



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: <https://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 4

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

4.1 Introduction

The effect of the setup of markets and the information available to traders has been studied intensively on allocative efficiency. Common examples of studied markets are the Call Market and the Continuous Double Auction (CDA). In these markets, experiments as in Cason and Friedman (1996) showed that subjects do not learn to optimise in a Bayesian sense, and hence are not fully rational. However, this paper showed that still a fast converge towards the equilibrium occurs, and hence the efficiency moves towards one quickly. This led to a large literature on simulations to model the boundedly rational behaviour of traders. In these models traders select their next strategy on the basis of the past trading history (Brock and Hommes, 1997, 1998), or by imitating other past strategies (Dawid, 1999). This branch of research distinguishes itself from the standard economic research in which traders are forward looking and hence select a rational equilibrium. One approach to model boundedly rational behaviour is to use learning algorithms, as this does not impose any strict assumptions on the behaviour of traders. The literature uses learning algorithms under full information about

trading history, to describe the boundedly behaviour of traders and to show that this may still lead to full efficiency.

In this chapter we study the impact of the information about trading history that is available to traders. In 2002 the New York Stock Exchange introduced the OpenBook system, which opened the content of the limit order book to the public. This allowed experienced traders to use a full history of orders submission, instead of solely knowledge of global market statistics as under the former ClosedBook system. This decision resulted in decreased price volatility and increased liquidity in the market as shown in Boehmer et al. (2005). A number of papers describe the effect of this OpenBook system in simplified markets. In a Continuous Double Auction model, Ladley and Pellizzari (2014) show that the information of the OpenBook is useless. We use the Individual Evolutionary Learning (IEL) algorithm, introduced by Arifovic and Ledyard (2003, 2007) to model the boundedly rational learning behaviour of agents in a Call Market model, in a multi-period Continuous Double Auction that models the common stock exchanges. This individual learning algorithm allows traders to select their strategy depending on the, hypothetical, performance in the previous period. Arifovic and Ledyard (2004) show that this algorithm performs better than other learning rules, and that it suits well in environments with continuous or large strategy spaces. Under this OpenBook system traders can directly determine the hypothetical performance of a strategy, assuming that 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.

The effect of the OpenBook system is analysed in a simple Call Market by Arifovic and Ledyard (2007). In a Call Market, at the end of each period a market clearing price is computed, at which price agents can trade. Their paper shows that this IEL-algorithm captures the behaviour of subjects in experiments during the learning phase. In the OpenBook system agents can influence the market clearing price, and the offers converge towards an equilibrium value. In the former ClosedBook system it becomes more difficult to influence the mar-

ket clearing price and agents become pricetakers. The offers converge to the private valuations or costs of the agents. Both in experiments and simulations, Arifovic and Ledyard (2007) show that efficiency is higher in the ClosedBook system.

Anufriev et al. (2013) analyse the OpenBook system in a Continuous Double Auction. Agents place their offer when they enter the market and if possible trade with an existing agent. Otherwise the order is stored in the order book. Agents learn by using the IEL-algorithm, where the same hypothetical payoff functions as in Arifovic and Ledyard (2007) are used to value strategies based on their hypothetical performance in the previous period. Anufriev et al. (2013) find the same bidding behaviour in Open- and ClosedBook as in the latter paper. In the long-run, efficiency is similar in both designs and the price volatility is lower in the OpenBook system. In the formerly used ClosedBook system, where only information about past average prices is available, they proved divergence of bids and asks away from the equilibrium price range. This is the consequence of their choice for a payoff function that only distinguishes between offers below and above the average price of the previous period. As a result, investors trade with a high probability but may generate a very small profit. However, the latter paper states that "the specification (of the ClosedBook hypothetical foregone payoff function) is a strong assumption ... which may affect results of IEL".

The paper Fano et al. (2013) compares the Call Market and the Continuous Double Auction in a setting closely related to the ClosedBook system. The strategies of traders emerge over time by a 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. Thus only information contained in the ClosedBook system is used. Similar to Arifovic and Ledyard (2007) they find that traders in a Call Market become pricetakers and offer their valuation of cost. Contrary to Anufriev et al. (2013), in a Continuous Double Auction traders become pricemakers and offers converge towards the equilibrium price.

In this chapter we demonstrate in simulations that the specification of the hypothetical payoff functions indeed plays a crucial role in a Continuous Double Auction model under the IEL learning algorithm. We show, that when agents use more information to estimate the hypothetical foregone payoff of each possible offer, bids and asks tend to drift towards each other. In this approach investors learn to increase their expected profit by submitting an order that has a higher possible profit. The probability of trading will be lower in this situation but is outweighed by the increase in possible profit. Similarly to Fano et al. (2013), in this setting bids and asks will not converge to the redemption value, but to some equilibrium price.

In the learning phase of the Continuous Double Auction, we examine the effect of the Open-Book system in simulations. Furthermore we compare these results with the simulations in the Call Market performed by Arifovic and Ledyard (2007). Hence we study whether the comparison of efficiency between Open- and ClosedBook differs between both markets.

The results of the long-run simulations allow us to compare the effect of the difference between our ClosedBook hypothetical payoff function and the function used in Anufriev et al. (2013). We study whether the differences between Open- and ClosedBook, as indicated in Anufriev et al. (2013) still hold under the new ClosedBook hypothetical foregone payoff function. Moreover, 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.

This chapter is organised as follows. The Call Market and Continuous Double Auction models are described in Section 4.2. The renewed Individual Evolutionary Learning algorithm is explained in Section 4.3. The setup of the simulations as well as the parameters and methods used are described in Section 4.4. In Section 4.5 the simulations in the learning phase are compared with the Call Market results of Arifovic and Ledyard (2007). We compare the results in the long-run that are described in Section 4.6 with Anufriev et al. (2013). Robustness with respect to the number of units traded is studied in Section 4.7 and with respect to the number of traders in Section 4.8. Concluding remarks are given in Section 4.9.

4.2 Market setup

We describe the environments and the competitive equilibrium used to compare OpenBook (OP) and ClosedBook (CB). The Call Market and the Continuous Double Auction (CDA) are explained together with the benchmark environments we use.

4.2.1 The environments

Each environment is determined by a set of buyers and a set of sellers with their valuations and costs for the good, also referred to as redemption values. In each trading period $t \in \{1, \dots, T\}$ each of the buyers $b \in \{1, \dots, B\}$ desires to consume one unit of the good and each of the sellers $s \in \{1, \dots, S\}$ is endowed with one unit of the good. The buyers have the same valuation of V_b per unit in every period, sellers have fixed costs of C_s that only need to be paid when a transaction occurs. Agents have knowledge of their own redemption value, but not of the values or distribution of the other agents.

The environments are shown as a vector of valuations and a vector of costs. For example the environment $\{[0.9, 0.9], [0, 0.2, 0.4]\}$ consists of two buyers with valuation 0.9 and three sellers with costs of 0, 0.2 and 0.4. We mainly use two environments from Arifovic and Ledyard (2007) and Anufriev et al. (2013) in our analysis, in order to compare with these papers.

The demand and supply functions are determined from the valuations and costs of traders. The equilibrium quantity is denoted as q^* and the interval of equilibrium prices is given by $[p_L^*, p_H^*]$. The traders that can gain a positive profit in equilibrium are called intramarginal, whereas the traders that in equilibrium cannot make a positive profit and therefore will not trade are called extramarginal. The payoff of a buyer equals $U_b(p) = V_b - p$ if he traded at price p and zero otherwise. The payoff of a seller equals $U_s(p) = p - C_s$ after a trade at price p and zero otherwise.

The allocative value of a trading period is the sum of the payoffs of all agents. The allocative efficiency is defined as the ratio between the allocative value in a trading period and the maximal

allocative value. The market is fully efficient when during a period all intramarginal agents trade. The efficiency of a period can be lower when an extramarginal agent trades, or when intramarginals simply do not trade. Furthermore we study the average price, the price volatility and the number of transactions.

4.2.2 Call Market

A Call Market is often used to determine the opening- and closing prices at stock exchanges. In period t each buyer b submits a bid $b_{b,t}$ and each seller s submits an ask $a_{s,t}$. At the end of the period all offers are collected and a market clearing price is computed. This market clearing price is the midpoint of the range where supply equals demand. Each buyer that submitted a bid above this value, and each seller with an ask below this value will trade at the market clearing price.

In the ClosedBook system the market clearing price is publicly available after the period. Moreover, the own offer is known and thus traders know their own profit and own trading history. The market clearing price does not reveal the entire sequence of orders. Hence the limited information does not allow agents to determine precisely how they can influence their transaction price. In a multi-period setting all offers are converging towards the redemption values, since this increases the probability of trading and in principle does not influence the market clearing price. In this system, Arifovic and Ledyard (2007) show that agents behave as pricetakers under the Individual Evolutionary Learning algorithm.

In the OpenBook system not only the market clearing price is known, but also all offers become publicly available after the period. Full information is available and therefore agents can exactly calculate how they can influence the market clearing price, assuming that other agents submit the same offers. In a multi-period setting, offers will not converge towards the redemption values but to some equilibrium price and Arifovic and Ledyard (2007) conclude that agents behave as pricemakers.

4.2.3 Continuous Double Auction

A Continuous Double Auction model is used to describe the regular behaviour during the trading day at stock exchanges. In this model, buyers and sellers arrive in a random sequence during a trading period and submit their bid. Agents can select their order only once, before the period and hence unconditional on the state of the order book. The bid of buyer b and the ask of seller s in period t are denoted as $b_{b,t}$ and $a_{s,t}$. If an arriving order can be matched with an order from the book according to the price-time priority, the transaction takes place at the price of the order in the book and both orders are removed. If the arriving order cannot be matched, it is stored in the order book. At the end of the period the order book with unmatched orders is cleared. The latter is a strict assumption that is often made in Continuous Double Auction models as Anufriev et al. (2013). However we show that our results are robust with respect to larger markets, in which case this assumption is less important.

The prices in period t at which buyer b and seller s trade are denoted as $p_{b,t}$ and $p_{a,t}$. The payoff of a buyer equals $U_b(p) = V_b - p_{b,t}$ if he trades and zero otherwise. The payoff of a seller equals $U_s(p) = p_{s,t} - C_s$ after a trade and zero otherwise. We note that the payoff depends not only on the own offer, but also on the arrival sequence. The arrival sequence does not only determine the trading partner but also at which price the trade occurs.

In the ClosedBook system the average price of the last period is public information and each trader knows if his offer resulted in a trade and if so at which transaction price. Submitting an offer equal to the own redemption value will increase the probability of trading, but the trade may occur at the price of the own offer, yielding a profit of zero.

All offers and thus the average price are publicly available in the OpenBook system. Each agent can exactly determine what the payoff would have been for each possible offer, assuming that other agents submit the same offers and arrive in the same sequence. Agents learn to select the offer that has the highest expected payoff and the offers will converge to some equilibrium value.

4.3 Individual Evolutionary Learning algorithm

Agents learn to submit the most profitable offer by the Individual Evolutionary Learning (IEL) algorithm as introduced by Arifovic and Ledyard (2003, 2007). It is shown in the paper Arifovic and Ledyard (2004) that this algorithm performs better than other learning rules, and it suits well in environments with continuous or large strategy spaces.

At the beginning of the period every agent selects an offer from an individual pool of strategies probabilistically, depending on the hypothetical foregone payoff of these strategies. After the trading period, the new hypothetical payoffs of own strategies are determined assuming the same arrival sequence and offers of other traders. The pool of agents' strategies is also updated using these new hypothetical payoffs as follows. First, every strategy may mutate with a given probability. Second, the places in the new pool are filled by those strategies that have higher hypothetical payoffs. After every agent updated own pool, the new period starts and, again, every agent selects probabilistically a strategy.

Pool of strategies

Every trader has an individual pool of strategies, which is a subset of the continuous strategy space: the set $B_{b,t}$ of K bids $b_{b,t} \leq V_b$ for buyer b and the set $A_{s,t}$ of K asks $a_{s,t} \geq C_s$ for seller s . Hence the Individual Rationality constraint of Anufriev et al. (2013) is satisfied and traders cannot submit offers that could result in a negative profit. The offers are initially drawn from a uniform distribution on $[0, V_b]$ and $[C_s, 1]$ respectively. Even though this pool of strategies is updated every period, the number of strategies remains constant. However, it is possible that a single strategy starts to dominate the pool of strategies over time and occupies many positions in the pool.

Mutation

With a fixed small probability ρ a strategy mutates and otherwise remains the same. When a strategy mutates, a normally distributed variable with mean zero and a fixed variance is added to the old strategy.

Replication

At the end of the period and after possible mutations, the foregone payoff is calculated for each strategy while taking the strategies from others constant. Two strategies are randomly selected and compared on the basis of their hypothetical foregone payoff, after which the best performing strategy occupies a spot in the new pool of strategies. This procedure is repeated K times, in order to select a new strategy for every position in the pool. In the field of Genetic Algorithms this is denoted as a tournament selection process.

Hypothetical foregone payoff functions

After a period traders have knowledge about the profit of the strategy that they selected. However, it is vital for traders to value also the other strategies from the pool. Hence traders calculate, or estimate, the hypothetical payoff of other strategies in the previous period. This is done given the offers of traders and the sequence of order submission in the previous period. Hence traders do not take into account that others are also learning in between periods and thus behave boundedly rational. Moreover, similar to Arifovic and Ledyard (2007) and Anufriev et al. (2013) traders compare strategies solely on their performance in the previous period. It is unclear how a longer sampling period would affect the comparison between Closed- and OpenBook.

The order book and arrival sequence of the previous period are known in the OpenBook system and hence it is possible to calculate the hypothetical foregone payoffs. These can be determined exactly for every possible strategy, given the strategies of others and the arrival sequence from the previous period. We assume the same arrival sequence because of the computational problems this would yield for traders in large markets, as the number of computations will be multiplied by the factorial of the number of traders. For example, with only one buyer and one seller who in the previous period submitted ask $a_{s,t}$, the hypothetical foregone payoff of the buyers' strategy (b_i, n_i) is equal to $V_b - a_{s,t}$ when $n_i > n_{s,t}$ and $b_i \geq a_{s,t}$, equal to $V_b - b_i$ when $n_i < n_{s,t}$ and $b_i \geq a_{s,t}$ and zero otherwise. When $n_i = n_{s,t}$ one of the traders randomly arrives first, and the hypothetical foregone payoff of the strategy (b_i, n_i) equals $\frac{1}{2}(V_b - a_{s,t}) + \frac{1}{2}(V_b - b_i)$. The hypothetical foregone payoff functions are given by

$$U_{b,t}(b_i | \mathcal{J}_{b,t}^{OP}) = \begin{cases} V_b - p_{b,t}^*(b_i) & \left| \text{if bid } b_i \text{ resulted in a transaction at price } p_{b,t}^*(b_i) \right. \\ 0 & \left| \text{otherwise,} \right. \end{cases}$$

$$U_{s,t}(a_j | \mathcal{J}_{s,t}^{OP}) = \begin{cases} p_{s,t}^*(a_j) - C_s & \left| \text{if ask } a_j \text{ resulted in a transaction at price } p_{s,t}^*(a_j) \right. \\ 0 & \left| \text{otherwise.} \right. \end{cases}$$

In the ClosedBook system however, only the average price of the previous period is known and it is necessary to determine some estimated payoff for every strategy. The average price, P_t^{av} , is set to the previous average price if no trade occurred. In this setup it is only possible to estimate the probability that a different strategy would have resulted in a trade. We extend the paper of Anufriev et al. (2013) by introducing a different hypothetical foregone payoff function in the ClosedBook setting, that uses more of the available information.

Anufriev et al. (2013) proved convergence of submitted bids and asks towards the redemption values of the agents in the ClosedBook system. This is the consequence of their choice for a payoff function chosen to be the closest to the Call Market payoff function of Arifovic and Ledyard (2007). Anufriev et al. (2013) state that "this specification is a strong assumption" and that this "may effect results of IEL". This foregone payoff function solely distinguishes between offers below and above the average price of the previous period. The consequence is that investors trade with a high probability but may generate a very small profit. The ClosedBook hypothetical payoff function of Anufriev et al. (2013) is given by

$$U_{b,t}(b_i | \mathcal{J}_{b,t}^{CL}) = \begin{cases} V_b - P_t^{av} & \left| \text{if } b_i \geq P_t^{av} \right. \\ 0 & \left| \text{otherwise,} \right. \end{cases}$$

$$U_{s,t}(a_j | \mathcal{J}_{s,t}^{CL}) = \begin{cases} P_t^{av} - C_s & \left| \text{if } a_j \leq P_t^{av} \right. \\ 0 & \left| \text{otherwise.} \right. \end{cases}$$

Fig. 4.1 shows a simulation in the symmetric S5-environment that has valuations and costs $\{[1, 0.8, 0.6, 0.4, 0.2], [0.8, 0.6, 0.4, 0.2, 0]\}$ and consists of 5 buyers and 5 sellers, where the payoff function of Anufriev et al. (2013) is used. The figures can be described as follows. Part (a) and (b) show the average price, efficiency and number of transactions in every period for

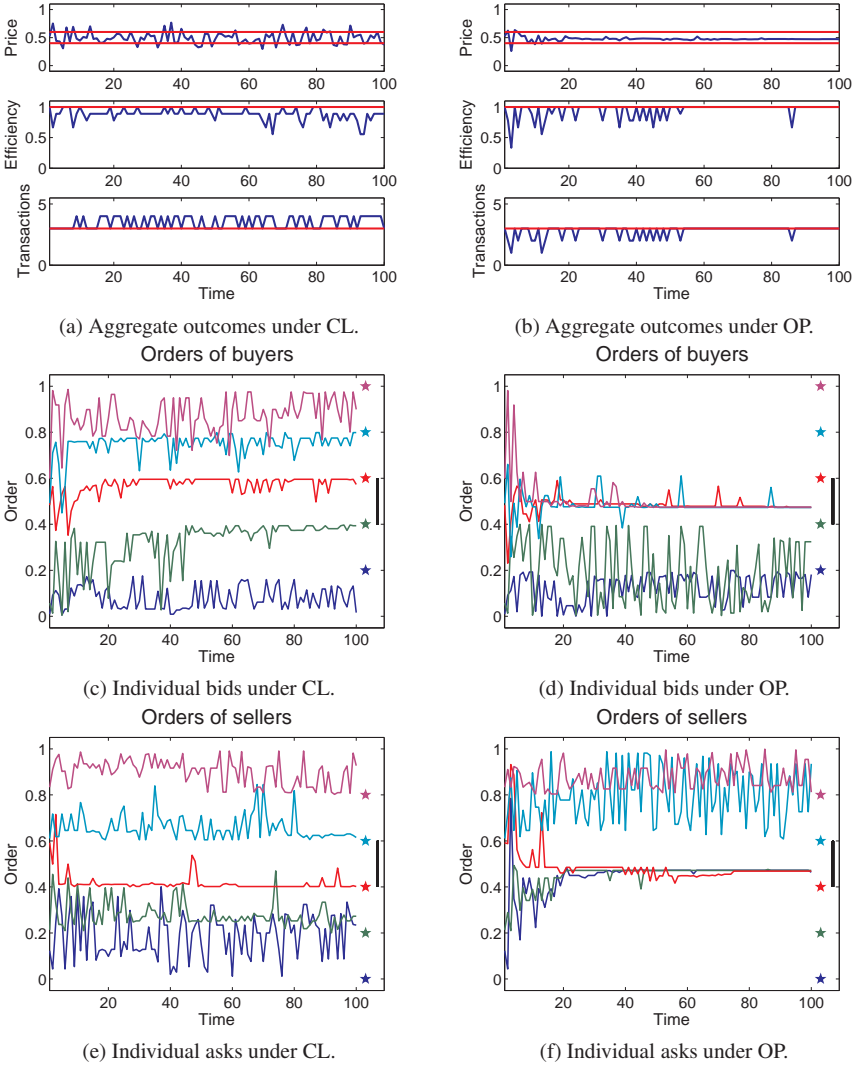


Figure 4.1: Dynamics in the S5-environment under the foregone payoff function of Anufriev et al. (2013). In the OpenBook system they observe that offers of intramarginal traders rapidly move towards the equilibrium price range. In the ClosedBook however, offers of traders move towards their valuation of cost. Extramarginal agents have more opportunities to trade, yielding a higher number of transactions and a slightly lower efficiency.

ClosedBook (CL) respectively OpenBook (OP). The horizontal lines indicate the equilibrium price range $0.4 - 0.6$, the equilibrium efficiency and the equilibrium quantity 3. The individual offers in the Open- and ClosedBook are shown in parts (c)-(f). The equilibrium price range is here indicated by the vertical line and the redemption values by the stars on the right. While in the OpenBook system offers of intramarginal traders converge towards the equilibrium price range, in ClosedBook they observe convergence of offers towards the redemption values, i.e. a divergence away from the equilibrium range of prices. As a result, extramarginal agents trade frequently and more transactions occur than the equilibrium number of 3. This leads to a lower efficiency and a higher price volatility. Under OpenBook mutation might prevent trades, hence the number of transactions and the efficiency are occasionally lower than the equilibrium value.

We now introduce a new ClosedBook hypothetical foregone payoff function, that uses more of the available information in the CDA. This payoff function uses more of the available information to derive hypothetical foregone payoffs. Both Arifovic and Ledyard (2007) and Anufriev et al. (2013) use a ClosedBook hypothetical foregone payoff function in which the main interest of agents is to trade. We will show that under the new ClosedBook payoff function intramarginal traders have a higher profit. We assume the following thought process of a buyer. A buyer who traded in the last period trades again with probability 1, if he submits a bid that is higher than the minimum of his last bid and the average price. A buyer who did not trade, trades if and only if he submits a bid above the maximum of his last bid and the average price. A bid that we assume to result in a trade, is matched with an ask price with the average price as estimated value. The trading price is equal to the bid with probability $\frac{1}{2}$ and with the same probability equal to P_t^{av} . A similar argument is used to construct the hypothetical payoff function for sellers. The new and old OpenBook hypothetical foregone payoff functions only coincide for a trader that traded at the average price. A buyer would under the old payoff function not distinguish between bids above the average price. Under the new payoff function a buyer prefers a lower bid, when he assumes that this also results in a trade. The hypothetical foregone payoff function is given by:

If the agent traded in the last period:

$$U_{b,t}(b_i | \mathcal{J}_{b,t}^{CL}) = \begin{cases} \frac{1}{2}(V_b - b_i) + \frac{1}{2}(V_b - P_t^{av}) & \text{if } b_i \geq \min(P_t^{av}, p_{b,t}) \\ 0 & \text{otherwise,} \end{cases}$$

$$U_{s,t}(a_j | \mathcal{J}_{s,t}^{CL}) = \begin{cases} \frac{1}{2}(P_t^{av} - C_s) + \frac{1}{2}(a_j - C_s) & \text{if } a_j \leq \max(P_t^{av}, p_{s,t}) \\ 0 & \text{otherwise.} \end{cases}$$

If the agent did not trade in the last period:

$$U_{b,t}(b_i | \mathcal{J}_{b,t}^{CL}) = \begin{cases} \frac{1}{2}(V_b - b_i) + \frac{1}{2}(V_b - P_t^{av}) & \text{if } b_i \geq \max(P_t^{av}, b_{b,t}) \\ 0 & \text{otherwise,} \end{cases}$$

$$U_{s,t}(a_j | \mathcal{J}_{s,t}^{CL}) = \begin{cases} \frac{1}{2}(P_t^{av} - C_s) + \frac{1}{2}(a_j - C_s) & \text{if } a_j \leq \min(P_t^{av}, a_{s,t}) \\ 0 & \text{otherwise.} \end{cases}$$

We observe the impact of the new ClosedBook hypothetical foregone payoff function in Fig. 4.2. Under ClosedBook, instead of a movement of offers towards the valuations and costs of traders as seen in Fig. 4.1, the offers move towards the equilibrium price range. The OpenBook hypothetical payoff function remains unchanged. After the learning phase offers occasionally fluctuate when traders use a different strategy due to mutation. These fluctuations may reduce efficiency and occur less frequently under the ClosedBook system. Hence the efficiency and number of transactions seem higher and the price volatility lower than under OpenBook.

Selection of a strategy from the pool

Initially, every strategy is equally likely to be chosen. In the next periods the probability that a certain strategy is selected is proportional to its hypothetical payoff in the previous period. After mutation has taken place, the probability that a buyer b selects strategy b_i for period $t + 1$ is given by

$$\pi_{b,t+1}(b_i) = \frac{U_{b,t}(b_i | \mathcal{J}_t)}{\sum_{i=1}^K U_{b,t}(b_i | \mathcal{J}_t)}.$$

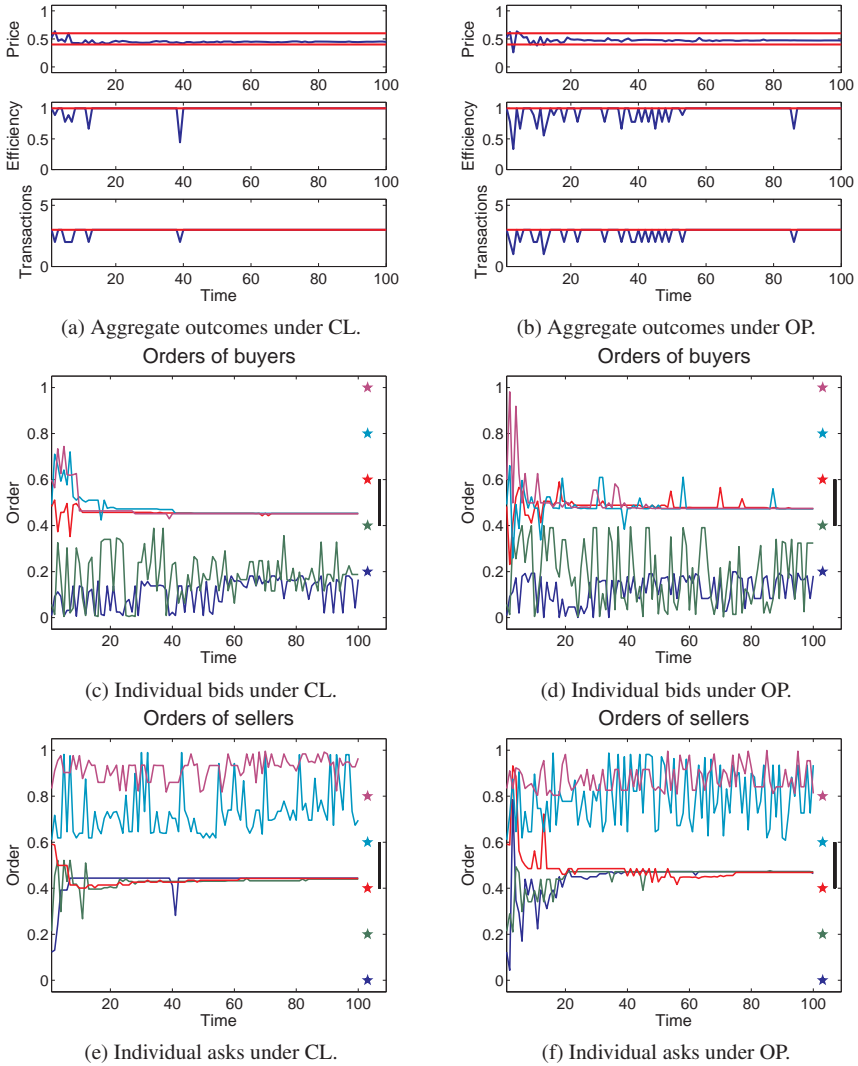


Figure 4.2: Dynamics in the S5-environment under the new ClosedBook payoff function. Where offers move towards valuations and costs in Fig. 4.1 under the ClosedBook payoff function of Anufriev et al. (2013), the new payoff function leads to a convergence of offers towards the equilibrium price range. The OpenBook payoff function is identical.

The variables used in this Individual Evolutionary Learning algorithm are the size of the individual pools, the probability and the distribution of mutation and the replication rate. In the next section the values of these variables are given, as well as the methods to compare average behaviour of Open- and ClosedBook.

4.4 Methodology

In the standard model we examine three environments used in Arifovic and Ledyard (2007) and Anufriev et al. (2013). The symmetric environment $\{[1, 0.8, 0.6, 0.4, 0.2], [0.8, 0.6, 0.4, 0.2, 0]\}$ which is denoted as S5 and $\{[1, 0.93, 0.92, 0.81, 0.5], [0.66, 0.55, 0.39, 0.39, 0.3]\}$ as the AL-environment, which both consist of 5 buyers and 5 sellers. The results of the latter environment will be shown in the appendices. The third environment is introduced by Gode and Sunder (1997) for its simplicity, which allows for obtaining intuition about the behaviour under the IEL-algorithm and the differences between Open- and ClosedBook. In this so called GS-environment we have 1 seller with cost 0, 1 buyer with valuation 1 and n buyers with valuation β . Of main concern is the GS-environment with 3 extramarginal buyers with valuation 0.5. We will find similar results for all environments and hence show a robustness with respect to the environment. The demand- and supply functions of these environments are shown in Fig. 4.3.




Figure 4.3: The main environments used: (a) The symmetric S5-environment $\{[1, 0.8, 0.6, 0.4, 0.2], [0.8, 0.6, 0.4, 0.2, 0]\}$ with 5 buyers and 5 sellers, (b) the AL-environment $\{[1, 0.93, 0.92, 0.81, 0.5], [0.66, 0.55, 0.39, 0.39, 0.3]\}$ with 5 buyers and 5 sellers and (c) the GS-environment $\{[1, 0.5, 0.5, 0.5], [0]\}$ with 4 buyers and 1 seller.

We study efficiency, average price, the price volatility and the number of transactions under the Individual Evolutionary Learning algorithm, where the allocative efficiency is defined as the ratio between the allocative value in a trading period and the maximal allocative value. This algorithm is studied under the parameters of Arifovic and Ledyard (2007), to be able to make a thorough comparison with their Call Market results. As a result our setup differs from Anufriev et al. (2013) where mutation is uniform. However, simulations have shown that our results are not affected by the distribution of this mutation. Every agent is given an individual pool of strategies, which is in principle of size $K = 100$. A strategy can mutate with a probability of 0.033. When a strategy mutates a normally distributed term with mean 0 and a standard deviation of 0.1 is added to the former strategy. When the mutated strategy lies outside the strategy space, a new normally distributed variable is drawn. In the replication phase K pairs are compared. All the averages are calculated over $S = 100$ random seeds.

We denote the periods 1 – 20 as the learning phase and moreover highlight learning by considering the periods 1 – 5 and 16 – 20 as well. With a small probability it can happen that no transaction takes place in a period during the learning phase, since we only use few agents. In that case the average price of the previous period remains, similar to real markets.

We will see that after the learning phase the market becomes quite stable and the offers and average price only fluctuate within a certain range. We denote this behaviour as the "equilibrium" phase. During the periods 101 – 200 we study the response of the learning algorithm to mutation.

4.5 Learning phase

The learning phase of a Call Market is studied in Arifovic and Ledyard (2007). They conclude that the IEL-algorithm with the parameters above explains their experiments quite well. In the OpenBook system agents behave as pricemakers and in the ClosedBook system as pricetakers. In their simulations and experiments the efficiency is higher in the ClosedBook system. In this

section we compare simulations of the Open- and ClosedBook in a Continuous Double Auction market during periods 1 – 20 and the subperiods 1 – 5 and 16 – 20. Moreover, during this learning phase we study whether the comparison of efficiency is identical to the comparison of Open- and ClosedBook in the Call Market.

4.5.1 Gode Sunder-environment

The Gode Sunder-environment, denoted as GS-environment, is simulated in Fig. 4.4 with 3 extramarginal buyers with valuation $\beta = 0.5$. In both settings we observe that the intramarginal traders seem to coordinate on offers above the valuation of the extramarginal buyers. However, during the first periods the asks from sellers are often quite low, which may result in a trade with an extramarginal buyer.

In Anufriev et al. (2013) the expected efficiency in the ClosedBook system is close to $\frac{1}{2} + \frac{1}{2}\beta$ as n goes to infinity, but under our payoff function the expected efficiency is close to the equilibrium value of 1 even for $n = 3$. In the new setting the agents will learn to select a strategy with a higher expected profit and the buyer will in general bid and the seller ask more than β , thus the efficiency will not be lowered due to transactions with extramarginal buyers.

4.5.2 S5- and AL-environments

Fig. 4.5, and Fig. 4.12 in the appendix, show the behaviour in simulations during the learning phase over periods 1 – 20 of the S5- en AL-environment. We notice that the orders of intramarginal traders converge fast towards the equilibrium price range. The initial pool of strategies is drawn uniformly and hence in the first periods the submitted offers are almost uniform. In these figures we observe that after five periods the orders are relatively close to each other and the standard deviation of individual offers has significantly dropped. The fast convergence originates from two effects. Strategies far away from the equilibrium price range are removed during the replication process, for example if the strategy with the lowest hypothetical foregone payoff occurs only once in the pool it cannot attain a spot in the updated pool. Second, the algorithm

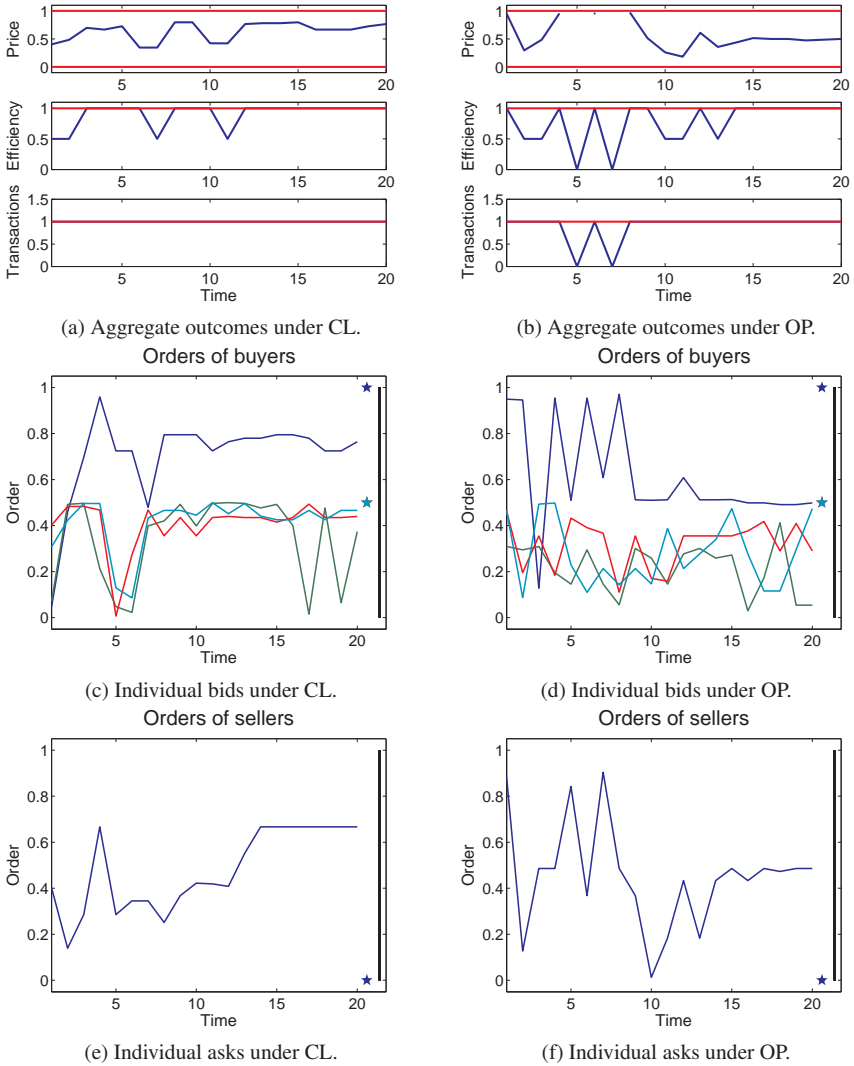


Figure 4.4: Learning phase dynamics in the GS-environment with 3 extramarginal buyers with valuation $\beta = 0.5$. Both in Open- and ClosedBook the intramarginal buyer learns to submit a bid above 0.5, such that the seller prefers to trade with this buyer. The seller increases his ask to make sure that he will trade with the intramarginal buyer and not with an extramarginal.

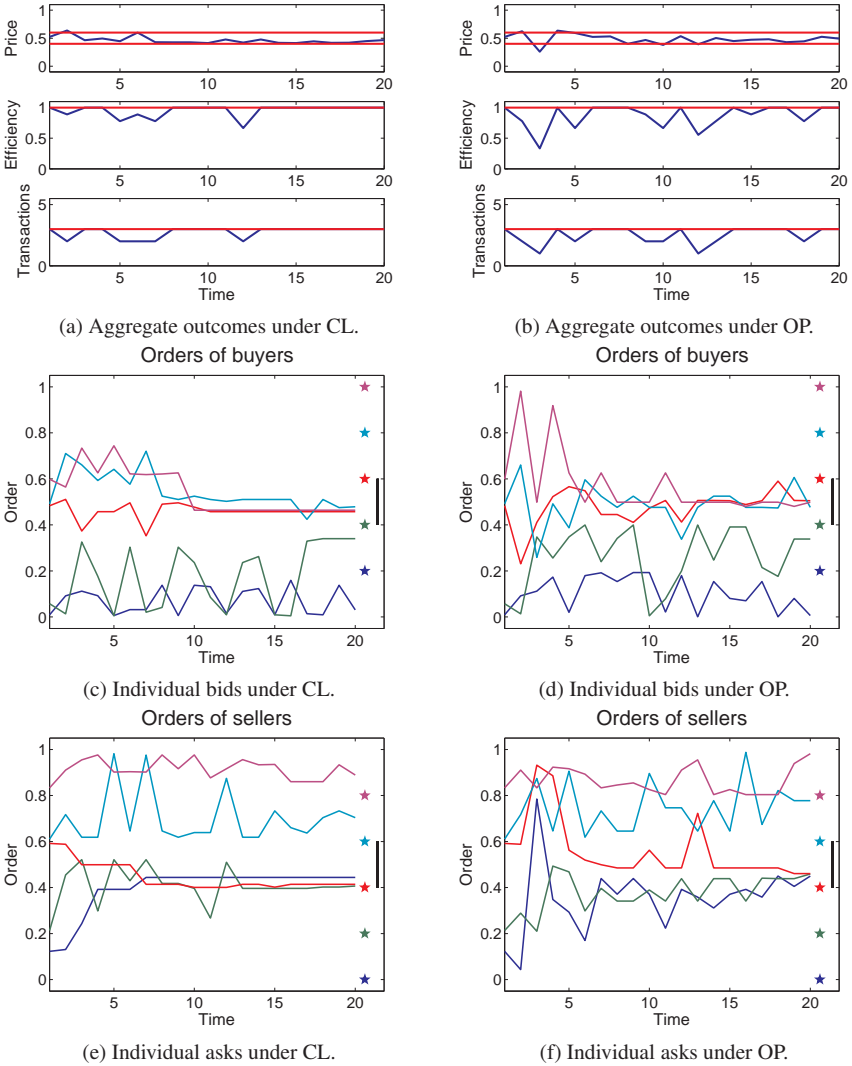


Figure 4.5: Learning phase dynamics in the symmetric S5-environment with 5 buyers and 5 sellers. In both settings a fast convergence occurs. After the first periods offers are relatively close to each other and these offers are less volatile.

selects more profitable offers, since the selection probability of a strategy is proportional to its hypothetical foregone payoff.

4.5.3 Comparison between Open- and ClosedBook

The IEL-algorithm is used to replicate the behaviour of agents. We simulated the different environments and averaged the efficiency, average price, price volatility and the number of transactions. The results of the simulations for Open- and ClosedBook are shown in Table 4.1 for the S5-environment and Table 4.5 in the appendix for the AL-environment, with the standard deviation between brackets. The t-values for comparing two means are given in Tables 4.2 and 4.6, the difference is significant at a level of 5% when the absolute t-value exceeds 1.96. These t-values are obtained by subtracting the OpenBook values from the ClosedBook values.

Of main interest is the question how the behaviour in the OpenBook system differs from the ClosedBook system. From the t-test for comparing two means, we can conclude that the efficiency is higher in the ClosedBook system, at a significance level of 5%. The average prices are similar in both systems and the price volatility is lower under ClosedBook. The number of transactions is higher in ClosedBook, but always lower than the equilibrium quantity. We show that this also holds for periods 1 – 5 and 16 – 20. We can now conclude that in the learning phase, the extra available information in the OpenBook setting leads to a higher price volatility and a lower efficiency and number of transactions. When under OpenBook some traders can increase their profit by submitting a more aggressive offer, they do not take into account that traders on the other side of the market do the same. This will often result in absence of trade when multiple traders are more aggressive and may lead to a higher price volatility when traders are matched with other tradingpartners. In the ClosedBook system traders will in such a case be less aggressive as the previous average price is used as a benchmark. Hence multiple traders may offer more aggressively, but this will less frequently result in absence of trade and a higher price volatility. This results in a lower expected efficiency and a higher price volatility in the OpenBook system.

Period:	CL: closed book			OP: open book		
	1-5	1-20	16-20	1-5	1-20	16-20
Efficiency	0.7771 (0.1109)	0.9031 (0.0491)	0.9562 (0.0562)	0.7304 (0.0900)	0.8433 (0.0477)	0.8887 (0.0695)
Price	0.5001 (0.0881)	0.5020 (0.0567)	0.5027 (0.0488)	0.4966 (0.0564)	0.4930 (0.0414)	0.4891 (0.0510)
Price Volat	0.0835 (0.0440)	0.0546 (0.0246)	0.0209 (0.0124)	0.1173 (0.0503)	0.0776 (0.0238)	0.0320 (0.0170)
Num transact	2.3260 (0.3299)	2.6860 (0.1536)	2.8200 (0.1938)	2.1680 (0.2737)	2.5010 (0.1616)	2.6120 (0.2324)

Table 4.1: Average outcomes during the learning phase in the symmetric S5-environment with 5 buyers and 5 sellers. We observe a clear learning effect, during the initial periods the efficiency and the number of transactions are relatively low but these increase fast during the learning phase.

Period:	T-values		
	1-5	1-20	16-20
Efficiency	3.27	8.74	7.55
Price	0.33	1.28	1.93
Price Volat	-5.06	-6.72	-5.28
Num transact	3.69	8.30	6.87

Table 4.2: T-values for testing the differences in average outcomes between ClosedBook and OpenBook during the learning phase in the symmetric S5-environment with 5 buyers and 5 sellers. The efficiency and number of transactions are significantly higher and the volatility significantly lower under ClosedBook at a 5% percent significance level. Average price is slightly higher under ClosedBook but not significantly.

4.5.4 Comparison with the Call Market

Arifovic and Ledyard (2007) find a higher efficiency in the ClosedBook system in various environments, a ClosedBook efficiency between 92% and 94% and an OpenBook efficiency between 77% and 90%. We used the same parameters in the simulations and, as in their paper, averaged over the first 20 periods of all environments. We find a ClosedBook efficiency between 86% and 90% and an OpenBook efficiency between 82% and 84%. The efficiency under ClosedBook is slightly higher in the Call Market and similar under OpenBook. We conclude that the differences in efficiency between Open- and ClosedBook are very similar for the Continuous Double Auction and the Call Market. Moreover, under ClosedBook the efficiency is higher in a Call Market than in a Continuous Double Auction.

4.6 Long-term behaviour

Arifovic and Ledyard (2007) do not consider the long-run in the Call Market, since the offers will converge to the redemption values in the ClosedBook system and mutation has little effect. The efficiency will almost always be equal to one. Arifovic and Ledyard (2007) find that under the OpenBook traders behave as pricemakers, and hence mutation does have a significant effect. Hence we argue that under OpenBook the efficiency is higher.

The long-term behaviour of agents in a Continuous Double Auction under the IEL-algorithm is studied by Anufriev et al. (2013). The foregone payoff function that they use is the same as in the Call Market model. The orders converge to the redemption values, as Anufriev et al. (2013) explain in their Result 1: "The strategy profile under which the pool of every trader consists of messages equal to his own valuation/cost is attractive under the ClosedBook treatment in the GS-environment".

In this section we present our results under the alternative ClosedBook foregone payoff function. We use normally instead of uniformly distributed mutation, but simulations have shown that this does not affect our results. We compare Open- and Closed book simulations and the

differences with the foregone payoff function from Anufriev et al. (2013). Under the foregone payoff function we introduced, also in the ClosedBook system orders will converge towards the equilibrium price range and this result does no longer hold.

4.6.1 GS-environment

Figs. 4.6 and 4.7 show the impact of the payoff function in the ClosedBook setting, in the GS-environment with 3 extramarginal buyers with valuation $\beta = 0.5$. In Fig. 4.6 the hypothetical foregone payoff function of Anufriev et al. (2013) is used. The intramarginal buyer will submit a bid close to 1 and the seller will submit an ask close to 0. The seller will often trade with an extramarginal buyer, giving him a low payoff. The seller can increase his expected profit by asking a higher price, without a decrease in the probability of trading. In Fig. 4.7 the new hypothetical payoff functions is used, where the agents use more of the available information, resulting in both in Open- and ClosedBook in some convergence towards an equilibrium price between 0.5 and 1.

4.6.2 S5- and AL-environments

Now we have studied the long-term behaviour in the simple GS-environment, we can formally consider the other environments. Figs. 4.8 and 4.13 show the S5- and the AL-environment. In these realisations the efficiency and the number of transactions are clearly higher in the ClosedBook system. In equilibrium the offers in the ClosedBook fluctuate less, as a result of the mutual IEL of all the agents, and in particular of their evaluation of mutations and their consequences. In the OpenBook system traders observe the entire trading sequence of the previous period and base their next strategy on this. If an agent traded at the price of the trading partner, slightly higher bids or lower asks have an identical hypothetical foregone payoff. However, in the ClosedBook system, the hypothetical foregone payoff function gives a lower profit for these slightly higher bids and lower asks. Hence mutated strategies are more often selected in the OpenBook system, which may reduce efficiency in the next period when other traders condition on this mutated strategy. The arrival sequence influences the strategy of traders and

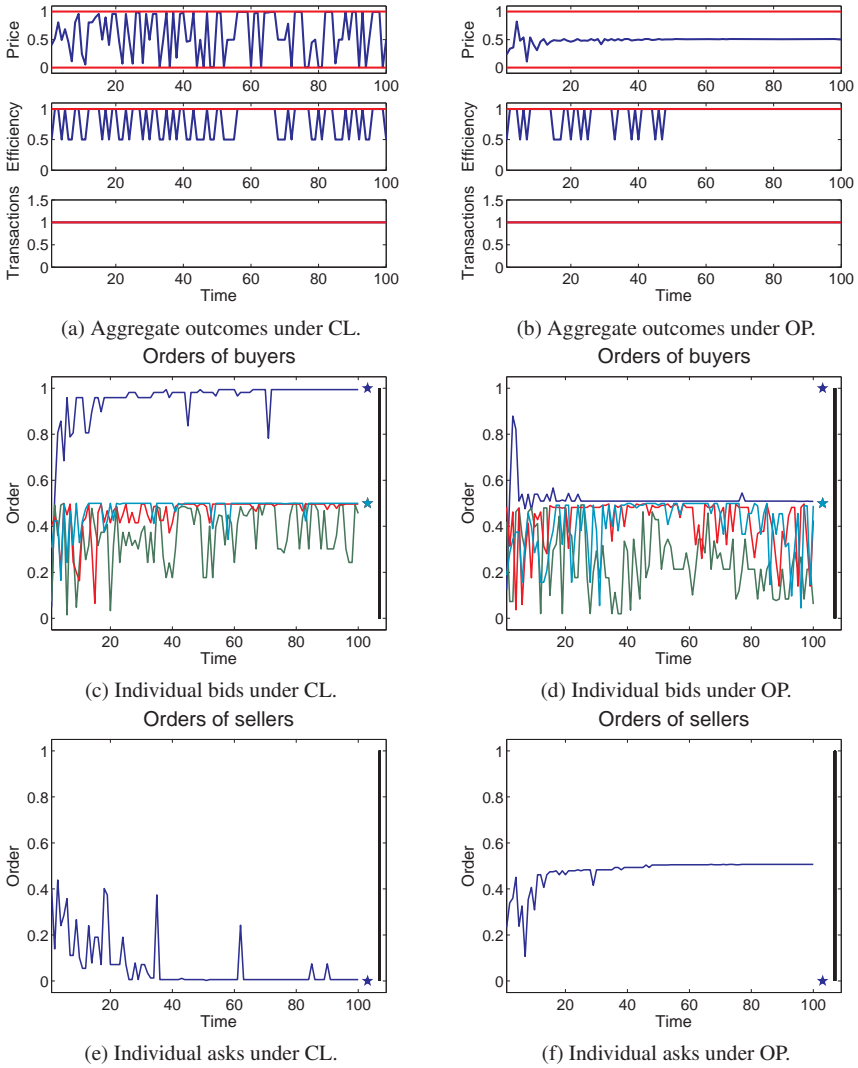








Figure 4.6: Long-term dynamics in the GS-environment with 3 extramarginal buyers with valuation $\beta = 0.5$ under the foregone payoff function of Anufriev et al. (2013). While in the OpenBook system the intramarginal buyer and seller coordinate on a price above 0.5, in the ClosedBook system the buyer bids close to 1 and the seller asks 0. Frequently an extramarginal buyer will trade which lowers the efficiency to 0.5. Hence the trading price is very volatile.

(a) Aggregate outcomes under CL.
Orders of buyers(c) Individual bids under CL.
Orders of sellers

(e) Individual asks under CL.

(b) Aggregate outcomes under OP.
Orders of buyers(d) Individual bids under OP.
Orders of sellers

(f) Individual asks under OP.

Figure 4.7: Long-term dynamics in the GS-environment with 3 extramarginal buyers with valuation $\beta = 0.5$ under the new foregone payoff function. The intramarginal buyer and seller coordinate on a price above 0.5, which ensures that they will trade with each other.

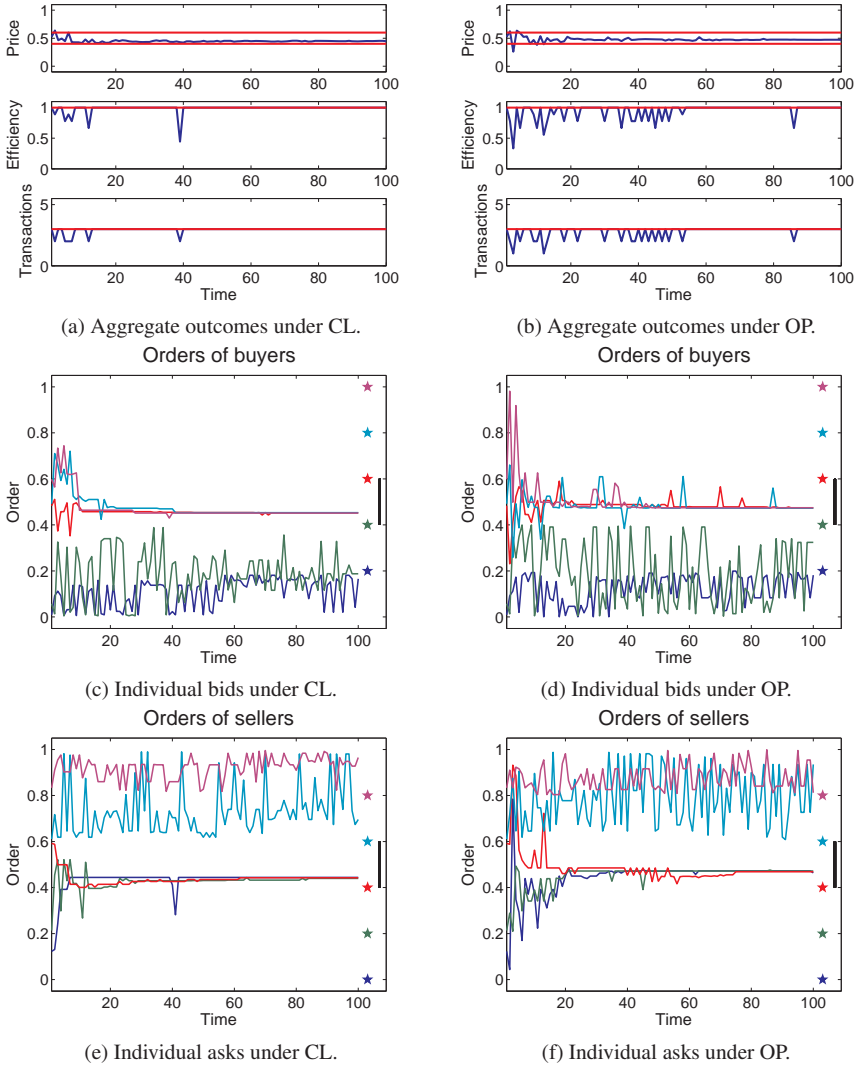


Figure 4.8: Long-term dynamics in the symmetric S5-environment with 5 buyers and 5 sellers. While under the OpenBook offers fluctuate frequently this occurs much less under ClosedBook. Hence the efficiency and the number of transactions are often lower under OpenBook.

they will occasionally offer relatively aggressive. This may lead to absence of trade if the arrival sequence in the next period is different.

4.6.3 Comparison between Closed- and OpenBook

We simulated the S5- and AL-environments and averaged the efficiency, average price, price volatility and the number of transactions, for different sizes K of the pools of strategies. In equilibrium we observe robustness with respect to the size of the pool of strategies. Anufriev et al. (2013) already showed robustness with regards to the probability of mutation; general statistics barely change as long as this probability is not too large. The results of the simulations for Open- and ClosedBook are shown in Tables 4.3 and 4.7 with the standard deviation between brackets. The t-values for comparing two means are given in Tables 4.4 and 4.8.

In all the environments the average efficiency is higher and the price volatility is lower in the ClosedBook setting. We do not observe a difference in average prices between Closed- and OpenBook. We observe a higher number of transactions for the ClosedBook setting in each environment, but this is in each case lower than the equilibrium number of transactions. All these results are significant at a level of 5%. Efficiency and number of transactions are higher when the size of the pool K increases, and the price volatility decreases. We conclude that in equilibrium, irrespective of the size of the pool of strategies, the extra available information in the OpenBook setting significantly leads to a higher price volatility and a lower efficiency and number of transactions. Hence we have shown robustness with respect to the size of the pool of strategies.

4.6.4 Comparison with the ClosedBook foregone payoff function in Anufriev et al. (2013).

The behaviour under the introduced ClosedBook hypothetical foregone payoff function is different than the behaviour in Anufriev et al. (2013). They find a divergence of offers; we find a convergence towards an equilibrium price. With the new hypothetical foregone payoff

	CL: closed book				OP: open book			
	$K = 10$	$K = 50$	$K = 100$	$K = 200$	$K = 10$	$K = 50$	$K = 100$	$K = 200$
Eff	0.9059 (0.0817)	0.9799 (0.0230)	0.9868 (0.0196)	0.9918 (0.0144)	0.8708 (0.0883)	0.9505 (0.0345)	0.9605 (0.0295)	0.9607 (0.0362)
Price	0.5035 (0.0503)	0.5042 (0.0405)	0.5030 (0.0432)	0.4986 (0.0442)	0.4973 (0.0535)	0.4993 (0.0412)	0.4966 (0.0407)	0.5031 (0.0426)
Vol	0.0227 (0.0114)	0.0090 (0.0060)	0.0066 (0.0050)	0.0054 (0.0053)	0.0252 (0.0120)	0.0130 (0.0046)	0.0122 (0.0045)	0.0117 (0.0048)
Trans	2.6176 (0.3305)	2.9362 (0.0724)	2.9559 (0.0587)	2.9723 (0.0435)	2.5377 (0.3310)	2.8461 (0.0989)	2.8714 (0.0881)	2.8862 (0.0977)

Table 4.3: Long-term average outcomes in the symmetric S5-environment with 5 buyers and 5 sellers. As the size K of the pool of strategies increases, so do the average efficiency and number of transactions, and the price volatility decreases.

	T-values			
	$K = 10$	$K = 50$	$K = 100$	$K = 200$
Eff	2.92	7.09	7.43	7.98
Price	0.84	0.85	1.08	-0.73
Vol	-1.51	-5.29	-8.32	-8.81
Trans	1.71	7.35	7.98	8.05

Table 4.4: T-values for testing the differences in long-term average outcomes between Closed-Book and OpenBook in the symmetric S5-environment with 5 buyers and 5 sellers. For every size K of the pool of strategies the average efficiency and number of transactions are significantly higher and the price volatility lower under ClosedBook. However, for a small size of the pool, $K = 10$, not all statistics are significant.

function, mutation has a smaller effect on the market under Closed Book. Hence, where they find a comparable efficiency and a higher price volatility in the ClosedBook system, under the new foregone payoff function the efficiency is higher and the price volatility lower under ClosedBook.

We argue that intramarginal traders are better off when they use the newly introduced hypothetical foregone payoff function. Efficiency and thus the total profit is significantly higher under the new payoff function. Moreover, we have shown that extramarginals trade less frequently, so that intramarginal traders receive a larger part of the total profit. Hence the intramarginal traders can increase their profit by using the new hypothetical foregone payoff function.

4.7 Multi-unit Continuous Double Auction market

So far we considered a market in which every buyer and every seller tries to trade a single unit of the good. In this section we extend this model by allowing agents to trade multiple units. Investors submit a strategy (b_i, n_i) respectively (a_j, n_j) , which not only consists of a bid or an ask price, but also the number of units they desire to trade. The more units a buyer already obtained in the period the lower he values an extra unit, and similarly for sellers. In this symmetric environment the 10 valuations for a single buyer are given by $\{[1, .95, .89, .82, .74, .63, .53, .42, .3, .17]\}$, where the first value denotes the valuation for the first obtained unit and so on. The costs for a seller are symmetric to these valuations. Fig. 4.9a shows the decreasing valuation function and the increasing cost function for 5 identical buyers and sellers.

In this symmetric environment there exist only equilibria in which all investors trade 7 units. Submitting a strategy for more units can result in a loss on the additional units. Mutation of the number of units can occur, in which case one unit is added or subtracted to the strategy with equal probability. In the ClosedBook setting the payoff function needs to be adjusted to consider in how many trades a certain strategy will result. The adjusted hypothetical payoff is calculated by multiplying the payoffs of the regular payoff function by the estimated number of units that would be traded. For buyers this number of units is as follows:

If the buyer traded in the previous period and the strategy consists of fewer units than in the previous period ($n_{i,t+1} < n_{i,t}$):

$$\left\{ \begin{array}{l} \text{all} \\ \# \text{ of transaction prices in the last period} < b_i \end{array} \right. \left| \begin{array}{l} \text{if } b_i \geq P_t^{av} \\ \text{otherwise.} \end{array} \right.$$

If the buyer traded in the previous period and the strategy consists of more units than in the previous period ($n_{i,t+1} \geq n_{i,t}$):

$$\left\{ \begin{array}{l} \text{all} \\ \# \text{ of transactions in the previous period} \\ \# \text{ of transaction prices in the last period} < b_i \end{array} \right. \left| \begin{array}{l} \text{if } b_i \geq \max(P_t^{av}, b_{b,t}) \\ \text{if } \max(P_t^{av}, b_{b,t}) > b_i \geq \min(P_t^{av}, b_{b,t}) \\ \text{otherwise.} \end{array} \right.$$

If the buyer did not trade in the previous period:

$$\left\{ \begin{array}{l} \text{all} \\ 0 \end{array} \right. \left| \begin{array}{l} \text{if } b_i > \max(P_t^{av}, b_{b,t}) \\ \text{otherwise.} \end{array} \right.$$

The symmetric environment in this extended market is shown in Figs. 4.9 and 4.10. Two-dimensional learning occurs as traders place an offer and a size of the offer. Both in Open- and ClosedBook these order move towards the equilibrium. We observe a robustness with respect to the number of units agents can trade. Mutation has a larger effect in the OpenBook system, resulting in a lower efficiency and number of trades and a higher price volatility.

A random environment where each investor trades 5 units in equilibrium is shown in Figs. 4.14 and 4.15 in the appendix. In this example, during early trading periods we observe a remarkable coordination on offers outside the equilibrium price, which is never observed before under IEL. Both in Open- and ClosedBook a disturbance after period 40 leads to coordination on offers closer to the equilibrium price. This may be the result of the relatively small equilibrium price range.

As shown in Tables 4.9-4.16 the comparisons between Open- and ClosedBook in the learning and in the equilibrium phase are slightly altered. The efficiency and number of transactions

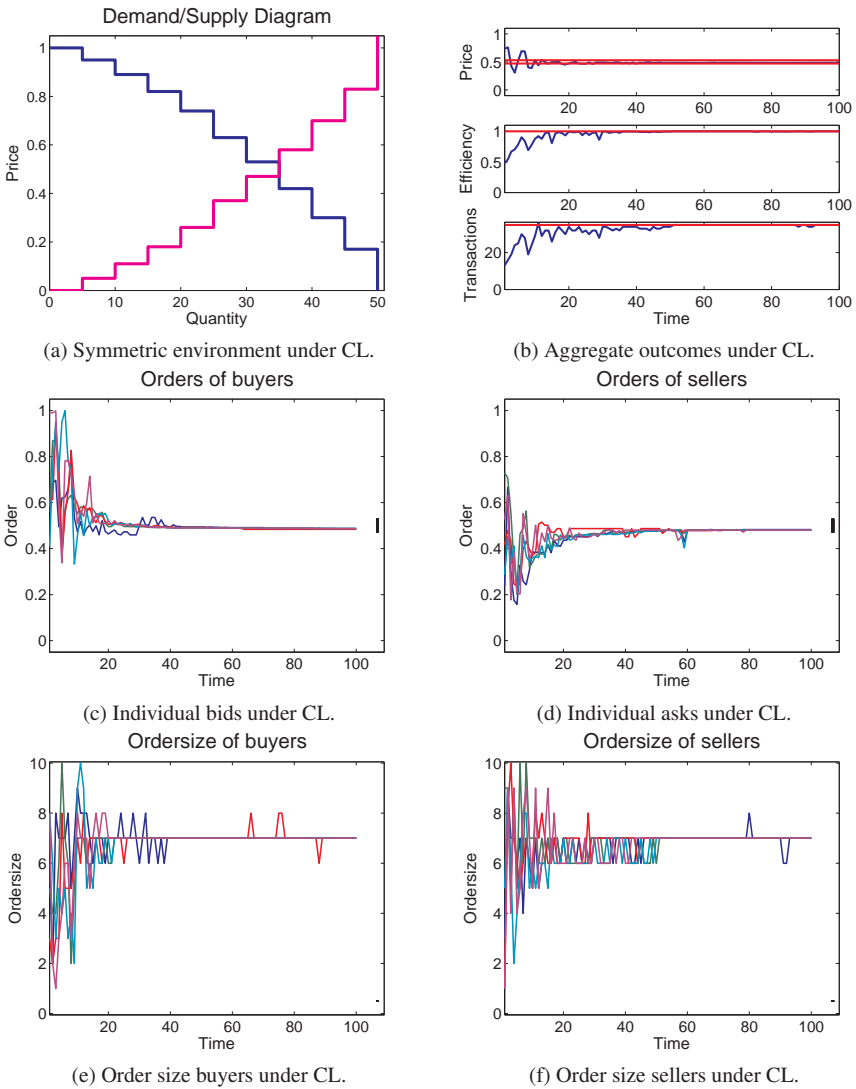


Figure 4.9: Long-term dynamics in the ClosedBook multi-unit symmetric environment with 5 buyers and 5 traders that can place an offer for a maximum of 10 units. The equilibrium offer is made for 7 units. Even though traders are required to make a two-dimensional decision, orders move towards the equilibrium. Mutation seems to have little effect after the learning phase.

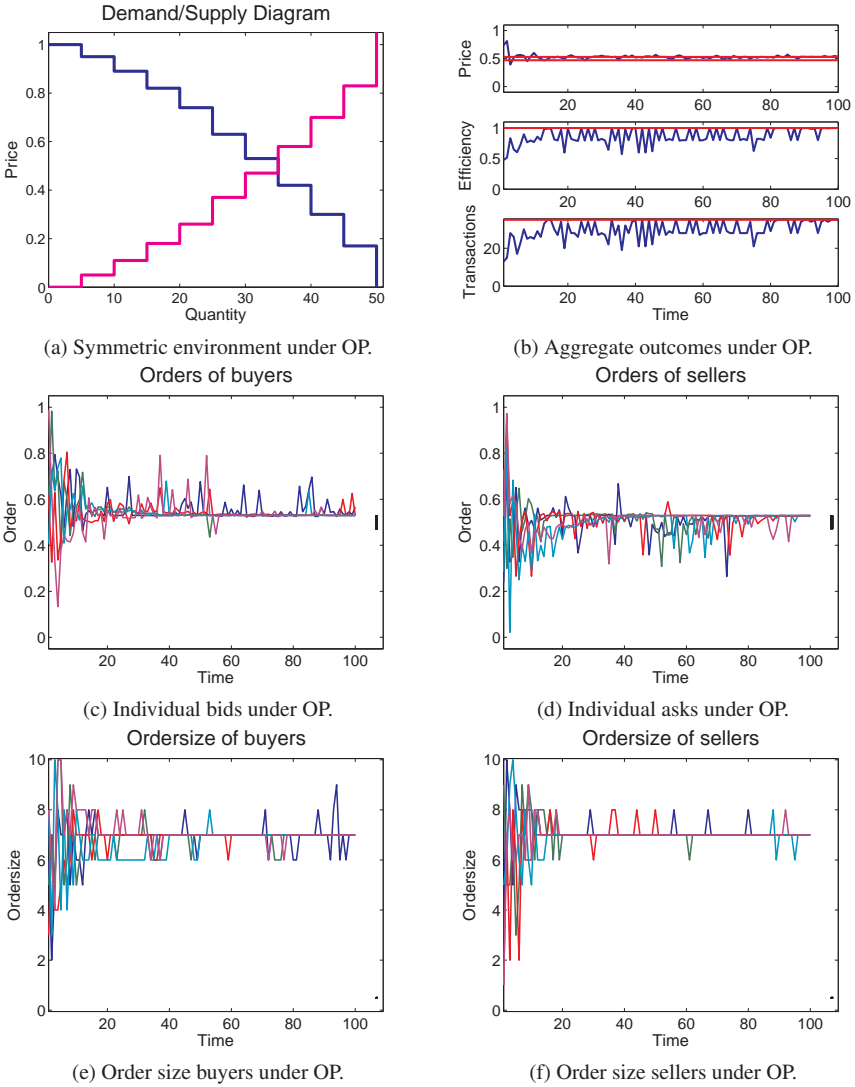


Figure 4.10: Long-term dynamics in the OpenBook multi-unit symmetric environment with 5 buyers and 5 sellers that can place an offer for a maximum of 10 units. The equilibrium offer is made for 7 units. Orders move towards the equilibrium quantity and price, but mutation has a larger effect than in the ClosedBook system. Efficiency and the number of transactions are often below the equilibrium value and the price volatility is higher.

remain higher under ClosedBook. However, the difference in price is now significantly larger under OpenBook. In the multi-unit symmetric environment the price volatility is significantly larger under OpenBook, but in the multi-unit random environment significantly lower. In the first environment offers converge to the equilibrium price range after the first 20 periods, but in the latter environment traders coordinate on offers outside the equilibrium price range as described above. The disturbance around period 40 increases the price volatility, which reverses the comparison between Open- and ClosedBook. Moreover, the average price is significantly larger under OpenBook, in the long run of this extended model. We find a robustness with respect to the number of units traded, but the IEL algorithm reacts slightly different when the equilibrium price range is relatively small.

4.8 Size of the market

Robustness with respect to the size of the pool of strategies and to the number of units that each investor prefers to trade is shown in the previous sections. In this section we will consider robustness with respect to the number of investors in the market.

To study the robustness with respect to the size of the market we increased the number of investors in the S5- and AL-environments. This results in a setting with 15 buyers and 15 sellers shown in Fig. 4.16 of the appendix and a setting with 25 buyers and 25 sellers in Fig. 4.11. Larger markets would be more realistic, but computationally not feasible. Although the efficiency increases and volatility decreases, the comparisons between Open- and ClosedBook remain in Tables 4.17-4.24. Our results are robust to the size of the market, both during the learning and the equilibrium phase.

4.9 Concluding Remarks

In this chapter we have studied the role of information about the trading history that is available to traders in a Continuous Double Auction market, when traders use the Individual Evolution-

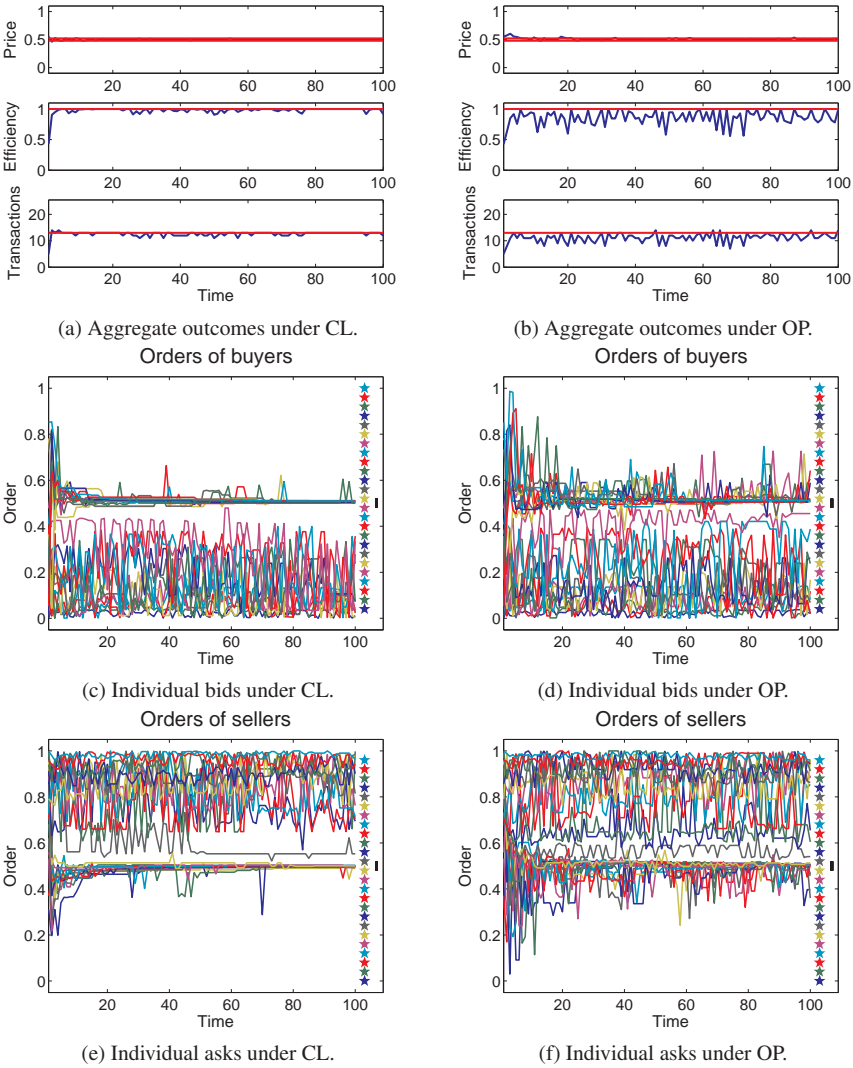


Figure 4.11: Long-term dynamics in the symmetric S25-environment with 25 buyers and 25 sellers. Both in the learning phase and long-term offers of intramarginal traders are less volatile under ClosedBook. Average efficiency and number of transactions are higher and price volatility lower under ClosedBook, showing a robustness with respect to the number of traders on both sides of the market.

ary Learning algorithm. In this learning algorithm traders select from a pool of strategies based on the, hypothetical, payoff in the previous period. In the ClosedBook system, where only information about past average prices is available, Anufriev et al. (2013) proved convergence of bids and offers towards the valuations and costs of agents. This is the consequence of their choice for a payoff function that only distinguishes between offers below and above the average price of the previous period, as in a Call Market. The consequence is that investors trade with a high probability but may generate a very small profit. We showed however, that when the payoff function uses more information to estimate the expected payoff of each possible offer, bids and offers tend to drift towards the equilibrium price range. In this approach investors learn to increase their expected profit by submitting an offer that has a higher possible profit. The probability of trading will be lower in this situation but is outweighed by the increase in possible profit. In this setting bids and asks will not diverge, but will converge towards some equilibrium price. These results are in line with Fano et al. (2013) who show that, in a setting closely related to the ClosedBook system, traders behave as pricetakers in a Call Market and as pricemakers in a Continuous Double Auction.

Both during the learning phase and in equilibrium we compared simulations of the Closed- and OpenBook treatments in different environments. In all the environments the efficiency and the number of transactions are significantly larger in the ClosedBook system. The number of transactions is in each case lower than the equilibrium number of transactions. In general, we did not observe a difference in average prices between Closed- and OpenBook. Moreover, we observed a significantly lower price volatility in the ClosedBook system. The cause is that agents in the OpenBook system are more tempted to try to influence their transaction price. We conclude that both during the learning phase and in equilibrium, the extra available information in the OpenBook treatment leads to a higher price volatility and a lower efficiency and number of transactions.

The efficiency found is comparable to efficiency in the Call Market from Arifovic and Ledyard (2007). The results differ however from the results under the payoff

function of Anufriev et al. (2013), which results in a comparable efficiency and a higher price volatility in the ClosedBook system. The behaviour in the ClosedBook system is also quite different under the new payoff function; instead of a divergence of offers some convergence occurs.


Some extensions to this Continuous Double Auction market are considered. We have shown that the above results are robust with respect to the number of units agents desire to trade and the size of the market. However, in the multi-unit random environment, price volatility is higher in the ClosedBook as a result of the relatively small equilibrium price range.

Both during the learning phase and in equilibrium, more information about the trading history leads to a higher price volatility and a lower efficiency and number of transactions. This is the result of the ClosedBook foregone hypothetical payoff function, that is introduced to estimate the expected profit in the previous period by using more of the available information in a Continuous Double Auction market. We conclude that it is optimal not to reveal the extra information of the OpenBook system.


Appendix A: Learning phase

In this appendix we consider the learning phase during the periods 1–20 for the AL-environment. This is done for the entire period, as well as the subperiods 1 – 5 and 16 – 20. An example is shown in Fig. 4.12 and averages in Table 4.5. Moreover, the t-values for testing the differences between ClosedBook and OpenBook are given in Table 4.6.


We observe a fast convergence of offers towards the equilibrium price range, both in the Open and ClosedBook system. In the entire time span and the subperiods, the efficiency and number of transactions are significantly higher and the price volatility significantly lower under ClosedBook. The average price does not significantly differ. Traders learn over time, and hence the efficiency and number of transactions increase and the price volatility decreases during the periods 1 – 20.




(a) Aggregate outcomes under CL.




(b) Aggregate outcomes under OP.



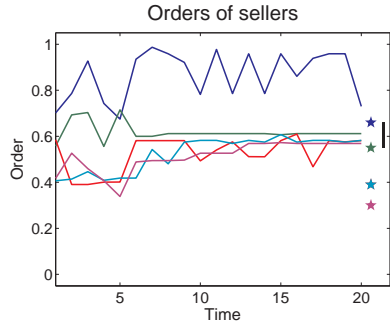
(c) Individual bids under CL.



(d) Individual bids under OP.



(e) Individual asks under CL.



(f) Individual asks under OP.

Figure 4.12: Learning phase dynamics in the AL-environment with 5 buyers and 5 sellers. Both in the Open- and ClosedBook system offers of traders move fast towards the equilibrium price range.

Period:	CL: closed book			OP: open book		
	1-5	1-20	16-20	1-5	1-20	16-20
Efficiency	0.7764 (0.1176)	0.9059 (0.0386)	0.9546 (0.0441)	0.7198 (0.0719)	0.8331 (0.0361)	0.8805 (0.0608)
Price	0.6328 (0.0703)	0.6316 (0.0490)	0.6297 (0.0429)	0.6396 (0.0488)	0.6436 (0.0324)	0.6459 (0.0394)
Price Volat	0.0648 (0.0313)	0.0416 (0.0147)	0.0163 (0.0078)	0.0819 (0.0309)	0.0552 (0.0140)	0.0226 (0.0119)
Num transact	3.2000 (0.4680)	3.6920 (0.1716)	3.8600 (0.1853)	2.9560 (0.2907)	3.4330 (0.1460)	3.6380 (0.2662)

Table 4.5: Average outcomes during the learning phase in the AL-environment with 5 buyers and 5 sellers. The efficiency and number of transactions are higher and the price volatility is lower under ClosedBook. A difference in average price is not observed. A learning effect occurs and efficiency and number of transactions increase over time and price volatility decreases.


Period:	T-values		
	1-5	1-20	16-20
Efficiency	4.11	13.77	9.87
Price	-0.79	-2.04	-2.78
Price Volat	-3.89	-6.70	-4.43
Num transact	4.43	11.50	6.84

Table 4.6: T-values for testing the differences in average outcomes between ClosedBook and OpenBook during the learning phase in the AL-environment with 5 buyers and 5 sellers. The efficiency and number of transactions are significantly higher and the price volatility is significantly lower under ClosedBook. A difference in average price is not observed.


Appendix B: Equilibrium phase

This appendix considers the long term behaviour of traders during the periods 101 – 200 for the AL-environment. This is done for different sizes K of the pool of strategies. Fig. 4.13 shows an example of the behaviour, with the averages given in Table 4.7. Moreover, the t-values for testing the differences between ClosedBook and OpenBook are given in Table 4.8.


In the example we observe a smaller effect of mutation under ClosedBook. Hence the equilibrium efficiency and number of trades are more often reached. The efficiency and number of transactions are significantly higher and the price volatility significantly lower in the Closed-Book system, irrespective of the size K of the pool of strategies. We do not observe a clear significant difference in average price.




(a) Aggregate outcomes under CL.



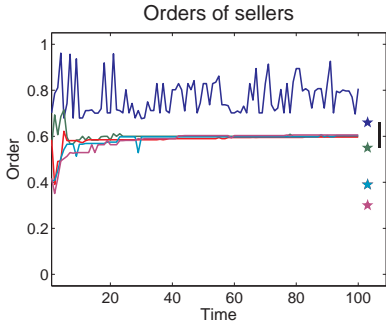
(b) Aggregate outcomes under OP.



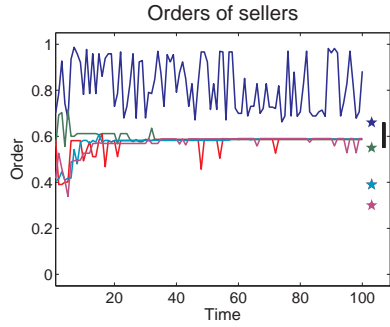
(c) Individual bids under CL.



(d) Individual bids under OP.



(e) Individual asks under CL.



(f) Individual asks under OP.

Figure 4.13: Long-term dynamics in the AL-environment with 5 buyers and 5 sellers. In both systems offers move fast towards the equilibrium price range. Under ClosedBook mutation seems to have a smaller effect and efficiency and the number of transactions more frequently attain the equilibrium value.

	CL: closed book				OP: open book			
	$K = 10$	$K = 50$	$K = 100$	$K = 200$	$K = 10$	$K = 50$	$K = 100$	$K = 200$
Eff	0.9197 (0.0584)	0.9745 (0.0219)	0.9806 (0.0210)	0.9885 (0.0151)	0.8646 (0.0647)	0.9229 (0.0319)	0.9259 (0.0334)	0.9267 (0.0347)
Price	0.6287 (0.0383)	0.6326 (0.0263)	0.6239 (0.0289)	0.6184 (0.0274)	0.6364 (0.0423)	0.6320 (0.0287)	0.6359 (0.0262)	0.6353 (0.0267)
Vol	0.0167 (0.0074)	0.0089 (0.0050)	0.0077 (0.0055)	0.0063 (0.0064)	0.0220 (0.0078)	0.0144 (0.0044)	0.0132 (0.0038)	0.0134 (0.0039)
Trans	3.6898 (0.3175)	3.9213 (0.0659)	3.9359 (0.0649)	3.9562 (0.0583)	3.5087 (0.2936)	3.7459 (0.1162)	3.7785 (0.1120)	3.7698 (0.1245)

Table 4.7: Long-term average outcomes in the AL-environment with 5 buyers and 5 sellers. The efficiency and number of transactions are higher and the price volatility is lower under ClosedBook. A clear difference in average price is not observed. These results are robust with respect to the size K of the pool of strategies.

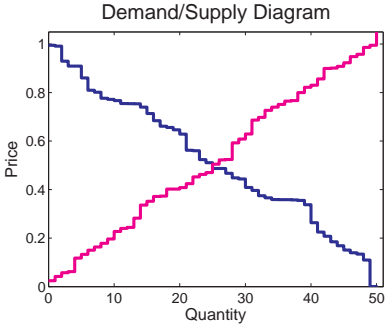
	T-values			
	$K = 10$	$K = 50$	$K = 100$	$K = 200$
Eff	6.32	13.34	13.86	16.33
Price	-1.35	0.15	-3.08	-4.42
Vol	-4.93	-8.26	-8.23	-9.47
Trans	4.19	13.13	12.16	13.56

Table 4.8: T-values for testing the differences in long-term average outcomes between Closed-Book and OpenBook in the AL-environment with 5 buyers and 5 sellers. The efficiency and number of transactions are significantly higher and the price volatility is significantly lower under ClosedBook. A clear significant difference in average price is not observed.


Appendix C: Multi-unit market

In this appendix we consider an extension of the model by allowing agents to trade multiple units. In the equilibrium of the random environment traders place an order for 5 units, shown in Figs. 4.14 and 4.15. The learning phase of both the symmetric and random environments is studied over the periods 1 – 20 and the subperiods 1 – 5 and 16 – 20. We present the average outcomes and the t-values for testing the differences between Closed- and OpenBook in Tables 4.9-4.12 for both the symmetric and random environment. Average outcomes and t-values for the long-term, during periods 101 – 200 and for different sizes K of the pool of strategies, are shown in Tables 4.13-4.16.


In the examples of the multi-unit random environment we observe that mutation plays a larger role under OpenBook, as full efficiency is often not obtained. During the learning phase the efficiency and number of transactions are higher under ClosedBook. The average price and price volatility do not show a significant difference. In the long-run efficiency and number of transactions remain higher under ClosedBook and the average price becomes significantly lower. The comparison between price volatility is different between both environments. In the symmetric environment the price volatility is significantly lower under ClosedBook, and in the random environment significantly higher. The latter is the effect of the coordination of offers outside the small equilibrium price range in the ClosedBook system.




(a) Symmetric environment under CL.




(b) Aggregate outcomes under CL.



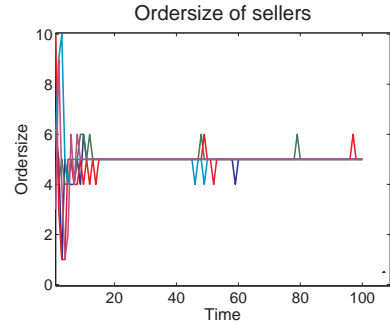
(c) Individual bids under CL.



(d) Individual asks under CL.




(e) Order size buyers under CL.




(f) Order size sellers under CL.


Figure 4.14: Long-term dynamics in the ClosedBook multi-unit random environment with 5 buyers and 5 sellers that can place an offer for a maximum of 10 units. The equilibrium offer is made for 5 units. Both size and offer converge. Traders coordinate their offers outside the relatively small equilibrium price range. This coordination is disturbed around period 40.




(a) Random environment under OP.




(b) Aggregate outcomes under OP.




(c) Individual bids under OP.



(d) Individual asks under OP.



(e) Order size buyers under OP.



(f) Order size sellers under OP.

Figure 4.15: Long-term dynamics in the OpenBook multi-unit random environment with 5 buyers and 5 sellers that can place an offer for a maximum of 10 units. The equilibrium offer is made for 5 units. Both size and offer converge to the equilibrium value. Mutation seems to play a larger role than under ClosedBook, as full efficiency is less frequently obtained.

Period:	CL: closed book			OP: open book		
	1-5	1-20	16-20	1-5	1-20	16-20
Efficiency	0.6895 (0.0654)	0.8657 (0.0322)	0.9516 (0.0336)	0.6423 (0.0604)	0.8109 (0.0311)	0.8878 (0.0525)
Price	0.4895 (0.0804)	0.4863 (0.0435)	0.4852 (0.0380)	0.5015 (0.0623)	0.5054 (0.0403)	0.5078 (0.0417)
Price Volat	0.1278 (0.0511)	0.0808 (0.0249)	0.0273 (0.0121)	0.1355 (0.0535)	0.0786 (0.0223)	0.0210 (0.0107)
Num transact	21.8240 (2.6538)	27.8625 (1.3104)	30.8560 (1.4896)	21.4420 (2.5747)	27.3770 (1.2164)	30.1460 (2.0175)

Table 4.9: Average outcomes during the learning phase in the multi-unit symmetric environment with 5 buyers and 5 sellers that can place an offer for a maximum of 10 units. The efficiency and number of transactions are higher under ClosedBook. A clear difference in average price and price volatility is not observed. A learning effect occurs and efficiency and number of transactions increase over time and price volatility decreases.

Period:	T-values		
	1-5	1-20	16-20
Efficiency	5.30	12.24	10.24
Price	-1.18	-3.22	-4.01
Price Volat	-1.04	0.66	3.90
Num transact	1.03	2.72	2.83

Table 4.10: T-values for testing the differences in average outcomes between ClosedBook and OpenBook during the learning phase in the multi-unit symmetric environment with 5 buyers and 5 sellers that can place an offer for a maximum of 10 units. The efficiency and number of transactions are significantly higher under ClosedBook. A clear significant difference in average price and price volatility is not observed.

Period:	CL: closed book			OP: open book		
	1-5	1-20	16-20	1-5	1-20	16-20
Efficiency	0.5395 (0.0883)	0.8201 (0.0444)	0.9614 (0.0413)	0.4850 (0.0675)	0.7593 (0.0352)	0.8889 (0.0544)
Price	0.4951 (0.0559)	0.4862 (0.0368)	0.4822 (0.0361)	0.5094 (0.0512)	0.5078 (0.0332)	0.5101 (0.0336)
Price Volat	0.0849 (0.0391)	0.0559 (0.0174)	0.0236 (0.0089)	0.0839 (0.0360)	0.0543 (0.0157)	0.0205 (0.0098)
Num transact	16.5240 (2.6932)	21.5830 (1.1873)	24.0320 (1.0307)	16.1740 (2.3771)	20.4600 (0.9495)	22.3260 (1.3516)

Table 4.11: Average outcomes during the learning phase in the multi-unit random environment with 5 buyers and 5 sellers that can place an offer for a maximum of 10 units. The efficiency and number of transactions are higher under ClosedBook. A clear difference in average price and price volatility is not observed. A learning effect occurs and efficiency and number of transactions increase over time and price volatility decreases.

Period:	T-values		
	1-5	1-20	16-20
Efficiency	4.90	10.73	10.61
Price	-1.89	-4.36	-5.66
Price Volat	0.19	0.68	2.34
Num transact	0.97	7.39	10.04

Table 4.12: T-values for testing the differences in average outcomes between ClosedBook and OpenBook during the learning phase in the multi-unit random environment with 5 buyers and 5 sellers that can place an offer for a maximum of 10 units. The efficiency and number of transactions are significantly higher and average price significantly lower under ClosedBook. A clear difference in price volatility is not observed.

	CL: closed book				OP: open book			
	$K = 10$	$K = 50$	$K = 100$	$K = 200$	$K = 10$	$K = 50$	$K = 100$	$K = 200$
Eff	0.9749 (0.0244)	0.9781 (0.0254)	0.9783 (0.0263)	0.9783 (0.0254)	0.9388 (0.0219)	0.9372 (0.0260)	0.9442 (0.0219)	0.9399 (0.0199)
Price	0.4721 (0.0171)	0.4764 (0.0195)	0.4795 (0.0232)	0.4766 (0.0189)	0.5038 (0.0307)	0.5040 (0.0376)	0.5040 (0.0324)	0.4991 (0.0312)
Vol	0.0077 (0.0042)	0.0055 (0.0044)	0.0054 (0.0055)	0.0055 (0.0049)	0.0120 (0.0037)	0.0110 (0.0036)	0.0102 (0.0039)	0.0106 (0.0038)
Trans	32.7328 (2.0657)	33.0609 (2.2211)	33.1029 (2.2171)	33.0880 (2.2109)	31.9787 (1.5151)	31.7267 (1.6777)	32.0048 (1.7313)	31.8543 (1.6213)

Table 4.13: Long-term average outcomes in the multi-unit symmetric environment with 5 buyers and 5 sellers that can place an offer for a maximum of 10 units. The efficiency and number of transactions are higher and the price volatility lower under ClosedBook. However, also the average price is lower under ClosedBook. These results are robust with respect to the size K of the pool of strategies.

	T-values			
	$K = 10$	$K = 50$	$K = 100$	$K = 200$
Eff	11.01	11.25	9.96	11.90
Price	-9.02	-6.52	-6.15	-6.17
Vol	-7.68	-9.67	-7.12	-8.22
Trans	2.94	4.79	3.90	4.50

Table 4.14: T-values for testing the differences in long-term average outcomes between ClosedBook and OpenBook in the multi-unit symmetric environment with 5 buyers and 5 sellers that can place an offer for a maximum of 10 units. The efficiency and number of transactions are significantly higher and the price volatility significantly lower under ClosedBook. However, also the average price is significantly lower under ClosedBook.

	CL: closed book				OP: open book			
	$K = 10$	$K = 50$	$K = 100$	$K = 200$	$K = 10$	$K = 50$	$K = 100$	$K = 200$
Eff	0.9524 (0.0267)	0.9470 (0.0310)	0.9517 (0.0294)	0.9517 (0.0250)	0.9306 (0.0251)	0.9391 (0.0286)	0.9410 (0.0267)	0.9426 (0.0271)
Price	0.4502 (0.0354)	0.4463 (0.0321)	0.4516 (0.0348)	0.4517 (0.0359)	0.4989 (0.0302)	0.5014 (0.0279)	0.4978 (0.0270)	0.5028 (0.0268)
Vol	0.0200 (0.0076)	0.0182 (0.0082)	0.0181 (0.0074)	0.0175 (0.0082)	0.0103 (0.0027)	0.0089 (0.0028)	0.0092 (0.0031)	0.0091 (0.0028)
Trans	23.8183 (0.8207)	23.6129 (0.9692)	23.7963 (0.8932)	23.6853 (0.9102)	23.2905 (0.7143)	23.4898 (0.7415)	23.5340 (0.7692)	23.5920 (0.7354)

Table 4.15: Long-term average outcomes in the multi-unit random environment with 5 buyers and 5 sellers that can place an offer for a maximum of 10 units. The efficiency, number of transactions and price volatility are higher and the average price lower under ClosedBook. Price volatility is higher under ClosedBook due to the disturbances between subsequent coordinations of offers outside the equilibrium price. These results are robust with respect to the size K of the pool of strategies.


	T-values			
	$K = 10$	$K = 50$	$K = 100$	$K = 200$
Eff	6.44	1.87	2.69	2.47
Price	-10.47	-12.96	-10.49	-11.41
Vol	12.03	10.73	11.09	9.69
Trans	4.85	1.01	2.23	0.80

Table 4.16: T-values for testing the differences in long-term average outcomes between ClosedBook and OpenBook in the multi-unit random environment with 5 buyers and 5 sellers that can place an offer for a maximum of 10 units. The efficiency, number of transactions and price volatility are significantly higher and the average price significantly lower under ClosedBook.


Appendix D: Size of the market

This appendix considers larger markets that consist of 15 buyers and 15 sellers or 25 buyers and 25 sellers. An example of the S15-environment is shown in Fig. 4.16. Both for the learning phase and the long-term we present average outcomes and the t-values for testing the differences between ClosedBook and OpenBook. These are given in Tables 4.17-4.24 for both the S15- and the S25-environment. The learning phase is studied over the periods 1 – 20 and the subperiods 1 – 5 and 16 – 20 and the long-term during periods 101 – 200 for different sizes K of the pool of strategies.


In the example under ClosedBook full efficiency is often attained, but only rarely in the OpenBook system. Hence we can conclude that the results shown in the S5-environment are robust with respect to the size of the market. The efficiency and the number of transactions are significantly higher under ClosedBook, and the price volatility significantly lower. The average price does not differ significantly between Open- and ClosedBook.




(a) Aggregate outcomes under CL.




(b) Aggregate outcomes under OP.




(c) Individual bids under CL.



(d) Individual bids under OP.



(e) Individual asks under CL.



(f) Individual asks under OP.

Figure 4.16: Long-term dynamics in the symmetric S15-environment with 15 buyers and 15 sellers. With more traders in the market offers again move towards the equilibrium price range. Where under ClosedBook mutation incidentally leads to a decreased efficiency, this occurs under OpenBook in most periods.

Period:	CL: closed book			OP: open book		
	1-5	1-20	16-20	1-5	1-20	16-20
Efficiency	0.8246 (0.0415)	0.9258 (0.0207)	0.9664 (0.0288)	0.7514 (0.0490)	0.8401 (0.0231)	0.8744 (0.0464)
Price	0.5048 (0.0379)	0.5026 (0.0219)	0.5008 (0.0176)	0.5015 (0.0314)	0.5003 (0.0193)	0.4995 (0.0230)
Price Volat	0.0549 (0.0197)	0.0336 (0.0094)	0.0125 (0.0056)	0.0574 (0.0242)	0.0406 (0.0106)	0.0199 (0.0093)
Num transact	6.7700 (0.4552)	7.3650 (0.2222)	7.5920 (0.2953)	6.1060 (0.4278)	6.7280 (0.2015)	6.9000 (0.4045)

Table 4.17: Average outcomes during the learning phase in the symmetric S15-environment with 15 buyers and 15 sellers. The efficiency and number of transactions are higher and price volatility lower under ClosedBook. A difference in average price is not observed. A learning effect occurs and efficiency and number of transactions increase over time and price volatility decreases.

Period:	T-values		
	1-5	1-20	16-20
Efficiency	11.40	27.63	16.85
Price	0.67	0.79	0.45
Price Volat	-0.80	-4.94	-6.82
Num transact	10.63	21.24	13.82

Table 4.18: T-values for testing the differences in average outcomes between ClosedBook and OpenBook during the learning phase in the S15-environment with 15 buyers and 15 sellers. The efficiency and number of transactions are significantly higher and price volatility significantly lower under ClosedBook. A significant difference in average price is not observed.

Period:	CL: closed book			OP: open book		
	1-5	1-20	16-20	1-5	1-20	16-20
Efficiency	0.8417 (0.0362)	0.9356 (0.0172)	0.9725 (0.0192)	0.7640 (0.0398)	0.8387 (0.0211)	0.8652 (0.0406)
Price	0.5032 (0.0290)	0.5028 (0.0158)	0.5031 (0.0116)	0.5017 (0.0274)	0.5007 (0.0148)	0.4992 (0.0155)
Price Volat	0.0433 (0.0169)	0.0263 (0.0080)	0.0098 (0.0040)	0.0395 (0.0185)	0.0293 (0.0075)	0.0142 (0.0061)
Num transact	11.3120 (0.5638)	12.1510 (0.2832)	12.4600 (0.3510)	10.0800 (0.5895)	10.9110 (0.2804)	11.1380 (0.5208)

Table 4.19: Average outcomes during the learning phase in the symmetric S25-environment with 25 buyers and 25 sellers. The efficiency and number of transactions are higher and price volatility lower under ClosedBook. A difference in average price is not observed. A learning effect occurs and efficiency and number of transactions increase over time and price volatility decreases.

Period:	T-values		
	1-5	1-20	16-20
Efficiency	14.44	35.60	23.89
Price	0.38	0.97	2.01
Price Volat	1.52	-2.74	-6.03
Num transact	15.10	31.11	21.05

Table 4.20: T-values for testing the differences in average outcomes between ClosedBook and OpenBook during the learning phase in the S25-environment with 25 buyers and 25 sellers. The efficiency and number of transactions are significantly higher and price volatility significantly lower under ClosedBook. A significant difference in average price is not observed.

	CL: closed book				OP: open book			
	$K = 10$	$K = 50$	$K = 100$	$K = 200$	$K = 10$	$K = 50$	$K = 100$	$K = 200$
Eff	0.9340 (0.0316)	0.9791 (0.0111)	0.9827 (0.0117)	0.9844 (0.0125)	0.8636 (0.0337)	0.9103 (0.0215)	0.9146 (0.0195)	0.9196 (0.0245)
Price	0.4996 (0.0240)	0.5010 (0.0129)	0.5007 (0.0143)	0.5010 (0.0154)	0.5015 (0.0225)	0.5038 (0.0148)	0.5014 (0.0137)	0.5012 (0.0164)
Vol	0.0112 (0.0038)	0.0052 (0.0018)	0.0052 (0.0022)	0.0047 (0.0024)	0.0181 (0.0043)	0.0107 (0.0025)	0.0104 (0.0019)	0.0100 (0.0018)
Trans	7.1764 (0.3919)	7.7946 (0.1082)	7.8386 (0.0989)	7.8599 (0.1054)	6.6505 (0.3692)	7.2267 (0.1677)	7.3048 (0.1502)	7.3640 (0.1624)

Table 4.21: Long-term average outcomes in the S15-environment with 15 buyers and 15 sellers. The efficiency and number of transactions are higher and price volatility lower under Closed-Book. A difference in average price is not observed. These results are robust with respect to the size K of the pool of strategies.

	T-values			
	$K = 10$	$K = 50$	$K = 100$	$K = 200$
Eff	15.24	28.69	29.95	23.56
Price	-0.58	-1.43	-0.35	-0.09
Vol	-12.02	-17.85	-17.89	-17.67
Trans	9.77	28.46	29.68	25.61

Table 4.22: T-values for testing the differences in long-term average outcomes between Closed-Book and OpenBook in the S15-environment with 15 buyers and 15 sellers. The efficiency and number of transactions are significantly higher and price volatility significantly lower under ClosedBook. A significant difference in average price is not observed.

	CL: closed book				OP: open book			
	$K = 10$	$K = 50$	$K = 100$	$K = 200$	$K = 10$	$K = 50$	$K = 100$	$K = 200$
Eff	0.9456 (0.0214)	0.9800 (0.0086)	0.9859 (0.0079)	0.9864 (0.0080)	0.8647 (0.0305)	0.9046 (0.0170)	0.9057 (0.0148)	0.9052 (0.0174)
Price	0.5021 (0.0186)	0.5030 (0.0096)	0.5022 (0.0085)	0.5011 (0.0094)	0.5010 (0.0174)	0.5016 (0.0083)	0.5002 (0.0080)	0.5002 (0.0088)
Vol	0.0083 (0.0020)	0.0043 (0.0011)	0.0037 (0.0014)	0.0035 (0.0013)	0.0145 (0.0033)	0.0085 (0.0015)	0.0085 (0.0013)	0.0085 (0.0013)
Trans	11.7892 (0.4615)	12.6831 (0.1521)	12.7884 (0.1274)	12.7947 (0.1190)	10.9008 (0.4903)	11.7102 (0.2301)	11.7794 (0.1869)	11.8101 (0.2145)

Table 4.23: Long-term average outcomes in the S25-environment with 25 buyers and 25 sellers. The efficiency and number of transactions are higher and price volatility lower under Closed-Book. A difference in average price is not observed. These results are robust with respect to the size K of the pool of strategies.

	T-values			
	$K = 10$	$K = 50$	$K = 100$	$K = 200$
Eff	21.71	39.58	47.81	42.40
Price	0.43	1.10	1.71	0.70
Vol	-16.07	-22.58	-25.12	-27.20
Trans	13.19	35.27	44.61	40.14

Table 4.24: T-values for testing the differences in long-term average outcomes between Closed-Book and OpenBook in the S25-environment with 25 buyers and 25 sellers. The efficiency and number of transactions are significantly higher and price volatility significantly lower under ClosedBook. A significant difference in average price is not observed.