

Network effects in double auction markets with automated traders

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

Simon Parsons
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

December 18, 2008

Abstract

Many electronic markets are linked together into larger “network markets” where the links reflect constraints on traders. These constraints mean that a choice to trade in one market limits the trader’s choice of other markets to use. This kind of network market is important because so many basic products, including gas, water, and electricity, are traded in such markets, and yet it has been little studied until now. This paper studies networks of double auction markets populated with automated traders, concentrating on the effects of different network topologies. We find that different topologies confer significant economic benefits to individual markets, suggesting that these network effects can be an important consideration in market design.

1 Introduction

An *auction*, according to [8], 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 [13], 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 [16], in structuring stock or futures exchanges [8], and, despite the current credit crisis, are the basis of a vast volume of trade in electronic markets.

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 Mercantile Exchange (CME). 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 are common 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 can be listed on both the National Stock Exchange (NSE) and the Bombay Stock Exchange (BSE). The interactions between such exchanges can be complex, 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 [29], or when unfulfilled orders on the CME overflowed onto the NYSE during the global stock market crash of 1987 [17]. This kind of interaction between markets has not been widely studied, least of all when the markets are populated by automated traders.

One multiple market scenario that is particularly interesting is that of *network markets*, markets in which individual markets are linked together into larger markets, where the links between markets reflect constraints on traders in the markets. Network markets are important because so many basic products, including gas [15], water, and electricity, are traded in such markets — the products proceed through a series of transactions at different locations from producer to final consumer, and the need to convey the product through a complex transportation network provides the constraints. Our specific focus in this paper is to examine the differences between network markets with different topologies.

2 Background

Double auctions have been extensively studied using agent-based methods. Gode and Sunder [10] were the first to use multi-agent simulations in this way, testing the hypothesis, suggested by [30], that the form of the market has more bearing on obtaining efficient allocation than the intelligence of traders in that market. [10] introduced a “zero-intelligence” trading strategy (denoted ZI-C) — which involves making offers at random under the constraint that they don’t lead to loss-making trades — and showed that agents using this strategy could generate high efficiency. Indeed, such agents come close enough to the performance of human traders that Gode and Sunder claimed that trader intelligence is not necessary.

This position was attacked by Cliff [6], who showed that if supply and demand are asymmetric, the average transaction prices of ZI-C traders can vary significantly from the theoretical equilibrium. 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, 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 [28], the reinforcement learning *Roth-Erev* approach (RE) [27] and the *Gjerstad-Dickhaut* approach (GD) [9] — and the performance of these algorithms evaluated under various market conditions. Despite the high performance of GD traders, research into automated trading mechanisms has continued. Recent examples of new approaches include those described in [11, 24, 33, 35].

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 [37] and [21] both of which show that the same set of trading strategies can have markedly different behaviors in different auction mechanisms.

As mentioned above, there has been little work within agent-based computational economics [34] on multiple connected markets, but what little has been carried out has studied a broad range of scenarios. [3] uses agent-based methods to examine the effects of linked markets on financial crises, while [38] looks at the effect of different trade routes on price convergence. [18, 19] study the bull-whip effect [14]¹ in supply chains. In addition, some initial results on multiple auctions that compete for traders were presented in [22] and the design of such auctions is the focus of the TAC Market Design competition analyzed in [23]. The work we report here further extends the use of agent-based computational economics to study groups of connected markets.

3 Experimental Setup

The aim of this work was to investigate the effect on market performance of different topological connections between markets. In the context of the double auction markets

¹Where small fluctuations in supply in one market can have an effect that magnifies through the network.

that we consider, these connections might reflect a number of different constraints. For example, they might reflect the physical layout of market makers on a trading floor, or they might reflect affiliations between electronic markets, or they might reflect the relationship between the time-zones in which different markets operate.

3.1 Software

To experiment with multiple markets, we used JCAT [12, 20], the platform that supports the TAC Market Design Competition [5]. JCAT provides the ability to run multiple double auction markets populated by traders that use a variety of trading strategies. 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. In addition, JCAT allows traders to move between markets at the end of a day, and over the course of many days they learn which market they perform best in.

In JCAT there are no restrictions on the movement of traders. To study network effects, we extended JCAT to restrict the movement of traders. In particular, our extension allows us to specify which markets a given market is connected to. At the end of every day that a trader spends in that market, the trader has a choice of remaining in that market or moving to any of the markets to which there are connections. The decision mechanism employed by the traders to make this choice is discussed below. In our experiments, market connections have four topologies:

- Fully connected. Each market is connected to every other market.
- Ring. Each market is connected to exactly two other markets. This is what [38] calls a “local connected network”.
- Chain structure. All but two of the markets are connected to two other markets as in the ring. The remaining pair form the ends of the chain and are connected to exactly one market.
- Star structure. One market is connected to every other market. There are no other connections between markets. This is the network topology studied in [26].

These topologies are illustrated in Figure 1.

3.2 Traders

In JCAT markets, traders have two tasks. One is to decide how to make offers. The mechanism they use to do this is their *trading strategy*. The other task is to choose the market to make offers in. The mechanism for doing this is their *market selection*




Figure 1: The different topologies we consider. Each node is a market, each arc a connection between markets. (a) fully connected, (b) ring, (c) chain, (d) star.

strategy. We studied markets in which all the traders used the same trading strategy, and considered two such strategies:

- Gode and Sunder’s zero intelligence strategy ZI-C [10]; and
- Cliff’s zero intelligence plus (ZIP) strategy [6].

The reason for picking the first of these is that given by [21, 36], 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 is that it is typical of the behavior of automated traders, rapidly converging to equilibrium in a single market.

The market selection strategy is a standard model taken from JCAT. Traders treat the choice of market as an n -armed bandit problem that they solve using an ϵ -greedy exploration policy [32]. Using this approach the behavior of the agents is controlled by the parameters ϵ . 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 available markets otherwise. We choose ϵ to take a constant value of 0.1. [22] suggests that market selection behavior is rather insensitive to the parameters we choose here, and we choose ϵ to remain constant so that any convergence of traders to markets is due to traders picking markets that work for them rather than being forced by a reduction in their tendency to explore.

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. Private values are set, just as in [6] to form perfect “staircase” supply and demand curves, with every buyer having a unique private value from the set $\{\$50, \$54, \$58 \dots, \$246, \$250\}$. Sellers are allocated values in the same way. A given trader has the same private value for all goods that it trades throughout the entire experiment. All of our experiments used 100 traders, divided into 50 buyers and 50 sellers. Initially they are equally distributed between the markets, and subsequently use their market selection strategy to pick the market to operate in.

3.3 Markets

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

- Registration fees, charges made by the market for entering the market. We set this to a low constant value (\$0.5) for every market following [23] which suggests that such a fee is effective in motivating extra-marginal traders to move between markets thus preventing their inertia from distorting results.
- Profit fees, charges made by the market on the bid/ask spread of any transactions they execute. The name arose the bid/ask spread is the transaction surplus, and with the $k = 0.5$ rule that is usually used in JCAT for allocating the surplus, the amount charged by this fee is thus directly related to the profit realized by both agents.

Unlike previous work that used JCAT to investigate multiple market scenarios [22], we used a simple, non-adaptive scheme for the profit fees. We used two variations:

- A 5% profit charge on all markets.
- A fixed 5% charge on some markets, a fixed 10%, 15%, 20% or 25% charge on the other markets.

All of our experiments we run five markets connected as described above, and we used both CDA and CH markets, both of which are provided in JCAT.

3.4 Experiments

We ran experiments that tested all the different combinations discussed above. That is we ran experiments for CH and CDA markets using each of the four different topologies, both the pricing schemes described above, and carried out each of these sets of experiments for traders that all used the ZI-C strategy and traders that all the ZIP strategy. Each of these experiments was run for 600 trading days, with each day being split into 50 0.5-second-long rounds. We repeated each experiment 50 times and the results that we give are averages across those 50 runs.

4 Results

There are several different sets of results that we consider.

4.1 Allocative efficiency

The first results to consider are those in Table 1 which measures the allocative efficiency of sets of markets of different topologies. In particular what they measure is what [1] calls the “global efficiency”, the ratio of the sum of profit made in all of the markets to the equilibrium profit that would be made in a hypothetical market that contained all the traders. Given the distribution of traders across different markets, it is clear that it is harder for trading strategies to achieve 100% efficiency in multiple markets than in one single market that contains the same traders, and this explains the efficiency values that we see. The global efficiency values in Table 1, which show no significant differences, tell us that the various topologies have no effect on the efficiency of the markets.

			Chain	Ring	Star	F.C.
ZIC	CDA	Efficiency	95.49	95.42	95.75	95.38
		Stdev.	0.30	0.25	0.2204	0.16
	CH	Efficiency	96.61	96.51	96.81	96.56
		Stdev.	0.25	0.19	0.15	0.13
ZIP	CDA	Efficiency	95.50	95.33	95.68	95.05
		Stdev.	0.24	0.19	0.22	0.17
	CH	Efficiency	96.86	96.77	96.96	96.54
		Stdev.	0.24	0.17	0.19	0.15

Table 1: The global efficiencies of sets of market with different connection topologies from left to right, chain, ring, star and fully connected networks. The table gives results for markets using both ZI-C and ZIP traders, and for both CDA and CH markets.




Figure 2: The total number of traders that move at the end of a given trading day, ZIP traders, fully connected CDA market. The x-axis gives the trading day, the y-axis gives the number of traders.

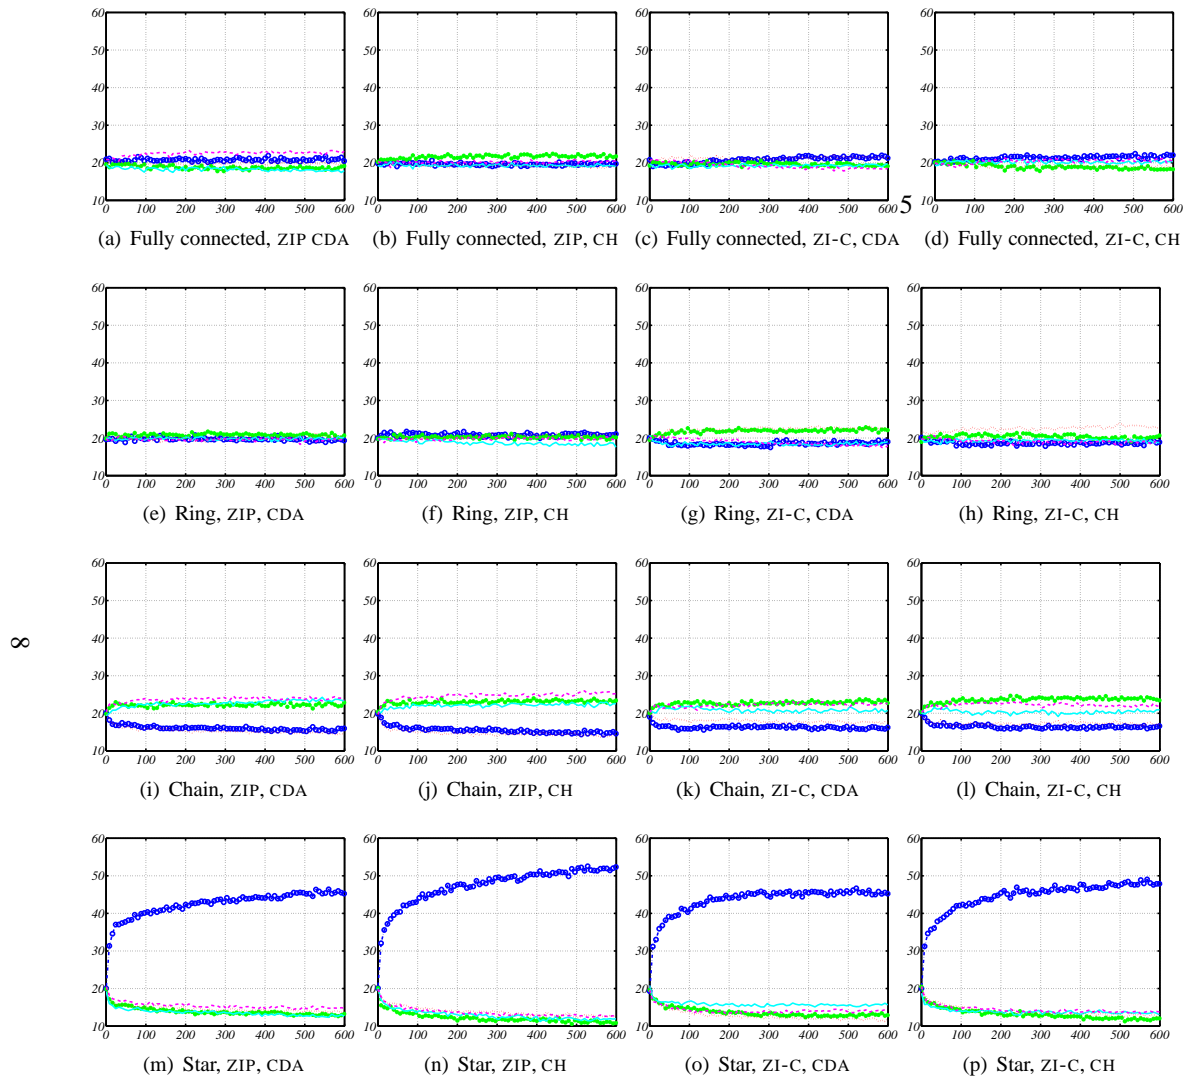


Figure 3: The number of traders in multiple connected CDA and CH markets with different connection topologies on each trading day. The x axis gives the trading days, the y axis the number of traders in each of the five markets. The left-hand graphs are results for CDA markets and right-hand graphs are results for CH markets. All markets make the same charges. In the chain markets, the lines marked with blue open circles and red dots give the numbers for the markets at the end of the chain, and for the star markets, the line marked with blue open circles gives the numbers for the market at the center.

				M0	M1	M2	M3	M4
Star	CDA	ZIC	Number of traders	43.67	13.65	15.82	14.14	12.72
			Stdev.	11.89	7.88	8.25	8.38	7.23
	CH	ZIP	Number of traders	42.50	13.90	13.57	15.42	14.61
			Stdev.	9.08	5.13	5.19	5.57	4.76
		ZIC	Number of traders	44.71	13.16	13.83	14.40	13.89
			Stdev.	5.70	2.68	3.01	3.03	3.90
	ZIP	Number of traders	47.41	12.14	12.92	13.60	13.93	
		Stdev.	8.44	3.32	3.07	4.40	4.58	
Chain	CDA	ZIC	Number of traders	16.24	22.74	20.88	22.21	17.93
			Stdev.	6.63	8.86	9.86	9.67	7.45
	CH	ZIP	Number of traders	16.10	22.22	22.66	23.72	15.31
			Stdev.	4.91	5.52	7.08	5.89	4.45
		ZIC	Number of traders	16.45	23.66	20.33	22.28	17.28
			Stdev.	4.82	6.67	5.78	6.03	4.67
	ZIP	Number of traders	15.50	23.02	22.10	24.76	14.63	
		Stdev.	4.80	6.32	7.01	6.31	4.66	

Table 2: The number of traders in each market for star and chain configurations. In the star market, M0 is the hub, the market at the center. In the chain markets, markets M0 and M4 are the markets at the end of the chain. All markets make the same charges.

4.2 Trader distribution

When we look at the distribution of traders across markets, however, some significant differences do show up. Figure 3 shows the number of traders in each market each trading day for the experiments involving both ZIP and ZI-C traders where each market makes the same registration charge and the same 5% charge on profits. These results show that the distribution of traders across markets is in fully-connected markets, Figures 3(a)–3(d), show no significant differences. Since markets are identical under our experimental conditions, this is exactly what we would expect. Markets organized as a ring, again identical, similarly have an equal distribution of traders (Figures 3(e)–3(h)).

Chain markets don't have the symmetry of fully-connected and ring-structured markets, and this shows up in the distribution of traders. As Figures 3(i)–3(l) show, markets at the end of the chain end up with fewer traders than the fully connected or ring markets. The effect of the loss of symmetry is even more marked in star markets, Here, as shown in Figures 3(m)–3(n), the hub market in the star collects many more traders than the otherwise identical markets that are connected to it.

The results are similar for CDA and CH markets. The only difference is that the extent of the variation between the chain ends and middle and between the star center and edges are slightly more marked in the CH markets. As mentioned above, the fact that the structural differences hold for markets with ZI-C traders suggests that these results are somewhat independent of the strategy that traders use to select offers.

The graphs of Figure 3 don't make it easy to decide what differences are significant




Figure 4: The number of traders in multiple CDA and CH markets with chain and star topologies where different levels of charging are applied to different markets. Results for CDA markets are given in (a)–(h), results for CH are given in (i)–(p). The results for chain markets are given in (a)–(d) and (i)–(l). Results for star markets are given in (e)–(h) and (m)–(p). In the chain markets, the line marked by blue open circles and the line marked by red dots give the numbers for the markets at the end of the chain, where charges are held fixed to 5% and the charges for the remaining markets are as given on the figures. In the star markets, the line marketed by blue open circles gives the numbers for the market at the center, the charge on the center market is as given on the figures while the charges made by the other markets are 5%.

so we show the actual trader numbers after the 600th trading day (that is at the end of the experiment) in Table 2. In the chain markets, the markets at the ends of the chain are M0 and M4. T-tests reveal that the numbers of traders in these markets are significantly different from the numbers of traders in markets M1, M2 and M3 at the 95% level for both CDA and CH whether or not the traders are ZI-C or ZIP. In the star markets, the market at the hub of the star is M0. T-tests show that the number of traders in this market is significantly different from that in all other markets at the 95% level for CDA and CH markets whether or not the traders are ZI-C or ZIP.

4.3 Speed of convergence

One thing that does show up in the graphs of Figure 3 is the fact that trader numbers seem to converge exponentially, suggesting that the markets reach some form of equilibrium. This convergence is illustrated by Figure 2 for ZIP traders and fully-connected markets. Results for other experiments are similar. An interesting question, then, is whether the network structure has any effect on the speed of convergence. Since, as described above, the market selection mechanism we are using will mean that we always have some number of traders still moving at the end of each trading day, we can't determine equilibrium by looking for the point at which all traders stop moving. Instead we need to find a way to estimate the speed of convergence.

Smith and Williams [31] solve a similar problem, that of deciding the speed of approach of offers to equilibrium, using regression. However they are dealing with the results of a human-subject experiment, and so only have a single piece of data at each time point. In contrast we have 50 pieces of data at each time point, and as Figure 2 shows, although there is a clear trend in the data, there is considerable variance from day to day. As a result, instead of using the approach in [31], we did the following. Starting from the 50th day, we calculated the 10-day moving average of the number of moving traders. The first value in this sequence is the average of the number of moving traders from day 41 through day 50, the next value is this average for day 42 through day 51, and so forth, until the final value, which is the average for day 591 through day 600. The repetitions of each experiment mean that we have 50 samples of this average at each day. We now consider each set of samples as a population and perform a t-test between each population and the population for day 600. We do this in order, starting with the population that ends with day 50, then the population that ends with day 51 and so on. We can then identify earliest day for which the test fails to reject the null hypothesis at the 0.05 level, meaning we can't say with 95% confidence that the two sets of samples are different, and we take this day as a measure of how quickly the markets converge. Table 3 shows the results of this test for CDA markets.

We see that for both ZIP and ZI-C traders, the disparity in the chain markets shows up early, with convergence happening faster than in the ring or star topologies. In other words, the end markets quickly get starved of traders. The star market converges more slowly than the chain. This is perhaps because the disparity ends up being much greater — it has further to go before it converges. In any case, it converges faster than the ring market, and for both ZIP and ZI-C traders, it converges before 400 days. Looking at the differences between markets using ZIP and ZI-C traders, we see that in general the ZIP markets converge faster. Given that the randomness of the offers made by ZI-C

ZIC	Fully connected	347
	Ring	460
	Chain	362
	Star	389
ZIP	Fully connected	397
	Ring	374
	Chain	205
	Star	317

Table 3: Our measure of convergence for markets with ZI-C and ZIP traders and different market connection topologies — the earliest day on which the distribution of traders is not significantly different from the distribution on the final trading day. We only give the results for CDA markets.

traders will tend to add noise to the learning of the market selection strategy, this is perhaps not surprising. However we are currently at a loss to explain why ZIP traders in fully-connected markets take so long to converge — with ZI-C traders, the fully-connected market converges fastest of all, rather as one would expect since, with every market accessible from every other market, exploration of the set of markets should be quickest.

4.4 Charges versus structure

The results so far show that connecting identical markets in different ways can lead to significantly different behaviors. In particular connecting markets using star and chain networks mean that certain markets become privileged in terms of the number of traders they attract. With all markets making similar charges, this will lead to the privileged markets making greater profits. Since the results from [22] show that higher charging markets will, over time, attract less traders (to the extent that eventually they make less profit than lower charging markets) a natural question is to what extent the privileged markets in the star and chain networks can exploit their situation with higher charges.

To answer this question, we looked at the results for our experiments with different (fixed) charges on profit. In particular, we looked at chain markets where the end markets in the chain charged 5%, and the middle three markets charged progressively higher percentages of the profit. We also looked at star markets where the hub made progressively higher charges on profit while the other markets charged a steady 5%. These results are shown in Figure 4 — we only have room here for the results of markets involving ZIP traders. As before, results for ZI-C traders are very similar. These show that, for both chain and star markets, the advantages inherent in the structure are eroded by as charges are increased, but remain strong even when the privileged markets set a 25% profit charge.

5 Related work

A number of authors have looked at the properties of network markets. [26] studies a three-node star network with a uniform-price double auction at each node. The same authors [25] report experiments using a 9-node gas network that, in addition to buyers and sellers, also includes pipeline owners, and in [15] study another small gas market. A further small network model, including just two markets, is the basis of the study in [4] into the effects of cheating (that is, either not paying for goods, or failing to deliver goods that have been paid for) and [7] investigates how a 6-node railway network responds to two different pricing mechanisms. While these markets are similar to those in our study, the investigation dealt with markets with human traders.

As mentioned above, agent-based methods have been used before in the context of network markets. [3] used such methods to examine the effects of linked markets on financial crises, while [18, 19] consider the behavior of supply chains.² This work all studies smaller sets of markets than we have considered.

The agent-based studies in [2] and [38] are larger but consider a set of connection topologies that overlap with, but does not contain, the set we consider. Both [2] and [38] deal with “local” networks, which are equivalent to our ring, as well as small-world networks, which we don’t consider. However, neither considers the chain or star topologies, which are the most interesting of the topologies we looked at, and neither study considers traders that move between markets.

6 Conclusions

This paper has examined the effect of different connection topologies on network markets in which the constituent markets are double auctions and the connections denote the allowed movements of traders between markets. Traders used either ZI-C or ZIP strategies, and markets were either CH or CDA. We looked at the behavior of four different topologies — fully connected, ring, chain and star — and considered the overall allocative efficiency, the distribution of traders and the speed of convergence. We found that topology had no significant effect on efficiency, but has a significant effect on the distribution of traders in star and chain markets, and hence on the liquidity of the components of those markets. Chain markets also seem to be quicker to converge than ring or star markets, while star markets are quicker to converge than ring markets (results for the convergence of fully connected markets are ambiguous). We also found that the topological effect on trader distribution could be offset by making suitable changes to charges imposed by markets.

We believe that these results will help to inform the design of network markets. With this in mind, we are currently working to analyze the performance of network markets with different topologies — in particular small-world, random and scale-free topologies — and to handle larger sets of markets than we considered here.

²The TAC supply chain competition also studies supply chains, but comes at it from the perspective of individual traders rather than from the perspective of overall market performance.

Acknowledgments

This work was partially funded by 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] K. Cai, J. Niu, and S. Parsons. On the economic effects of competition between double auction markets. In *Proceedings of the 10th Workshop on Agent-Mediated Electronic Commerce*, Estoril, Portugal, 2008.
- [2] A. Cassar. Coordination and cooperation in local, random and small world networks: Experimental evidence. *Games and Economic Behavior*, 58(2):209–230, February 2007.
- [3] 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.
- [4] A. Cassar, D. Friedman, and P. H. Schneider. Cheating in markets: A laboratory experiment. Working paper, Learning and Experimental Economics Projects of Santa Cruz (LEEPS) Laboratory, 2006.
- [5] <http://www.marketbasedcontrol.com/>.
- [6] D. Cliff. Minimal-intelligence agents for bargaining behaviours in market-based environments. Technical Report HPL-97-91, Hewlett-Packard Research Laboratories, 1997.
- [7] J. Cox, T. Offerman, M. Olson, and A. Schram. Competition For versus On the Rails: A laboratory experiment. *International Economic Review*, 43(3):709–736, August 2002.
- [8] 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.
- [9] S. Gjerstad and J. Dickhaut. Price formation in double auctions. *Games and Economic Behavior*, 22:1–29, 1998.
- [10] 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.
- [11] M. He, H. Leung, and N. R. Jennings. A fuzzy logic-based bidding strategy for autonomous agents in continuous double auctions. *IEEE Transactions on Knowledge and Data Engineering*, 15(6):1345–1363, 2003.

- [12] <http://jcat.sourceforge.net/>.
- [13] P. Klemperer. How (not) to run auctions: The European 3G telecom auctions. *European Economic Review*, 46:829–845, 2002.
- [14] H. L. Lee, V. Padmanabhan, and S. Whang. Information distortion in a supply chain: The bullwhip effect. *Management Science*, 43(4):546–558, 1997.
- [15] K. A. McCabe, S. J. Rassenti, and V. L. Smith. Auction design for composite goods. *Journal of Economic Behavior and Organization*, 14:127–149, 1990.
- [16] J. McMillan. *Reinventing the Bazaar: A Natural History of Markets*. W. W. Norton & Company, 2003.
- [17] 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.
- [18] 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 Management*, Troyes, France, 2006.
- [19] 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.
- [20] 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, 2008. Demo Paper.
- [21] 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.
- [22] 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.
- [23] 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.
- [24] S. Park, E. H. Durfee, and W. P. Birmingham. Use of Markov chains to design an agent bidding strategy for continuous double auctions. *Journal of Artificial Intelligence Research*, 22:175–214, 2004.

- [25] S. Rassenti, S. Reynolds, and V. Smith. Cotenancy and competition in an experimental auction market for natural gas pipeline networks. *Economic Theory*, 4(1):41–65, 1994.
- [26] S. Rassenti and V. Smith. Deregulating electric power: Market design issues and experiments in designing competitive electricity markets. In *International Series in Operations Research and Management Science*, volume 13, pages 105–120. Kluwer, Dordrecht, 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. Shah and S. Thomas. David and Goliath: Displacing a primary market. *Global Financial Markets*, 1(1):14–21, 18th June 2000.
- [30] V. L. Smith. An experimental study of competitive market behaviour. *Journal of Political Economy*, 70(2):111–137, April 1962.
- [31] 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.
- [32] R. S. Sutton and A. G. Barto. *Reinforcement learning: An introduction*. MIT Press, Cambridge, MA, 1998.
- [33] G. Tesauro and D. Bredin. High-performance bidding agents for the continuous double auction. In *Proceedings of the 1st International Conference on Autonomous Agents and Multiagent Systems*, New York, NY, 2002. ACM.
- [34] L. Tesfatsion. Agent-based computational economics: Growing economies from the bottom up. *Artificial Life*, 8(1):55–82, 2002.
- [35] P. Vytelingum, R. K. Dash, E. David, and N. R. Jennings. A risk-based bidding strategy for continuous double auctions. In *Proceedings of the 16th European Conference on Artificial Intelligence*, pages 79–83, Valencia, Spain, 2004.
- [36] V. Walia, A. Bye, and D. Cliff. Evolving market design in zero-intelligence trader markets. Technical Report HPL-2002-290, Hewlett-Packard Research Laboratories, Bristol, England, 2003.
- [37] 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.
- [38] A. Wilhite. Bilateral trade and 'small-world' networks. *Computational Economics*, 18:49–64, 2001.