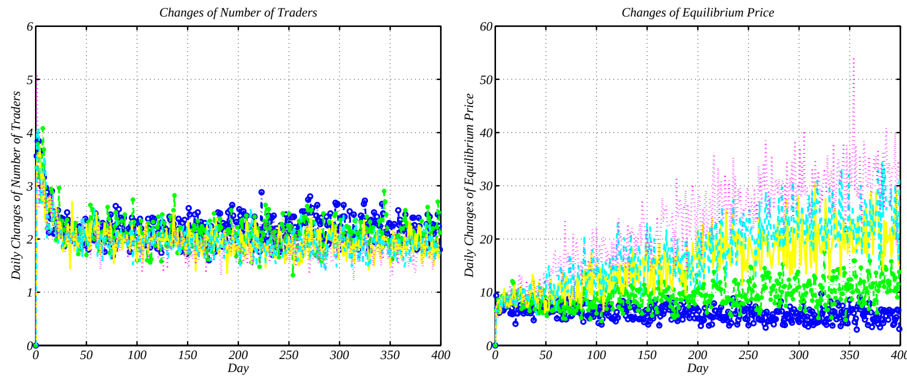

(a) Number of traders changing market each day

(b) Change in equilibrium price each day

Fig. 1: How individual markets change over time. Mobile ZIP traders in CH markets that use ZI to set charges on trader profits.

way that markets change over time. All the other experiments have very similar results, and the results parallel those we reported in [12].

Figure 1 (a) shows the number of traders leaving each of the five markets at the end of each day. 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. Over the first 50 days, the amount of "churn" falls steadily and eventually the movement between markets stabilizes and settles to a constant value. However, because the market selection strategy always keeps exploring, on average each market still has two traders leave each day. (On average, the same number of traders also enter).

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, in the multiple market case, this does not happen. Figure 1(b), which plots the daily *change* in equilibrium price in each market, 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 period. However, we do see a certain level of stability emerge — 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 ranges from around $10 in M0 to around $30 in M4

In case these results suggest that there is no overall pattern, consider Figure 2. This plots the global values of efficiency and the coefficient of convergence for the same experiment as in Figure 1. 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, the overall picture has settled down after around 200 days.

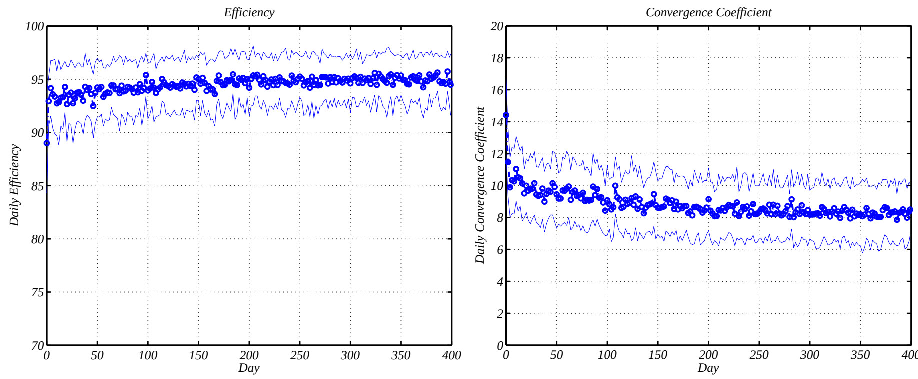

(a) Global efficiency (b) Global coefficient of convergence

Fig. 2: How individual markets change over time. Mobile ZIP traders in CH markets that use ZI to set charges on trader profits. The plots show average value and standard deviation.

Having sketched the overall behavior of the markets in our experiments, the main results of this paper are given in Tables 1–3. These give, for each of the experiments outlined above, the efficiencies of markets M0 to M4, the global efficiency, and the efficiency of a single market containing all the traders. This latter differs from the global measure in that the actual trader profits are obtained in the single market rather than in the individual markets (while the theoretical profit is the same in both cases). The values of the efficiency 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 [17], the efficiency of the single market of 100 traders is 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.

Some other interesting points emerge. First, looking just at the global 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

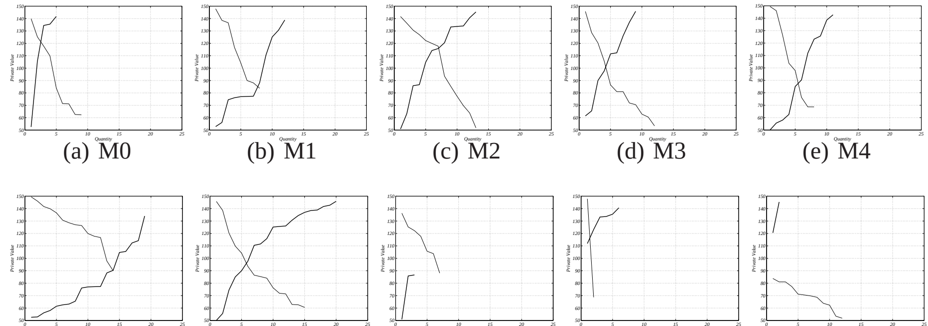

(a) $\overline{\text{M}}$0 (b) $\overline{\text{M}}$1 (c) $\overline{\text{M}}$2 (d) $\overline{\text{M}}$3 (e) $\overline{\text{M}}$4

(f) $\overline{\text{M}}$0 (g) $\overline{\text{M}}$1 (h) $\overline{\text{M}}$2 (i) $\overline{\text{M}}$3 (j) $\overline{\text{M}}$4

Fig. 3: Typical final day supply and demand curves for the fixed charging CDA markets (a)–(e) with stationary traders and (f)–(j) with mobile traders.

do not charge[3]. 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.

5 Discussion

An explanation for the effects that we see is provided by Figure 3. 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 the 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 [24] 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

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

Table 1: Market allocative efficiency for 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	⋆***88.86***
			11.96	*20.07*	*27.04*	*29.85*	*30.99*	*3.49*	*2.05*
		ZI	87.15	80.88	78.36	66.25	60.49	85.54	⋆***87.58***
			12.17	*20.65*	*22.89*	*30.09*	*32.26*	*3.31*	*2.35*
		Free	78.80	76.37	78.27	79.36	78.24	85.66	⋆***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	⋆***88.86***
			11.11	*9.19*	*9.19*	*8.25*	*8.40*	*5.80*	*2.05*
		ZI	82.38	84.70	81.65	80.51	81.51	77.18	⋆***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	⋆***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	⋆**81.99**
			20.01	*24.85*	*30.87*	*30.87*	*31.30*	*3.20*	*2.99*
		ZI	***87.41***	79.55	69.29	60.43	57.81	⋆**83.52**	81.30
			7.17	*16.02*	*25.13*	*29.86*	*30.89*	*2.94*	*2.63*
		Free	74.58	76.38	71.83	72.94	77.90	83.72	⋆***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	⋆**81.99**
			8.47	*8.85*	*7.63*	*8.72*	*8.72*	*5.11*	*2.99*
		ZI	79.78	81.08	***81.72***	78.62	77.69	76.09	⋆**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	⋆***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 ⋆ 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 ZI 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: Market allocative efficiency for 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*
		ZI	96.79	96.91	96.63	94.43	93.49	94.51	$\star$***98.55***
			7.15	*5.34*	*5.74*	*12.68*	*14.30*	*2.62*	*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*
		ZI	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*
		ZI	98.56	97.89	96.88	96.65	94.17	96.74	$\star$***99.68***
			4.71	*7.50*	*10.07*	*10.86*	*16.45*	*2.34*	*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*
		ZI	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 ZI 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: Market allocative efficiency for 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*
		ZI	***89.94***	89.20	79.69	70.43	66.90	86.55	⋆**87.39**
			2.93	*8.41*	*29.40*	*38.49*	*41.10*	*3.21*	*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*
		ZI	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*
		ZI	98.86	97.71	95.74	92.48	87.57	95.83	⋆***99.42***
			2.30	*6.76*	*12.84*	*20.95*	*28.84*	*3.28*	*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*
		ZI	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 ZI 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.

will, by definition, generate a larger profit), the segregation that we observe will lead to increased efficiency. In addition, as we observed in [13], 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 efficiencies are rather low 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 given the random allocation of private values to traders). We attribute this to churn. When a trader moves from one market to another, any learning it underwent in the old market 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.

6 Conclusions

The main conclusion of this paper is that while dividing traders into multiple markets leads to a loss of efficiency, this loss is reduced when traders are allowed to move between markets in search of greater profits, and this movement is encouraged by the imposition of fees on the traders. 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 further discourage agents from remaining in markets that are unprofitable for them.

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; 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 The authors are grateful for financial support received from the National Science Foundation under grant NSF IIS-0329037 *Tools and Techniques for Automated Mechanism Design*, and from the UK EPSRC under grant GR/T10657/01 *Market Based Control of Complex Computational Systems*.

References

1. `http://www.marketbasedcontrol.com/`.

2. D. Cliff. Minimal-intelligence agents for bargaining behaviours in market-based environments. Technical Report HPL-97-91, Hewlett-Packard Research Laboratories, 1997.
3. D. Cliff and J. Bruten. Less than human: Simple adaptive trading agents for CDA markets. Technical Report HP-97-155, Hewlett-Packard Research Laboratories, Bristol, England, 1997.
4. 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.
5. S. Gjerstad and J. Dickhaut. Price formation in double auctions. *Games and Economic Behavior*, 22:1–29, 1998.
6. 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.
7. `http://jcat.sourceforge.net/`.
8. P. Klemperer. How (not) to run auctions: The European 3G telecom auctions. *European Economic Review*, 46:829–845, 2002.
9. J. McMillan. *Reinventing the Bazaar: A Natural History of Markets*. W. W. Norton & Company, 2003.
10. 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.
11. 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.
12. 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.
13. 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.
14. 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.
15. 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.
16. J. Rust, J. H. Miller, and R. Palmer. Characterizing effective trading strategies. *Journal of Economic Dynamics and Control*, 18:61–96, 1994.
17. 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.
18. A. Shah and S. Thomas. David and Goliath: Displacing a primary market. *Global Financial Markets*, 1(1):14–21, 18th June 2000.
19. V. L. Smith. An experimental study of competitive market behaviour. *Journal of Political Economy*, 70(2):111–137, April 1962.
20. V. L. Smith. Experimental auction markets and the Walrasian hypothesis. *The Journal of Political Economy*, 73(4):387–393, August 1965.
21. R. S. Sutton and A. G. Barto. *Reinforcement learning: An introduction*. MIT Press, Cambridge, MA, 1998.

22. 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.
23. 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.
24. W. Zhan and D. Friedman. Markups in double auction markets. Technical report, LEEPS, Department of Economics, University of Santa Cruz, 2005.