The effects of market structure on a heterogeneous evolving population of traders

Dan Ladley[1] and Seth Bullock[2]

[1] Leeds University Business School,
University of Leeds,
UK
danl@comp.leeds.ac.uk
[2] School of Electronics and Computer Science,
University of Southampton,
UK
sgb@ecs.soton.ac.uk

Abstract. The majority of market theory is only concerned with centralised markets. In this paper, we consider a market that is distributed over a network, allowing us to characterise spatially (or temporally) separated markets. The effect of this modification on the behaviour of a market with a heterogeneous population of traders, under selection through a genetic algorithm, is examined. It is demonstrated that better-connected traders are able to make more profit than less connected traders and that this is due to a difference in the number of possible trading opportunities and not due to informational inequalities. A learning rule that had previously been demonstrated to profitably exploit network structure for a homogeneous population is shown to confer no advantage when selection is applied to a heterogeneous population of traders. It is also shown that better-connected traders adopt more aggressive market strategies in order to extract more surplus from the market.

1 Introduction

Understanding the centralised market has been one of the key aims of economic research for many years. Both the behaviour of the market and the traders within it have been intensely scrutinised in order to determine how they operate. Analytical studies, (e.g. [1]), experimental studies, (e.g. [2]) and empirical analysis (e.g. [3]) have all been employed in this attempt.

In addition to analytical, empirical and experimental results, the use of simulation and more recently multi-agent simulation[4] has become increasingly important [5–10]. Multi-agent approaches have enabled the relationships between trader micro-behaviour and market phenomena to be modelled, which is often analytically intractable and experimentally time consuming. In virtually all of these micro studies, the market is assumed to occupy a single location. All bids and offers are submitted in the same place, where all others may see and respond to them. Not all markets, however, are like this. Retail markets, for instance, are spatially embedded and consequently impose costs in terms of the time and effort that it takes to visit other traders and acquire information. As a consequence of this, it is usually impossible for a trader to visit *all* possible partners. Instead, the trader will probably restrict information gathering to nearest neighbours, or key operators in the market. In this case the market no longer has a central location to which information is submitted and, as a result, different traders within the market may have access to different histories of bids and offers.

It is not only spatially embedded markets that may limit the ubiquity of market information. Traders in a financial market have ready access to all trading information. However, in this case the shear quantity of information may segregate the market. The traders incur very little cost in gathering information, instead the main cost is that of analysis. Analysing information takes time, meaning that it may be impossible for a single trader to study and accurately respond to all of the information in the market in a fast enough manner. Traders are therefore likely to ignore some of the information available and fail to take it into account when making decisions. In effect the trader will not be hearing some of the information even though it is available in principle. One possible consequence of this is to focus the attention of traders on a small subset of market products, leading to specialisation.

There is, however, an important difference between these cases. Although a market may be segregated in terms of information flow, trade is not as restricted as it is in the spatially extended case. In either of these cases, however, assumptions about centralisation of market processes no longer hold. Different traders within the market have access to different histories of bids and shouts and, potentially, a propensity to deal with particular partners rather than others. These problems aren't necessarily limited to human traders. It is possible to conceive of markets that are sufficiently large, fast-moving, and complex that even computer programs would find it inefficient to analyse all information present, or consider trading with every agent in the market. Recently models have started to appear that examine these types of problems. For instance [11, 12] both examine trading scenarios that take place across networks, similarly [13–15], amongst others, consider the connected problem of trade network formation.

This paper aims to investigate the valuation of information within distributed markets. As has previously been described, traders in these markets will have access to different information sources and therefore different pictures of the market state. This will be particularly apparent if some traders are more connected than others, i.e., they have more information sources and/or trading partners. These better-connected traders are, on average, likely to have a better understanding of the market than those traders who are less well connected.

The effect of this imbalance is important because to some extent the degree to which a trader is connected can be altered by the trader itself. It is well known that resources must be expended to gather information and that properly analysing information takes time. In many situations it is possible for a trader to change the proportion of its resources dedicated to gathering and analysing information, however, it is important to know under which circumstances to do this.

In previous work [16] we have examined markets where both trade and information flow are restricted in a manner represented by an explicit, fixed network of possible agent-agent interactions. The network governed which agents were able to communicate with each other and, therefore, which agents were able to trade with each other.

Importantly, this network was not complete (fully connected), i.e., some traders within the market could not communicate directly with others.

In this initial work we wished to gain an understanding of the value of information in a simple separated market so the market network was fixed. Traders were not permitted to change their connections during the simulation. In future we hope to develop this system so as to better understand the circumstances in which it is favourable to change connectivity. The market used for these simulations was very simple, it was not designed to reflect the intricacies of any particular distributed market. Instead it was designed to provide general insight into the valuation of information in separated markets. The results found could be applied to any markets where information cannot flow freely. This includes retail markets, OTC markets, and many others.

It was found that traders who were more heavily connected had a valuation that was significantly closer to the theoretical equilibrium price of the market. The better-connected a trader was, the more information sources it would have and so the more accurate an opinion it could generate. The quality of information a trader possessed was, therefore, directly related to its connectivity.[3] In order to exploit this knowledge a simple modification was made to each trading agent's learning rule so that they weighted information according to the difference between the sender's connectivity and their own. As a result a trader would place more weight on information it heard from traders more connected than itself, and less on information from traders less connected. It was shown that on average a population of traders gained an advantage from using this rule, and that this advantage was enjoyed mostly by the least connected individuals in the population. In fact, the most connected individuals suffered a slight drop in performance when adopting this rather crude fixed learning rule.

This experiment assumed a homogeneous population. All traders within the population had parameters drawn from the same distribution and so behaved in a very similar manner. It is not difficult to argue that traders with different connectivities might perform better by adopting different strategies in order to exploit their position within the market network. For instance, we would expect that if the most highly connected individuals described above had had the choice, they would have chosen not to employ the learning rule that disadvantaged them, whereas those that were least connected may have chosen to use the rule more strongly in order to gain more benefit.

In order to allow such a heterogeneous population of traders it is necessary to individually specify each trader's parameters. One-way to do this is to hand tune every trader to find its optimal parameter set. However, the optimal parameter set for a particular trader is likely to depend on the parameter sets of the other traders within the market. In order to solve this problem it was decided to compete trading strategies against each other. A co-evolutionary genetic algorithm was designed to allow the evolution of competitive strategy sets.

[3] Although in that case information quality was based on a traders connectivity there was no reason why it could not be decided by other factors in a real-world market, such as a company's reputation or size or the previous history of information received.

2 Method

This section will first describe the structure and function of the markets that will be investigated, before detailing the traders that will populate them. It will then go on to describe the topological learning rule first introduced in [16] before finally describing the set-up of the co-evolutionary genetic algorithm.

2.1 Network Generation

Trading networks were constructed in which nodes represented traders and edges represented bi-directional communication channels. There are many possible network configurations that could be investigated for their effect on market performance, including lattices, Erdős-Rényi random graphs, small worlds, and graphs resulting from preferential attachment. This paper will focus on the latter class of networks since they exhibit some interesting properties, including the presence of well-connected "hubs", that have an intuitive appeal in terms of real-world markets, where it would be expected that certain major shops or investment banks would be much better-connected than individual shoppers or investors in their respective markets.

We employ an existing preferential attachment scheme [17]. A network of N unconnected nodes is gradually populated with Nm bi-directional edges. In random order, each node is consulted, and allocated an edge linking it to a second node chosen according to probabilities calculated as $p_i = (n_i + \delta)^P$. Here, P is the exponent of preferential attachment and remains constant, n is the node's current degree (number of edges), and δ is a small constant (0.1 for all results reported here) that ensures unconnected nodes have a non-zero probability of gaining a neighbour. Self-connections and multiple connections between the same pair of nodes were not allowed. All probabilities, p_i, were updated after every edge was added. After m cycles through the population, the network was complete. Note that every node will have a minimum of m edges, and a maximum of $N - 1$.

Markets explored here have a relatively high preferential exponent of $P = 1.0$ in order to generate networks that display a wide range of degrees. For all results reported here, $m = 10$. Initial tests showed that if m was significantly less than this value, the market failed to converge as few traders were able to trade with their limited number of neighbours i.e. it was separated.

2.2 Market Mechanism

The market mechanism operates in discrete time. Each time period, one active agent (one who is still able to trade) is selected at random to make an offer or a bid. The other agents in the market may only respond to that shout during that time period either to make a trade or ignore it. Once that time period has elapsed the shout is removed. Second, we limit an agent's ability to trade such that they are only able to make offers to, or accept bids from, their network neighbours. Each market was simulated for a fixed number of time steps.

The inspiration for this market mechanism came from the work of Gode and Sunder [5] and Cliff and Bruten [7]. Both of these cases investigated the effects of the continuous double auction mechanism on market behaviour. [4] The aim of this paper is to gain a similar understanding of the effect of the trade network. It is hoped that a relatively simple trade mechanism such as this will allow the effect of the network to be more easily identified and isolated.

2.3 Trading Agents

Here, the ZIP trading algorithm is used to govern trader behaviour. ZIP, or Zero Intelligence Plus, traders were created by Cliff and Bruten [7] in response to work by Gode and Sunder [5], who created the "Zero Intelligence" trading algorithm in some of the first agent-based market simulations. The Zero Intelligence algorithm was designed to be the simplest possible algorithm that would allow trade to occur in a market. Two types of Zero Intelligence trader were introduced. The first, unconstrained traders (ZI-U), choose shout prices at random from a uniform distribution across the whole range of possible prices permitted, disregarding any limit prices. It was found that markets populated by these traders exhibited none of the normal properties associated with markets, such as convergence to the equilibrium price. The second type of zero intelligence traders (ZI-C) were *constrained* in the range of prices that could be shouted. Shout prices were again drawn at random from a uniform distribution, however, this distribution was now constrained by a trader's limit price. In the case of sellers, shouts were constrained to be greater than the limit price, while in the case of buyers, shouts had to be less than the specified limit price. Importantly, markets populated by traders using this algorithm were shown to behave analogously to real markets in that they converged to the theoretical equilibrium price [5]. This was interpreted as indicating that the market mechanism itself was the most significant factor in market behaviour, and that the design of the trading algorithm was not as important. Cliff and Bruten [7], however, showed this to be incorrect, demonstrating that the convergence observed during each trading period was an artifact of the supply and demand schedules used by Gode and Sunder. They demonstrated that, for a certain type of supply and demand schedule that was close to symmetric, the probability distribution of likely ZI-C bids and offers would result in convergence to the mean price. They then performed simulations to verify these results with a broader range of supply and demand schedules. For non-symmetric schedules, markets populated by ZI-C traders failed to converge, or converged to a non-market-equilibrium value.

The ZIP trader differs from the ZI-C trader in that it learns from the market. Each ZIP trader has a profit margin associated with its limit price. In the case of buyers, the profit margin is the amount by which they wish to undercut their limit price to make a trade, and in the case of sellers, it is the amount by which they wish to exceed their limit price. When a ZIP trader shouts, the price is constrained by its limit price and profit margin. The trader uses the market's response to its activity (and the observable activity of others) to update its profit margin. For instance, buyers observe the bids made on the

[4] In our case the distributed nature of the market and the impossibility to maintain a single market price has meant that we can no longer describe the mechanism as a double auction.

market and whether or not they are accepted and adjust their profit margin accordingly. The ZIP algorithm employs the Widrow-Hoff learning rule with momentum [18] to adapt these profit margins throughout each trader's lifetime, maximising for each trader the possibility of making a profitable trade (for full details of this algorithm, see [7]). This learning rule allows the traders to rapidly converge on the optimal price, while the momentum term allows blips in the market to be ignored. Unlike ZI-C, ZIP traders are capable of finding the market equilibrium under a wide range of supply and demand schedules.

Here each ZIP trader was initialised with a random profit margin drawn from a uniform distribution $[A_1, A_1 + A_2]$. Each trader was also initialised with a random learning rate drawn from a uniform distribution $[B_1, B_1 + B_2]$ and random momentum value drawn from a uniform distribution $[C_1, C_1 + C_2]$. The target price for the Widrow-Huff learning rule was calculated from $T = Fq + G$ where q is the shouted price and F is a value drawn from a uniform distribution $[1.0, 1.0 + D_R]$ for buyers and $[1.0 - D_R, 1.0]$ for sellers. G is a random variable drawn from a uniform distribution $[0.0, D_A]$ for buyers and $[-D_A, 0.0]$ for sellers. Thus, F provides a small relative perturbation to the target price and G provides a small absolute perturbation.

2.4 Topological Learning Rule

The standard ZIP learning rule makes no distinction between the information it receives from different individuals. It was demonstrated [16], however, that there is a relationship between trader connectivity and accuracy of valuation. The following modification to the standard ZIP Learning rule was devised to take this imbalance into account.

The Widrow-Hoff rule currently includes a fixed learning rate which influences how quickly a trader is able to learn. In order for the traders to take account of information quality, the learning rate was modified so that instead of being fixed, the value would be calculated for each piece of information received. This alteration results in ZIP traders placing more weight on information obtained from well-connected individuals than from less well-connected individuals.

The function $f(s, r)$ was defined, where s and r are the sender and recipient of a piece of information (a shout).

$$f(s,r) = \begin{cases} 0.3 + \frac{0.2 \log \frac{E(s)}{E(r)}}{\log(R_{max})} : E(s) \geq E(r) \\ \\ 0.3 - \frac{0.2 \log \frac{E(r)}{E(s)}}{\log(R_{max})} : E(s) < E(r) \end{cases}$$

The function, E, gives the number of neighbours (degree) of a trader, and R_{max} is the largest ratio of edges between two adjacent traders within the market. M is the mid-point of the function and Q controls the range of values that it can take. This enhanced learning rule weights information according to relative connectivity within the market, i.e., the ratio of the sender's connectivity to the recipient's connectivity determines the learning rate. When the sender is more highly connected than the receiver the information received is more likely to be accurate and so more adaptation occurs. When the receiver is more connected, the receiver's current picture of the market state is likely to be more accurate than the senders and so less adaptation occurs. The value is normalised by the maximum ratio present in the market to prevent unnatural learning rates. Connectivity ratios are log-scaled to ensure that learning rate adaptation is sensitive to the small differences in connectivity that characterise most sender-recipient pairs in a network generated by a preferential attachment process (where there will be only a few very well-connected individuals).

The Widrow-Hoff "delta" learning rule was modified by removing the learning rate and replacing it with the function $G(s, r)$:

$$G(s, r) = \alpha f(s, r) + (1 - \alpha)L$$

Where L was the original Widrow-Hoff learning rate. This function allowed simple control of how much importance the trading strategy placed on the enhanced rule.

2.5 Genetic Algorithm

In this experiment it was desirable for different trading strategies to compete against each other in order to examine the selection of trading strategies. A co-evolutionary system was designed in order to do this. As was noted by Cliff [19] the behaviour of a ZIP trader is governed by eight real valued parameters that may be expressed as a vector V: $V = [A_1, A_2, B_1, B_2, C_1, C_2, D_R, D_A]$. In the case of enhanced ZIP traders it was necessary to introduce three new parameters that controlled the function of the topological learning rule, therefore, the vector used was W: $W = [A_1, A_2, B_1, B_2, C_1, C_2, D_R, D_A, M, Q, \alpha]$ These parameters were used to form a real valued genotype in which each parameter was bounded to lie between zero and one.

From previous results [16] we know that a traders connectivity affects its profitability within the market. As a consequence a traders connectivity may also affect its optimal strategy, therefore, it was desirable for traders with different connectivities to be able to evolve their own strategies independently. In order to do this it was necessary to maintain multiple populations of genotypes. One method of doing this would be to have one population of traders for each possible connectivity, i.e. a population for traders with 99 connections, a population for traders with 98 connections etc. There is a problem with this system, although there are many examples of traders with few connections in the networks there are relatively few examples of traders with large numbers of connections. It would have required a prohibitively large number of trials in each generation to evaluate each strategy's fitness accurately. In order to avoid this problem each market was broken up into a fixed number of groups, G_N (20 in all experiments reported here), sorted by connectivity. The N/G_N most connected individuals formed one group, the next N/G_N most connected individuals formed the second group and so on. Previous results [16] showed that traders with similar connectivities tended to achieve similar results and so could possibly employ the same strategy. This justified the formation of a set number of populations that would each contribute the same number of traders to each experiment.

To populate the groups G_n populations were formed each of size S_p (in all experiments reported here $S_p = 25$), each population corresponded to a particular group. In each trial N/G_N members of each population were chosen at random to form each group. The members of these groups were then added to the network in the appropriate places. Every member of each population participated in T_n trials each generation (in all experiments reported here $T_n = 40$). Due to the randomness present in the market and allocation of limit prices it was necessary to assess the traders fitness multiple times in order to attain a meaningful estimation. Strategy fitness was the average profit extracted by a strategy over all trials in that generation. The average profitability has an intuitive appeal, as in the real-world profitable strategies are more likely to survive and be copied. Standard roulette wheel selection with fitness proportionate weighting was used to select individuals for entry into the next generation. Mutation occurred at every locus of a selected genotype with probability P_m ($P_m = 0.05$ in all experiments presented here). Mutation consisted of a perturbation of the locus by a value drawn from a uniform distribution $(-0.05, 0.05)$. If the mutated value was greater than one or less than zero then the mutation was discarded and the original value used. In accordance with the method employed by [19], single point crossover was performed with probability P_c ($P_c = 0.3$ in all experiments presented here).

3 Results

Evolution occurred over 1000 generations. Each market was populated by 100 ZIP traders. Each trader was randomly allocated a limit price in the range $[1.00, 2.00]$, and either the ability to buy one unit or sell one unit of an unnamed indivisible commodity. Each market simulation lasted for 400 time steps. Markets were constrained by networks, constructed as described above, with $P = 1.0$ and $m = 10$, and all markets operated through the market mechanism described above. At the start of the experiment genotypes were initialised with parameters randomly chosen from the uniform distribution$[0, 1]$.

Figure 1 shows the average fitness of individuals within five different populations averaged over 24 experimental runs. The left figure shows the result of evolving standard ZIP traders, the right figure shows the result of evolving the enhanced ZIP traders. In both cases, strong trading strategies are quickly found by all populations (within approximately the first 40 generations). From this point onwards, however, the fitness levels of the populations remain approximately constant.

It should be noted that although the absolute fitness does not change after the first 40 generations this might not mean that the strategies are not continuing to adapt. Fitness is measured by the average amount of profit a trader makes. In these markets, however, the amount of profit available is fixed (though there is some small variation depending on the random distribution of limit prices). In order for one population of traders to increase its fitness it is necessary for it to become more profitable relative to another population. This is difficult to do as other population are simultaneously attempting to adapt their strategies to do the same. As in many co-evolutionary settings the trading strategies are continually adapting against each other and cancelling out each other's advantages. So although the fitness's may appear constant, the strategies may still be moving and changing in the strategy space [20].

In both experiments, corresponding populations attain similar fitness's. This indicates that some populations may have inherent advantages within the market and that

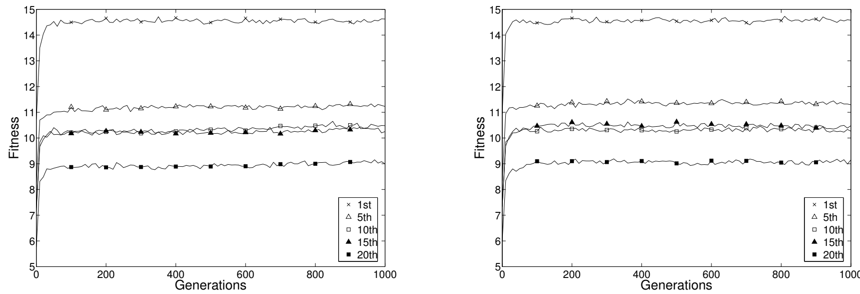

Fig. 1. Absolute fitness averaged over twenty four experiments for five of twenty populations ranked in decreasing order of connectivity for (left) standard ZIP traders, and (right) enhanced ZIP traders using a learning rule adapted to exploit market topology information.

the traders are not able compensate for these differences. It appears that the more heavily connected populations (low numbered) are able to exploit their connectivity advantage and extract the same amount of profit from the market in both cases.

Figure 2 deals with the deviation of the traders valuations over the length of a market experiment. These results were obtained by performing 10,000 market experiments using the final populations from each of the twenty-four genetic experiments. At each time-step the deviation from the equilibrium price of each trader's valuation was measured (traders that had already traded were not included in this measure). In previous work it was demonstrated that traders who were more heavily connected had valuations that were closer to the equilibrium price than those who were weakly connected. In this case, however, all traders quickly converge to equally good approximations of the equilibrium price. The addition of the topological learning rule does not appear to have any effect on the ability of traders to identify the equilibrium price. The convergence occurs at the same speed, and to the same level, both with and without the topological learning rule. (Note, this measure will never converge to zero as some traders have limit prices beyond the equilibrium price which bounds their valuation away from it).

Figure 3 shows the average learning rate for each of the final populations. In the case of the standard ZIP traders the learning rate is inversely proportional to the connectivity (r value < 0.01), i.e. more connected individuals have a lower learning rate than the less connected individuals. In the case of the enhanced ZIP traders the learning rate appears to remain approximately constant across the populations. The enhanced ZIP traders have on average a higher learning rate than the standard ZIP traders.

Figure 3 shows the average initial profit margin for each of the populations. The range indicated on the graph is the range from which each trader's initial profit margin is drawn. This value is then adapted throughout the course of the experiment as shouts are heard. In both cases there is a positive correlation with connectivity (r value < 0.01). i.e. the more connected a trader the higher the initial profit margin. The value of the initial profit margin does not appear to depend on the use of the topological learning rule.

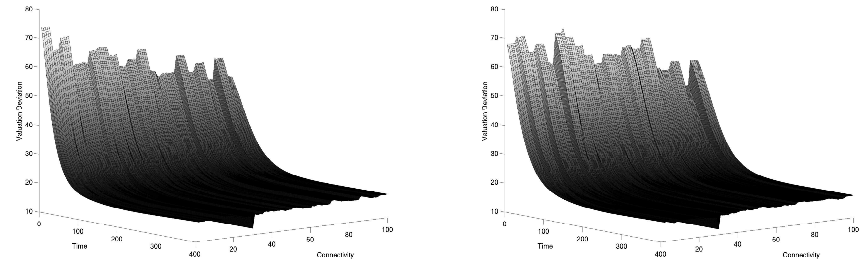

Fig. 2. Absolute deviation from optimum price averaged over 10000 runs for the final populations of each of 24 experiments. Traders ranked in decreasing order of connectivity for (left) standard ZIP traders, and (right) enhanced ZIP traders using a learning rule adapted to exploit market topology information.

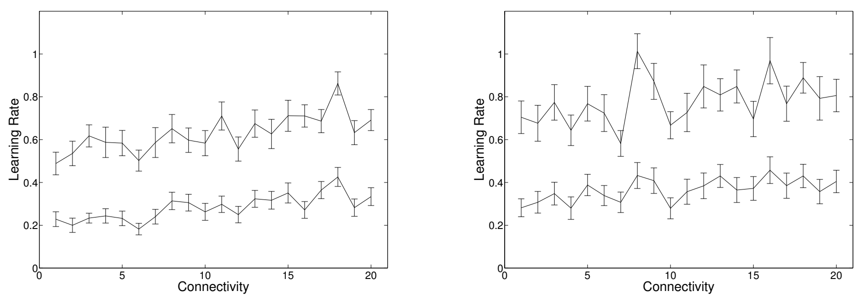

Fig. 3. Learning Rate for all members of the final populations of twenty four experiments, sorted in decreasing order of connectivity. (Left) Standard ZIP traders, and (right) enhanced ZIP traders using a learning rule adapted to exploit market topology information.

Figure 5 shows the average weighting factor for the topological learning rule for each of the populations. The graph shows that this remains approximately constant across populations. No population exploits the rule more than any other. It should be noted that the midpoint (M) and range (Q) of the rule remain approximately constant over all populations at 0.35 and 0.45.

In both sets of experiments the remaining parameters were approximately constant across populations. The momentum parameters $B_1 = 0.35$ and $B_2 = 0.70$ and the perturbation parameters $C_A = 0.35$ and $C_R = 0.3$.

4 Discussion

This paper aimed to investigating effects of diverse trading strategies on trading behaviour in a structured market. Previous work had examined markets with a homogeneous population of traders [16]. In this paper this limitation was removed in order to

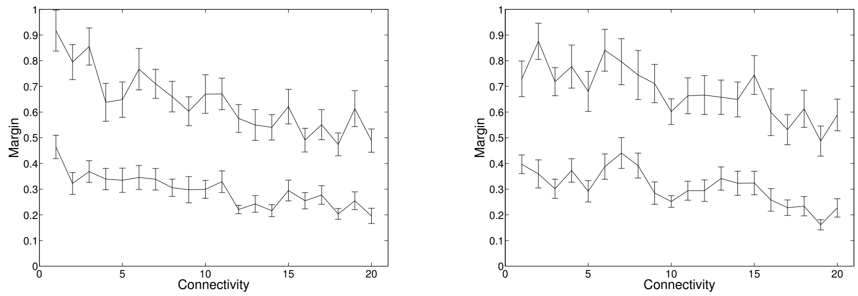

Fig. 4. Initial profit margin for all members of the final populations of twenty four experiments, sorted in decreasing order of connectivity. (Left) Standard ZIP traders, and (right) enhanced ZIP traders using a learning rule adapted to exploit market topology information.

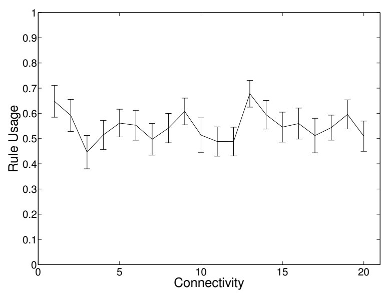

Fig. 5. Weighting factor for the topological learning rule adapted to exploit market topology information for the final populations of twenty four experiments, sorted in decreasing order of connectivity.

allow a more realistic representation of a market where individual traders may develop their own strategies based on their circumstances and environment.

The first finding of this paper was that, traders who are better-connected on average start with higher initial profit margins. By having a higher initial profit margin the more connected traders are adopting more *aggressive* market strategies. A higher profit margin means the well-connected traders demand more from their trading partners and as a result will probably take a larger cut of the profit from the trade. They are effectively able to charge their trading partner a premium for the right to trade with them. How are they able to do this? There is no advantage for the partner in trading with the well-connected individual, the unit of goods bought or sold has the same value and no reputation is gained through trading with the well-connected individuals.

It is the market position that is exploited by the well-connected traders in order to increase their profits. A better-connected individual has many potential partners but it will only trade with one of them. Once it has traded all other traders are left with one less potential partner. If a trader is quick and agrees to the disadvantageous terms it is able to reliably make a trade and extract some (small) profit. If, however, it does not then it takes a chance on finding another partner who is more generous, or not finding any partner and so making no profit. The extent to which a well-connected individual may do this is governed by its connectivity, the better-connected an individual, the more likely it is that another trader will take the unreasonable terms and trade. Therefore, the better-connected an individual is the higher it can set its initial price and still expect to make a trade.

This paper demonstrated that the way in which traders learn is affected by their connectivity. In the case of standard ZIP traders the results (3, left), clearly show that learning rate is inversely proportional to connectivity. The less well connected a trader is the more it learns from each piece of information. This seems to be intuitively correct, if a trader receives a large amount of information it is possible to average over them all and place less weight on any individual piece. If information is relatively sparse then the trader must place more importance on each piece heard.

Figure 3 (right) shows the learning rates of the enhanced ZIP traders. It would appear from this graph that the enhanced ZIP traders do not follow the same trend, as there is virtually no slope present. When the effect of the topological learning rule is included, however, a slope appears. Figure 5 shows that the weighting of the learning rule remains approximately constant across all populations, in addition, the midpoint and range of the function also remains constant. The effect of this rule, however, is not the same for all populations. Traders in the more connected populations are more likely to hear information from a traders who is less connected than they are, i.e. those from a lower numbered population, and vice versa for those in less connected populations. This effect becomes larger towards extremes. When the effect of the topological learning rule is added to the fixed learning rate, a similar pattern is observed to that seen in the standard ZIP traders. Although the learning rule does not improve the trader's performance, it is used as an easy way to correctly shape the learning function.

The results also show that, in a market populated by standard ZIP traders, those traders who possess more connections are able to make more profit than those who have less connections, as demonstrated by their higher fitness (figure 1 left). The fitness remains almost unchanged with the addition of the topological learning rule (figure 1, right).

This is a very surprising result. Previous experiments [16] had suggested that the addition of the topological learning rule allowed less well connected traders to reduce the informational advantage of better-connected traders. As a result the performances gap between the best and least well-connected individuals could be narrowed. Accordingly it was expected that with the addition of the topological learning rule the less well-connected populations would gain a higher fitness and the better-connected populations a lower fitness than before. In appears, however, that this was not the case. The topological learning rule had no effect on the ability of traders to extract surplus from the market.

Figure 2 shows parameter sets are evolved such that after a small number of time steps, all traders, on average, have a valuation that is equally close to the equilibrium price. In previous experiments, using markets populated by homogeneous traders, those with more connections had a valuation significantly closer to the equilibrium price than those with fewer connections. The enhanced ZIP traders produce an almost identical result. They converge to a similar level in a similar amount of time.

The significance of this result should not be understated. It demonstrates that a very simple trading strategy, ZIP, is able to evolve to perform well in a structured market with limited information. Previously, in a homogeneous population, it was demonstrated that simple traders could use information quality in order to increase the accuracy of their valuations. In this case, a heterogeneous adaptive population was able to adapt a few simple parameters such that the deviation from the equilibrium price of the shouts was approximately equal across all individuals (figure 2). As a result information inequalities were no longer visible in the shouts and so could no longer be exploited. By adapting their parameters the ZIP traders were able to remove the effect of information inequalities from the market, all that remained was the effect of the market structure itself.

This fact has important consequences. If the structure of the market (the number of possible partners a trader has) is the only remaining factor that is unequal between traders then it must be this that causes inequalities in profits. Since this factor cannot be affected by the trading strategy then the design of more sophisticated trading algorithms may not be able to mitigate this effect. In this paper it was shown that the addition of an topological learning rule, that was previously demonstrated to be effective in structured markets such as this, had no effect on the fitness or valuation deviations. In other words the standard ZIP strategy was sufficiently advanced that it could find the competitive market equilibrium in these separated markets and that the addition of a more complex strategy could not improve on the result found for any group.

This leads to two possible hypotheses. First, in this simple trading scenario more sophisticated trading algorithms may not be able to significantly outperform the standard ZIP algorithm. Second, that no population of competitive traders (those that do not voluntarily give up possible profit) will be able significantly improve the performance of one population relative to another. In order to test these hypotheses more experiments must be performed. In particular experiments that pit the standard ZIP algorithm and the enhanced algorithm against each other and against other trading systems in the same market.

5 Conclusion

This paper has demonstrated that it is possible to evolve simple trading strategies to function well in structured markets. It has demonstrated that by tuning a few simple parameters it is possible for the traders to quickly remove any informational imbalances present within the market. Any differences in profits that remain are then solely due to differences in the number of possible trading partners that each trader has. It was also demonstrated that the addition of a more complex trading strategy that had proved to be effective in structured markets populated by homogeneous traders, had no effect on the distribution of profit within the market. The more advanced trading strategy was not able to mitigate the imbalances due to the market structure. The inherent imbalances were shown to be exploited by the better-connected traders, allowing more aggressive trading strategies to be employed successfully. This was primarily due to a larger number of potential trading partners.

These results also have important consequences for real markets. The fact that simple trading strategies may be tuned to remove informational imbalances indicates that this may also be true in real markets. The markets used in this experiment were simple, however, they do capture important features of real commodity markets i.e. there are shouts and trades that specify prices for goods of a known quality and volume. Some real markets, such as financial markets, are not entirely dissimilar from this, though it is accepted that real markets posses more commodities and more information sources. This work suggests that in heterogeneous populations of self-interested adaptive traders, such as those found in the real-world, it is possible to ignore informational advantages as a result of market connectivity. After a short amount of time, barring the effect of private knowledge, all traders should have an equally good valuation of the commodity. Given that maintaining trading connections probably has some cost, what then is the advantage of having multiple connections? The advantage comes from having more possible trading partners. The more trading partners a trader has the more aggressive it can be in its trading strategy and as a result the more profit it can make. So even though all traders may know the fundamental value of a commodity the structure of the trade network itself allows some traders to extract a higher price for that good than should theoretically be possible. In other words traders can exploit their market position to extract more surplus from a market than theory suggests they should.

References

1. DeLong, J.B., Shleifer, A., Summers, L.H., Waldmann, R.J.: Noise trader risk in financial markets. Journal of Political Economy **98** (1990) 703–738
2. Smith, V.L.: An experimental study of competitive market behaviour. Journal of Political Economy **70** (1962) 111–137
3. Bouchaud, J.P., Mezard, M., Potters, M.: Statistical properties of stock order books: empirical results and models. Quantitative Finance **2** (2002) 251–256
4. Axtell, R.: Economics as distributed computation. In Deguchi, H., Takadama, K., Terano, T., eds.: Meeting the Challenge of Soical Problems via Agent-Based Simulation, Springer-Verlag (2003)
5. Gode, D.K., Sunder, S.: Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. Journal of Political Economy **101** (1993) 119–37
6. Arthur, W.B., Holland, J.H., LeBaron, B., Palmer, R., Tayler, P.: Asset pricing under endogenous expectations in an artificial stock market. In Arthur, W.B., Durlauf, S.N., Lane, D.A., eds.: The economy as a complex evolving system II. Volume XXVII of Santa Fe Institute Studies in the Sciences of Complexity. Addison-Wesley (1997) 15–44
7. Cliff, D., Bruten, J.: Minimal-intelligence agents for bargaining behaviors in market-based environments. Technical Report HPL-97-91, HP Labs (1997)
8. Gjerstad, S., Dickhaut, J.: Price formation in double auctions. Games and Economic Behavior **22** (1998) 1–29
9. Westerhoff, F.H.: Multiasset market dynamics. Macroeconomic Dynamics **8** (2004) 596–616
10. Farmer, J.D., Patelli, P., Zovko, I.I.: The predictive power of zero intelligence in financial markets. Proceedings of the National Academy of Science **102** (2005) 2254–2259

11. Wilhite, A.: Bilateral trade and ‘small-world’ networks. Computational Economics **18** (2001) 49–64
12. Bell, A.M.: Bilateral trading on a network: a simulation study. In: Working Notes: Artificial Societies and Computational Markets. (1998)
13. Kirman, A.P., Vriend, N.J.: Evolving market structure: An ace model of price dispersion and loyalty. Journal of Economic Dynamics and Control **25** (2001) 459–502
14. Howitt, P., Clower, R.: The emergence of economic organization. Journal of Economic Behavior & Organization **41** (2000) 55–84
15. Pingle, M., Tesfatsion, L.S.: Evolution of worker-employer networks and behaviors under alternative unemployment benefits: An agent-based computational study. Technical Report 12327, Iowa State University, Department of Economics (2005)
16. Ladley, D., Bullock, S.: Who to listen to: Exploiting information quality in a zip-agent market. In: To appear. (2005)
17. Noble, J., Davy, S., Franks, D.W.: Effects of the topology of social networks on information transmission. In Schaal, S., Ijspeert, A.J., Billard, A., Vijayakumar, S., Hallam, J., Meyer, J.A., eds.: Eighth International Conference on Simulation of Adaptive Behavior, MIT Press, Cambridge, MA (2004) 395–404
18. Widrow, B., Hoff, M.E.: Adaptive switching circuits. IRE WESCON Convention record **4** (1960) 96–104
19. Cliff, D.: Evolution of market mechanism through a continuous space of auction types. In: Computational Intelligence in Financial Engineering. (2002)
20. Cliff, D., Miller, G.F.: Tracking the red queen: Measurements of adaptive progress in co-evolutionary simulations. In: Proceedings of the Third European Conference on Advances in Artificial Life, London, UK, Springer-Verlag (1995) 200–218