



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 5

Timing under Individual Evolutionary Learning in a Continuous Double Auction

5.1 Introduction

In many agent-based models of order-driven financial markets traders submit their order at a random moment during a trading period and are required to make a one-dimensional decision; to choose a bid or ask price as in LiCalzi and Pellizzari (2006) or to forecast a future price as in Chiarella and Iori (2002). However, in a Continuous Double Auction (CDA) the moment of order submission plays a crucial role; submitting at the end of the period will yield a lower probability of trading, submitting at the beginning of the period will most likely result in a trade at the own submitted price which yields a lower profit. Allowing traders to submit their order at their preferred moment may influence these effects as traders may decide to condition their offer on the moment of submission. In agent-based models learning is often used to avoid making extreme assumptions about the rationality of traders and to select between multiple equilibria. With non-random timing learning of agents becomes multidimensional; not only learning about the offer but also learning about the timing of submission is of great importance.

In this chapter we introduce learning about the timing of order submission in an agent-based model. Traders also learn about the offer that they submit, and hence we extend the Individual Evolutionary Learning (IEL) algorithm used in Arifovic and Ledyard (2003, 2007) and in Anufriev et al. (2013) to a multidimensional version and allow for contemporaneous learning about the moment of submission and about the submitted orders. In the IEL algorithm traders select from a pool of possible strategies. After a trading period the hypothetical payoff is calculated for every possible strategy and some strategies are replaced with randomly modified strategies. Adopting the IEL algorithm to incorporate the decision about timing, we study the distribution of preferred submission moments, the interrelation between these moments and the submitted orders, and also the impact of the size of the market on the timing of submission and the offers of traders. In simulations we find that the distribution of the submission moments highly depends on the size of the market.

Starting from early contributions it is common that investors in agent-based models make a decision about the price of the order, but not about its timing. It is typically assumed that they submit their one unit orders at a random moment in the trading period and that between periods their learning is only one-dimensional: buyers learn which bid to submit and sellers learn which price to ask. Sometimes agents directly learn bids and asks, sometimes their bids and asks depend on the expectations and learning is over the space of prediction rules. For examples of the former approach, see LiCalzi and Pellizzari (2006, 2007) who compare efficiency in the CDA with other market protocols such as the call market, under boundedly rational resp. zero intelligent agents; and Bottazzi et al. (2005) who focus on the properties of price time series under different trading protocols. In the market protocols with sequential trade these papers assume that agents arrive in a random sequence. For examples of the latter approach see Chiarella and Iori (2002) who study properties of asset pricing under Continuous Double Auctions and other mechanisms in a model with heterogeneous expectations, Yamamoto and LeBaron (2010) who study the number of order splits and Anufriev and Panchenko (2009) who study the switching between forecasting rules. Again, in all simulations under Continuous Double Auctions, agents submit orders in a random sequence.

Arifovic and Ledyard (2003, 2007) introduced the Individual Evolutionary Learning algorithm to model the boundedly rational learning behaviour of agents in a Call Market model. Anufriev et al. (2013) use IEL in a Continuous Double Auction and compare efficiency under full and no information about the history of orders. Furthermore the latter paper studies the GS-environment from Gode and Sunder (1993, 1997) in the case where traders have zero intelligence and submit every possible offer with equal probability. In Chapter 4 we have shown that the results of Anufriev et al. (2013) depend on the hypothetical foregone payoff function that is chosen, under no information about the order history. This chapter extends the model in Anufriev et al. (2013) by considering multidimensional learning in which traders also learn about the moment of order submission.

An important feature of IEL is that it is essentially a backward-looking learning process. This approach contrasts with the standard economic approach where optimising agents make their decision and use all information rationally. In Friedman (1991) traders can submit orders at any moment of time and can also improve their outstanding orders. Traders regard other's orders as random, update their beliefs about the order distribution using Bayes' formula and submit orders on the basis of their updated distribution. The classical financial literature contains many studies on limit and market orders. Rosu (2009) and Parlour and Seppi (2008) provide a survey about the theoretical research on limit and market orders under random arrival of traders. The surveys Gould et al. (2013b) and Hachmeister (2007) discuss the main theoretical, experimental and empirical papers on limit orders of informed and uninformed traders. Bloomfield et al. (2005) performed an experiment on the choice between limit and market orders by informed and uninformed traders over time. As the period advances, uninformed traders use more market orders and informed traders more limit orders. Bae et al. (2003) and Chung et al. (1999) empirically consider the number of limit and market orders during a trading day and their relation with spread, order size and price volatility. Biais et al. (1995) determine the empirical distribution of large and small trades, orders and cancelations during a trading day. These papers find a U-shaped distribution of orders during a day. They explain this finding as motivated by the desire of traders to perform price discovery in the beginning of the day and

react to events during the closing of the exchange. At the end of the day traders desire to unwind their positions. In this chapter there is no modelling of the news process; information is equal for all traders and does not evolve during periods. Rather we are interested in how traders with given valuation and cost are able to find their trading moment and strategy during the period.

The distribution of submission moments is studied in a benchmark environment under full information about trading history. We find that under the IEL-algorithm investors in a medium size market learn to submit their order around the middle of the trading period to avoid a lower trading probability or lower profit. Moreover, we observe an increasing bid function and decreasing ask function over time, similar to Fano and Pellizzari (2011). We show that the size of the market and competition between traders influence the distribution of submission moments. Furthermore of interest are the placed offer, efficiency, profit and the probability of trading as a function of the submission moment. General market statistics are compared with the setup in Anufriev et al. (2013) where traders submit orders at random moments.

The organisation of this chapter is as follows. The model and the trading mechanism are described in Section 5.2, followed by the extended Individual Evolutionary Learning algorithm in Section 5.3 and the methodology used. The distribution of the preferred moment of submitting and its relation to the bid and ask are described in a benchmark environment in Section 5.4. The impact of the size of the market is considered in Section 5.5. In Section 5.6 we study submission moments and their relation with offers as the amount of competition changes. The Gode-Sunder environment is attractive for its simplicity and studied in Section 5.7. Finally, Section 5.8 concludes.

5.2 Market setup

We describe the environments and the trading mechanism in which we study the simultaneous decision about the time of order submission and the submitted offer. Each trader buys or sells in a Continuous Double Auction market one unit of the good and has to decide the moment of

submission and the offered price.

5.2.1 The environments

Each environment is determined by a set of buyers and a set of sellers with their redemption values for the good. In each trading period $t \in \{1, \dots, T\}$ each of the buyers $b \in \{1, \dots, B\}$ likes 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. Such a trading period consists of the time moments $\{0, 1, \dots, 100\}$. The buyers have a fixed valuation of V_b per unit, sellers have fixed costs of C_s that only needs to be paid in case of a transaction. Agents know their own redemption value, but not the values of the other agents.

We will denote the environments by vectors of valuations and costs. For instance, $\{[1, 1], [0, 0.1, 0.2]\}$ denotes an environment with two buyers having identical valuations 1 and 1 and three sellers with costs 0, 0.1 and 0.2. The supply and demand functions of the benchmark environment $\{[1, 1, 1, 1, 1], [0, 0, 0, 0, 0]\}$ and the main symmetric environment $\{[1, 0.85, 0.7, 0.55, 0.4], [0.6, 0.45, 0.3, 0.15, 0]\}$ are shown in Fig. 5.1.

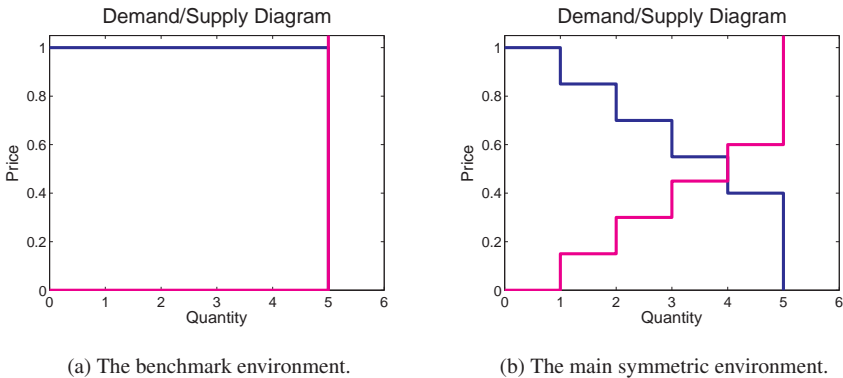


Figure 5.1: The demand and supply functions of the main environments used, the benchmark environment $\{[1, 1, 1, 1, 1], [0, 0, 0, 0, 0]\}$ and the main symmetric environment $\{[1, 0.85, 0.7, 0.55, 0.4], [0.6, 0.45, 0.3, 0.15, 0]\}$.

Based on the valuations and costs the demand and supply functions can be determined. We denote the equilibrium quantity by q^* and the interval of equilibrium prices by $[p_L^*, p_H^*]$. The traders that can gain a positive profit in equilibrium are denoted as intramarginals, whereas the traders that cannot make a positive profit and therefore will not trade in equilibrium are called extramarginals. 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.

5.2.2 Continuous Double Auction

A Continuous Double Auction model is used to describe the regular behaviour at stock exchanges. During a trading period buyers and sellers arrive at their preferred moment and immediately submit their order. The bid of buyer b and the ask of seller s in trading period t are denoted as $b_{b,t}$ and $a_{s,t}$ and their arrival times as $n_{b,t}$ and $n_{s,t}$. If an arriving order can be matched with the best order from the book, the transaction takes place at the price of the order in the book. 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.

If in period t buyer b traded, the transaction price of this buyer is denoted as $p_{b,t}$. Similarly the transaction price of seller s is denoted as $p_{s,t}$. Hence the payoff of a buyer equals $U_{b,t}(p) = V_b - p_{b,t}$ if he traded and zero otherwise. The payoff of a seller equals $U_{s,t}(p) = p_{s,t} - C_s$ after a trade and zero otherwise. We note that the payoff depends not only on the own offer, but also on the trading sequence. For example, if there are only one buyer and one seller, given their offers $b_{b,t} > a_{s,t}$ a buyer will get higher payoff if he will submit his order after the seller, as this will yield a transaction price equal to the ask of the seller. In this chapter we will focus on learning of traders about their timing of submitting the order and about the price of submission. That is why we assume that no order can be cancelled and restrict the traders to buy or sell only one unit of the good. The effects of learning about cancellation and size of the order are left for further research.

5.2.3 Nash equilibria

In the extended model where traders are required to make a two-dimensional decision, a multiplicity of possible long run outcomes may exist. Let us consider a one-period version of our model, where valuations and costs are common knowledge. In this one-period model traders are required to select a strategy consisting of a submission moment and an order only once. Then irrespective of the environment, one set of Nash equilibria exists in which every intramarginal trader submits the same offer in the equilibrium price range, at any possible arrival moment. Hence in this equilibrium the timing of order submission is of no importance. Trivially this constitutes an equilibrium. If a trader adjust its offer price this will result in a lower profit if the offer is more conservative or in absence of trade if more aggressive. However, after a deviation of any offer the arrival moment does play an important role. It is optimal for traders on the other side of the market to arrive at moment $n = 100$ if the deviation leads to a more conservative offer. A more aggressive offer leads to absence of trade for one of the traders on the other side of the market. Hence it is optimal for traders on the other side to arrive at moment $n = 0$.

Furthermore other Nash equilibria may exist in the one-period model, depending on the environment. A trivial example consists of one buyer and two sellers, such that the sellers attempt to outcompete each other. An example of a Nash equilibrium in which not all traders submit the same offer consists of the two sellers submitting an ask price of 0 and the buyer submitting a bid price b at a later moment than both sellers.

The Nash equilibria of this one-period model are possible long run outcomes of the multi-period model used in this chapter. Under the Individual Evolutionary Learning algorithm we find that the offers of traders converge towards the equilibrium price range, such as in the first Nash equilibrium of the one-period model. However, the timing of order submission is of importance; traders learn the optimal submission moment when offer prices of intramarginal traders are not all identical.

5.3 Individual Evolutionary Learning algorithm

Agents learn which strategy to select by a multidimensional version of the Individual Evolutionary Learning algorithm as introduced by Arifovic and Ledyard (2003). Every agent can choose from a set of strategies, which consists of an offer and a submission moment. These strategies might mutate from time to time to allow for some sort of experimentation. Based on how these strategies would have performed in the last trading period, some strategies are replaced by better performing ones. At the beginning of the next period one strategy is selected with a probability proportional to the hypothetical foregone payoff.

Past offers and arrival moments and thus the average price are publicly available. This setup is comparable with the OpenBook setting of Chapter 4. Therefore, after the trading period each agent can determine exactly what his payoff would have been for each possible strategy, assuming no changes in the behaviour of other agents. Agents learn to select the strategy that has the highest hypothetical payoff in the previous period.

Pool of strategies

Every trader has an individual set of strategies: the set $B_{b,t}$ of K randomly drawn pairs of bids and arrival moments (b_i, n_i) for buyer b and the set $A_{s,t}$ of K randomly drawn pairs of asks and arrival moments (a_j, n_j) for seller s . Offers are initially drawn from a uniform distribution on $[0, V_b]$ and $[C_s, 1]$ respectively and the arrival moments are uniformly drawn from the set $\{0, 1, 2, \dots, 100\}$. The periods correspond to days in reality and the arrival moments to the time period of a trading day. Note that the set of arrival moments is larger than the number of traders, which in most environments equals 10. This is done to prevent a large random component in the sequence of submission. When two or more traders decide to submit their order at the same moment in time the orders are handled randomly. Expanding the set of arrival moments reduces this random component.

Traders observe not only the sequence in which traders arrived in the last period, but also their actual moment of submitting. Under the OpenBook system introduced in Chapter 4, the full

order book is shown to traders in the New York Stock Exchange. Hence the actual submission moments are known and not only the sequence of submission. Under this assumption agents evaluate the hypothetical payoff in the previous period on the basis of the moment of submission of other traders. This assumption is important as the following example illustrates. If there are 10 traders and trader i desires to submit in place 9 and the submitted moments of others are for example $\{1, 1, 1, 1, 1, 1, 1, 1, 3\}$ (trader i excluded), he prefers to submit at moment 2. This results in submitting the order at place 9, *ceteris paribus*. We compare with the setup where only the sequence of submissions is known and traders select at which position in the sequence they prefer to arrive, from the set $[1, 2, \dots, 10]$. Calculating foregone payoffs is more difficult as it is uncertain in which place in the sequence a certain submission moment results. If the trader in this example prefers to maximise the probability of submitting in place 9, he would submit moment 9 and *ceteris paribus* arrive at place 10.

Mutation

A part of a strategy mutates with a fixed small probability ρ . A normally distributed variable with mean zero is added to the part of the strategy that mutates while the other part may remain unchanged. The mutated arrival moment is rounded to the nearest integer $n \in \{0, 1, 2, \dots, 100\}$. The variance of the normally distributed variable depends on which part of the strategy mutates. This distribution is truncated; when the mutated strategy lies outside the strategy space, a new normally distributed variable is drawn.

Replication

After the trading period has ended and some strategies have possibly mutated, the foregone payoffs are calculated for each strategy while taking the chosen strategies of others from this period constant. Replication consists of a comparison of two strategies, randomly selected from the pool of strategies. The strategy with the highest foregone hypothetical payoff will obtain a place in the updated pool of strategies. This is repeated K times to fill the entire updated pool.

Hypothetical foregone payoff functions

Calculating the foregone payoffs is a straightforward task under full information about the trading history. Since the entire order book and the arrival moments of others in the last period are known, the foregone payoff can precisely be determined for every possible strategy, given that others remain to use the same strategies.

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 $b_i \geq a_{s,t}$, and is 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 in general given by

$$U_{b,t}(b_i, n_i) = \begin{cases} V_b - p_{b,t}^*(b_i, n_i) & \text{if strategy } (b_i, n_i) \text{ resulted in a trade at price } p_{b,t}^*(b_i, n_i) \\ 0 & \text{otherwise,} \end{cases}$$

$$U_{s,t}(a_j, n_j) = \begin{cases} p_{s,t}^*(a_j, n_j) - C_s & \text{if strategy } (a_j, n_j) \text{ resulted in a trade at price } p_{s,t}^*(a_j, n_j) \\ 0 & \text{otherwise.} \end{cases}$$

Selection of a strategy from the pool

After the hypothetical foregone payoffs are determined each trader has to select a strategy for the next trading period. The probability that a certain strategy is selected is proportional to its hypothetical foregone payoff. In the first period every strategy is equally likely to be chosen. For a buyer b_i the probability of selecting strategy (b_i, n_i) for period $t + 1$ is given by

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

This Individual Evolutionary Learning algorithm depends on some variables, such as the size of the individual pools, the probability and the distribution of mutation and the replication rate. The next section shows the values of these variables that are used in the simulations, as well as the characteristics used to describe the overall outcome in a trading period.

5.3.1 Methodology

The benchmark environment $\{[1, 1, 1, 1, 1], [0, 0, 0, 0, 0]\}$ is considered to show the basic results of IEL-learning regarding the moment of arrival. We simulate this environment with different numbers of traders to study the effect of the size of the market. Furthermore we will use symmetric environments which mainly consist of five buyers and five sellers to study the impact of changes in the amount of competition and the size of the equilibrium price range on the distribution of arrival moments. Some of these environments as the main symmetric environment $\{[1, 0.85, 0.7, 0.55, 0.4], [0.6, 0.45, 0.3, 0.15, 0]\}$ are introduced in Arifovic and Ledyard (2007) and Anufriev et al. (2013). We will show that the results are not robust with respect to the environment; when one side of the market is much larger the other side will extract their power and submit their order as late as possible.

We study the long-run distribution of arrival moments and the expected offer per submission moment. Also of interest are the *allocative efficiency*, which is the ratio between the allocative value in a trading period and the maximal possible allocative value. The allocative value of a trading period is the sum of the payoffs of all agents. It is fully efficient when all intramarginal investors trade during a period. Efficiency can be lower when an extramarginal investor trades, or when intramarginals do not trade at all. Furthermore we study the *average transaction price*, the *price volatility* and the *number of transactions*. All these characteristics are considered per trading period as well as per possible arrival moment and are compared with the one dimensional model.

The Individual Evolutionary Learning algorithm is used with most of the parameters of Arifovic and Ledyard (2007). Every agent is given an individual pool of strategies of size $K = 300$. A part of a strategy mutates with a probability of 0.033; in the case that the offer mutates a normally distributed term with mean 0 and a standard deviation of 0.1 is added to the offer, in the case that the arrival moment mutates a normally distributed term with mean 0 and a standard deviation of 10 is added to the arrival moment. The mutated offer is truncated on the bounds of the interval $[0, 1]$ and the mutated arrival moment is rounded to the nearest integer

$n \in \{0, 1, 2, \dots, 100\}$. When the mutated strategy lies outside the strategy space, a new normally distributed variable is drawn. This mutation differs from Anufriev et al. (2013), where a uniform distribution is used to form the new strategy. In the replication phase K pairs are compared.

All the *averages* are calculated over $S = 3000$ *random seeds*. The benchmark environment $\{[1, 1, 1, 1, 1], [0, 0, 0, 0, 0]\}$ shows that 3000 seeds is indeed sufficient. After a number of periods the market becomes more or less stable and the offers and average price only fluctuate within a certain range, mainly due to mutation. We denote this behaviour as an "*equilibrium*" in which the offers of intramarginals are close to the equilibrium price range and the agents choose the time to submit that showed to perform the best given these offers. All the results are averaged over periods 41 – 50 to avoid the random impact of the first learning periods. We show that the distribution of submission moments is stable after 40 periods and thus the impact of the first learning periods is negligible. This is done by conducting a two-sample Kolmogorov-Smirnov test on periods 39 and 40. The test statistic $D = 0.0022$ lies outside the critical region $D > 0.0159$ for $\alpha = 0.001$. Chapter 4 studies this impact by considering both the learning and the equilibrium phase.

5.4 Benchmark environment

In this section we focus on the distribution of arrival moments and its correlation with the chosen offer, in a benchmark environment. Important characteristics as efficiency, variance of transaction prices and volume are considered both per trading period as per possible arrival moment. We consider simulations of the basic environment with the following redemption values: $\{[1, 1, 1, 1, 1], [0, 0, 0, 0, 0]\}$. In this environment all buyers and all sellers are identical and buyers and sellers are symmetric. In Fig. 5.2 we show the distribution of arrival times and the expected offer per arrival moment for every trader individually, where buyers are represented by solid lines and sellers by dotted lines. Averages are shown in solid black lines. Bids of buyers are positively correlated with the moment of arrival, asks of sellers are negatively correlated. Buyers' valuations are shown in the offer function plotted by stars and sellers' costs by circles.

When multiple traders have the same valuation or cost, colours do not correspond to the colours of the offer functions, but are shown in black.




Figure 5.2: Distribution of learned submission moments (left) and offers as function of submission moment (right) in environment $\{[1, 1, 1, 1, 1], [0, 0, 0, 0, 0]\}$. Traders learn to submit their order more often during the middle of the period. The increasing bid function and decreasing ask function show that early submitting traders are more aggressive.

We derive more intuition about the distribution of submission moments by further investigating the environment with 5 buyers and 5 sellers. In Section 5.5 we study robustness with respect to the number of traders. This environment is attractive since it is an intermediate case. In the environment with one trader on either side of the market the probability of trading does not play a role, and as the number of traders converges to infinity the effect of timing on the expected profit from trade fades away. However, in this intermediate environment both effects play a significant role.

Moment of order submission

With respect to the preferred moment of submitting we observe in Fig. 5.2 that agents desire to submit their order in the middle part of the trading period. This illustrates the trade-off any individual trader faces: submitting the same offer earlier increases the probability that a trade will occur at the price of their own offer which results in a lower expected profit, submitting later decreases the probability of trading. A peak at the end of the period exists and we observe an

increasing bid function and decreasing ask function. The latter conclusion is also drawn in the paper Fano and Pellizzari (2011). Traders who submit their offer late, bid close to the valuation or ask close to the cost, which yields a high probability of trading. It makes sense for such a late offer to be submitted as late as possible. This ensures that the trade occurs at the preferable price of the other trader, which outweighs the minimal decrease in the probability of trading.

Traders learn which arrival moment performs best when the offers of intramarginal traders are not identical. Under the IEL algorithm, mutation is the main cause that strategies do not entirely converge. Due to mutation the trade-off between the probability of trading and the expected profit is of importance and the moment of submission plays a crucial role. This results in the distribution shown in Fig. 5.2. The micro-motives of this distribution are further investigated in Section 5.5.

Offer

Our main observation, with regards to the offer that agents submit in relation with the preferred moment of submitting, is that the earlier they intend to submit their order, the further bids are from the valuations and asks from the costs. If an agent prefers to submit at a late moment, he intends for a lower profit to increase his probability of trading. Thus we find a positive correlation between bids and time and negative between asks and time.

Characteristics per possible arrival moment are shown in Fig. 5.3. The average profit per transaction is shown in panel (a), conditional on the arrival moment. This excludes the instances where a trader arrives at that moment but does not trade. The next panel shows the probability of trading; defined as the number of trades divided by the number of arrivals at a given moment. Panel (c) shows the average profit over all the instances in which a trader arrives at that moment. Included in this average are the instances at which no trade occurs and a trader receives a profit of 0. This panel is the product of the first two panels; the probability of trading and the average profit of a transaction. Finally the standard deviation of transaction prices is shown for every arrival moment in panel (d).

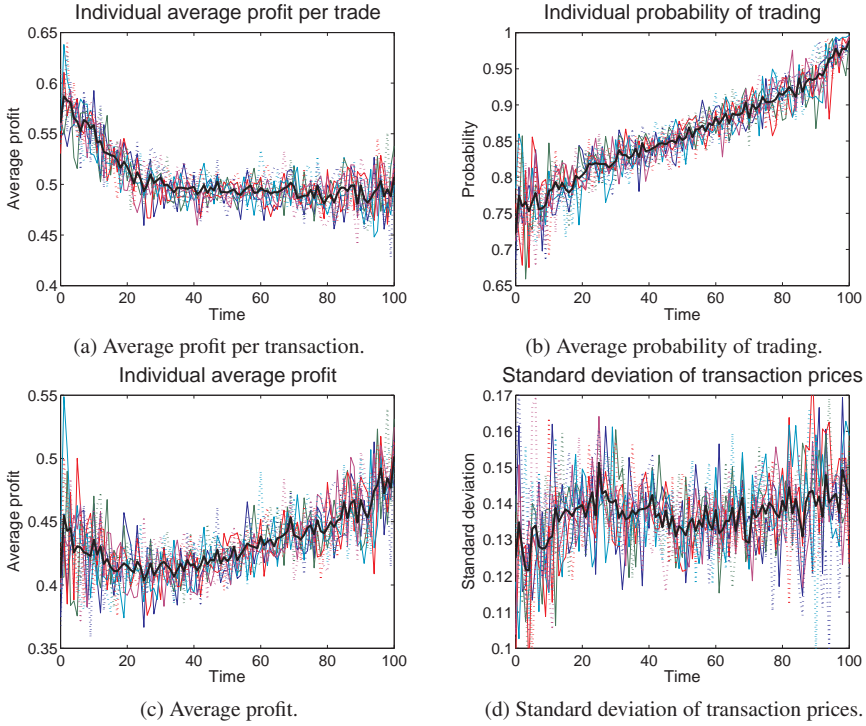


Figure 5.3: Characteristics environment $\{[1, 1, 1, 1, 1], [0, 0, 0, 0, 0]\}$ per possible arrival moment. The average profit per transaction is decreasing and the average probability of trading increasing. The average profit, which is the product of the two, is U-shaped. The standard deviation of transaction prices is increasing at the end of the period and is μ -shaped.

The average profit per transaction is decreasing and the probability of trading is increasing during the period. Early submitting traders submit a very aggressive offer, a bid below and a ask above one half. This will often not result in a trade, but if a trade occurs this will yield a very high profit. An early or late arrival leads to a higher average profit, but also a higher variance of transaction prices. A higher variance of the price at which the trade occurs (if late), or a higher variance because of more occasions where the agent does not trade (if early). In the replication process early and late strategies are therefore often removed from the pool and that is why

agents more often submit in the middle of the trading period. It is remarkable that in the IEL-algorithm traders learn to arrive more frequent at moments that yields a lower average profit. This may be the result of the learning algorithm, which only considers the profit in the previous period. It is likely that when the IEL-algorithm is modified in such a way that hypothetical profit is averaged over multiple periods, a reversal will occur and traders more frequently submit at the beginning or the end of the period.

Let us suppose that traders do not condition their offer on the moments of submitting, thus always submit the same bid or ask. In this case the probability of trading would be decreasing over time because fewer possible trading partners remain, whereas the expected profit per trade is increasing over time because it is more likely to occur at the preferable price of the other trader. However, the buyers' bid function is increasing and the sellers' ask function is decreasing. This results in a reversal of these two effects.

5.4.1 Knowledge of the submission moments

We compare the results in the benchmark environment with 5 buyers and 5 sellers to the setting where traders only observe the sequence of arrivals and the offers of traders, but not the actual submission moments. In the latter setting traders choose from the set $[1, 2, \dots, 10]$ in which position in the sequence they want to submit their order. We observe in Fig. 5.4 that knowledge of the actual arrival moment in addition to the sequence impacts the arrival moment distribution. This is confirmed with a two-sample Kolmogorov-Smirnov test. The test statistic $D = 0.0530$ lies in the critical region $D > 0.0051$ for $\alpha = 0.001$. The additional knowledge reduces the kurtosis and increases the peak at the end of the period.

5.4.2 Allowing the choice of submission moment

For the two-dimensional model with timing and the one-dimensional model without timing, we measure the overall results for efficiency, number of trades and volatility in the benchmark environment with 5 buyers and 5 sellers over trading periods 41 – 50. These characteristics and

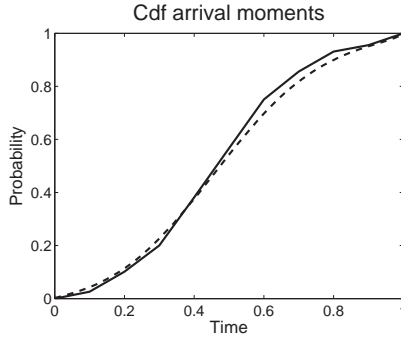


Figure 5.4: CDF of arrival times for knowledge of sequence (solid line) and actual arrival moments (dotted). The distribution of arrival moments is significantly altered when knowledge of the actual arrival moments of others in the previous period is added. This additional information reduces the kurtosis and increases the peak in the distribution of arrival moments at the end of the period.

	With timing	Without timing
Efficiency	0.8567 (0.1718)	0.8952 (0.0447)
Price Volatility	0.0208 (0.0101)	0.0177 (0.0077)
Number of transactions	4.2837 (0.8591)	4.3520 (0.1290)

Table 5.1: Average outcomes with and without timing in environment $\{[1, 1, 1, 1, 1], [0, 0, 0, 0, 0]\}$. The average efficiency and the average number of trades significantly decrease, and the average price volatility significantly increases when traders are allowed to submit orders at their preferred moment.



their standard deviations are shown in Table 5.1. Allowing traders to submit at their preferred moment has a negative effect; the average efficiency and the average number of transactions decrease, and the average price volatility increases. These comparisons are all significant at a significance level of 1%. It is optimal not to allow traders this extra decision, since it results in a lower efficiency and thus a lower profit.

5.5 Size of the market

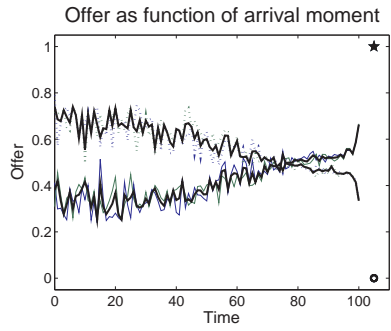

The size of the market plays a crucial role on the distribution of arrival moments. Two forces are important for the expected profit: the expected profit from a transaction and the probability of trading. In thinner markets the first has a larger impact than in thicker markets. In Fig. 5.5 we show the distribution of arrival moments and the average offer per trading moment depending on the number of traders. The results are given for 1, 2, 5, 10 and 15 traders on either side of the market.

Moment of order submission



In a market with only one buyer and one seller the moment of arrival does not affect the probability of trading. It is optimal to submit the order at moment 100, since submitting after the other trader results in a trade at the price of the other trader. In the simulations we find indeed that traders submit as late as possible. Submitting later strongly dominates submitting earlier in the IEL algorithm and hence in a market with one buyer and one seller the IEL algorithm selects submission moment 100. A similar distribution of submission moments is shown for two buyers and two sellers. When the size of the market increases, the probability of trading does play a role. Also the effect of the moment of arrival on the expected profit from trade decreases, since the probability that the transaction price equals the own offer tends towards one half for every arrival moment. With five traders on either side of the market traders submit around the middle of the period. The larger the size of the market, the earlier traders arrive. The simulations suggest that the moment of arrival will converge to zero as the size of the market converges to infinity, which would in the benchmark environment with infinitely many traders be optimal as the effect of the expected profit from trade disappears.



Dynamics in environment with 1 buyer with valuation 1 and 1 seller with cost 0.



Dynamics in environment with 2 buyers with valuation 1 and 2 sellers with cost 0.



Dynamics in environment with 5 buyers with valuation 1 and 5 sellers with cost 0.




Figure 5.5: Distribution of learned submission moments (left) and offers as function of submission moment (right), for 1, 2, 5, 10 and 15 traders on either side of the market. In the environment with one buyer and one seller both traders submit as late as possible, since the probability of trading does not depend on the submission moment. As the size of the market increases, this probability does play a role and traders submit earlier and earlier, and moreover bids increase and asks increase for every possible arrival moment.

Offer

Irrespective of the arrival moment traders on average submit a higher bid and a lower ask as the size of the market increases. Average bids and asks of early arriving traders are below respectively above one half, but the intersection points of the bid and ask function with one half approaches $n = 0$. For traders arriving late the average bids and asks are further away from one half and thus closer to their valuation and cost.

5.6 Competition

In this section we study how the distribution of arrival moments and the function of average offers is affected by different aspects of competition. Sets of environments are considered in which the size of one side of the market increases, competition to extramarginal traders increases, additional extramarginal traders enter or the range of equilibrium prices decreases.

5.6.1 Decreasing competition between buyers, increasing competition between sellers

Competition between buyers decreases and competition between sellers increases as more and more sellers enter, as illustrated in Fig. 5.6. These environments range from a large difference in size between both sides of the market, towards a symmetric environment.

Moment of order submission

When there is little competition between sellers they tend to trade late. After more and more sellers enter, the intramarginal sellers trade earlier and earlier. Intramarginal buyers trade early to outcompete extramarginals. In the final environment in Fig. 5.6 the distribution is symmetric and traders prefer not to trade too early. The extramarginal buyers submit every moment with the same probability in the first environment, and act more like the intramarginals as more sellers enter; in which case they have more opportunities to trade.

Offer

Sellers submit lower asks as more sellers enter; since their market power decreases they can be less aggressive. The intramarginal buyers submit a bid higher than the extramarginal buyers. When a seller enters, the buyer that becomes the most competitive extramarginal buyer significantly increases his bid. The other intramarginal buyers lower their bid, so that they slightly overbid the most competitive extramarginal buyer. The average offer of extramarginals is not affected by their moment of submitting their order.

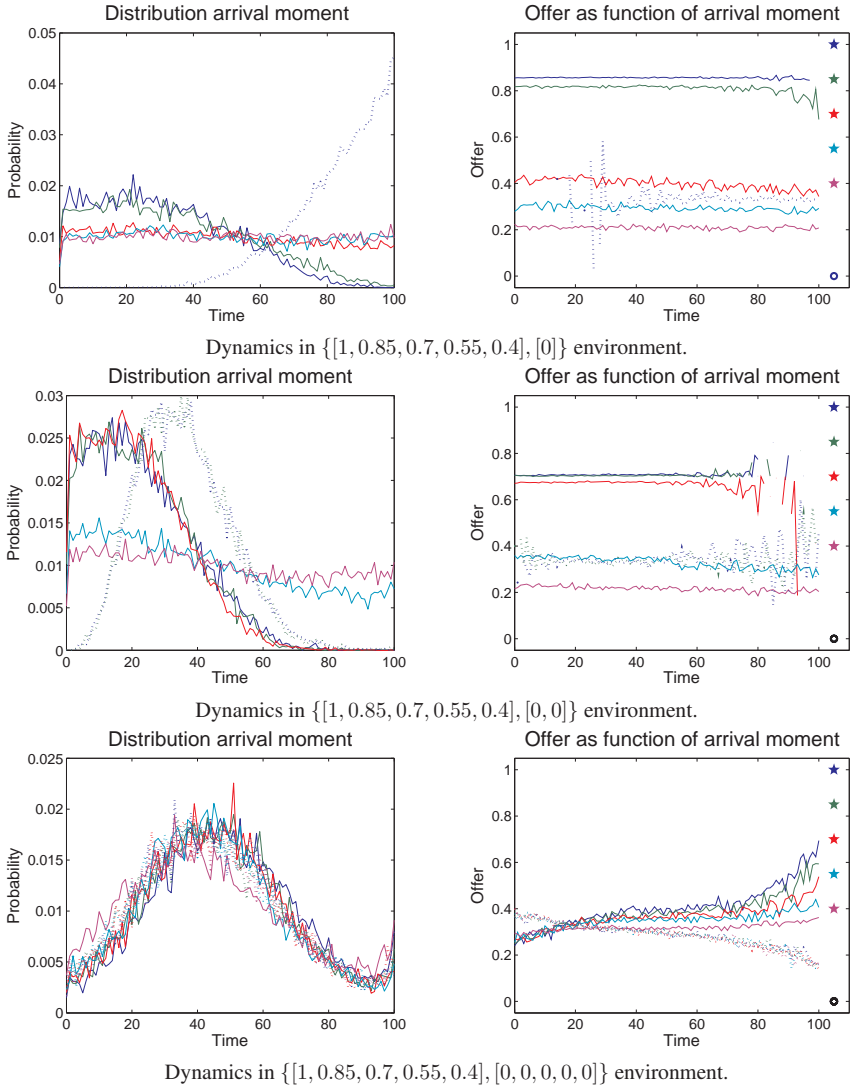


Figure 5.6: Distribution of learned submission moments (left) and offers as function of submission moment (right) with decreasing competition between buyers and increasing competition between sellers. With little competition between sellers, they tend to trade late in order to trade with the buyer who submitted the highest bid, and intramarginal buyers trade early. When more sellers are added to the environment the arrival moments move towards the middle of the period. Intramarginal buyers slightly overbid the most competitive extramarginal buyer.

5.6.2 Increasing competition to extramarginal traders

Competition between intramarginal traders and the two extramarginals increases in Fig. 5.7 as their valuation and cost become closer to the equilibrium price range.

Moment of order submission

When extramarginals in Fig. 5.7 have little opportunity to trade they have no clear preferred arrival moment. As they can compete more with intramarginal traders they prefer to trade earlier to increase their probability of trading. The intramarginals that set the range of equilibrium prices, and thus are the traders that face the most competition from extramarginals, trade earlier than the other intramarginals when competition to extramarginals is less. They behave more similar to the other intramarginals when competition increases.

Offer

When competition increases and the valuation and the cost of the extramarginals get closer to the equilibrium price range, they post less aggressive offers. As a result, the intramarginals that set the range of equilibrium prices also submit less aggressive offers to ensure that their offers are better than the offers from the extramarginals.

5.6.3 Extramarginal traders entering

In Fig. 5.8 we show environments with zero, one and two extramarginals on either side of the market. Over these environments the same intramarginals face an increasing probability of absence of trade due to competition to extramarginals.

Moment of order submission

Intramarginals tend to trade earlier when extramarginals enter. The intramarginals that set the range of equilibrium prices tend to trade even earlier than the rest of the intramarginals. The most competitive extramarginals in the last environment prefer to trade earlier to increase their probability of outcompeting an intramarginal.

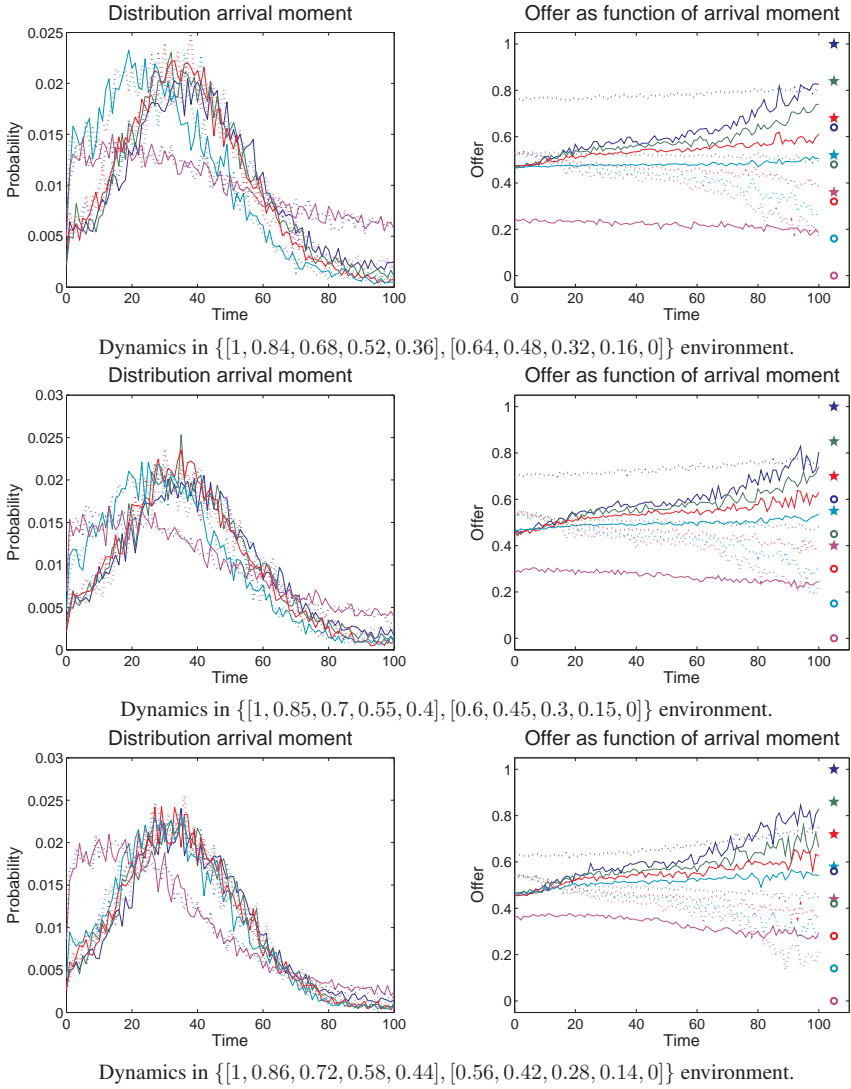


Figure 5.7: Distribution of learned submission moments (left) and offers as function of submission moment (right) with increasing competition to extramarginal traders. As competition to extramarginal traders increases, these traders learn to submit early and post less aggressive offers to increase their probability of trading. The weakest intramarginal traders submit later and submit less aggressive offers.

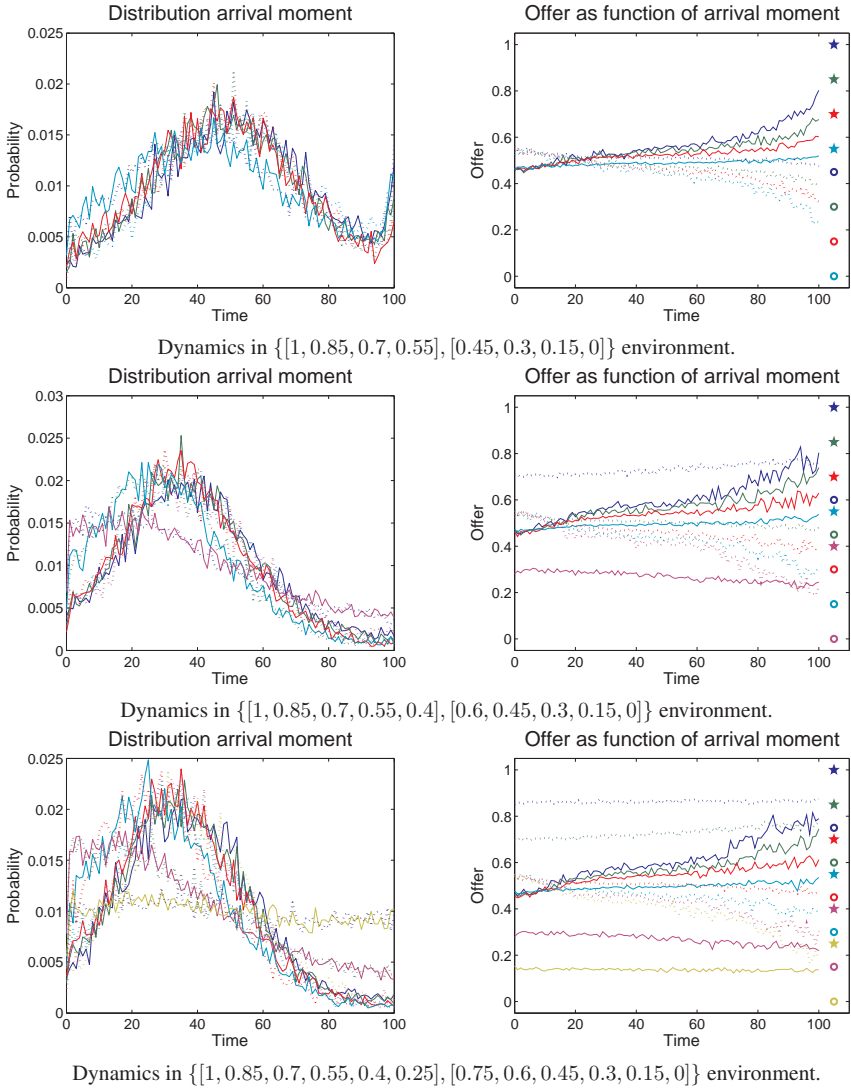


Figure 5.8: Distribution of learned submission moments (left) and offers as function of submission moment (right) as more intramarginal traders enter the market. This increasing competition forces intramarginal traders to submit their order earlier.

Offer

There is not a clear effect of extramarginal traders entering in Fig. 5.8 on the average offer of intramarginals. Hence the intramarginals solely respond to extra competition to extramarginals by trading earlier, not by altering their offer function. The extramarginal traders submit a similar offer for every arrival moment.

5.6.4 Decreasing range of equilibrium prices

Over the environments of Fig. 5.9, all valuations decrease by 0.02 and all costs increase, which reduces the equilibrium price range. Thus competition between intramarginal traders is increased and competition to extramarginals decreased.

Moment of order submission

As the range of equilibrium prices decreases in Fig. 5.9, the intramarginals that set the range of equilibrium prices tend to trade earlier. The remaining intramarginals alter their moment of arrival very limited. Extramarginals prefer to trade early, but this effect decreases as the equilibrium price range decreases and their valuation and cost are relatively further away from this range.

Offer

The offer functions of traders that set the equilibrium price range become more constant as the equilibrium price range decreases. The average offer function lies within the equilibrium price range. There is no clear impact of the increasing competition on the other intramarginals and the decreasing competition on extramarginals. However, the change in valuations and costs naturally causes the average bid to decrease and the average ask to increase.

5.7 Gode-Sunder environments

Gode and Sunder (1997) study the impact of extramarginals in the so-called GS-environments. These environments consist of one seller with cost 0, one buyer with valuation 1 and a set of ex-

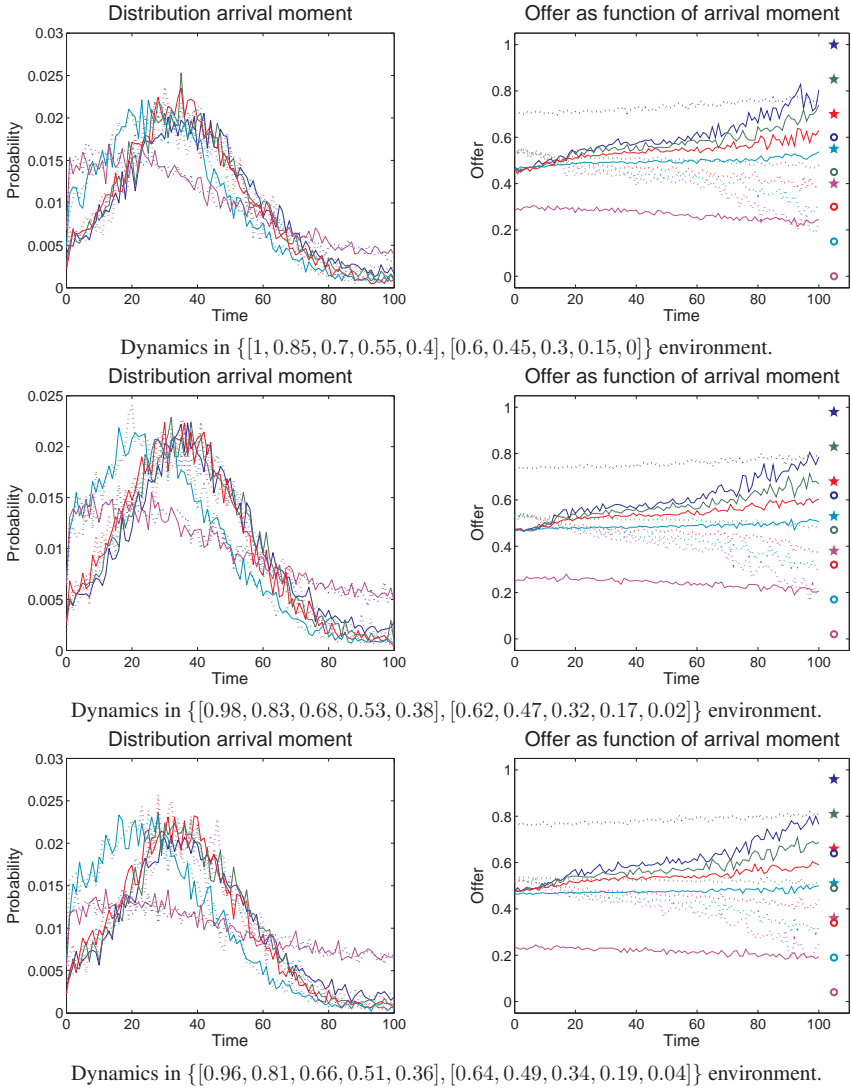


Figure 5.9: Distribution of learned submission moments (left) and offers as function of submission moment (right) when the range of equilibrium prices decreases. The least competitive intramarginals submit earlier and the extramarginal traders later.

tr marginal buyers with valuation β . In this analysis we consider three extramarginal buyers under different values of β . The demand and supply function of the GS-environment with $\beta = 0.5$ is shown in Fig. 5.10. Anufriev et al. (2013) determine efficiency in these GS-environments in a simple CDA. They show that the efficiency under the IEL-algorithm is very close to one and significantly larger than under Zero Intelligence. We investigate the impact of timing in the GS-environments in Fig. 5.11 and compare efficiency with Anufriev et al. (2013) in Fig. 5.12.




Figure 5.10: The demand and supply function of the GS-environment with 1 buyer with valuation 1, 1 seller with cost 0 and 3 extramarginal buyers with valuation $\beta = 0.5$.

Moment of order submission

The seller submits late and makes use of his market power to face the best possible bid and trade at the price of the buyer. As β increases the seller has a weaker incentive to enter late. The intramarginal buyer faces more competition and is forced to trade earlier. As extramarginals are competing more with the intramarginal buyer they tend to trade earlier. Arriving early can be used to outperform other buyers if the seller submits early.

Offer

The intramarginal buyer increases his bid as β increases to outbid other buyers. The seller increases his ask to trade more often with the intramarginal buyer; which yields a higher profit. Early arriving extramarginals relatively bid higher to outbid other early arriving buyers.

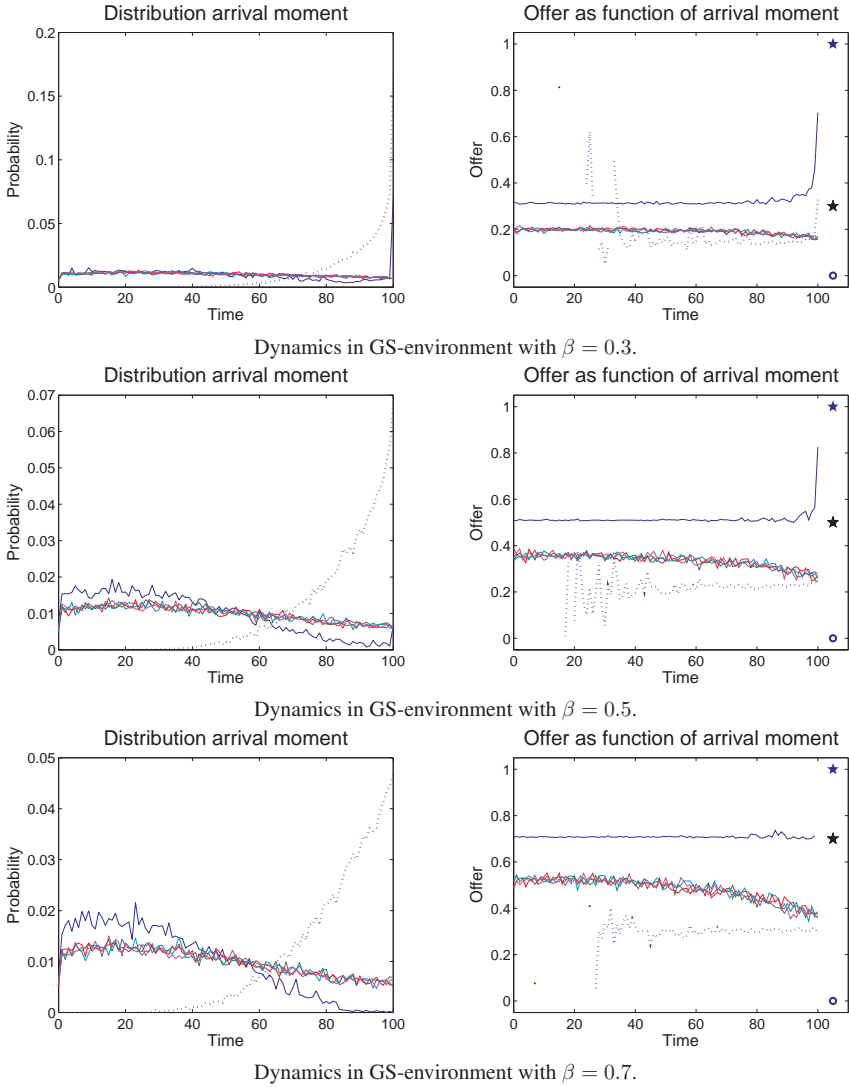


Figure 5.11: Distribution of learned submission moments (left) and offers as function of submission moment (right) in the GS-environment with 3 extramarginal buyers with valuation β . The seller submits his offer late in order to make use of his market power to trade against the best possible bid. As β increases the intramarginal buyer faces more competition, requiring him to increase his bid.

The efficiency function shows the same non-monotonicity as in Gode and Sunder (1997) and Anufriev et al. (2013). However, as stated earlier, with a more complex decision problem for traders that yields more freedom we find a lower efficiency than in Anufriev et al. (2013). The number of trades is increasing and the volatility decreasing in β .

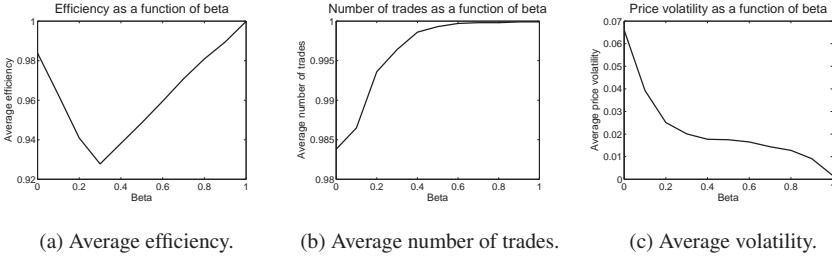


Figure 5.12: Characteristics of the GS-environments with 3 extramarginal buyers with valuation β . The efficiency as a function of β follows a U-shape similar to Gode and Sunder (1997) and Anufriev et al. (2013). The number of transactions is increasing and the price volatility decreasing in β .

5.8 Concluding Remarks

In the Continuous Double Auction the moment of submitting the order plays a crucial role; submitting at the end of the trading period may yield a lower probability of trading, submitting at the beginning of the period will most likely result in a trade at the own submitted price which yields a lower profit. This chapter is a step forward to a more complete model of learning in markets. Moreover, it is distinguished from other papers by the decision of traders. 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 the offer during the trading period. We showed that the size of the market and competition between traders influence this distribution.

The distribution of arrival moments is studied in a benchmark environment under full information about trading history. We found in simulations that under the Individual Evolutionary

Learning algorithm intramarginal traders learn to submit their order around the middle of the trading period. This result holds for a medium size market with a comparable number of traders on each side. If one side of the market is thinner it can extract more profit by submitting later. Our main observation with regards to the offer that agents submit, in relation with the preferred moment of submitting, is that the earlier they submit their order, the higher profit they aim for. If an agent submits his order at a late moment, he submits a conservative offer to increase his probability of trading. As a result, an early or late arrival results in a higher expected profit. However, an early arrival increases the risk of not trading and a late arrival results in a higher price volatility. Therefore agents tend to trade more often in the middle of the period. This shows that in the IEL-algorithm traders learn not to submit risky strategies, resulting from the algorithm that considers only the performance of strategies in the previous period. A possible future research subject would be to adjust the IEL-algorithm in such a way that the average profit over multiple periods is considered. This may result in a more realistic outcome and traders may prefer to enter at the beginning or the end of the period.

Allowing traders to submit at their preferred moment has a negative effect; the expected efficiency and the expected number of trades decrease significantly and the expected price volatility significantly increases. Hence, allowing traders to make this extra decision results in a lower expected profit. It is optimal not to allow traders this extra decision.

When the size of the market increases, the probability of trading and the probability that trade occurs at the price of the own offer change. The larger the size of the market, the earlier traders submit their order. It appears that the moment of submission will converge to zero as the size of the market converges to infinity. Irrespective of the submission moment traders on average submit a higher bid and a lower ask as the size of the market increases. We conclude that the size of the market is of great importance to the distribution of submission moments.

Ceteris paribus as competition increases in some sense, the probability of trading decreases. Intramarginal traders are hence forced to trade earlier and to submit less aggressive offers to

cope with the increased competition, which increases the probability of trading. Extramarginal traders have a clearer preference for their submission moment when they have more opportunity to trade, in which case they submit early to outcompete others.

We found that under the Individual Evolutionary Learning algorithm investors in a medium size Continuous Double Auction market learn to submit their order around the middle of the trading period to avoid a lower trading probability or profit. The earlier traders submit the more aggressive offer they submit and thus aim for a higher profit. In a large market the latter effect reduces and traders submit earlier and earlier. Moreover, we have shown how the distribution of submission moments and the expected offer as a function of the submission moment change with the amount of competition and the size of the market.