Social networking and information diffusion in automated markets

Martin Chapman[1], Gareth Tyson[1], Katie Atkinson[2], Michael Luck[1], and Peter McBurney[1]

[1] Department of Informatics, King's College London, Strand, London, UK
{martin.chapman, gareth.tyson, michael.luck, peter.mcburney}@kcl.ac.uk
[2] Department of Computer Science, University of Liverpool, UK
katie@liverpool.ac.uk

Abstract. To what extent do networks of influence between market traders impact upon their individual performance and the performance of the specialists in which they operate? Such a question underpins the content of this study, as an investigation is conducted using the JCAT double auction market simulation platform, developed as a part of the CAT Market Design Tournament. Modifications to the JCAT platform allow for influential networks to be established between traders, across which they transmit information about their trading experiences to their connected peers. Receiving traders then use this information (which is the product of an n-armed bandit selection algorithm) to guide their own selection of market specialist and trading strategy. These modifications give rise to a sequence of experimental tests, the documented results of which provide an answer to the question phrased above. Analysis of the results shows the benefits of taking advice as a collective and demonstrates the properties of the communities which emerge as a result of engaging in widespread social contact.

Keywords: JCAT, CAT, agent-based trading, social networking, agent-based simulation, automated trading.

1 Introduction

Empirical ethnographies of financial trading firms and hedge funds have shown that traders are typically influenced by other people when making trading decisions or deciding strategies, e.g. [1, 6]. These others may be fellow employees or superiors of the same company, or outsiders, including direct competitors, passing on tips or advice. If such interactions only happened once, one could imagine that they would be subject to manipulation by the advice-giver, for example, talking up the value of a company stock in which the advice giver has an undisclosed stake. But most participants in financial markets interact with one another repeatedly, and known manipulators risk being shunned in future interactions, and thus losing opportunities to receive future information. From an evolutionary game theory perspective, then, it makes sense even for self-interested utility-maximizing participants to be honest in these interactions, even with the provision of information to competitors.

While there has been considerable attention paid in recent years to modelling social influences on trading behaviours, dating back at least to [10], few researchers have considered trader interactions in a context where traders may choose between competing marketplaces. The only work we know is by Dumesny et. al [3], which considers how traders may detect deceptive information received from others in their social networks by means of a trust model. Our focus is different. We are concerned here with two key research questions:

– To what extent, if any, does reliance on a social network improve trader performance?
– To what extent, if any, does trader performance depend on the topology of the social network of traders?

The answers to these questions are not necessarily obvious or straightforward. For example, although it may seem that the network topology would be important for trader performance, a simulation study of adoption decisions for new information technology found that business network topologies had little influence on overall adoption patterns in some circumstances [12].

To explore these two questions, we modified the `JCAT` platform, a simulation platform for agent-based trading in competing markets[3], to allow for social networking between traders. Specifically, we integrated a set of algorithms that allow autonomous software traders to dynamically receive, reconcile, and judge advice from each other before deciding where to trade and what to quote. Through these modifications, we explored the impact of sharing historical information regarding strategies on trader performance. Following the modifications, we then performed a comprehensive set of computer simulations, looking at several metrics relevant to our research questions.

This paper is structured as follows. Section 2 first presents some background information on the operation of the `CAT` Market Design Tournament, the platform of which our study is based. Following this, in Section 3, we present the core algorithms used to integrate social awareness into our traders. Section 4 then describes the design of the simulation experiments undertaken using the modified platform, with the key results arising from these simulations being presented in Section 5. Finally, Section 6 concludes the paper with a brief discussion of possible future work.

2 The CAT Tournament

The `CAT` Market Design Tournament is an international competition established to promote research into automated and adaptive market mechanism design. `JCAT` was created in 2006 by a team from the Universities of Liverpool and Southampton, UK, and Brooklyn College, USA, to support the Tournament, which has been held annually since 2007.[4] Entrants to the competition design and implement software entities, called *specialists*, which operate as matchmakers between automated software traders. Specialists are competing marketplaces, in the same way that, say, the New York Stock

[3] See `http://jcat.sourceforge.net/`.
[4] See `www.marketbasedcontrol.com`.

Exchange, NASDAQ, the Paris Bourse, etc, are all competing markets. The traders are adaptive software agents provided by the Tournament organisers, which buy or sell a single generic commodity. Each trading day, each trader first selects a specialist to trade through on that day, and then decides what level of shout to bid or ask in that specialist. For the shout decision, each trader is endowed with one of four standard automated trading strategies.[5] For the research reported in this paper, only one trading strategy was used, namely an implementation of the Gjerstad and Dickhaut algorithm [5], choosing an offer based on an estimation of utility.

Each trading day in the CAT Tournament begins with specialists deciding on and posting their fee levels for all traders to see. The fee structure and fee levels inform the traders' decisions as to which specialist to select that day. Entrants to the Tournament are scored on three equally-weighted criteria each trading day: share of all profits earnt by specialists that day; share of traders attracted to their marketplace that day; and the specialist's transaction success rate, which is the percentage of traders that were successfully matched with counter-parties by the specialist. These daily scores are summed over the Tournament, with the winner being that specialist having the highest cumulative score across all scored days. There is a growing literature on the strategies of specialists, for example, [14, 16, 18], and on how traders may best choose between competing markets, e.g., [4, 13].

3 Enabling social interaction between autonomous entities

3.1 Overview

In accordance with the apparent benefits of social connection presented above, this section presents several novel approaches to communication in a multi-agent system such as the JCAT platform.

Although conceptually social interaction can simply be understood by means of point to point connection, in practice, achieving worthwhile cross-entity communication is a multi-tiered process. First, traders must obtain the information to share; in the context of competing marketplaces, one would imagine the most sought after information to be the relative benefits of trading within each environment itself. With this in mind, our social entities exchange opinions on their historical strategies; that is, how well, in retrospect, their previous day's strategy caused them to perform with their last specialist. This includes their specialist choice and the exchange of a finer granularity of information, namely trading techniques themselves or the precise formation of market shouts.

Our social networks are established as restricted static subsets of the trading body, defined under the syntax of an adjacency matrix. This static nature is representative of a typical trading social network in which information will only be exchanged between previously known trusted parties. Once each trader is aware of their connected peers, a specific algorithm is employed to assess any potentially conflicting advice, the product of which is a recommendation for the invoking trader based on the opinions of the peer most likely to produce success via imitation.

[5] See [2] for details of these strategies.

As a final step to the information exchange process — employing advice — traders choose whether to use the experience of their peers with a given probability T, their own strategy at GD, or make a random selection at R. These probabilities are defined according to a given mentality or *profile*; credulous, sceptical or non-bias. The fact that traders therefore do not simply take advice verbatim, produces a secondary logical topology within which edges represent the transfer of advice.

Once advice is taken and a profit (or otherwise) made as a result of using it, our approach attributes this gain back to the original purveyor of the advice as (a surrogate for) utility in order to establish a subset of agents providing the strongest advice. Hence the intention is to develop a system which facilitates the comparison of conflicting advice and attempts to establish those traders with whom others are most likely to *correlate*. That is, those traders whose advice is most generalisable to an individual agent.

3.2 Trader experiences as the precursor to giving advice

Before traders can actively direct one another towards particular specialists they first need to be endowed with the ability to gauge their own historical performance. In order to capture this notion of personal experience, traders are given a specific success metric; that of trading above a proportion of their given trade entitlement within a particular specialist. If traders are able to transact above specifically half of their trade entitlement, they deem their experiences positive; if they can only attain a level below this threshold, their experiences are negative. Algorithm 1 specifies the operation of such a procedure. Based on an event-driven architecture, it stores the details of the buyer and seller traders involved in this transaction and sets their *rating* of the specialist in which they are transacting according to how many overall matches they obtain. This rating is based upon a buyer or seller being able to match half of its entitled trade units.

Algorithm 1 Invoked whenever a transaction occurs.

Require: TRANSACTIONEXECUTEDEVENT
Ensure: $numTraderMatches \leftarrow$ updated, $seller \leftarrow$ experience,
$buyer \leftarrow$ experience

$seller \leftarrow$ SELLEROBJECT
$buyer \leftarrow$ BUYEROBJECT
$numBuyerMatches \leftarrow numBuyerMatches + 1$
$numSellerMatches \leftarrow numSellerMatches + 1$
if $numTraderMatches(buyer) > buyer.tradeEntitlements()/2$ **then**
    $buyer.setExperience(true)$
else if $numTraderMatches(seller) > seller.tradeEntitlements()/2$ **then**
    $seller.setExperience(true)$
end if

3.3 Giving advice over network topologies

With this system in place, the potential to exchange such information becomes apparent. Using a variant upon existing mechanisms, Algorithm 2 seeks to select a trader who, itself, holds an informed historical opinion on how beneficial its previous actions have been. Alongside this, the set of available traders is limited to a fixed social subset using the information contained in an adjacency matrix, thus restricting the body of entities from whom advice is available. Such a state of affairs is analogous to real world scenarios, in which individuals tend to hold core social groups.

In practice, the algorithm consists of two key parts. The first establishes those traders with whom the invoking agent is linked. Note that the social channels provided by these links are static. The second selects a trader from whom to take information and returns this information as a *recommendation* after storing a record of the exchange. This second algorithm results in a dynamic 'logical' topology, which represents the flow of good advice between agents taking place over the underlying static social topology. The content of a recommendation is either a `<market, advice>` pair or a `<shout, advice>` pair; where to trade and how to trade respectively.

Algorithm 2 This method is called *before* each trader makes a market decision.

Require: The *traderId* of the invoking trader
Ensure: $recommendation \leftarrow (History(recommendedTrader), experience)$

Let $row[] \leftarrow AdjacencyMatrix(traderId)$ {Store this trader's row from the adjacency matrix.}
Let $linkedTraders < integer >$ be a new array of integers
for $i = 0$ to $row.length$ **do**
  if ValueInCell $= 1$ **then**
    $linkedTraders \leftarrow i$ {If the value in the matrix is a 1 (they are linked), store the trader at that position as a potential trader to receive advice from.}
  end if
end for
$recommendedTraderId \leftarrow adviceTaken$ {Taking a list of $linkedTraders$ store *one* of these traders as the trader to take advice from.}
REFERRER $\leftarrow$ recommendedTraderId
REFEREE $\leftarrow traderId$
REFEREE.$setReferrer$(REFERRER)
if $adviceTaken = \emptyset$ **then**
  return ERROR
else
  return $recommendation \leftarrow History.get(formalTraderIds.get(recommendedTraderId)), experience$ {Return the historical record of the linked trader recommended and an opinion on the merit of this choice, to the invoking trader.}
end if

3.4 Correlation System and n-armed bandit

Expanding upon this notion of trader selection mentioned above, the next stage in achieving social interaction involves solving the conflicts apparent in Line 8 of Algorithm 2. Although Line 8 abstractly selects an advising trader, with a trader being open to an $n$ary number of connections, the potential for conflicting experiences is high. This presents the need for a specific selection technique. Whilst the possibility of majority voting could be a device considered, a mechanism which weights traders according to a given metric would provide a more accurate choice. The statistical n-armed bandit approach provides an efficient solution in this case [17]. Using existing platform mechanisms, we employ a technique which attributes to each possibility (each 'advisor') utility gain estimates based on previous results, and approaches a decision amongst these utilities in a greedy manner; that is to say, a *greedy* strategy is employed which consistently *exploits* the highest estimate as opposed to *exploring*. We express the pure greedy strategy as follows:

$$Q_t(a^*) = max_q Q_t(a) \tag{1}$$

whereby $Q_t(a)$ is any given play (choice) and $a*$ is the chosen greedy action.

Despite the classification of this strategy as greedy, it does introduce a factor *epsilon* (ϵ) by which a random option may be explored thus classifying it as ϵ-$greedy$. In this way, at times, such an algorithm may make an explorative decision based upon an unexpected *play*.

In addition to this, we also establish the notion of *potential correlation* as an input into the n-armed decision process. Extracting profit from the previously self-contained records of each trader, we attribute it back to the advising trader upon whose advice this trader is currently acting. The success of others using an advising trader's advice is therefore attributed back to that specific advising agent and seen to be indicative of how transferable one trader's actions are to the rest of its social subset.

4 Experimental Setup

To study the effects of the above techniques on an agent populous, we implemented them into the `JCAT` platform and ran of series of full, local `CAT` game simulations. To vary our experimental state, for each network topology and trader type, we ran 5 of these simulations, creating 80 in total.

4.1 Traders

For the purposes of our experiments, we chose a sample size of 20 trading agents with *types* applied uniformly across the population. This size we felt was justifiable in relation to the standard number of traders in a normal `JCAT` simulation. Types we define to be of three specific categories, each denoting how a trader approaches social interaction. **Credulous** traders are susceptible to the advice of others and almost always willing to act on it; **sceptical** traders are unlikely to take the advice of others and more trusting of their inherent strategies, and **non-bias** traders are equally as likely to take advice as to ignore it. Theoretically, we define these through the variable probability $P[T \cup GD \cup R\ ]$; taking advice, using an inherent strategy and using a random strategy respectively. Table 1 formalises this probability distribution whilst Algorithm 3 presents an implementation.

Action / Profile	Non-bias	Sceptical	Credulous
Use networked trader advice	49%	4%	94%
Use underlying strategy	49%	94%	4%
Select a random specialist	2%	2%	2%

Table 1. Trader profiles.

Algorithm 3 Selects, for an individual agent, a market in which to trade

Require: $P[T], P[GD], P[R] \leftarrow$ game parameters $TraderProfile, T, GD, R$
Ensure: $courseOfAction \leftarrow choice$

```
rand ← (DISTRIBUTION.NEXTRANDOM)
if rand ≤ P[R] then
    courseOfAction ← R
else if P[R] > rand ≤ P[GD] then
    courseOfAction ← GD
else if (P[GD]) > rand ≤ 1.0 then
    courseOfAction ← T
else
    THROW ERROR
end if
```

4.2 Social Networks

In addition to the above, we implemented five standard static network structures within which our traders operate: random, fully connected, small-world [9], ring and hierarchical, whilst holding the use of no social network as a control benchmark. As discussed, these networks were simulated under the formation of adjacency matrix files and varied in co-ordination with the trader types above to provide more insightful analysis into the impact of social structures.

4.3 Specialists

Our trader population was provided with access to four specialists within our simulations, three of which were successful entrants to previous real games, namely `jackaroo` (Winner 2009), `PersianCAT` (Winner 2008), and `PoleCAT` (Standard entrant, 2010)[6]. The final specialist, `MetroCAT`, we describe as experimental as it was engineered for research based on an intimate knowledge of the game structure.

[6] Retrieved from `http://www.sics.se/tac/showagents.php`.

4.4 Measures

Although a large number of response variables were recorded from our simulations those most pertinent to the subsequent analysis are listed in Table 2.

Metric	Description
Trader Profit	The difference between a trader's valuation and the amount centrally paid or received.
Specialist Profit	Profit from all fees charged that trading day.
Advice taken and from whom	A record, for each trader, of who they are taking advice from across a given period.

Table 2. Recorded variables.

Notably, variables were to be recorded on average per trader (or specialist), per day and then per game using the following metric:

$$\alpha = \sum_{m=0}^{t} m \to \sum \frac{\alpha}{t} \times d \tag{2}$$

in which we record a moving average for a group of traders t, for our simulations, over a particular metric m, across the 30 trading days, d, of a simulation.

5 Results

5.1 Trader and market performance

The histograms in Figures 1 and 2 demonstrate the differences in aggregate trader and aggregate market profit that arose between the configured simulation environment with its realised social extensions and the environment without. Here, profit is used as the dominant indicator of both trader and specialist performance.

Figure 1 shows the comparative profits of traders who communicate and are able to use each other's advice, against the profits of those traders who work independently. Similarly, Figure 2 depicts the ability of specialists to make a profit, again, when traders work together and when they work alone. Depicted on the far right of both graphs are the profits made when no networking strategy is used (traders are independent). Immediately, one can observe that a much lower profit level is obtained for both the traders and the markets in which they trade when advice is not available.

Using this information and returning to the focus of our original research questions, it can be stated that, indeed, social influence has an impact on both the performance of the autonomous traders within the platform and the markets in which they transact.

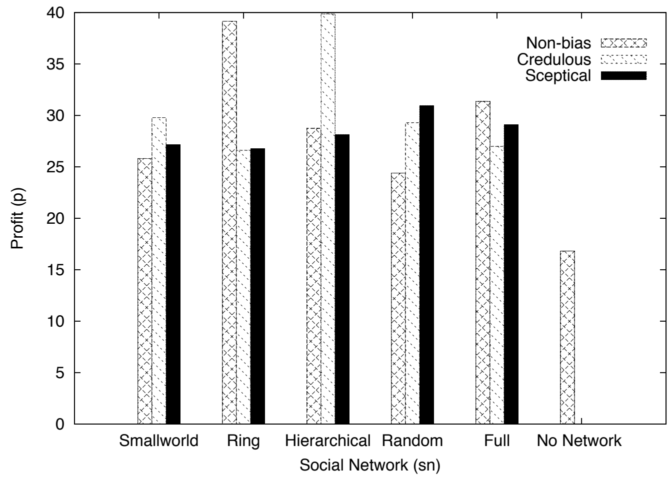

Fig. 1. How aggregate trader profits are affected by the introduction of social networking. 'Social' traders are displayed by type, grouped by network and compared to traders who operate with no social network.

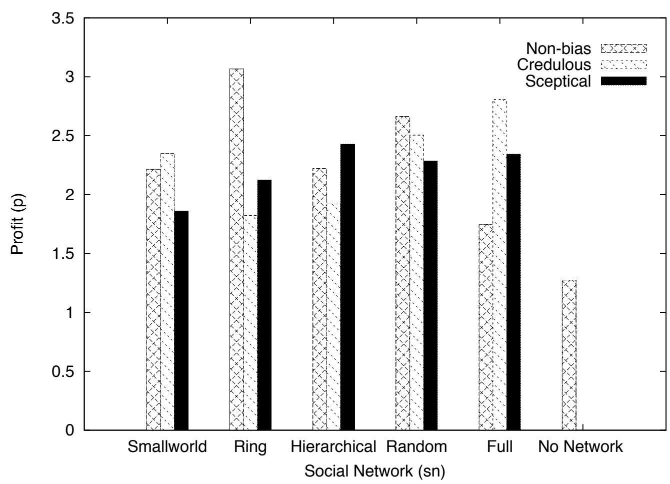

Fig. 2. How aggregate market profits are affected by the introduction of social networking.

The extent of this impact is so positive that it creates a marked statistical difference in the financial performance of both these entities within the simulation environment (c.f. Table 3).

Measure	Informed Social Strategy	No Social Network
Mean	29.60	21.96
Variance	90.25	44.80
Standard Deviation	9.50	6.69
Mean Difference	7.64	
Standard Error of Difference	1.34	
Degrees of Freedom	148	
t-statistic	5.70	
p-value	< 0.001	

Table 3. Statistical Test of Difference of mean values over the metric given in (2)

Intuitively, given the Algorithm in Section 3.3, this state of affairs makes sense. If trader T has access to $max\{T'_0 \ ... \ T'_n\}$ whereby each member of this set is a trader from T's network marked with a utility value representing their correlation to the rest of their social circles, then (assuming advice is taken) one would expect T's strategy to be optimal. In other words, having access to the experiences of others clearly allows traders to combine more informed market and shout choices, and thus potentially make a higher profit. In turn, due to the fee setup of a `CAT` game [2], the by-product of this is that specialists may also make a higher profit as they attain a greater portion of higher transaction levels.

However, we do make further comment on the differences in profit between traders and specialists, particularly under the hierarchical network. This seems to suggest some conflict of interest in both the gain and retention of profit within the market paradigm; when traders succeed, the level of profit made by specialists is reduced dramatically. By framing specialists and traders as conflicting agents, such a state of affairs makes sense. This is a beneficial result in context, as the focus of a `CAT` tournament is to increase the complexity of automated markets and the challenges they face in becoming successful entities. Therefore, by increasing the capabilities of agents, we change the dynamics of the market place, and so increase this challenge by forcing specialists to adapt.

5.2 Clustering, hubs and minority powers

By focussing on our agents as a group of intelligent social entities, we can make some comment on their patterns of behaviour as a collective. Although the initial network structures used in our platform are static and thus the communication channels open to individual traders fixed, due to the utilisation of our correlation system, the formation of virtual network clusters around certain trader individuals can be noted. In other words, if one imagines the utilisation of an advice channel as the formation of a secondary network edge superimposed onto an existing one, the clustering coefficients of some of the individual agents become orthogonal to the others. An exemplary logical topology can be seen in Figure 3.

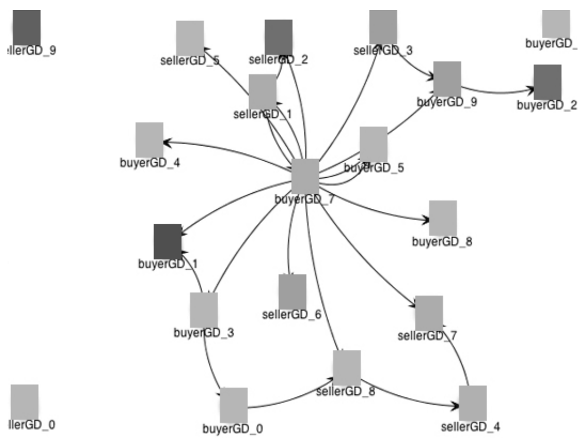

Fig. 3. The typical attraction of traders to specific local hubs. Here, *buyerGD7* has been weighted significantly by our matching algorithm and thus is highly connected in the logical topology. Traders are coloured according to market selection.

We describe these very 'popular' individuals (e.g. buyerGD7 in Figure 3) as *opinion leaders* [11] or power nodes, as their high level of connectivity makes a significant and, by the notion of preferential attachment, growing subset of the trader population, highly susceptible to their recommendations. Whether this almost reputation-like by-product of introducing utility as a measure of correlation is positive or not, is obviously therefore dependent on whether entities quote their historical results truthfully. Although, in our model the danger of traders purporting their poor performance as copyable (effectively lying) is not a possibility[7], injecting this potential into the network could be an interesting focus of future work. Such work is likely to be grounded in the trust and reputation literature, e.g. [7].

Despite the apparent influence of a small number of nodes in a network, Figure 4 tells us that in fact this dominant position does bring with it a large amount of volatility. Plotting each trader against the number of socially linked colleagues taking their advice, we can observe the fluctuations in the most reputable trader to be huge, as the consensus as to whom it is best to take advice from continuously shifts. Whilst traders seem to be easily admitted into their social hubs, the length of time they spend in such a role is notably limited.

We attribute the changeability of leadership in our networks to the minority power phenomenon [8] whereby opinions of a few eventually have a large effect on the majority. The behaviour of these outliers [19] is usually tangential to the performance of the opinion leaders and thus we observe a constant contest for influence in the network.

[7] Hence using static trusted social networks.

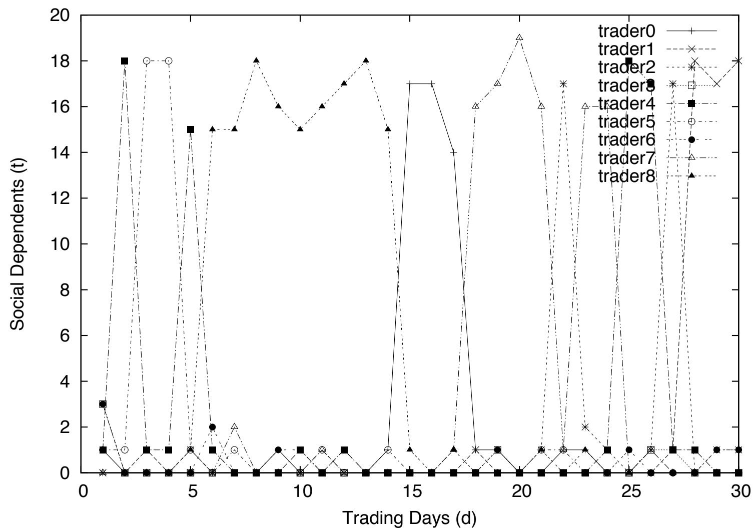

Fig. 4. A graphical representation of the tumultuous 'reign' of (a subset of) individual traders. Interestingly the choice of whom to take advice from seems to be restricted to an individual reputable group. *Game State: fully-connected, credulous.*

The performance of these outliers and their influence is therefore not necessarily negligible as we observe consistently high profit levels throughout our simulations. It would appear therefore that, in fact, exploring away from the consensus of the collective is altruistically beneficial.

5.3 Initial insights on the importance of topology

Our second research question brings us into a discussion about the relative importance of the underlying topology employed to facilitate social interaction. Interestingly, it can be observed that, in our simulations, topology actually has little effect on the ability for traders to disseminate advice (c.f. Figure 5). On the one hand, it is possible to attribute this to the limited size of the simulations.[8] However, we believe these results allow us to tentatively state that, in accordance with [12], topology is not a critical factor in the transfer of advice. Instead, of more importance is the topology of social trust that progressively evolves on top of this fixed underlying structure. Despite this, small benefits can be gained by tuning the underlying network to better facilitate dissemination; the hierarchical topology, for instance, does gain marginally better performance, likely due to its ability to scalably disseminate information.

[8] With simulations of 20 nodes, in practice, information dissemination can occur quickly over any topology.

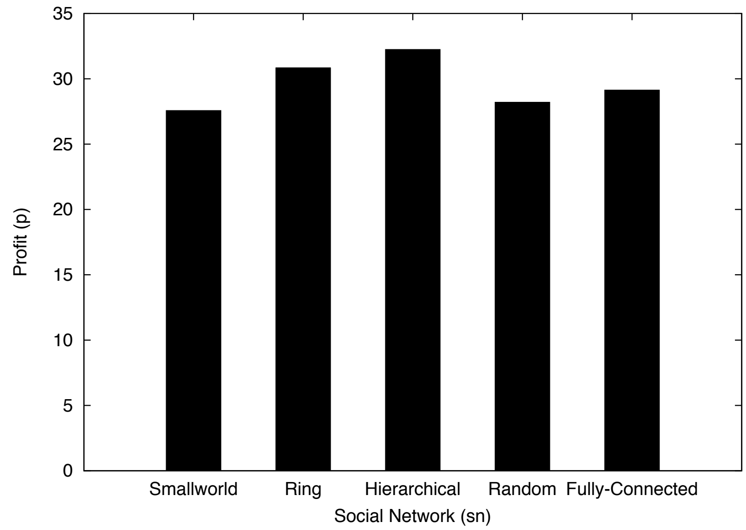

Fig. 5. The varying profit levels of traders exhibited between different **network topologies.**

6 Conclusions

In this paper we have presented preliminary research on the social influence of traders on each other's performance in automated trading systems. Although our experimental setup may require further alterations, our studies, undertaken using the `JCAT` platform, have shown that reliance on a social network improves both trader performance, and the profit performance of the market as a whole.

In future work, we are considering a more fined grained analysis of alternate social network topologies so as to better understand the dynamics of influence and correlation. In particular, we would like to explore the ability of traders to dynamically manipulate their directly-connected network through the information passed to it (thereby, introducing the potential for receiving information through new non-trusted agents). A related topic of interest is the extent to which social networks of traders facilitate or inhibit the co-evolution of segmentation that we have observed in `CAT` games and elsewhere, as traders and specialists co-self-organise into clusters or segments, discussed, for example, in [15]. Further, to expand on some of our findings, we consider it vital to look at larger population sizes to study the flow of information.

Acknowledgements

This work arose from the first author's dissertation at the University of Liverpool. We are grateful for assistance from and discussions with Guan Gui, Timothy Miller, Jinzhong Niu, Edward Robinson and Rahul Savani.

The development of the `JCAT` platform was funded by the UK EPSRC, through the *Market Based Control of Complex Computational Systems* project (GR/T10657/01) and this support is gratefully acknowledged.

References

1. D. Beunza and D. Stark, *Tools of the trade: the socio-technology of arbitrage in a Wall Street trading room*, Living in a Material World: Economic Sociology Meets Science and Technology Studies (T. Pinch and R. Swedborg, eds.), MIT Press, Cambridge, MA, USA, 2008, pp. 253–290.
2. K. Cai, E. Gerding, P. McBurney, J. Niu, S. Parsons, and S. Phelps, *Overview of CAT: A market design competition (version 2.0)*, Technical Report ULCS-09-005, Department of Computer Science, University of Liverpool, Liverpool, UK, 2009, Version 2.0.
3. J. Dumesny, T. Miller, M. Kirley, and L. Sonenberg, *TATM: a trust mechanism for social traders in double auctions*, Advances in Artificial Intelligence: Proceedings of 24th Australasian Joint Conference on Artificial Intelligence (AI 11) (Berlin, Germany) (D. Wang and M. Reynolds, eds.), vol. LNAI 7106, Springer, 2011, pp. 402–411.
4. K. Ganchev, Y. Nevmyvaka, M. Kearns, and J. Wortman Vaughan, *Censored exploration and the Dark Pool Problem*, Communications of the ACM **53** (2010), no. 5, 99–107.
5. S. Gjerstad and J. Dickhaut, *Price formation in double auctions*, Games and Econic Behaviour **22** (1998), 1–29.
6. I. Hardie and D. MacKenzie, *Assembling an economic actor: the* agencement *of a Hedge Fund*, The Sociological Review **55** (2007), no. 1, 57–80.
7. T. D. Huynh, N. R. Jennings, and N. R. Shadbolt, *Fire: An integrated trust and reputation model for open multi-agent systems*, ECAI **16** (2004), 18–22.
8. M. Kearns, S. Judd, J. Tan, and J. Wortman, *Behavioral experiments on biased voting in networks*, Proceedings of the National Academy of Sciences of the United States of America **106** (2009), no. 5, 1347–1352.
9. J. M. Kleinberg, *Navigation in a small world*, Nature **406** (2000), no. 2, 845.
10. T. Lux, *Herd behaviour, bubbles and crashes*, The Economic Journal **105** (1995), no. 431, 881–896.
11. K. Mahdi, S. Torabi, and M. Safar, *Diffusion and Reverse Diffusion Processes in Social Networks: Analysis using the Degree of Diffusion alpha*, Proceedings of the 3rd IEEE International Conference on Ubi-media Computing, U-Media (US), May 2010.
12. J. McKean, H. Shorter, M. Luck, P. McBurney, and S. Willmott, *Technology diffusion: the case of agent technologies*, AAMAS **17** (2008), no. 3, 372–396.
13. T. Miller and J. Niu, *An assessment of strategies for choosing between competitive marketplaces*, Electronic-Commerce Research and Applications **11** (2012), no. 1, 14–23.
14. J. Niu, K. Cai, S. Parsons, P. McBurney, and E. Gerding, *What the 2007 TAC Market Design Game tells us about effective auction mechanisms*, AAMAS **21** (2010), no. 2, 172–203.
15. E. Robinson, P. McBurney, and X. Yao, *Market niching in multi-attribute computational resource allocation systems*, Proceedings of the 13th International Conference on Electronic Commerce (ICEC 2011) (Liverpool, August 2011) (W. Kettler, K. R. Lang, and K. J. Lee, eds.), 2011.
16. B. Shi, E. Gerding, P. Vytelingum, and N. Jennings, *An equilibrium analysis of market selection strategies and fee strategies in competing double auction marketplaces*, Autonomous Agents and Multi-Agent Systems (2012), *(Forthcoming)*.
17. R. S. Sutton and A. G. Barto, *An n-armed bandit problem*, Reinforcement Learning: An introduction (Bradford, ed.), The MIT Press, Cambridge, Massachusetts, 2005, pp. 2.1–2.2.
18. P. Vytelingum, I. A. Vetsikas, B. Shi, and N. R. Jennings, *IAMwildCAT: The winning strategy for the TAC Market Design Competition*, Proceedings of 18th European Conference on Artificial Intelligence (Patras, Greece), 2008, pp. 428–434.
19. J. Zhou, C. Zhou, J. Wang, L. Qi, and W. Huang, *Detection of clusters and outlying nodes in spatial networks*, Geoinfomatics 2008 and Joint Conference on GIS Built Environment **7147** (2008), no. 1, 1–8.