

# The Strategic Exploitation of Limited Information and Opportunity in Networked Markets

Dan Ladley · Seth Bullock

Accepted: 2 April 2008 / Published online: 24 April 2008  
© Springer Science+Business Media, LLC. 2008

**Abstract** This paper studies the effect of constraining interactions within a market. A model is analysed in which boundedly rational agents trade with and gather information from their neighbours within a trade network. It is demonstrated that a trader's ability to profit and to identify the equilibrium price is positively correlated with its degree of connectivity within the market. Where traders differ in their number of potential trading partners, well-connected traders are found to benefit from aggressive trading behaviour. Where information propagation is constrained by the topology of the trade network, connectedness affects the nature of the strategies employed.

**Keywords** Trade network · Agent-based computational economics · Information · Strategy

**JEL Classifications** D85 · D83 · D40

## 1 Introduction

Economics research has demonstrated the power of markets for allocating goods amongst many individuals with heterogeneous allocations and beliefs. The majority of this work has focused on markets in which all traders interact through a centralised body or auctioneer, however, many are not organised in this manner. For instance,

---

D. Ladley (✉)  
Leeds University Business School, University of Leeds, Leeds LS2 9JT, UK  
e-mail: danl@comp.leeds.ac.uk

S. Bullock  
School of Electronics and Computer Science, University of Southampton, Southampton SO17 1BJ,  
UK  
e-mail: sgb@ecs.soton.ac.uk

markets are frequently spatially distributed with each trader only interacting directly with a small proportion of the other market participants. This paper investigates these decentralised markets; it aims to examine the effect of limited trading opportunities and information availability on the behaviour of individuals and of the market as a whole. It will examine how a trader's ability to make profit is influenced by their location within a trade network and how trader strategy must be adapted to cope with this constraint.

The Walrasian auctioneer oversees the classic example of a centralised market; information is complete, there are no transaction costs and decisions are taken centrally. Such idealised markets are attractive for work in economics and related fields as they combine many favourable elements including General Equilibrium through perfect competition. Work using this approach has generated many powerful results regarding the behaviour of markets and the traders within them, justifying its use and central role. In certain circumstances, however, objections to the idealisations involved may be raised.

The first and most significant is the issue of complexity (Scarf 1973). Work by Papadimitriou (1994) has shown that calculating the Walrasian equilibrium is NP complete. Hence, as the number of market participants or goods increases, the time required for some central executive to complete the necessary computation increases exponentially. Consequently, except for very specialised applications, the Walrasian system is rarely observed empirically.

Instead, many markets operate using systems that tend to approximate the Walrasian equilibrium whilst avoiding the impractical computational costs. Double auctions and the order book systems employed in many financial markets do this by distributing the computation of price and allocation of assets. These markets are highly efficient whilst also allowing traders to respond rapidly to changing information (Smith 1962). Market efficiency may be further increased at the cost of speed of trade by the use of batch auctions (Gode and Sunder 1997).

While distributed, the above market mechanisms are still centralised in that they operate at a single location in which trade is conducted and through which all traders communicate. This, however, does not have to be the case. It is possible for markets to be decentralised, i.e., for traders to communicate and trade with each other through private, pairwise interactions. Axtell (2005) demonstrates the potential of this type of market structure. He shows that a decentralised exchange can produce stable allocations which are computationally much more desirable than Walrasian general equilibrium but which are not in the core. In Axtell's model although the market is decentralised in terms of its computation it is still *integrated* in the sense that all traders may hypothetically interact and communicate with all others. Many real markets are a step beyond this. They are frequently decentralised in terms of their computation and *segregated* either in terms of trading opportunities (traders may only interact with a subset of others within the market) or in terms of information (not all traders have access to the same information).

Segregation in terms of trading opportunities may occur if traders only know about a certain subset of the potential trading partners within the market, or will only be willing to trade with a certain fraction of them. This may be dispositional, but may also be due to the inherent nature of the interacting entities within the market.

For instance, in labour markets referrals play an important role (Tassier and Menczer *in press*). Or, alternatively, transaction costs may make it impractical for certain traders to interact. This is especially frequent when distance plays a role in the market (see, e.g., Epstein and Axtell 1996, for a model in which trade is spatially constrained). As a consequence a trader may limit themselves to trading with a subset of the market and consequently risk not achieving the best possible interaction.

Markets may also be segregated in terms of information. Shoppers in retail markets incur costs in terms of the time and effort required to visit multiple shops in order to acquire information prior to trading. For this reason, a trader will frequently restrict information gathering and trading to nearest neighbours or key operators in a market and as a consequence may not execute the optimal trade.<sup>1</sup> Even in markets where all information is freely available the costs of analysing information may limit the amount of information that is actually considered, leading to different traders having different information sets (Brock and Hommes 1997). For instance, traders within a financial market have virtually free access to all market information. However, the sheer volume of information within the market prevents traders from analysing it all. Consequently, traders tend to specialise in a small subset of commodities for which they can study all of the available information. As a consequence, just like a retail shopper, each trader only interacts with a limited subset of the market.

The segregation of markets and the resulting less-than-fully-integrated nature of trader interactions allow heterogeneity's to exist. There may be differences in terms of who a trader can effectively trade with and in terms of the information each trader can effectively gather and process. In many markets, traders may expend resources to change their market situation, e.g., visit more shops or analyse more information. As a consequence, the dynamics of these markets and the behaviour of the traders within them are potentially very complex.

## 2 Previous Work

The behaviour and efficiency of decentralised markets is strongly affected by how the traders are segregated (Wilhite 2001). Within models investigating these types of markets it is important to be able to specify which traders can interact. Possibly the most natural way to do this is to situate the market on a network. Traders may interact with others to which they are connected and each trader is only connected to a subset of all others. Evstigneev and Taksar (1995, 2002) provide results demonstrating the existence of equilibria and the variation of equilibrium prices subject to changes in trader behaviour within such a setting. Several papers demonstrate the importance of the structure of the trading network in governing market behaviour (Evstigneev and Taksar 1994; Falbo and Grassi 2004; Tassier and Menczer *in press*). The network has been shown to affect the number of trades required in order to reach equilibrium (Bell 1998), the amount of searching for trade partners that is required (Wilhite 2001) and the results of negotiations between market participants (Markovsky *et al.* 1988;

---

<sup>1</sup> Restricted information gathering differs from segregated trading opportunities. A trader may be prepared to trade with many more traders than they are prepared to visit to acquire information.

Lovaglia et al. 1995). Together these papers have shown that a non-centralised market may converge to the theoretical optimum distribution of commodities, although the speed at which this happens is highly influenced by the structure of the network underlying the market. Networks with relatively high characteristic path lengths,<sup>2</sup> for instance lattices, converged much more slowly than those with short characteristic path lengths, for instance small world networks. Small world networks, those with few links but also low characteristic path length (Watts and Strogatz 1998), have the added benefit of requiring relatively little expenditure on searching for trading partners compared to more densely connected networks.

The formation of these networks has received significant attention. Key to this work is the process of traders establishing mutually beneficial relationships between each other based on their own preferences. Over time these preferences lead to the formation of stable networks over which trade occurs. This process has been examined in several settings, for example labour markets (Tefatsion 2001), buyer seller relationships (Kirman and Vriend 2001) and the emergence of money and shops (Howitt and Clower 2000) (see Vriend 2006, for a review).

This paper will focus on the effects of market segregation on market dynamics. We aim to understand how boundedly rational traders cope with the limited information set implied by a segregated market and how this affects their trading strategy and the market's ability to converge.<sup>3</sup>

The effect of market structure on the dissemination of information and how traders make use of this information is of central concern. Previous studies have shown in a relatively simple setting how better connected traders are able to exploit their greater access to trading opportunities to make higher profits by trading with more individuals. In contrast, in this paper we wish to examine the “market power” provided to a trader by their position, i.e., how a trader is able to leverage their greater access to the market to make the best possible trades. To this end we examine a model of trade in which every trader is able to buy or sell a single unit of a commodity with the aim of obtaining as high return as possible.

In order to be able to understand how market segregation affects individual traders it is necessary to employ a model that captures processes of information exchange and learning at the level of individual traders. In several of the above-mentioned models, pairs of traders on the network directly interrogate each others utility functions and calculate an optimal allocation of goods. There is no role for information transfer or strategy since both traders know exactly how each other values the commodity and this valuation does not change throughout the lifetime of the market. This approach is therefore inappropriate for investigating how individual trading behaviour might be influenced by constraints on information flow and trade. Instead we employ a model in which each trader is allowed to learn individually using a learning rule.

---

<sup>2</sup> A network's characteristic path length is the mean of the shortest distances that separate all pairs of network nodes.

<sup>3</sup> By the term “segregated market” we mean a market that is less-than-fully integrated, i.e., one in which not every trader is connected to every other trader. Here we will mostly be concerned with markets that, while segregated in this sense, are not *fragmented*, i.e., are not composed of multiple isolated sub-markets.

The trade networks investigated within this paper are fixed. They do not change during the duration of the experiment. As Wilhite states in relation to changes in market topology “if this evolution is slow relative to economic activity taking place on the network the characteristics of the network will have a larger effect on economic decisions than those decisions will have upon the network” (Wilhite 2006). In this case we are focusing on a relatively short period during which the equilibrium price of the market is identified and during which information is most significant, consequently it is justified to model a fixed network in this case.

### 3 Model

We employ an agent-based model of trader interaction in which the actions of each trader are governed by individual behavioural rules. Throughout the simulation traders react to, and may learn from, the market conditions they find themselves in. Through analysis of the success of the differing strategies we gain insight into the nature and dynamics of trader behaviour in segregated markets.

#### 3.1 Market Structure

In order to examine the effect of market segregation trader interactions were constrained by embedding the market on a network where traders occupied nodes and edges represented bi-directional trade and/or communication channels. Many different network configurations can be investigated, e.g., lattices, Erdős-Rényi random graphs, small worlds, etc.<sup>4</sup> In this paper we concentrate on network structures resulting from *preferential attachment* (Barabasi and Albert 1999), a process that grows networks incrementally by tending to attach new nodes to existing well-connected nodes. The graphs that result exhibit power-law degree distributions and positive assortedness (Newman 2003). Such graphs are intuitively attractive as structures for markets as they share many properties with real world social networks. Importantly, they contain well connected network “hubs”, i.e., major players within the market who are well connected to many other traders. Whereas such hubs tend *not* to be directly connected within technological networks such as the Internet, power grids, etc., they do tend to be directly connected within social networks (and those resulting from preferential attachment schemes) since influential individuals tend to influence one another directly.

An existing preferential attachment scheme is employed here (Noble et al. 2004). A network of  $N$  unconnected nodes is gradually populated with  $Nc$  edges. In random order, each node is consulted, and allocated a bi-directional edge linking it to a second node chosen according to probabilities calculated as  $b_i \propto (E(i) + \epsilon)^E$ . Here,  $E$  is the exponent of preferential attachment and remains constant,  $E(i)$  is node  $i$ 's current degree (number of edges), and  $\epsilon$  is a small constant (0.1 for all results reported here) that ensures unconnected nodes have a non-zero probability of gaining a neighbour.

<sup>4</sup> Previous work (Bell 1998; Wilhite 2001; Falbo and Grassi 2004) has demonstrated some of the properties of markets set on these different structures.

Self-connections and multiple connections between the same pair of nodes are not allowed. All probabilities,  $b_i$ , are updated after every edge is added. After  $c$  cycles through the population, the network is complete. Note that every node will have a minimum of  $c$  edges, and a maximum of  $N - 1$ .

Markets explored here have a relatively high preferential exponent of  $E = 1.0$  in order to generate networks that display a wide range of degrees. For all results reported here,  $c = 10$ . Initial tests showed that if  $c$  was significantly less than this value, the market failed to converge as few traders were able to trade with their limited number of neighbours. With significantly more edges the behaviour of the market approaches that of a completely connected centralised market.

### 3.2 Market Mechanism

The market mechanism for these experiments was chosen for its simplicity. It is sufficient to capture the essentials of information exchange and trade without introducing additional complexity. Each simulation is run for a fixed number of discrete time steps. At the start of a simulation every trader within the market is issued, at random, with either one unit of an indivisible nameless commodity to sell or the right to buy one unit of the nameless commodity. Each of these traders is also issued an integer reservation price, chosen with uniform probability from the integer distribution  $[P_{\min}, P_{\max}]$ . For buyers, this reservation price is the maximum value that they can pay for a unit of the commodity. For sellers, it is the minimum value they can accept for their unit of the commodity. At each step, a trader who has yet to trade is selected at random to make a shout. The other traders to which the shouter is connected in the trade network are then polled in random order for responses to the shout. The market mechanism operates on a simple “take it or leave it” basis. If the value of the shout is acceptable to a listener then a trade is executed at the shouted price, if not the next agent is polled. If no trader accepts the shout then no trade is made.

### 3.3 Trader Behaviour

The traders employed within these experiments employ the “ZIP” learning rule governing their price-determining behaviour (Cliff and Bruten 1997). This learning rule allows the traders to learn from the information present within the market and make informed trade decisions whilst also being simple, thus avoiding unnecessary complexity.

The ZIP learning rule attempts to maximise the amount of profit made by the trader based on information it hears from other market participants. Each trader maintains a profit margin,  $m$  bounded to lie in  $[0, 1]$ , associated with its fixed reservation price. For buyers, this determines the amount by which they wish to undercut their reservation price when they trade. For sellers, it determines the amount by which they wish to exceed their reservation price. The trader’s valuation of the commodity is  $p$ . This is the value they shout if requested and the least favourable price they will accept from another trader to trade at. It is defined as  $p = \lambda - m(p_{\max} - p_{\min})$  for buyers and  $p = \lambda + m(p_{\max} - p_{\min})$  for sellers, where  $\lambda$  is a trader’s reservation price.

Each time a trader within the market makes a shout, connected traders adapt their profit margin in the following manner. If the trader is a buyer and they hear a shout at price  $q$  which is accepted and if  $p \geq q$  (a trader successfully trades the commodity at a lower price) the buyer raises their profit margin (decreases their price). If the shout was an offer and  $p \leq q$  (a seller offered the commodity at a higher price than the buyer valued it) the buyer lowers their profit margin (increases their price). If the last shout was a bid which was not accepted then if  $p \leq q$  (a trader bid more than the buyers valuation and it was not accepted) the buyer lowers their profit margin (increases their price). These rules are defined analogously for sellers.

Traders adjust their profit margin,  $m$ , using the Widrow–Hoff learning rule with momentum (Widrow and Hoff 1960). The traders the learning rule is defined at time  $t$  as

$$m(t + 1) = (p(t) + F(t))/(\lambda - 1)$$

Here  $p(t)$  is the trader's valuation as defined above and  $F(t)$  is defined as

$$F(t + 1) = \gamma F(t) + (1 - \gamma)\delta(t)$$

where

$$\delta(t) = \beta(\tau(t) - p(t))$$

$\gamma$  is the momentum term and  $\tau(t)$  is the target or desired output price and is calculated as

$$\tau(t) = R(t)q(t) + A(t)$$

Here  $R(t)$  is a small relative perturbation drawn from a uniform distribution [1.0, 1.05] for buyers and [0.95, 1.0] for sellers,  $A(t)$  is a small absolute perturbation drawn from a uniform distribution [0.0, 0.05] for buyers and [-0.05, 0.0] for sellers, and  $q(t)$  is the shouted price.

### 3.4 Parameter Optimisation

Each traders parameter values are tailored to its position within a network. The initial behaviour of a trader using the ZIP learning rule is governed by three real-valued parameters, learning rate, momentum and initial profit margin, which may be expressed as a vector  $V$ :  $V = [\beta, \gamma, m]$ . While Cliff (2003) employs a more sophisticated genetic algorithm to search the strategy space, this model significantly simplifies the trader strategy space allowing investigation with a simpler hill-climbing search algorithm.

In order to allow parameter values to adapt to different degrees of trader connectedness, separate populations were maintained for each location within the market. For a market of  $N$  traders,  $N$  populations of size four were maintained. Members of the first population were assessed in the most well-connected market location

(i.e., the network node with highest degree), while members of the  $N$ th population were always assessed in the least well-connected market location (i.e., the network node with lowest degree). In general, each of the  $N$  populations was associated with a specific rank of network node degree.

In each trial a market network was generated as described in the previous section, and one member of each population was chosen at random and added to the network in the appropriate place. Every member of each population participated in 20 trials each generation. Due to the random allocation of reservation prices it was necessary to assess trader strategy fitness multiple times in order to obtain a meaningful estimation. Strategy fitness was calculated as the average profit extracted by a strategy over all trials in that generation. After each set of trials the fittest individual in each population was selected to be copied into the next generation. The remaining three places in each population were then filled by taking this fittest individual and mutating each of its three parameters to produce three new strategies each differing from the original in one parameter. Mutation consisted of perturbing one parameter 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.

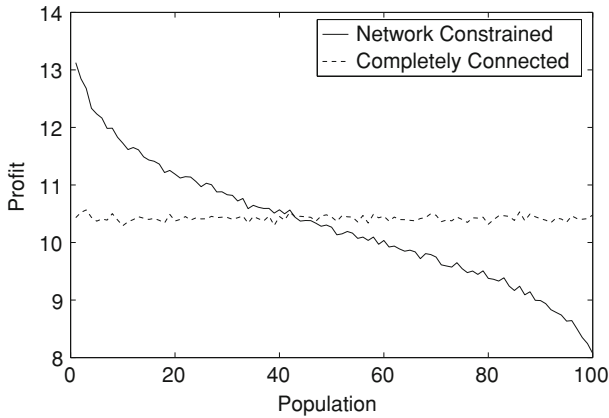
## 4 Results

Experiments were performed using markets populated by 100 traders. Each trader was randomly allocated a reservation price in the range  $[0, 100]$  ( $p_{\min} = 0$  and  $p_{\max} = 100$ ), 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. During each simulation run, both trade and information flow were constrained by a network constructed as described above with  $E = 1.0$  and  $c = 10$ . Selection occurred over 1,000 generations with 100 populations each of size 4. We contrast these result with those obtained under the same conditions for a market in which trade and information flow are not constrained.

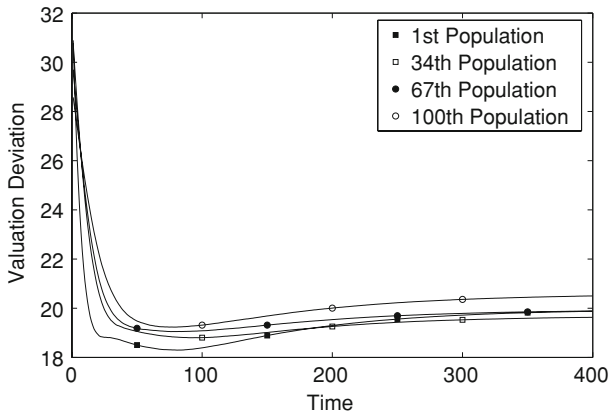
Figure 1 shows the average profit (fitness) of each population in the final generation, ranked in order of decreasing connectivity. As expected, all traders within the fully connected market make approximately the same amount of profit. By contrast, this is not the case in the segregated market. In line with the results of [Wilhite \(2001\)](#), the better connected traders make significantly more profit. The segregated markets are far more sparsely connected than a fully connected market, with their better-connected traders having at most the same number of connections as every trader within the fully connected market. However, since these traders are *relatively* well-connected compared to their competitors, they are able to leverage their market position to find better trading opportunities and so make more advantageous trades.

Figure 2 demonstrates that, over time, traders' valuations approach the equilibrium price (as calculated from the reservation prices of the traders), regardless of their connectivity. The traders do not all converge at the same rate, however. Those with most connections converge fastest. Traders who are better connected receive information more frequently and so have a better impression of the state of the market and are better able to accurately judge the equilibrium price. Consequently the better connected





**Fig. 1** Profit obtained by traders in a segregated market (traders ordered by decreasing connectivity) and a fully connected market



**Fig. 2** Absolute deviation from equilibrium price over the lifetime of a market, averaged over 10,000 markets, for populations of traders ranked in order of decreasing connectivity (i.e., the first rank corresponds to the population of strategies optimised to occupy the highest degree node in the trade network; the hundredth rank corresponds to the population optimised to occupy the lowest degree node in the trade network)

traders are able to make larger profits at the start of the simulation by exploiting their greater number of connections and more accurate valuation of the commodity. As time progresses all traders converge towards the equilibrium price and the initial advantage of the better connected traders in making larger profits is reduced.

Importantly, the final deviation from the equilibrium price is not zero. There are several reasons for this. Firstly, traders are allocated units which are untradable at the equilibrium price on approximately 50% of occasions, i.e., half of the buy reservation prices are below equilibrium and half of the sell prices are above. Consequently traders issued with these prices will never achieve the equilibrium price. When trade is segregated by a trade network this situation is exacerbated. Some traders will be unable to trade

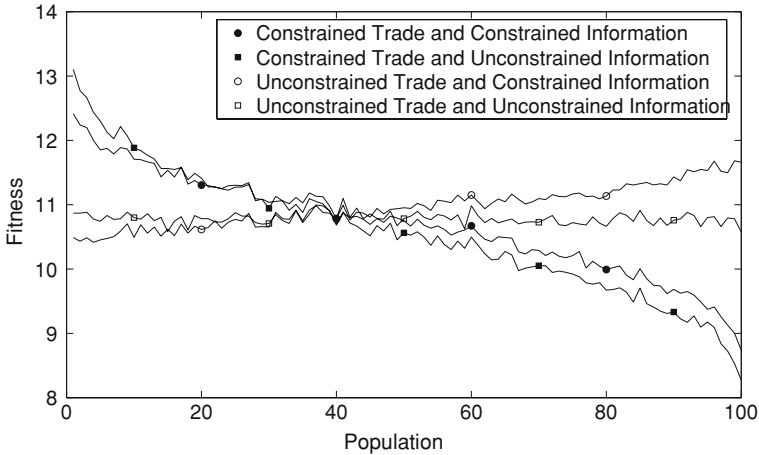
despite the presence of willing partners in the market as a whole, because they will not have a *neighbour* prepared to trade with them. As a result, assuming all traders converge towards the equilibrium price up to the point at which they expect to make no profit, the half of the traders with tradable units will have a deviation of zero while the remaining half will have a deviation of  $\frac{1}{2}(p_{\max} - (p_{\max} + p_{\min})/2) = \frac{1}{2}(100 - 50) = 25$ , and so the average deviation across both groups of traders is  $(0 + 25)/2$ . Subject to the above assumption 12.5 points of the deviation is due to this effect and should be considered the baseline. The remainder of the deviation is due to traders trading early and no longer modifying their valuation and the informational constraints of the market, i.e., if all of a trader's neighbours have traded or if the trader only has neighbours who are unable to reach the equilibrium price they may receive inaccurate information resulting in them being unable to identify the equilibrium price.

## 5 Distinguishing Trade from Information Flow

In order to understand the effect of the network on trade and the importance of information flow in a segregated market, the effects of trade and information flow across a market were separated. Four scenarios were investigated and compared: (i) fully-connected trade network with fully-connected information network; (ii) segregated trade network with fully-connected information network; (iii) fully-connected trade network with segregated information network; and (iv) trade and information flow segregated by the same network. In cases where a network is described as fully-connected all traders are connected to all others and the specified type(s) of interaction can occur between all pairs of individuals—the market is effectively centralised. In the cases where the market is described as segregated, a network is generated via preferential attachment as described above and the specified type(s) of interaction can only occur between traders that are network neighbours.

The first case, both aspects completely connected, is the equivalent of a small, centralised market where all traders can communicate and trade with each other and the amount of information is sufficiently small that all traders can process it effectively. This case also serves as the baseline case against which the others may be compared. The second case is similar to some financial markets where it is possible to gather all of the information about all trades that occur, yet it is only possible to trade with some of the traders present. For instance, an individual investor may only be able to trade with a small subset of brokers who are willing to deal with small customers. The third case has some similarities with the football transfer market. Theoretically any club can deal with any other yet the details of trades between other clubs may not be known. The fourth case is similar to a spatially embedded market such as a retail market where every trader can only communicate and trade with a subset of the other traders within the market. In the case of a retail market this restriction is due to geography and the size of the market. The details of the effects of the information and trade networks are described in the next two sections.

Figure 3 shows the average fitness of individuals in the final generation averaged over 100 experimental runs. The effect of both the information and trade networks may be clearly seen. The presence or absence of the trade network has a clear and

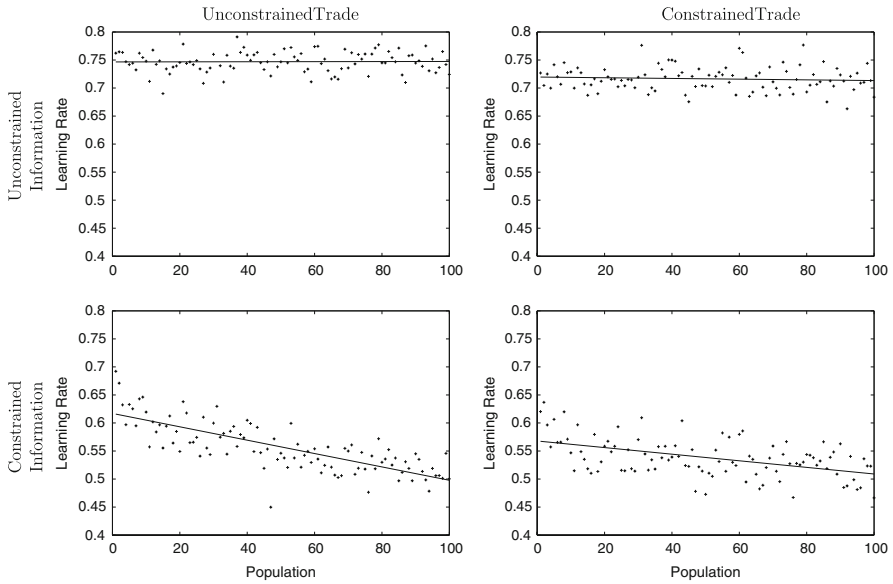


**Fig. 3** Average fitness of traders in the final generation averaged over 100 runs for various market conditions

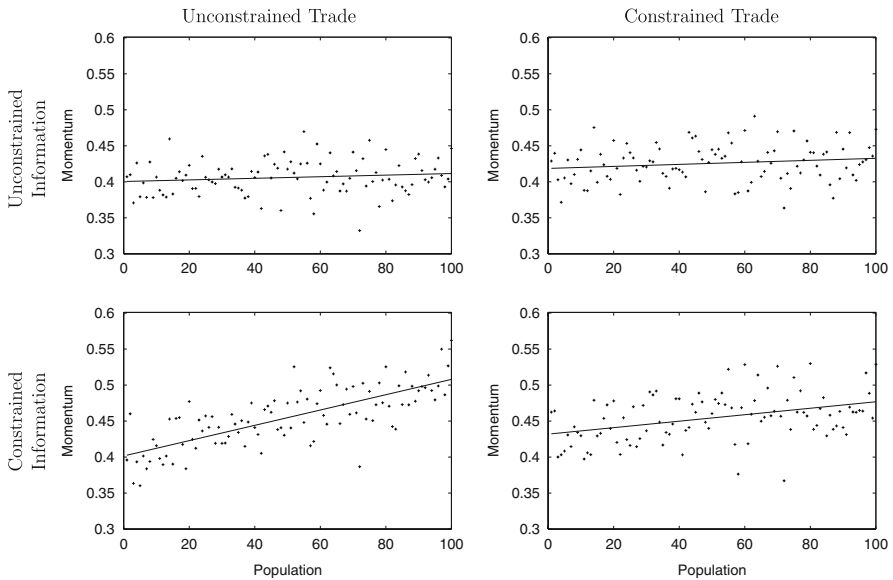
intuitive effect on trader fitness. When trade is constrained to take place across the network, those traders who are best connected are substantially fitter. This is primarily because the more connections a trader has the more likely they are to be able to find a suitable trade partner with which to make an advantageous trade.

It can be seen that constraining information leads to an increase in fitness for those less well connected individuals and a decrease in fitness for the most well connected individuals. It is easiest to understand this result when considered together with those for the momentum and learning terms which control traders' reactions to information. The learning rate and momentum terms associated with the trading strategies are shown in Figs. 4 and 5. It can be seen that when information flow is unconstrained, whether or not trade is constrained, both the learning rate and momentum terms are approximately constant across the optimised populations. In contrast, when information flow is restricted by the network there are obvious trends: momentum values are negatively correlated with trader connectivity, while learning rate values become positively correlated with connectivity. In both cases, these correlations are stronger when trade is unconstrained.

The learning rate and momentum parameters control the learning process of each individual trader. When information flow is unconstrained and consequently all individuals receive the same information it is perhaps not unexpected that all traders deal with this situation in the same way. Understanding the reason for the correlations between learning parameters and connectivity when information is constrained is slightly more complex. First, it is important to remember that the "accuracy" of information changes throughout the life of the market. As was shown in Fig. 2, initially valuations and therefore shouted prices are far from the equilibrium valuation. However, this situation changes rapidly as valuations quickly become more accurate. By definition, when information flow is constrained by a network, those traders who are better connected receive more information than those who are less connected. Well-connected traders initially receive large amounts of information the content of



**Fig. 4** Learning rates of traders in the final generation of each of 100 populations averaged over 100 runs for four market conditions. Populations are ordered from most to least connected



**Fig. 5** Momentum values of traders in the final generation of each of 100 generations averaged over 100 runs for four market conditions. Populations are ordered from most to least connected

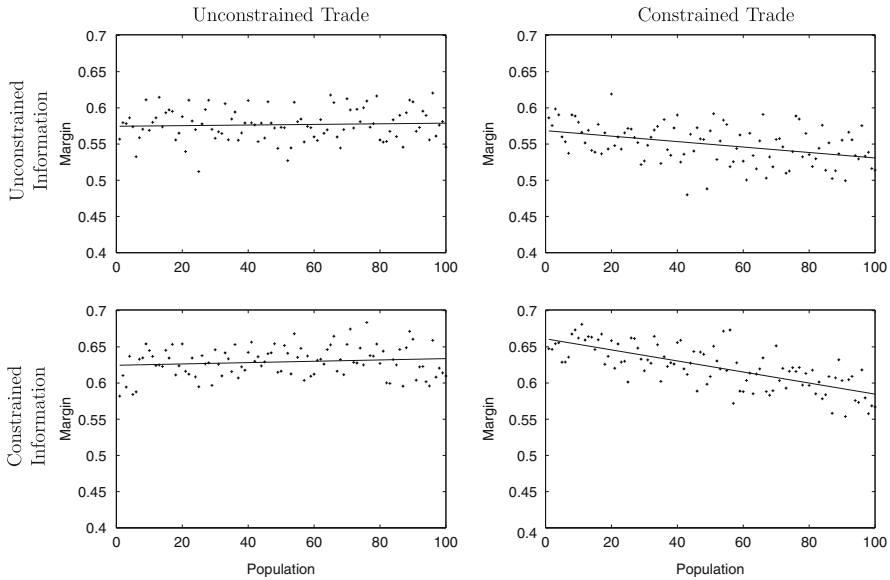
which is rapidly changing. Therefore, it is important that these traders are able to respond and quickly adapt to this changing content in order to allow them to benefit from advantageous prices and avoid being seduced by poor trades due to out of date information. Consequently these traders require a high learning rate to allow them to keep abreast of the market situation and be well suited to it at all points in time. In markets where information is unconstrained, all traders are forced to have high learning rates to respond to this rapid influx of information. Similarly these traders have a low momentum so that they focus on the most recent information and pay less attention to events in the past as the market state may have changed.

Less well connected traders receive information more slowly and consequently are able to adopt a different strategy. Rather than trying to keep abreast of every market action they instead attempt to gain a picture of the market by aggregating information over time. Because these traders receive information less frequently and adopt a relatively low learning rate, they are able to weather much of the initial period of high volatility in the market. The relatively high momentum value allows them to gain an accurate picture of the competitive price by averaging over the more stable mid and late-market prices whilst allowing the trader to ignore misleading shouts. This strategy is successful as evidenced by Fig. 3. When trade is unconstrained, any mistakes in valuation are penalised quickly due to all traders having more potential partners. In order to compensate for this more radical strategies are adopted resulting in the increased trend in the learning and momentum parameters. Segregating information flow confers a relative advantage on less well connected traders (note that this is only a relative advantage as there is only a fixed amount of profit available within the market and so an increase in profit for one group necessitates a decrease in profit for another). This rather counter intuitive result indicates that without the ability to exploit a greater number of connections in order to make more trades, the challenge of processing a greater volume of possibly misleading information may make it preferable to concentrate on fewer information sources.

Figure 6 shows the average initial margin of traders in the experiments. It can be seen from these graphs that when trade is constrained by the network, well-connected traders tend to have a higher initial margin than the less well-connected traders. This means that at the start of each experiment traders in such markets find it advantageous to ask a relatively high price for their commodity. Figure 3 indicates that they achieve increased profits from this strategy. In effect a good market position allows traders to adopt more aggressive market strategies. The presence or absence of the information network affects all traders equally. When information is constrained traders have a higher initial margin. This is because information takes longer to propagate within the market and so it takes longer for a trader to arrive at an accurate valuation of the commodity. As a result, traders initially demand more in order to offset the risk of, and avoid prematurely entering, a disadvantageous trade.

## 6 Exploiting Neighbourhood Information

It can be seen from the previous section that both the information network and trade network provide advantages to those who can exploit knowledge of their own connectivity.



**Fig. 6** Initial profit margins of traders in the final generation of each of 100 populations averaged over 100 runs for various market conditions. Populations are ordered from most to least connected

However, until now traders have made no distinction between the information they receive from different neighbours. Here, by modifying the standard learning rule, we allow traders to act based upon both their position within the network(s) and that of their neighbours.

There are at least two ways in which traders might distinguish between different kinds of neighbour. A trader might decide to pay more attention to information gleaned from certain neighbours (perhaps the more well connected ones). Alternatively better connected traders may attempt to exploit their connectivity advantage by demanding more favourable prices from less well connected traders who have less choice.

In order for traders to vary their strategy based on the connectivity of their neighbours the function  $g(s, r)$  was defined, where  $s$  and  $r$  are the sender and recipient of a shout:

$$g(s, r) = \begin{cases} \frac{\log \frac{E(s)}{E(r)}}{\log(R_{\max})} & : E(s) \geq E(r) \\ -\frac{\log \frac{E(r)}{E(s)}}{\log(R_{\max})} & : E(s) < E(r) \end{cases}$$

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.

In order for traders to adapt the weight they place on information from different sources the Widrow–Hoff learning rule was modified so that the learning rate  $\beta$  was dependant on trader connectivity.

$$\delta(t) = \begin{cases} (\beta + \alpha(1 - \beta)g(s, r))(\tau(t) - p(t)) & : E(s) \geq E(r) \\ (\beta + \alpha\beta g(s, r))(\tau(t) - p(t)) & : E(s) < E(r) \end{cases}$$

Within this function, the parameter  $\alpha$  allows control of how much importance the trading strategy places on the topological rule. If  $\alpha$  is set equal to zero the topological information is ignored and the unmodified learning rule is employed.  $\alpha$  may range from  $-1$  to  $1$ . Negative values weight information from less well connected traders more heavily, positive values do the converse. The terms  $\beta$  and  $(1 - \beta)$  directly before the function  $g(s, r)$  scales the changes in the learning rule so that the learning rate remains in the range  $[0, 1]$ .

Traders may also adapt the profit margin they demand from other traders.<sup>5</sup> For standard traders, the price  $p$  at which the trader values the asset is defined as  $p = \lambda - m(p_{\max} - p_{\min})$  for buyers and  $p = \lambda + m(p_{\max} - p_{\min})$  for sellers, where  $\lambda$  is the trader's reservation price.

We modify this situation with the introduction of a parameter,  $\theta$ , and the function  $g(s, r)$ , where  $s$  and  $r$  are the sender and recipient of a shout.

For buyers:

$$p = \lambda - m(1 + \theta g(s, r))(p_{\max} - p_{\min})$$

For sellers:

$$p = \lambda + m(1 + \theta g(s, r))(p_{\max} - p_{\min})$$

Like  $\alpha$ ,  $\theta$  is allowed to vary between  $-1$  and  $1$ , where negative values lead to higher margins being demanded from less well connected traders and positive values lead to higher margins being demanded from better connected traders, a value of zero leads to the rule having no effect.

## 6.1 Results

We wish to compare the effect of these two rules on the behaviour of traders. First, does the possession of either rule confer an advantage on a single trader located within a market populated by standard ZIP traders? Answering such a question requires care since a trader with variable learning rate, say, might enjoy an advantage in such a market merely by virtue of achieving a learning rate that is high with respect to all individuals within the market. Such an advantage would have been achieved without exploiting any topological information about the trader's neighbours. In order to control for this potential artifact, initially we compare the performance of the two topological rules in a rather artificial market constructed such that the modified trader being assessed has learning rates and margins that are guaranteed, on average, to equal those of unmodified

<sup>5</sup> In order to preserve market structure it is not used when traders are making a shout as shouts are made publicly to all of the traders neighbours (potentially of different connectivities) and therefore must be of one value.

traders. Here we report the market structure without comment. See Appendix A for a full account.

Two equal-sized groups of 51 standard traders are constructed. The first group form a fully connected clique (i.e., each trader has 50 neighbours). The second group forms a minimally connected ring (i.e., each trader has two neighbours). The modified trader possessing one of the two topological learning rules is added to the market by assigning five connections to random agents within each of the two groups (i.e., the trader has degree 10). Like any other trader, this agent adapts its valuation based on the information it hears and will accept favourable offers, but it will never make any shouts. In this way, we can measure the traders profit and its ability to estimate the market's equilibrium price.

In each of 10,000 market runs, traders were initialised with a random profit margin drawn from a uniform distribution [0.2, 0.8], a learning rate of 0.5 as discussed in Appendix A, a random momentum value drawn from a uniform distribution [0.2, 0.8]. These parameters were chosen to allow the traders a wide range of possible behaviours and are based on those used in previous studies (Cliff and Bruten 1997).<sup>6</sup>

Experiments on a range of values of  $\alpha$  show that for positive values of the weighting factor traders make more profit than those not using topological information. By weighting information from the better connected traders more highly than that received from less well connected traders the monitored trader is able to make higher profits. The reason for this is shown in Fig. 7. As can be seen, the valuation of a trader using the modified topological learning rule converges faster and to a significantly lower asymptotic value than the trader using the standard rule (*t*-test,  $p < 0.0001$ ). Hence, the topological learning rule has a positive effect on traders' understanding of the market conditions and hence their valuation of the asset.<sup>7</sup> This result is promising, since it indicates that weighting information according to its source may increase a trader's ability to identify the correct value of the commodity and make increased profits.


Similarly, a complementary set of experiments demonstrate that for small negative values ( $-0.5 \leq \theta \leq 0$ ) of the weighting parameter  $\theta$  the margin weighting rule leads to increased profits over a trader not using the rule. The increased performance is achieved for negative values of the weighting parameter, indicating that the trader finds it beneficial to charge higher prices to traders with less connections, in effect exploiting their weak position in the market.

We now extend this examination to consider trade across more realistic preferential attachment networks. Figure 8 illustrates selected traders in markets composed of (i) naive traders (i.e., those with randomised parameter values drawn from the standard distributions) and (ii) those in which the traders' parameters were optimised. Figures report the difference between the performance of a standard trader and a trader in

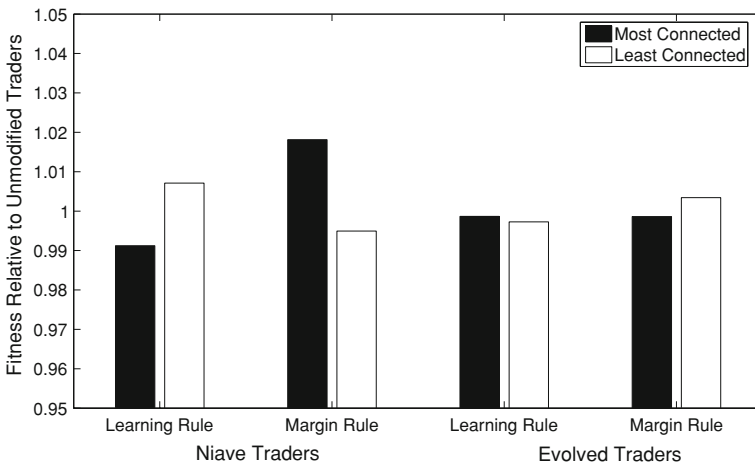
<sup>6</sup> Note that the previously optimised parameters were inappropriate as they were not optimised for this market setup.

<sup>7</sup> The deviation from the equilibrium price can be seen to increase slightly after it reaches its minimum. The reason for this is that towards the end of the experiment the traders with the reservation prices that allow trade have traded leaving only those who have worse prices. Consequently the information from these traders is less informative leading to an increase in deviation from the equilibrium price.





**Fig. 7** Absolute deviation from equilibrium price averaged over 10,000 runs for the monitored traders in markets designed to control for learning rate, for the topological learning rule,  $\alpha = 0.5$



**Fig. 8** Fitness of traders employing the learning or margin rules relative to those not employing them. Fitness calculated from 10,000 repetitions of each condition: best connected fitness averaged over 10 best connected individuals in each market, least connected fitness averaged over fitness of 10 least connected individuals in each market. For learning rule,  $\alpha = 0.5$ ; for margin rule  $\theta = -0.5$

the same market location with the ability to adapt either their learning rate or their profit margin based on market topology information.

From Fig. 8 it can be seen that when pitted against naive trader behaviour the use of both the margin and learning rate adaptation rules lead to increased profits for adaptive traders of either high or low connectivity. In the case of learning rate adaptation, less well connected traders are able to gain better information about the market and so make more profit. When margin adaptation is employed, better connected traders are able to exploit their larger number of connections more effectively to make more profit.

However, Fig. 8 also shows that when all traders in the market are permitted to adapt their parameters to their market location, including whether and how they employ

the adaptive rule being examined, the adaptive rule confers very little advantage. In fact, the optimised values of  $\theta$  and  $\alpha$  are all very close to zero indicating that the rules were not used. This indicates that a trader can exploit knowledge of the relationship between a neighbour's connectivity and their own only if that neighbour is not able to make strategic use of the same information, i.e., the inherent advantage due to market location that was demonstrated in Fig. 3 can neither be exaggerated nor attenuated by attending to topology in order to set learning rates or margins.

The explanation for this is hinted at by Fig. 2. Well-connected traders initially enjoy a more accurate estimate of the market price, but this advantage diminishes over the course of the market. If poorly-connected traders are prepared to trade early, they risk being exploited by their better-connected and better-informed neighbours. By evolving strategies that tend to prevent early trades, poorly-connected traders protect themselves against the volatility in the early period of trading. After this period, all traders will have tended to acquire the same accurate estimation of market price. Since both topological learning rules rely on exploiting *disparity* between well-connected and poorly-connected traders, their utility tends to be extinguished by the cautious optimised trading strategies of less well-connected traders.

## 7 Conclusion

This paper has investigated the effect of market structure on market behaviour and trader strategies. We have demonstrated that the assumption that markets are homogeneous and centralised has a significant effect on market behaviour. We can see that when this assumption is relaxed by imposing a network structure on the market, a trader's position within the market has a noticeable effect on their performance and strategy. Traders who are "centrally located" within the market have a clear advantage over more peripheral traders due to their many trading opportunities. However, they do have to integrate large quantities of potentially inaccurate and misleading information from neighbouring traders in order to gain an accurate picture of the market. This problem is acute. When all traders are free to trade with one another, but some traders are better connected in terms of access to market information, such traders are actually disadvantaged. This result shows clearly that market structure can be important and that assumptions regarding market behaviour derived from completely connected centralised markets do not necessarily hold in distributed settings.

When traders are given the chance to tailor their strategy to their position within the market, we see that it is advantageous for traders to adopt strategies based on their connectivity. For instance, better connected traders tend to be more aggressive and reactionary, while less connected traders tend to be more cautious. The simple topological learning rules for learning rate and profit margin presented in this paper demonstrated that traders are able to make use of market structure information to positively affect their performance in markets composed of naive traders. However, it was also shown that when traders' strategies are tailored to their market location, no trader may gain any benefit from these rules. As with many evolutionary scenarios, the environment in which a trader operates, i.e., the collection of other trading strategies in the market, has a clear impact on strategy choice.

This paper investigated the effect of one market structure on trader strategy. There are potentially many alternative market structures that could be investigated. In this case we examined markets formed through a preferential attachment scheme, however, there are many alternatives with different local or global structures which should be considered. For instance the graph generation algorithm we considered does not encourage clustering as found in social networks. Traders engaging in tripartite negotiations are likely to behave in different ways to those only dealing with bipartite interactions, though in this case the relatively simple traders presented in this paper do not have the capacity to exploit this knowledge. Moreover, here the market structure was fixed. This is rarely the case in reality, where market structure changes over time as traders generate new relationships and end old ones based on their value and profitability. By allowing the networks to be formed endogenously there is scope for further investigating the relationship between individual strategy, network formation and learning. The market model employed here was relatively simple and focused on the initial period of the market's operation until it reaches equilibrium. In particular, every trader only had a single unit of the commodity to buy or sell. The next step is to combine this work with models such as that of [Wilhite \(2001\)](#) to examine the interaction of volume effects and the market power effects. By allowing traders to trade multiple times we expect the advantage of better connected traders to be enhanced as they use their many connections to gain additional benefit. In the scenarios explored here, favourable market position confers an advantage in terms of ability to find a profitable trade but not in terms of quantities of trades. There is also the possibility of investigating the optimality of strategy in a more general framework. The use of genetic programming may be one possible method for this. [Lensberg and Schenk-Hoppé \(2007\)](#) point the way with a paper in which theoretically optimal investment strategies for multiple financial assets are evolved.

## Appendix A

Here we explain how to construct a market topology with which to control for the learning rate or margin of a trader employing a topological learning rule. In the case of the standard traders, the learning rate  $\beta = 0.5$ . When the topological learning rule is employed, it is possible to specify the degree of each of the neighbours of the monitored trader such that its average learning rate,  $\bar{\beta}_n$  is also equal to, e.g., 0.5. Given that the learning rates of traders using the topological learning rule lies within the range  $[0.0, 1.0]$ , that the monitored trader has an equal number of neighbours in each of two groups of standard traders, and that within each group all traders share the same connectivity, we can write an expression for  $\bar{\beta}_n$ .

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

Here  $E_{\max}$  is the connectivity of standard traders in the fully connected group,  $E_{\min}$  is the connectivity of standard traders in the weakly connected ring, and  $E(n)$  is

the number of connections to the monitored trader. For  $\beta = 0.5$  this equation can be rewritten:

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

The above approach may be applied, *mutatis mutandis*, for the margin parameter.

## References

- Axtell, R. (2005). The complexity of exchange. *Economic Journal*, 115(504), F193–F210.
- Barabasi, A.-L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, 286, 509.
- Bell, A. M. (1998). Bilateral trading on a network: A simulation study. In: Working Notes: Artificial Societies and Computational Markets, Autonomous Agents '98 Workshop, Minneapolis, St. Paul, pp. 31–36.
- Brock, W. A., & Hommes, C. H. (1997). A rational route to randomness. *Econometrica*, 65, 1059–1096.
- Cliff, D. (2003). Explorations in evolutionary design of online auction market mechanism. *Journal of Electronic Commerce Research and Application*, 2(2), 162–175.
- Cliff, D., & Bruten, J. (1997). Minimal-intelligence agents for bargaining behaviors in market-based environments. Technical Report HPL-97-91, HP Labs.
- Epstein, J. M., & Axtell, R. (1996). *Growing artificial societies: Social science from the bottom up*. Washington, DC: Brookings Institute Press; Cambridge, MA: MIT Press.
- Evstigneev, I. V., & Taksar, M. (1994). Stochastic equilibria on graphs, I. *Journal of Mathematical Economics*, 23(5), 401–433.
- Evstigneev, I. V., & Taksar, M. (1995). Stochastic equilibria on graphs, II. *Journal of Mathematical Economics*, 24(4), 383–406.
- Evstigneev, I. V., & Taksar, M. (2002). Equilibrium states of random economies with locally interacting agents and solutions to stochastic variational inequalities in  $\langle L_1, L_\infty \rangle$ . *Annals of Operations Research*, 114(1–4), 145–165.
- Falbo, P., & Grassi, R. (2004). Equilibrium prices on a financial graph. *Computational Economics*, 24(2), 117–157.
- Gode, D. K., & Sunder, S. (1997). What makes markets allocationally efficient? *The Quarterly Journal of Economics*, 112(2), 603–630.
- Howitt, P., & Clower, R. (2000). The emergence of economic organization. *Journal of Economic Behavior and Organization*, 41(1), 55–84.
- Kirman, A. P., & Vriend, N. J. (2001). Evolving market structure: An ACE model of price dispersion and loyalty. *Journal of Economic Dynamics and Control*, 25(3–4), 459–502.
- Lensberg, T., & Schenk-Hoppé, K. R. (2007). On the evolution of investment strategies and the Kelly rule – a Darwinian approach. *Review of Finance*, 11, 25–50.
- Lovaglia, M. J., Skvoretz, J., Willer, D., & Markovsky, B. (1995). Negotiated exchange in social networks. *Social Forces*, 74(1), 123–155.
- Markovsky, B., Willer, D., & Patton, T. (1988). Power relations in exchange networks. *American Sociological Review*, 53(2), 220–236.
- Newman, M. E. J. (2003). The structure and function of complex networks. *SIAM Review*, 45(2), 167–256.
- Noble, J., Davy, S., & Franks, D. W. (2004). Effects of the topology of social networks on information transmission. In S. Schaal, A. J. Ijspeert, A. Billard, S. Vijayakumar, J. Hallam, & J.-A. Meyer (Eds.), *Eighth International Conference on Simulation of Adaptive Behavior* (pp. 395–404). Cambridge, MA: MIT Press.
- Papadimitriou, C. H. (1994). On the complexity of the parity argument and other inefficient proofs of existence. *Journal of Computer and System Sciences*, 48(3), 498–532.
- Scarf, H. E. (1973). *Computation of economic equilibrium*. New Haven, CT: Yale University Press (with T. Hansen).
- Smith, V. L. (1962). An experimental study of competitive market behaviour. *Journal of Political Economy*, 70, 111–137.
- Tassier, T., & Menczer, F. (in press). Social network structure, segregation, and equality in a labour market with referral hiring. *Journal of Economic Behaviour and Organisation*.

- Tesfatsion, L. (2001). Structure, behavior, and market power in an evolutionary labor market with adaptive search. *Journal of Economic Dynamics and Control*, 25(3–4), 419–457.
- Vriend, N. J. (2006). ACE models of endogenous interactions. In: L. Tesfatsion & K. L. Judd (Eds.), *Handbook of computational economics*, Vol. 2, Chap. 21 (pp. 1047–1079). Elsevier.
- Watts, D., & Strogatz, S. (1998). Collective dynamics of “small-world” networks. *Nature*, 393, 440–442.
- Widrow, B., & Hoff, M. E. (1960). Adaptive switching circuits. *IRE WESCON Convention Record*, 4, 96–104.
- Wilhite, A. (2001). Bilateral trade and ‘small-world’ networks. *Computational Economics*, 18, 49–64.
- Wilhite, A. (2006). Economic activity on fixed networks. In L. Tesfatsion & K. L. Judd (Eds.), *Handbook of computational economics*, Vol. 2, Chap. 20 (pp. 1013–1045). Elsevier.