

Network effects in double auction markets with automated traders

Kai Cai¹, Jinzhong Niu¹, 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 & Info. Science, Brooklyn College, City University of New York
2900 Bedford Avenue, Brooklyn, NY 11210, USA. parsons@sci.brooklyn.cuny.edu

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 the topology has a significant effect on the equilibrium behavior of the set of markets.

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 recession, 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, especially 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. We describe some experiments using network markets where the nodes in the network are double auction markets, traders can move between the markets, and the connections between markets are limitations on such moves. These experiments identify whether network topology has a significant effect on the steady state behavior of a set of connected markets and the speed with which the set of markets converges to that steady state. We see this work as a first step towards understanding the relationship between market topology and performance. Our long-term goal is to be able to use our understanding of this relationship to engineer network markets with appropriate properties.

2 Background

2.1 Double auctions

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 has been proposed — including those that took part in the Santa Fe double auction tournament [28], the reinforcement learning *Roth-Erev* approach (RE) [26] and the *Gjer-stad-Dickhaut* approach (GD) [9] — and the performance of these algorithms has been evaluated under various market conditions. Despite the high performance of GD traders, research into automated trading mechanisms has continued.

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 mechanism plays an important role in determining the outcome of an auction, and this is further borne out by the work of [35] and [21], both of which show that the same set of trading strategies can have markedly different behaviors in different auction mechanisms. This leads us to anticipate that in a set of connected markets the way that the markets are connected will also have an effect on the behavior of the markets.

2.2 Methodology

The basis of our approach comes from Smith [31] via Gode and Sunder [10] and then Cliff [6]. We follow these authors in having all traders, whether human or machine, be chosen to be either buyers or sellers. No trader can both buy and sell in the same experiment. On any given day, each seller is given some number of indivisible goods that they are allowed to exchange for money, and is given a value for each good — the trader’s *limit price* or *private value*. A typical restriction, which we adopt, is that no seller may sell a good for less than its private value. Buyers have a similar private value for a number of goods, but rather than goods, they are given an allocation of money which they may exchange for goods. No buyer is permitted to pay more than the private value for any good.

These conditions are what Smith [31] calls “conditions of normal supply and demand”, the conditions in which the flow of goods through the market is at equilibrium and each day sellers bring to market the same goods that cost the same to produce, and buyers look to buy the same goods at the same price. The aim of our experiments is to identify what this equilibrium would be, and to allow us to find the equilibrium point — bearing in mind that there is a certain amount of learning going on that will take time to converge — we repeat the same trading conditions day after day, allowing our trading agents to recall the outcomes of trade on the previous day and trading multiple goods to speed convergence to equilibrium. Despite this, the slow convergence of the learning³

³ Which we can attribute to the movement of traders between markets since we know that the trading strategies we use converge in a few days at most in single market experiments.

means that to get close to a steady state we run our experiments for 600 trading days under identical conditions with each day allowing for multiple rounds of trading.

Clearly, this is not a realistic model. There is no existing market in which the same set of traders will continue to trade with the same limit prices for more than a year of trading without some price shock altering prices or traders entering and leaving the market. The model is not intended to be realistic in this sense. The model is just intended to tell us about the steady state, and we know from the literature that introducing price shocks [9] and permitting traders to enter or leave the market [22] just slows convergence to the steady state.

Our justification for working with such a simplified model is that we see our work as fitting within the “class-of-models” approach, due to Sutton [18, 32]. According to Sutton, the aim of modelling economic systems is rarely to model a real market, but is to model an abstraction from a real market that captures the behavior of a whole class of markets — exactly those which are the instantiations of the abstract model. In this work we are trying to see what the steady state behavior is in all sets of competing markets, both those with price shocks and those without, both those in which traders enter and leave, and those that don’t. To do that we look first at the most abstract market. We can take the results of our shock-free and fixed-trader experiments and use them to predict the results of removing these restrictions, and in the future we can investigate whether these predictions are true. This approach, of course, ties in with Rubinstein’s suggestion [27] that economic modelling be used to help sharpen our economic intuitions about complex phenomena as well as being used to predict the behavior of real systems⁴.

3 Experimental Setup

3.1 Software

To experiment with multiple markets, we used JCAT [11], 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

⁴ [27] presents four purposes for economic modelling in general: to predict behavior; to guide decision-making by economic agents or policy-makers; to sharpen the intuition of economists when studying complex phenomena; and to establish linkages between theoretical economic concepts and everyday thinking.




Fig. 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.

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 (1) Fully connected. Each market is connected to every other market. (2) Ring. Each market is connected to exactly two other markets. This is what [36] calls a “local connected network”. (3) 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. (4) Star structure. One market is connected to every other market. There are no other connections between markets. This is the network topology studied in [25]. These topologies are illustrated in Fig. 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 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 [34], 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.

In this work we use the standard market selection strategy used by JCAT. Traders treat the choice of market as an n -armed bandit problem that they solve using an ϵ -greedy exploration policy [33]. Using this approach, a trader chooses what it estimates to be the best available market, in terms of its average daily trading profit in that market on previous days, with probability $1 - \epsilon$, for some ϵ , and chooses one of the remaining available markets with equal probability otherwise. We choose ϵ to take a constant value of 0.1. Our previous work 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 because 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 fee, placing a 5% profit charge on all markets. In 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 Hypotheses

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 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.

In any case, we would expect that, as in [12], the topology of the relationships to have an effect on market behavior. In a model where traders move between markets, we would expect that placing different restrictions on movement between markets would lead to differences in the ease with which traders can explore the space of markets and then reach their preferred market, affecting the time it takes the set of markets to reach their steady state. In addition, we might expect that these different restrictions might lead to the steady state favoring some markets over other. These considerations give us two hypotheses that we will test:




Fig. 2. How the markets change over time. (a) shows the total number of traders that move at the end of a given trading day, (b) shows the average transaction price each day for a set of five fully connected CDA markets with ZIP traders. The x-axis gives the trading day, the y-axis gives (a) the number of traders, (b) the transaction price.

1. The topology of the network market will affect the speed with which the set of markets reaches its steady-state configuration; and
2. The topology of the network will have a significant effect on the steady state configuration of the set of markets.

Note that in discussing these hypotheses, we find it helpful to distinguish the fact that some of the topologies we consider — the star and chain — are *asymmetric* in the sense that traders in some markets are more restricted in the markets that they can move to as opposed to the *symmetric* ring and fully-connected markets where, in terms of connections, all markets are equal.

3.5 Experiments

To test these hypotheses, 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 trading strategies described above and both the market selection strategies. 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.

In order to assess the effect of the different topologies on the convergence of the markets, we looked at the number of traders that moved each day. The market selection strategy picks a random market with probability ϵ , so there will always be some movement of traders, but we would expect to see the number of traders decreasing from an initial high to a steady state, and the speed with which the steady state is reached is one way to measure how quickly the system of traders and markets converges.

To identify any differences between the steady state configurations of different market topologies we looked at two things — the number of traders in each market, and the efficiency of each market. The number of traders in each market gives us some idea of the preference that traders have for markets, and any time that there is an uneven distribution it is an indication that from the traders' point of view differences in market

ZIC	CDA	Fully connected	141.25	ZIP	CDA	Fully connected	142.91
		Ring	107.48			Ring	108.90
		Chain	93.47			Chain	95.61
		Star	93.34			Star	98.44
CH		Fully connected	184.56	CH		Fully connected	155.75
		Ring	143.95			Ring	120.73
		Chain	125.43			Chain	109.11
		Star	127.68			Star	113.43

Table 1. The average number of traders moving each day for the different topologies.

topologies have an effect. Efficiency, of course, is a standard measure of market behavior, and will indicate whether differences in the market topologies have an effect on the performance of the set of markets as a whole.

4 Results

4.1 Speed of convergence

When we look at the movement of traders between markets it is clear from Fig. 2 (a) that the markets make an exponential approach to the steady state (these results are for ZIP traders and fully connected markets, but the results for other experiments are similar). This is despite the fact that the average transaction price in each market is, like that shown in Fig. 2 (b), far from steady⁵. Since, as described above, the market selection strategies 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.

To do this we borrowed from the usual measure of the convergence of a market to equilibrium [31]. To compute this measure, Smith's *alpha* as it is known, we compute the average deviation between the price of each transaction and the equilibrium price suggested by theory. Here, we look at the number of traders moving each day and compute the average difference from the number we would expect if the only cause of trader movement was the ϵ in the market selection strategy (which would mean that, on average, 10% of the traders would move each day). Markets that are faster to converge to the steady state will have lower values of this difference. These results are shown in Table 1 and show that there is a clear difference between the speeds with which the different topologies converge. In particular, the asymmetric topologies converge much faster than the symmetric topologies.

⁵ Because the transaction price is a function of the traders in a market (in particular it depends on both their private values and when and how they choose to bid), it changes as traders move between markets. Since the reward gained by traders is a function of the transaction price the dynamics are more complex than those of a set of n-armed bandit learners converging to static rewards.

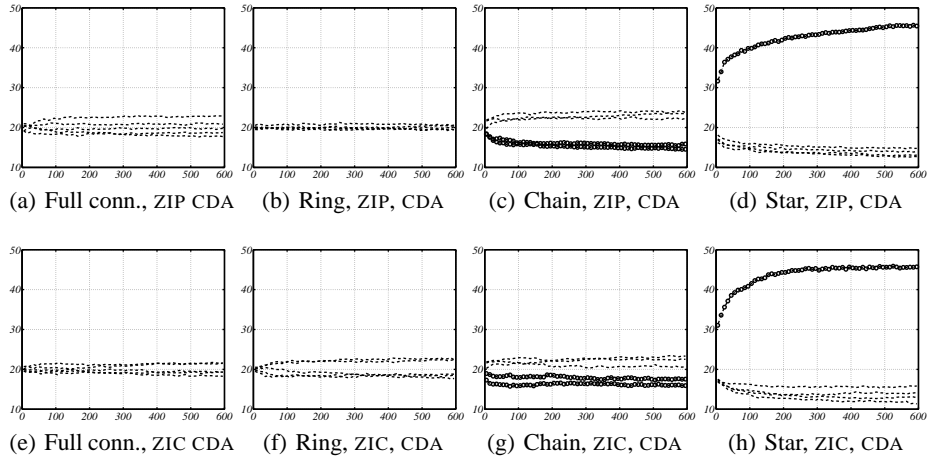


Fig. 3. The number of traders in multiple connected CDA markets with different connection topologies on each trading day. The traders in (a)–(d) use the ZIP strategy, those in (e)–(h) use the ZIC strategy. The x axis gives the trading days, the y axis the number of traders in each of the five markets. In the chain markets, the dark lines give the numbers for the markets at the end of the chain, and for the star markets, the dark line gives the numbers for the market at the center. All other markets are marked with dashed lines.

4.2 Trader distribution

To examine the steady-state for differences due to connection topology, we looked at the number of traders in each market. Figure 3 shows this for each day of the experiment for both ZIC and ZIP traders in CDA markets (the other experiments give very similar results). The graphs in the figure show that the distribution of traders in fully-connected (Fig. 3(a), Fig. 3(e)) and ring (Fig. 3(b), Fig. 3(f)) markets is pretty uniform.

Chain markets, however, don't have the same symmetry, and this shows up in the distribution of traders. As Fig. 3(c) and Fig. 3(g) show, markets at the end of the chain end up with fewer traders than the markets in the middle of the chain. The effect of the loss of symmetry is even more marked in star markets, Here, as shown in Figures 3(d) and 3(h) the hub market in the star collects many more traders than the otherwise identical markets that are connected to it.

The graphs of Fig. 3 don't make it easy to decide what differences are significant so we show the actual trader numbers after the 600th trading day (that is at the end of the experiment) in Table 2. This includes the results of all the experiments on star and chain markets, not just those from Fig. 3 (the ones from the figure are in the first and third rows of the table). 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. This holds for both CDA and CH markets whether 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 again for both CDA and CH markets for ZI-C and ZIP traders.

	Star					Chain				
CDA ZIC No. of traders	43.67	13.65	15.82	14.14	12.72	16.24	22.74	20.88	22.21	17.93
Stdev.	11.89	7.88	8.25	8.38	7.23	6.63	8.86	9.86	9.67	7.45
ZIP No. of traders	42.50	13.90	13.57	15.42	14.61	16.10	22.22	22.66	23.72	15.31
Stdev.	9.08	5.13	5.19	5.57	4.76	4.91	5.52	7.08	5.89	4.45
CH ZIC No. of traders	44.71	13.16	13.83	14.40	13.89	16.45	23.66	20.33	22.28	17.28
Stdev.	5.70	2.68	3.01	3.03	3.90	4.82	6.67	5.78	6.03	4.67
ZIP No. of traders	47.41	12.14	12.92	13.60	13.93	15.50	23.02	22.10	24.76	14.63
Stdev.	8.44	3.32	3.07	4.40	4.58	4.80	6.32	7.01	6.31	4.66

Table 2. The number of traders in each market for star and chain configurations for both market selection strategies. In the star configuration, 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. In the star configuration the number of traders in M0 is significantly greater than that in all the other markets with 95% confidence in all cases and in the chain markets the number of traders in M0 and M4 is significantly smaller than in all the other markets with 95% confidence in all cases.

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

Table 3. 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.

4.3 Allocative efficiency

The final results to consider are those in Table 3 which measures the allocative efficiency of sets of markets of different topologies. In particular what they measure is what we call “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 market containing all the traders.

Pairwise t-tests on the efficiency values in Table 3 reveal that there are differences between the efficiencies obtained with different configurations that are significant at the 95% level. In all the experiments the symmetric markets are significantly less efficient than the asymmetric markets. In all of the experiments except the CH with ZIC traders, fully-connected markets are less efficient than ring markets, ring markets are less efficient than chain markets, and chain markets are less efficient than star markets — all of these differences being significant at the 95% level. A possible explanation for this may be the fact that the asymmetric markets tend to concentrate traders in particular markets but results from our prior work [22] (on the effect of allowing traders to move in fully connected markets) suggests that such effects are only a partial explanation.

Note that the efficiency results we report for ZIP traders are somewhat lower than are reported for such traders in single markets (and are lower than the results we have

obtained for the same implementation of ZIP in a single market, results which are similar to those seen in the literature). We believe that there are a couple of reasons for this. First, we are computing efficiency as the surplus obtained divided by the surplus that would be obtained were all the traders in one market and that market traded at theoretical equilibrium. It is easy to see that it is possible to match traders in such a way that individual markets are efficient, but the combined surplus will fall below what would be possible if all traders were in one market and that is what we believe is happening here. (ZIP achieves higher efficiency when the efficiency is computed in a more conventional fashion.). Second, traders are constantly moving between markets, which means that the equilibrium point of all the markets is constantly changing (recall the transaction prices of Fig. 2 (b)). We know from [6] that ZIP takes several trading days to identify market equilibrium, and since this is changing every day, ZIP is always playing catch-up. Naturally this will mean it is less than completely efficient. (When traders are constrained not to move, the efficiency of ZIP improves.)

4.4 Discussion

The aim of this work was to test the hypotheses that:

1. The topology of the network market will affect the speed with which the set of markets reaches its steady-state configuration; and
2. The topology of the network will have a significant effect on the steady state configuration of the set of markets.

The results in Table 1 suggest that the first of these hypotheses is correct — for most of the experiments that we carried out, the time we estimate it takes the set of markets to converge varies considerably from topology to topology.

To address the second hypothesis, we measured both the number of traders in each market and the overall efficiency of the set of markets. When we looked at the number of traders (Table 2), it was clear that many more traders congregated in the central market of the star configuration and many fewer traders choose the end markets of the chain configuration, and pairwise t-tests confirmed that the differences are statistically significant. This suggests that the second hypothesis is correct. This suggestion is supported by looking at the efficiency of different sets of markets (Table 3) where we find that sets of markets with different topologies have significantly different efficiencies.

5 Related work

While network markets have not been studied in the same detail as single markets, there is a growing body of work to consider. [25], for example, describes a study of a three-node star network with a uniform-price double auction at each node. The same authors [24] 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 investigations all dealt with markets with human traders.

Agent-based methods were used by [3] to examine the effects of linked markets on financial crises, while [19, 20] consider the behavior of supply chains⁶. This work all studies smaller sets of markets than we have considered. The agent-based studies in [2] and [36] are larger but consider a set of connection topologies that overlap with, but does not contain, the set we consider. Both [2] and [36] deal with networks equivalent to our ring (their term is “local”) as well as small-world networks, which we don’t consider. Neither looks at chain or star topologies, the most interesting of the topologies we looked at, and neither study considers traders that move between markets.

The most closely related research we know of is [12], [37] and [14]. Judd and Kearns [12] describe experiments with human traders that clearly show that restrictions on who is allowed to trade with who — restrictions that are somewhat different from those imposed in our work — have a significant effect on market clearing performance. Wilhite [37], though mainly concentrating on results from network versions of the Prisoner’s dilemma, describes agent-based experiments in the same kind of scenario as studied in [12] with similar results. Ladley and Bullock [14] looked at networked markets of ZIP traders and showed that differences in topology affected an agent’s ability to make a profit. Like the results reported here, all of this work helps us to understand different aspects of the effect of network topology on market performance.

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. This work is the first systematic study of the effects of network topology on a set of double auction markets.

Traders in our experiments used either ZI-C or ZIP strategies, and markets were either CHs or CDAs. We looked at the behavior of four different topologies — fully connected, ring, chain and star — and considered the speed with which markets converge to a steady state, the distribution of traders across markets in the steady state, and the overall allocative efficiency in the steady state. We found that for all of these aspects, the connection topology can have a significant effect. In particular, the asymmetric topologies, chain and star, lead to an unequal distribution of traders, and in most cases an overall increase in efficiency of the markets.

Our main conclusion that topology affects steady state behavior is in line with previous work on network markets [12, 37]. In addition, since our results are consistent across different trading strategies (including the minimally rational ZI-C) and different market selection strategies, we believe that they will prove to be robust across other variants of our experimental scenario. 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.

⁶ The 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 NSF IIS-0329037 , and EP-SRC GR/T10657/01. We are grateful for use of the computational facility at the CUNY Graduate Center and to the reviewers whose comments helped us to make many improvements to the paper

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. <http://jcat.sourceforge.net/>.
12. J. S. Judd and M. Kearns. Behavioral experiments in networked trade. In *Proceedings of the ACM Conference on Electronic Commerce*, pages 150–159, Chicago, IL, July 2008.
13. P. Klemperer. How (not) to run auctions: The European 3G telecom auctions. *European Economic Review*, 46:829–845, 2002.
14. D. Ladley and S. Bullock. Who to listen to: Exploiting information quality in a ZIP-agent market. In *Proceedings of the Workshop on Trading Agent Design and Analysis*, 2005.
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. M. Morgan. Symposium on *Marshall's Tendencies*: 1 how models help economists to know. *Economics and Philosophy*, 18:5–16, 2002.
19. 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.

20. 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.
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 the dynamic behavior of traders in multiple markets. In *Proceedings of the Workshop on Trading Agent Design and Analysis*, 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. 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.
25. 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.
26. 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.
27. A. Rubinstein. *Modeling Bounded Rationality*. Zeuthen Lecture Book Series. MIT Press, Cambridge, MA, USA, 1998.
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. An experimental study of competitive market behaviour. *The Journal of Political Economy*, 70(2):111–137, April 1962.
32. J. Sutton. *Marshall's Tendencies: What Can Economists Know?* MIT Press, Cambridge, MA, 2000.
33. R. S. Sutton and A. G. Barto. *Reinforcement learning: An introduction*. MIT Press, Cambridge, MA, 1998.
34. 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.
35. 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.
36. A. Wilhite. Bilateral trade and 'small-world' networks. *Computational Economics*, 18:49–64, 2001.
37. A. Wilhite. Economic activity on fixed networks. In L. Tesfatsion and K. L. Judd, editors, *Handbook of Computational Economics*, volume 2, pages 1013–1045. Elsevier B. V., 2006.