



UvA-DARE (Digital Academic Repository)

Essays on markets over random networks and learning in Continuous Double Auctions

van de Leur, M.C.W.

[Link to publication](#)

Citation for published version (APA):

van de Leur, M. C. W. (2014). Essays on markets over random networks and learning in Continuous Double Auctions

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <http://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Chapter 1

Introduction and Thesis Outline

The design of a market and the information that is available before traders make their decision largely influence traders' behaviour and the efficiency of the market. For example the OpenBook system as introduced in 2002 by the New York Stock Exchange, opened the content of the limit order book to the public. This allows for a change in behaviour of traders, who can now condition their strategy on the full history of orders. We study whether a market design with more information, such as the OpenBook system, is preferable in terms of efficiency. More information benefits traders with a high market power and hurts others, but it is unclear whether the total profit in the market and thus efficiency will increase. We consider boundedly rational behaviour of traders and the resulting efficiency depending on the available information in the market design, to study what information should be made available to traders. In the markets examined in this dissertation traders are truthful, or behave boundedly rational. In the first case, traders offer their valuation for the asset or ask their cost, which is in general not rational. In the latter case traders are boundedly rational by only considering linear strategies or by using a learning algorithm that is based on the hypothetical payoff of strategies in the previous period. Boundedly rational behaviour is commonly modelled by putting mild restrictions on the strategy of traders or by learning algorithms. Such algorithms are used in agent-based models of financial markets since they do not impose strict assumptions on the behaviour of traders or their strategy space, and are considered in the second part of this dissertation. An underexposed type of market is a market in which trade occurs over a network, where the network structure is

not entirely known to traders. An example is the spot foreign exchange market which is modelled in the first part of this dissertation by imposing mild restrictions on the strategy function of traders.

1.1 Network theory

Network theory is applicable to many research fields besides pure mathematics. In neuroscience, biological networks of the neural system are considered. In sociology networks are applied for instance to social media and relational connections. A common example in computer science is the use of networks in Google's PageRank and in operational research directed networks are used for transportation problems. In economic theory the banking crisis has led to a large literature on banking networks.

The seminal papers of Erdős and Rényi (1960, 1961) have introduced a mathematical theory on random graphs, often referred to as Erdős-Rényi graphs. We consider vertices in the network as traders and edges as links between traders. In these graphs traders are linked with an equal probability, independently of other connections. Erdős and Rényi derive phase transitions for infinitely many traders. During these phase transitions the structure of the network changes abruptly. The most surprising result of Erdős and Rényi occurs when the expected number of links per trader crosses the threshold value of one half. During this phase transition the structure of the graph changes from a collection of mainly isolated spanning trees to a network that contains a giant component of positive measure. Such a spanning tree connects a subset of traders of the graph but does not contain any cycle. Alon and Spencer (2008), Bollobás (1982) and Janson et al. (2000) summarise the work in the field of random graphs.

Markets over networks have been studied in various settings and trading mechanisms. In these markets trade may only occur between linked traders. The literature has in common that there is full knowledge of the network structure when traders determine their strategy. Spulber (2006) and Kranton and Minehart (2001) consider a market in which sellers jointly raise their ask

until supply equals demand and trade occurs, known as simultaneously ascending-bid auctions. In Corominas-Bosch (2004) and Chatterjee and Dutta (1998) traders submit an offer side by side, which can be accepted or rejected by traders on the other side of the market. It is shown in Corominas-Bosch (2004) that the network can be split into different subgraphs in which the short side of the market extracts all the surplus, when all buyers have the same valuation and sellers the same cost. Intermediaries that act strategically and extract surplus are added in Easley and Kleinberg (2010) and Blume et al. (2009). In a market over a network, the power of a trader is measured in Calvó-Armengol (2001) on the basis of the number of linked traders and their links. The market power of a trader is higher when linked to more traders and when the linked traders have fewer links themselves. Moreover, a branch of network theory in economics and sociology studies the formation of links in a network, starting from Jackson and Wolinsky (1996). However, entirely random graphs are very seldomly studied in economic theory. These random graphs are important since they allow for studies on the effect of information about the network structure that is available to traders.

1.2 Learning algorithms

Learning algorithms are used in economic theory to model boundedly rational behaviour of traders. These algorithms are attractive because they do not make strict assumptions on the behaviour. For instance in reinforcement learning traders may learn to select the optimal strategy without having knowledge of the equilibrium. Genetic algorithms are developed in game theory for cobweb and overlapping generations models. In genetic algorithms every period a new generation of individuals is generated, depending on the fitness or profit of individuals in the previous period. Many agent-based models use learning to avoid making extreme assumptions about the rationality or strategies of traders. For example, the Individual Evolutionary Learning (IEL) algorithm is introduced in Arifovic and Ledyard (2003, 2007) to model the boundedly rational learning behaviour of agents in a Call Market model. In this learning algorithm traders learn to select from a pool of strategies, based on the hypothetical payoffs in the previous period. Moreover, this learning algorithm is used in a Continuous Double

Auction in Anufriev et al. (2013) to compare efficiency under full and no information about the history of others' strategies. Anufriev et al. (2013) also study the GS-environment from Gode and Sunder (1993, 1997) under the assumption that traders have zero intelligence and submit every possible offer with equal probability.

The introduction of the OpenBook system in 2002 by the New York Stock Exchange allows for studies on the effect of the information that is available to traders. This OpenBook system opened the content of the limit order book to the public, which allows experienced traders to use a full history of orders submission, instead of solely knowledge of global market statistics as under the former ClosedBook system. Boehmer et al. (2005) empirically show that this led to a decrease in price volatility and an increase in liquidity. The opening and closing of stock exchanges can be modelled with a Call Market. For such Call Markets Arifovic and Ledyard (2007) analyse experiments and simulations under the IEL algorithm, in which traders select strategies on the basis of their hypothetical performance in the previous period. Under the OpenBook system traders can directly determine the hypothetical performance of a strategy, assuming that other traders would have behaved the same. Under the ClosedBook system however, traders have to make additional assumptions to estimate the hypothetical foregone payoff of selecting another strategy. Arifovic and Ledyard (2007) show that in the OpenBook system agents try to influence the market clearing price. Agents behave as price makers and offers converge towards an equilibrium price. However, in the ClosedBook system traders learn to become pricetakers and offers diverge away from the equilibrium price range.

Anufriev et al. (2013) analyse the effect of the OpenBook system in a Continuous Double Auction. Agents enter the market and trade with an existing agent if possible. Otherwise their offers are stored in the order book until trade occurs with newly arriving traders or the book is emptied. In the IEL-algorithm the same hypothetical payoff functions as in Arifovic and Ledyard (2007) are used to value strategies. Anufriev et al. (2013) find the same bidding behaviour in a Continuous Double Auction as in the latter paper. They conclude that

in the long-run, efficiency is similar in both designs and the price volatility is lower in the OpenBook system. Under this hypothetical payoff function in the formerly used ClosedBook system, where only information about past average prices is available, Anufriev et al. (2013) proved divergence of bids and asks away from the equilibrium price range. This results from the chosen ClosedBook hypothetical foregone payoff function, which only distinguishes between orders below and above the average price of the previous period. As a consequence investors trade with a high probability but may generate a very small profit. Anufriev et al. (2013) state however that "the specification (of the ClosedBook hypothetical foregone payoff function) is a strong assumption ... which may affect (their) results of IEL". Contrary to the latter paper, Fano et al. (2013) use a genetic algorithm in a setting closely related to the ClosedBook system, and show that traders behave as pricemakers and thus offers converge towards the equilibrium price. In this genetic algorithm, traders with the same valuation are compared on the basis of their average profit over some evaluation window, after which individuals with a low average profit take on strategies of better performing agents.

Starting from early contributions it is common in many agent-based models of order-driven financial markets that traders submit their order at a random moment during a trading session. Moreover, they are often required to make a one-dimensional decision, namely to choose a bid or ask price as in LiCalzi and Pellizzari (2006) or to forecast a future price as in Brock and Hommes (1997, 1998). For example, LiCalzi and Pellizzari (2006, 2007) compare efficiency in a Continuous Double Auction with other market protocols such as the Call Market, under boundedly rational respectively zero intelligent agents that arrive in a random sequence. Chiarella and Iori (2002) as well as Yamamoto and LeBaron (2010) use traders that submit their order at a random moment and use simple rules to make predictions about future prices, similar to Brock and Hommes (1998). In the classical financial literature many studies focus on limit and market orders. The surveys Gould et al. (2013a) and Hachmeister (2007) discuss the main theoretical, experimental and empirical papers on limit orders of informed and uninformed traders. Bae et al. (2003) and Biais et al. (1995) empirically find that the number of orders during a day follows a U-shaped distribution. Their reasoning behind this distribution is

that at the beginning of the day traders desire to perform price discovery and react to events that occurred during the closing of the exchange and at the end of the day traders desire to unwind their positions. With the use of learning algorithms such as Individual Evolutionary Learning, agent-based models of financial markets can be extended to allow traders to submit their order at a chosen moment during a period. This requires an extension of the learning algorithm, in which traders are required to make a two-dimensional decision.

1.3 Dissertation outline

This dissertation consists of 4 chapters after this introduction. These chapters consider the effect of the available information in the market design on expected efficiency, in markets over networks when we assume that traders use linear markup strategies and in Continuous Double Auctions when traders use the Individual Evolutionary Learning algorithm to select their strategy. The first two chapters study efficiency in markets over random networks; in infinitely large markets when we assume that traders behave truthfully and in thin markets under boundedly rational behaviour of traders. The last two chapters consider the Individual Evolutionary Learning algorithm in Continuous Double Auctions. We introduce a new hypothetical foregone payoff function under no information about the history of others' actions and moreover extend the model by requiring traders to make the additional decision of choosing the timing of order submission.

Efficiency in Large Markets over Random Erdős-Rényi Networks

Chapter 2 follows Erdős and Rényi and derives phase transitions of bipartite graphs, depending on the probability of a link. Links are realised with the same probability, independently of each other. We find a similar transition of the bipartite graph, when the expected number of links per trader crosses the value one: the graph consists of many small isolated spanning trees below the threshold value and contains a giant component after the threshold. A market over random bipartite graphs with infinitely many traders is considered in the second part of this chapter. Agents desire to trade one unit and we assume that every trade yields the same surplus.

We study the restrictions of the network on the maximal efficiency, which can be calculated as the maximal expected number of trades divided by the number of traders on the thin side of the market, under identical valuations and costs of traders. The problem of finding the maximal number of trades is known as the Maximum Matching problem, studied for instance in Mucha and Sankowski (2004) and in West (1999). We derive bounds on expected efficiency as a function of the probability of a link, and improve these bounds for the range where the graph contains mainly spanning trees. An algorithm is introduced to construct all spanning trees and we determine the distribution of the degree of the vertices in a spanning tree.

Information and Efficiency in Thin Markets over Random Networks

A thin Erdős-Rényi market with two buyers and two sellers is considered in Chapter 3. Similar to the model of the spot foreign exchange market studied in Gould et al. (2013a), trades occur over links of the network. In contrast to their model we assume that links are realised with the same probability and independently of each other. Traders receive information about the network structure and behave strategically. We compare the equilibrium configurations for three nested information sets about the network structure; no, partial and full information. Under no information traders do not receive information about the realisation of links, but only the probability that a link is realised. The existence of one's links is given under partial information, as well as the probability of links of other traders. Under full information the entire network structure is revealed. We consider the effect of the amount of information on the allocative efficiency. This work shows that this effect is not only *non-monotonic*, but that a reversal of this non-monotonicity occurs when we switch from complete to incomplete information about traders' valuations. Contrary to Corominas-Bosch (2004), we show that under partial information about the network structure, or under incomplete information about valuations and costs, not all the surplus is necessarily extracted. Under complete information about valuations and costs, partial information about the network structure is weakly dominated. Under incomplete information about valuations and costs, we restrict attention to linear markup and markdown strategies. This type of strategies is introduced in Zhan and Friedman (2007) and a symmetric version is derived in Cervone et al. (2009). Myerson and Satterthwaite (1983) and

Chatterjee and Samuelson (1983) show for bilateral trading that Nash equilibrium strategies are monotone and piecewise linear transformations of valuations into offers. For the subset of linear markup strategies, partial information about the network structure strongly dominates full and no information.

On the role of Information under Individual Evolutionary Learning in a Continuous Double Auction

In Chapter 4 we demonstrate through simulations that the specification of the hypothetical foregone payoff functions indeed plays a crucial role in a Continuous Double Auction model under the IEL learning algorithm, as suggested by Anufriev et al. (2013). Traders use the payoff function to estimate how other strategies would have performed in the previous period. Under their hypothetical foregone payoff function bids and asks diverge away from the equilibrium price range in the ClosedBook system. This work, jointly with Mikhail Anufriev, Jasmina Arifovic and Valentyn Panchenko, introduces a new foregone payoff function, that uses more information to estimate the hypothetical foregone payoff of each possible offer, which results in bids and offers drifting towards an equilibrium price similar to Fano et al. (2013). Under this payoff function investors learn to increase their expected profit by submitting an order that has a higher possible profit. This results in a lower probability of trading, but this effect is outweighed by an increase in possible profit from trade. First we perform simulations during the learning phase of a Continuous Double Auction, to study the effect of the OpenBook system. We compare with the results of the simulations in the Call Market performed by Arifovic and Ledyard (2007), by comparing efficiency between both markets. Second, we examine the effect of the OpenBook system during long-run simulations. This allows for a comparison of the new ClosedBook hypothetical foregone payoff function with the function used in Anufriev et al. (2013). Thirdly we show robustness of our results with respect to the size of the market and the number of units a trader desires to buy or sell. As indicated in Anufriev et al. (2013) the specification of the hypothetical foregone payoff function indeed plays a crucial role and largely affects their main results.

Timing under Individual Evolutionary Learning in a Continuous Double Auction

In Chapter 5 we extend the IEL-algorithm used in Arifovic and Ledyard (2003, 2007) and in Anufriev et al. (2013) by introducing learning about the timing of order submission. In this joint work with Mikhail Anufriev, traders submit a multidimensional strategy which allows for contemporaneous learning about the submitted order and the moment of submission. In a benchmark environment with complete information about the trading history in the previous period, we study the distribution of submission moments under the extended IEL algorithm and the interrelation between the submission moments and the orders. This chapter is a step forward to a more complete model of learning in markets and is distinguished from previous research by the decision traders are required to make. Instead of a one-dimensional decision traders are required to make a two-dimensional decision; which bid or ask to submit and when to submit this offer during the trading session. We show that traders in medium size markets learn to submit around the middle of the trading session to avoid a lower profit or trading probability. Moreover, we consider the impact of competition and the size of the market on the timing of the submission. We conclude that the size of the market highly influences the preferred arrival moment. We show that the effect of the extra decision that traders are required to make is negative, by comparing general market statistics with Anufriev et al. (2013), where traders submit at a random moment during the trading period.