

Who to listen to: Exploiting information quality in a ZIP-agent market

Dan Ladley and Seth Bullock

School of Computing

University of Leeds

LS2 9JT

Leeds

UK

danl@comp.leeds.ac.uk seth@comp.leeds.ac.uk

Abstract

Market theory is often concerned only with centralised markets. In this paper, we consider a market that is distributed over a network, allowing us to characterise spatially (or temporally) segregated markets. The effect of this modification on the behaviour of a market populated by simple trading agents was examined. It was demonstrated that an agent's ability to identify the optimum market price is positively correlated with its network connectivity. A better connected agent receives more information and, as a result, is better able to judge the market state. The ZIP trading agent algorithm is modified in light of this result. Simulations reveal that trading agents which take account of the quality of the information that they receive are better able to identify the optimum price within a market.

1 Introduction

The study of the centralised market has been one of the key areas of economic research for many years. There have been many attempts to understand the behaviour of markets and that of the traders within them. These attempts range from analytical studies, (e.g., DeLong *et al.*, 1990), to experiments on real subjects, (e.g., the studies of Smith, 1962).

In addition to analytical and experimental results, the use of simulation has become increasingly important [Gode and Sunder, 1993; Cliff and Bruten, 1997; Gjerstad and Dickhaut, 1998; Das *et al.*, 2001; Farmer *et al.*, 2005]. In particular simulation has allowed the modelling of trader micro-behaviour, which would be 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 which may limit the ubiquity of market information. Traders in a financial market have ready access to all trading information. However, in this case the sheer 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 within the market. 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 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 Wilhite [2001] and Bell [1998] have both examined trading scenarios that take place across networks.

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 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. This paper will first examine the advantage this inequality provides to the better connected traders along with the effect this has

on trading within the market.

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. This paper will begin to investigate this question. It will consider a market where both trade and information flow are restricted in a manner represented by an explicit, fixed network of possible agent-agent interactions. The network will govern which agents are able to communicate with each other and, therefore, which agents are able to trade with each other. Importantly, this network will not be complete (fully connected), i.e., some traders within the market will not be able to communicate directly with others. In this initial work we wish to gain an understanding of the value of information in a simple segmented market so the market network is fixed. Traders are 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 is very simple, it is not designed to reflex the intricacies of any particular distributed market in particular. Instead it is designed to provide general insight into the valuation of information in segmented markets. The results found within this paper could be applied to any markets where information cannot flow freely. This includes retail markets, OTC markets, and many others.

This model will differ from previous work in that it will model the micro-behaviour of the traders. In both of the previous studies mentioned above [Wilhite, 2001; Bell, 1998], trade between agents was abstract. When two agents were chosen to trade, their utility functions were examined and an allocation of resources was calculated such that the utilities of both agents were increased. In this study, we will use a well-established trading agent algorithm to investigate the effect of the market constraints on the ability of agents to identify the optimum price. In addition, an attempt will be made to modify the trading agent algorithm to better cope with, or exploit, this situation.

2 Method

This section will first describe the structure and function of the markets that will be investigated, before detailing the trading agents that will populate them.

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 which 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”, which have an intuitive appeal in terms of real world markets, where it would be expected that certain major investment banks would be much better connected than individual investors.

An existing preferential attachment scheme is employed here [Noble *et al.*, 2004]. A network of N unconnected nodes is gradually populated with Nm 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_i is 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 $P = 1.5$ 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 agents were able to trade with their limited number of neighbours.

2.2 Auction Dynamics

The market functioned according to an adapted continuous double auction mechanism. The standard continuous double auction allows buyers and sellers to submit bids to the rest of the market for consideration at any time. First, in order to simplify the implementation of a continuous double auction on a network, we adopted the system presented by Cliff and Bruten [1997]. In their simulations, the auction mechanism acts in discrete time and has no order book. 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 are then polled in random order for responses to the shout. If the response and the shout cross then a trade is executed at the first shouted price, if not the next agent is polled. If no trader accepts the shout then 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.

2.3 Trading Agents

Here, the ZIP trading algorithm is used to govern agent behaviour. ZIP, or Zero Intelligence Plus, agents were created by Cliff and Bruten [1997] in response to work by Gode and Sunder [1993], 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 an agent'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 [Gode and Sunder, 1993]. 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 [1997], 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 agent differs from the ZI-C agent 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 agent 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 market and whether they are accepted or not and adjust their profit margin accordingly (for full details of this algorithm, see Cliff and Bruten, 1997). The ZIP algorithm employs the Widrow-Hoff learning rule with momentum [Widrow and Hoff, 1960] to adapt these profit margins throughout each agent's lifetime—maximising for each agent the possibility of making a profitable trade. This learning rule allows the agents to rapidly converge on the optimal price, while the momentum term allows blips in the market to be ignored. Unlike ZI-C, ZIP agents are capable of finding the market equilibrium under a wide range of supply and demand schedules.

Here, each ZIP agent was initialised with a random profit margin drawn from a uniform distribution $[0.05, 0.35]$. Each agent was also initialised with a random learning rate drawn from a uniform distribution $[0.1, 0.5]$ and random momentum value drawn from a uniform distribution $[0.2, 0.8]$.

3 Initial Results

Experiments were performed using markets populated by 100 ZIP traders. Each agent was randomly allocated a limit price

in the range $[100, 200]$, 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 $P = 1.5$ and $m = 10$, and all markets operated through the continuous double auction mechanism.

Figure 1(left) shows the price deviation from the theoretical optimum averaged over forty thousand repetitions. Each agent's valuation was obtained at each time step of each repetition, and the average calculated. Notice that timeseries are shown for agents with connectivity *rank* ranging from 1st (most well-connected) to 100th (least well-connected). Over time, the average price shouted by all agents, regardless of connectivity, approaches the equilibrium price. This is to be expected, as it is a fundamental property of markets that they tend to converge to equilibrium. The agents do not all converge at the same rate, however. Those agents who have most connections converge fastest. Agents who are more connected receive more frequent information and so have a better impression of the state of the market. They are, therefore, better able to accurately judge the equilibrium price.

Agents converge on a market price that deviates from the equilibrium price. This is due to the allocation of supply and demand. As the market converges, it will become increasingly difficult for agents who have been allocated limit prices beyond the market's theoretical equilibrium price to find partners prepared to trade with them. Since agents cannot alter their limit price, and are not prepared to trade at a value below it, some will effectively price themselves out of the market. Indeed, some agents will be unable to trade despite the presence of willing partners in the market as a whole, because they will not have a *neighbour* prepared to trade with them.

3.1 Extension

As described, traders in the market assume that all of the information that they hear is of equal quality, regardless of its source, i.e., the ZIP learning rule makes no distinction between the information it receives from different individuals. It has been demonstrated, however, that there is a relationship between agent connectivity and accuracy of valuation. How, therefore, could the traders take advantage of this fact?

In reality, it is known that some sources of information are of better quality than others. This may be for a number of reasons, including market experience, quality of information sources, reputation, or size. For reasons such as these, people are more likely to trust information about financial markets obtained from a market trader than information from a pub landlord. Alternatively they are more likely to trust the manager of a large importer than a market trader for information about the fruit market.

In order to incorporate this factor into the model, the traders were modified to *weight* the quality of information received. Information judged to be of high quality was weighted strongly when adapting profit margins. There are many (possibly sophisticated) ways to evaluate the quality of information received by an agent, especially as each agent has multiple sources of information each of which may have sent messages several times in the past. It could be possible to construct an algorithm that determined a profit margin by





Figure 1: Absolute deviation from optimum price averaged over 40000 runs for agents ranked in decreasing order of connectivity for (left) standard ZIP agents, and (right) ZIP agents with a learning rule adapted to exploit market topology information.

comparing the most recent bid received to all previous bids received from that and other traders, and the relative information quality of those bids. The complexity of such an algorithm, however, would seem to be inappropriate for a model of this simplicity at this stage. Here, a very simple method for evaluating the quality of information was implemented, based on the results presented above.

ZIP agents adapt their price using the Widrow-Hoff rule every time they hear a shout from another trader in the market. This rule includes a learning rate which influences how quickly an agent is able to learn. Currently a fixed learning rate is assigned to each individual in the simulation from a uniform distribution $[0.1, 0.5]$.

In order for the agents 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 agents placing more weight on information obtained from well-connected individuals than from less well-connected individuals.

The Widrow-Hoff “delta” learning rule was modified by removing the learning rate and replacing it with the function $f(s, r)$, 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 an agent, and R_{max} is the largest ratio of edges between two adjacent agents within the market. This adaptive learning rate 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 in order to ensure that the learning rate remains within the same bounds as standard ZIP traders. 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).

4 Results

Figure 1(right) shows the results obtained with the modified learning rule. All other parameters are the same as the previous scenario. As before the deviation of the valuations decreases over time. Again the most connected agents converge more quickly than the least connected agents. Figure 2(left) affords an easier comparison between the two studies. The least well-connected agents converge more quickly when using the modified learning rule than when using a fixed learning rate. At all times they have a lower deviation from the optimum price than agents using a fixed learning rate. By the end of the market they are significantly closer to the optimum price than those using fixed learning rates (t-test, $p < 0.0001$). Over longer experiments they may, however, eventually converge to the same value. The convergence of the better-connected agents is very similar, with or without the presence of the modified learning rule, although convergence is slightly retarded in the former case. By the end of the market, however, both groups have attained very similar values. It must be remembered that, as a consequence of the preferential attachment scheme that generates the market network, the distribution of agent connectivities exhibits a power law. Poorly-connected individuals vastly outnumber well-connected agents. As a result, even if the adaptive learning rate does significantly retard the convergence of well-connected agents, only a very small number of traders will suffer (although these agents may in some sense be of above average “importance”).





Figure 2: Absolute deviation from optimum price averaged over 40000 runs for (left) the most and least connected agent in each experimental condition, and (right) for the single monitored agents in markets designed to control for learning rate.

Before discussing these results further, a slight bias introduced by the adaptive learning rate scheme must be dealt with. Although the adaptive learning rate is constrained to lie within the same bounds that constrain regular ZIP agents, the average learning rate employed by the adaptive ZIP agents is higher than that of standard ZIP traders. Recall that there are a greater number of weakly connected agents than strongly connected agents. In the case of the weakly connected agents, nearly all of their neighbours will be at least as well-connected, if not better-connected. This means that the typical learning rate employed by an agent will rarely be below the population mean. As a result, the average learning rate of these agents is increased, so faster convergence is not necessarily surprising. In order to demonstrate that the modified rule has an effect on convergence above that which would be expected to result from a simple increase in average learning rate, a further study was designed.

Two equal-sized groups of standard ZIP agents are initialised, the first group forming a completely connected clique, while the second forms a minimally connected ring. The quality of information being exchanged in the first group should, therefore, tend to be much higher than that being trafficked in the second group. A final modified ZIP trader with an equal number of connections to random agents within each group is added to the market. This agent does not make any shouts, nor respond to shouts. It simply adapts its valuation based on the information it hears. This study was run under two conditions: using a standard ZIP algorithm for the final agent, or using the adaptive learning rule instead. In the second condition, the connections from the final agent to the rest of the market are ignored by the adaptive learning rate rule, as they are never used to convey information to the market.

In the second condition, it is possible to set the number of connections from the final agent such that its average learning rate is equal to the average learning rate of a standard ZIP agent (0.3 in the studies reported here). The learning rates of the modified ZIP traders lie within the range $[0.1, 0.5]$. The following equation for the average adaptive learning rate, $\bar{\eta}(n)$, for a modified ZIP trader with n neighbours may be

written given that all agents within each of the two groups share the same connectivity and the final agent has an equal number of connections to each group.

$$\bar{\eta}(n) = \frac{1}{2} \left(0.3 + 0.2 \frac{\log \left(\frac{E_{max}}{E(n)} \right)}{\log(R_{max})} + 0.3 - 0.2 \frac{\log \left(\frac{E(n)}{E_{min}} \right)}{\log(R_{max})} \right)$$

Where E_{max} is the connectivity of agents in the fully connected group and E_{min} is the connectivity of agents in the weakly connected ring, and n is the number of connections of the final agent. If we require this average to be the standard ZIP average of 0.3 this equation can be solved to give:

$$n = \sqrt{E_{max} E_{min}}$$

Here, we employ a fully-connected clique of fifty-one agents and a ring of fifty-one agents, giving $E_{max} = 50$ and $E_{min} = 2$, requiring $n = 10$ connections in total, or five connections to each population.

The results of 40000 repetitions are shown in figure 2(right). As can be seen, the modified ZIP agent using an adaptive learning rate converges faster and to a significantly lower asymptotic value than the standard ZIP agent (t-test, $p < 0.0001$). Hence, the adaptive learning rate rule has a positive effect on convergence beyond simply increasing the learning rate.

5 Discussion

In this paper we wished to explore the effect on convergence of market structure in terms of trader connectivity. The result obtained in the first part of this paper clearly show that the more connected an agent is, the faster it is able to converge, and the closer it is able to get to the optimum price. In the short term the more well-connected traders receive more information, and so are able to adapt faster. In the longer term, this greater volume of information means that they have a better overall picture of the market, and so may evaluate the optimum price more accurately.

As a consequence, the source of a shout has an impact on the quality of the information obtained. This was demonstrated by factoring information quality into the Widrow-Hoff adaptation rule via an adaptive learning rate. Our results show that agents who employ this strategy value the commodity more accurately in the majority of cases.

The fact that a small number of well-connected agents do worse by adopting this adaptive learning rate rule demonstrates that it is not of universal utility. The reason for this may be related to who these individuals are connected to. In the market structures studied here, for the vast majority of the time, highly connected agents will receive their information from less well-connected individuals. Since such information is judged (by the learning rule) to be of relatively low quality, less attention is paid to it. This may be a good decision in the long run. However, when the market opens this proves to be detrimental. At the beginning of a market every agent has a random valuation. Therefore, it would pay to attend to any information, even if it originates from poorly connected individuals. Whereas the most poorly-connected individuals do just this, the most well-connected agents tend to trust their initial valuation to a greater extent. This results in the slower initial convergence seen in the most well-connected agents. One might imagine that since these well-connected, but initially misguided individuals have a very significant influence on their neighbours, they might retard the market's convergence as a whole. Results suggest that this does not occur for the topologies considered here, but one could imagine market structures in which the hubs are so large and scarce that they could disturb the market for some time. As it stands, the simple adaptive learning rate rule could obviously be modified to better suit the more well-connected individuals, or, alternatively, separate rules could be used.

The adaptive learning rate rule used in this paper is very simple. It was chosen in order to demonstrate that the quality of an information source could be an important factor in making trading decisions, and that this could be used by trading agents to improve their valuations. It is not an attempt to provide an optimal rule. There are many other possible factors which could be incorporated in order to make this rule more sophisticated.

As it currently stands this rule has several weaknesses. The most obvious is that it relies on information which may not be publicly known. First, it uses the maximum connectivity ratio present within the market in order to normalise the rate of change. This is necessary in order to ensure that learning rates were scaled to fall within the same range as that employed by standard ZIP agents. Second, whenever an agent adapts, it uses the connectivity of the shouter in order to determine how much attention should be paid to the shout. In real situations it seems unlikely that an agent would have access to either of these kinds of information. In order for an agent to know the maximum connectivity ratio, it would be necessary for the agent to know how the whole of the market was structured. For the simulations reported here, this information is easy to obtain. However, in real markets it seems highly unlikely that this information would be available. The only probable way for an agent to know the whole market structure is to be in contact with every agent in the market.

If this were possible, then it is likely that other agents would also be completely connected, and an individual agent's connectivity would cease to be an issue. The connectivity of an individual trader may be even more difficult to obtain. In the studies reported here, every agent knows its own connectivity and that of its neighbours. In reality it is unlikely that the agent would possess this second piece of information, or perhaps even the first. The only realistic way to determine how well-connected another agent is, is to obtain the information from the agent directly. As has been suggested, however, it seems unlikely that agents would want to give this information away, as our results suggest that it is valuable.

This begs the question, if it is impossible to obtain the necessary connectivity information in reality, what practical significance do the results presented here have? Although it may be impossible to obtain exact connectivity information, it is not impossible to generate an evaluation of the information quality of ones trading partners in a real market. Moreover, information quality can be measured via indicators other than connectivity. In the case of human markets, traders may implicitly evaluate many aspects of their trading partners when deciding the significance of a piece of information. These aspects may include estimates of the partner's size (is it an investment bank with many traders or is it an sole trader), reputation, market position, market experience, etc. As yet, it may be difficult for an artificial trader operating in a real market, i.e., the stock market, to do this, due to the complexity of the information processing involved. It could be possible, however, for such automatic trading algorithms to be provided with human-generated estimates of known trading partners. The problem becomes even more complex in the case of markets solely populated by artificial trading agents. However, as mentioned above, if sufficient information is available it may be possible to analyse the time series of shouts by a particular agent in order to estimate the quality of its information. There is obviously a great deal of further work to be done in this area, particularly in examining the effect of different evaluation rules and the effect of different market structures on market (and agent) behaviour.

This paper has been more concerned with the performance of the market as a whole rather than the situation of any one group of traders. There are, however, many interesting questions which can be asked about the ways in which information may be used. In particular it would be interesting to consider how the better connected traders could exploit their advantageous positions to make a larger profit and also how they could exploit the knowledge that other agents gain an advantage by considering information quality. In order to properly understand these issues, however, it may be necessary to make the market more sophisticated. Currently, traders may only trade once. This effectively limits the ways in which traders can exploit information because as soon as a trader makes a trade they are effectively removed from the market. As a consequence some areas of the market may become stagnant as all available trades are made. This could be remedied by the introduction of a continuous flow of buy and sell orders entering the market. Allowing traders to interact multiple times and to develop more sophisticated strategies, whilst preventing the market from stagnating.

In this paper only a small (but significant) difference between traders using the new strategy and those not using it is demonstrated. Two points should be noted, firstly this rule was not chosen as an optimal rule for increasing valuation accuracy, instead it was chosen for its simplicity in demonstrating a point. Secondly the market employed in these experiments is by its nature “one-shot” in that all traders only trade once for one unit of the commodity. This naturally limits the opportunities for making profits, in particular in eliminates reselling. In real markets, however, this is not the case. Traders in real markets often trade many times for large volumes of products. A small increase in accuracy when dealing with large volumes may make a significant difference to the profit obtained. In the extreme case, foreign exchange markets have a turn over in the region of one trillion dollars a day. Even very small increases in valuation accuracy in contexts such as these can result in huge increases in profits.

6 Conclusion

This paper has demonstrated that simple markets populated by simple trading agents may function despite trading constraints represented by an explicit, fixed network of possible agent-agent interactions. We have not attempted to classify or analyse the effect of topology in general. However, for one particular type of market, some interesting and encouraging results have been shown. We have demonstrated that more well-connected agents have an informational advantage within a market. Having more neighbours ensures that a trading agent has a better picture of the market and, as a result, a better valuation of the commodity being traded. By making a simple modification to the ZIP trading algorithm to take account of this observation, improved performance was demonstrated for all but the most well-connected individuals.

In a more general sense, these results show that if it is possible to estimate the quality of a trader’s knowledge, it may be beneficial to factor this information into the way in which the trading agent learns. This result isn’t necessarily restricted to trading agents. In general, in any market with incomplete information flow, it may be beneficial to pay more attention to the most significant players within the market. However, care is necessary. Currently the ZIP trading strategy does not allow the more well-connected agents to exploit their informational advantage. However, this could be changed fairly easily. It is not difficult to imagine methods that allow simple trading agents to exploit informational advantages in order to make profit. It is, however, difficult to imagine ways in which to modify trading strategies in order to prevent this occurring in network based markets. The only solution to this problem may be to design market systems which minimise informational asymmetries. This may require research into more sophisticated market mechanisms to replace the continuous double auction. Alternatively, it may require the addition of completely new processes to the market. For instance, a process analogous to the financial press may allow agents to regularly gain an overview of the behaviour of the entire market.

References

- [Bell, 1998] Ann Maria Bell. Bilateral trading on a network: a simulation study. In *Working Notes: Artificial Societies and Computational Markets*, 1998.
- [Cliff and Bruten, 1997] Dave Cliff and Janet Bruten. Minimal-intelligence agents for bargaining behaviors in market-based environments. Technical Report HPL-97-91, Hewlett-Packard Labs, UK, 1997.
- [Das *et al.*, 2001] Rajarshi Das, James E. Hanson, Jeffrey O. Kephart, and Gerald Tesauro. Agent human interactions in the continuous double auction. In Bernhard Nebel, editor, *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 1169–1187. Morgan Kaufman, San Francisco, CA, 2001.
- [DeLong *et al.*, 1990] J. Bradford DeLong, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann. Noise trader risk in financial markets. *Journal of Political Economy*, 98(4):703–738, 1990.
- [Farmer *et al.*, 2005] J. D. Farmer, P. Patelli, and I. I. Zovko. The predictive power of zero intelligence in financial markets. *Proceedings of the National Academy of Science*, 102(6):2254–2259, 2005.
- [Gjerstad and Dickhaut, 1998] Steven Gjerstad and John Dickhaut. Price formation in double auctions. *Games and Economic Behavior*, 22(1):1–29, 1998.
- [Gode and Sunder, 1993] Dhananjay K. Gode and Shyam Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1):119–37, 1993.
- [Noble *et al.*, 2004] J. Noble, S. Davy, and D. W. Franks. Effects of the topology of social networks on information transmission. In S. Schaal, A. J. Ijspeert, A. Billard, S. Vijayakumar, J. Hallam, and J.-A. Meyer, editors, *Eighth International Conference on Simulation of Adaptive Behavior*, pages 395–404. MIT Press, Cambridge, MA, 2004.
- [Smith, 1962] Vernon L. Smith. An experimental study of competitive market behaviour. *Journal of Political Economy*, 70:111–137, 1962.
- [Widrow and Hoff, 1960] B. Widrow and M. E. Hoff. Adaptive switching circuits. *IRE WESCON Convention record*, 4:96–104, 1960.
- [Wilhite, 2001] Allen Wilhite. Bilateral trade and ‘small-world’ networks. *Computational Economics*, 18:49–64, 2001.