Timing under Individual Evolutionary Learning in a Continuous Double Auction*

Michiel van de Leur [a,†]   Mikhail Anufriev [b,‡]

[a] *CeNDEF, Amsterdam School of Economics, University of Amsterdam, Valckenierstraat 65-67, NL-1018 XE Amsterdam, Netherlands*
[b] *Economics Discipline Group, Business School, University of Technology Sydney PO Box 123, Broadway, NSW 2007, Australia*

Abstract

The moment of order submission plays an important role for the trading outcome in a Continuous Double Auction; submitting an offer at the beginning of the trading period may yield a lower profit as the trade is likely to be settled at the own offered price, whereas late offers result in a lower probability of trading. This timing problem makes the order submission strategy more difficult. We extend the behavioural model of Individual Evolutionary Learning to incorporate the timing problem and study the limiting distribution of submission moments and the resulting offer function that maps submission moments to offers. We find that traders learn to condition their orders on the submission moment, and that this behaviour exacerbates efficiency loss from learning. Traders in a medium size market learn to submit the order around the middle of the period, but in a large market tend to submit their orders earlier and earlier. If traders evaluate profitability of their strategies over longer history, orders are submitted later, but the same effect of market size is found.

Keywords: Bounded rationality; Individual evolutionary learning; Agent-based models; Moment of order submission; Order-driven market.

JEL Classification: D83, D44, C63.

*We thank Jasmina Arifovic, Valentyn Panchenko and participants of the Computing in Economics and Finance conference 2013 at Vancouver, Canada, for their helpful comments and useful suggestions and comments. Michiel van de Leur acknowledges financial support by the European Commission in the framework of the European Doctorate in Economics-Erasmus Mundus (EDE-EM). Mikhail Anufriev acknowledges financial support from the Australian Research Council through Discovery Project DP140103501.

† *Phone:* +31 20 525 4125, *E-mail:* M.C.W.vandeLeur@uva.nl

‡ *Phone:* +61 2 9514 4078; *Fax:* +61 2 9514 7711, *E-mail:* Mikhail.Anufriev@uts.edu.au.

1 Introduction

Most of the automated trading systems of modern financial markets are operated under the Continuous Double Auction protocol. This protocol allows asynchronous trading when the unexecuted orders are stored in an order book. The moment of order submission may affect both profitability of trade and aggregate properties of the market such as price volatility even in the absence of information concerns. A trader who decided which limit order to submit faces the following *timing trade-off*. Offers submitted at the beginning of the trading period tend to be stored in the book, and hence are executed at the submitted price. Submitting early thus forsakes the possibility of getting a higher profit. On the other hand, offers submitted at the very end of the trading period may not be executed at all. Submitting late may thus result in getting zero profit. This timing problem is intertwined with the basic issue of choosing the optimal order strategy, making the strategical problem of order submission multi-dimensional.

Despite the importance of the timing problem, in many agent-based models of order-driven financial markets[1] traders submit their orders at random moments during a trading period. This assumption simplifies the traders' behaviour, as they are left then with a one-dimensional decision, such as choosing a bid or ask price (as in LiCalzi and Pellizzari, 2006) or forecasting a future price and condition its offer on this forecast (as in Chiarella and Iori, 2002). On the other hand, it is plausible that if traders are allowed to submit their orders at preferred moments of the trading session, their other decisions will change, as traders may condition their bids/asks or expectations on the moment of submission. Thus multi-dimensionality of the problem of order submission translates into multi-dimensionality of learning where learning both about the orders and about the timing of submission is important.

In this paper we contribute to the growing literature on agent-based modeling of financial market by studying the issue of timing of order submission in the markets organised as Continuous Double Auction. In our model traders learn about which bid/ask offers to submit and also about when to submit them during the trading session. To model this behaviour, we extend the Individual Evolutionary Learning (IEL) algorithm which was introduced by Jasmina Arifovic and John Ledyard in 2003 and published as Arifovic and Ledyard (2011). The IEL is studied in Arifovic and Ledyard (2007) and Anufriev et al. (2013) in the context of financial markets for the call markets and CDA, respectively.[2] The IEL algorithm assumes that traders select their active strategies from a pool of potential strategies. After a trading period the hypothetical payoffs of all strategies are calculated using the past information and some strategies are

[1]For a review of such models, see Chakraborti et al. (2011). Agent-based methodology avoids making an extreme assumption of full rationality and models traders as boundedly rational who learn from the past. The approach of bounded rationality seems to be a good research strategy for the analysis of the order-driven markets, given its complexity (even without the timing issue) and the possibility to have multiple equilibria even in the simplest cases.

[2]Anufriev et al. (2013) focused on studying market allocative and informational efficiency by comparing full and no information about the history of orders, and comparing the IEL-algorithm with zero intelligent behaviour (see Gode and Sunder, 1993, 1997). This paper extends Anufriev et al. (2013) by considering learning about the moment of order submission (in the full information case).

replaced with randomly modified strategies having higher hypothetical payoffs. 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. Our simulations show that the distribution of the submission moments highly depends on the size of the market as well as on the size of the memory used to evaluate hypothetical payoffs.

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 orders 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 influences the distribution or submission moments. Extending the IEL-algorithm to allow for evaluating strategies over multiple past periods shifts the distribution of submission moments. Traders submit later and later, which on average results in the highest profit. We also discuss the resulting offers, market allocative 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.

An important feature of IEL is that it is a backward-looking learning algorithm. This approach should be contrasted with the standard economic approach of full rationality where optimising agents make their decisions and use all information rationally. However, also in the fully (or almost fully) rational setting the timing is often random, see the survey of Parlour and Seppi (2008). In Friedman (1991) traders can submit orders at any moment in 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. Under these assumptions their model predicts the occurrence and nature of the next transaction. In Roşu (2009) traders submit orders following a Poisson submission process. At the submission moment the order is determined from a trade-off between opportunity cost and future profit. Whereas implications of rationality on the timing problem are important and interesting, we stick here to the behavioural approach for several reasons. First, numerous experimental studies show that people have difficulty to behave rationally even in the most simple environments. Nevertheless, when the task is replicated in a fixed environment, agents still tend to reach a reasonable behaviour with good welfare consequences. Thus, to study realistic market outcomes one may not need to make extreme and unrealistic assumptions. Second, rational models are naturally restricted in their scope. Consequently many agent-based studies appeared recently[3], whereas many researchers call for more studies, e.g., Dawid (1999b), LeBaron (2001), Lux (2009), Hommes (2013).

Another stream of literature which is related to the problem we study is devoted to the

[3] See, in particular, Dawid (1999a), Bottazzi et al. (2005), Anufriev and Panchenko (2009), Ladley and Pellizzari (2012), Chiarella and Iori (2009), Chiarella et al. (2014) and Leal et al. (2015).

reaction of traders to events during a trading day. Biais et al. (1995) study the empirical distribution of large and small trades, orders and cancellations during a trading day. They find a U-shaped distribution of orders and 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 our paper there is no modeling of the process of events during trading days. Rather we are interested in how traders are able to find their trading moments and strategies in a market environment where other traders are doing the same. In fact, this is a typical setup in experiments with financial markets which can be used as a testing device for various theories. One common regularity found in the experiments is that the competitive equilibrium (if found) arises only after replications. Our model is also built in the same spirit: we assume that the types of traders (defined by their role in the market and by their valuations and costs) are fixed and do not evolve between trading periods. What evolves are trading strategies which are updated due to individual learning. Thus, in our model (as in many others) the current state of the order book cannot be used by traders, which resembles very active markets where the order book changes faster than traders can react.

As we mentioned before, the timing issue was largely ignored in the agent-based literature. It is, indeed, typical to assume that investors submit their one unit orders at random moments during the trading period and that between the periods their learning is only one-dimensional. In some models buyers learn which bid to submit and sellers learn which price to ask, and in some models the bids and asks depend on the expectations and the 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, 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, Fano et al. (2013) who use a Genetic Algorithm in which traders mimic well-performing strategies of others and Anufriev and Panchenko (2009) who study the switching between forecasting rules. In all simulations under Continuous Double Auctions, agents submit orders in a random sequence. Therefore, our goal in this paper is to relax this assumption and investigate the consequences of this relaxation.

The organisation of this paper is as follows. The market environment and the Individual Evolutionary Learning algorithm are described in Sections 2 and 3, respectively. The distribution of the preferred moment of submitting, its relation to the bids and asks, and market efficiency are studied in Section 4 for a benchmark environment, where the buyers and sellers occupy a symmetric position in the market. The impact of the size of the market as well as asymmetry are considered in Section 5. In Section 6 the IEL-algorithm is extended to multi-period learning. Finally, Section 7 concludes.

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 buyer is willing to buy and each seller is willing to sell at most one unit of the good. The market is organised as a Continuous Double Auction where traders can submit only limit orders and where the unexecuted orders are stored in an order book to be executed later according to the price-time priority. Traders have to decide the offered price and the moment of submission.

2.1 The environments

Each environment is determined by a set of B buyers and a set of S sellers with their valuations and costs. In each trading period each buyer would like to consume one unit of the good and each seller can deliver one unit of the good. Buyer $b \in \{1, \dots, B\}$ has a fixed valuation $V_b \geq 0$ and seller $s \in \{1, \dots, S\}$ has fixed costs $C_s \geq 0$ that need to be incurred only in case of a transaction. Agents know their own valuations (or costs), but not these values of the other agents. We denote the environments by vectors of valuations and costs.[4] In this paper we will use the following two environments: the *benchmark* environment $\{[1, \dots, 1], [0, \dots, 0]\}$ with B identical buyers and S identical sellers and the so-called *GS*-environment $\{[1, \beta, ..., \beta], [0]\}$, in which one intramarginal buyer has valuation 1, one seller has cost 0 and $B - 1$ extramarginal buyers have valuations $\beta \leq 1$. This environment was studied in Gode and Sunder (1997) for the Zero-Intelligent traders and in Anufriev et al. (2013) for the traders using IEL for the order submission.

Based on the valuations and costs, the demand and supply functions can be determined and the competitive equilibrium can be derived. We denote the equilibrium quantity by q^* and the interval of equilibrium prices by $[p^*_L, p^*_H]$ where $p^*_L \leq p^*_H$. The traders that trade in the equilibrium are referred to as intramarginals, whereas the traders that will not trade in the equilibrium are called extramarginals. In the competitive equilibrium with price $p^* \in [p^*_L, p^*_H]$, the payoff of intramarginal buyer b equals to $V_b - p^*$, the payoff of intramarginal seller equals to $p^* - C_s$, whereas extramarginal traders have zero payoff.

The supply and demand functions for the benchmark environment with $B = S = 5$ and the GS environment with $B = 5$ buyers are shown in Fig. 1.

[4] 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.

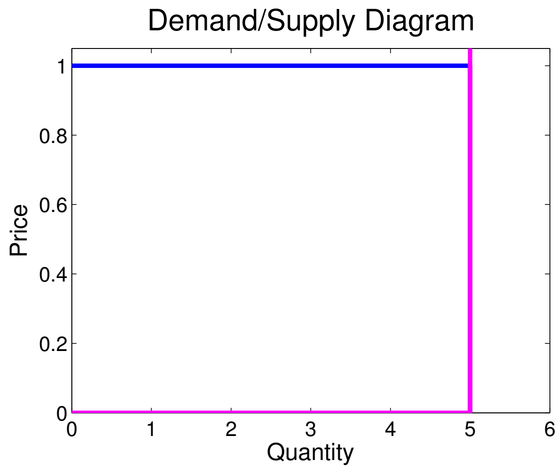

(a) The benchmark environment.

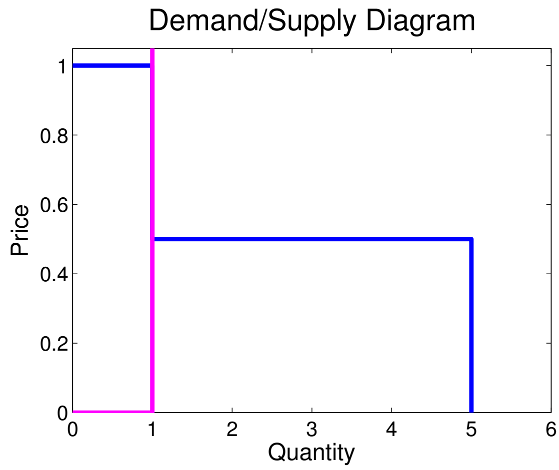

(b) The main symmetric environment.

Figure 1: The demand and supply functions of the main environments used in the paper, the benchmark environment with 5 buyers and 5 sellers, and the GS environment with 5 buyers and $\beta = 0.5$

2.2 Continuous Double Auction

A Continuous Double Auction model is often used to describe the asynchronous trading mechanism of stock exchanges. During trading period $t \in \{1, \ldots, T\}$, 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 $a_{s,t}$, respectively.

We discretise a trading period and split it into the *submission moments* $\{0, 1, \ldots, M\}$ when the traders can arrive to submit their orders. The set of submission moments of all traders determines the order in which traders arrive. When all traders nominate different submission moments, their order of submission is established in a natural way by sorting traders' submission moments: the trader with the smallest submission moment arrives first (and places her offer into the order book), the next order will be from the trader with the second smallest submission moment, and so on. An issue arises when a group of two or more traders decide to submit their orders at the same submission moment. Moreover, there might be several such groups, each with own submission moment for all traders of the same group. In this case, first, the submission order of the groups is determined by sorting their submission moments, and then the orders of the traders from the same group are arranged in a random order, uniformly over all possible queues.[5]

In simulations we set M sufficiently large to prevent congestion effects due to possible preference of traders to submit their orders in the same submission moment and, consequently,

[5]For example, if six traders choose submission moments 13, 4, 13, 46, 46 and 13, then the second trader arrives first, traders 1, 3 and 6 arrive next in a random order, and, finally, traders 4 and 5 arrive in a random order.

to avoid a large random component in the sequence of submission.[6] The submission moments of buyers and sellers in period t are denoted as $n_{b,t}$ and $n_{s,t}$, respectively. If a submitted order can be matched with the best order from the book, the transaction takes place at the price of the order stored in the book. If the arriving order cannot be matched, it is stored in the book. At the end of the trading 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 $p_{s,t}$. The payoff of a buyer during the trading session equals to $U_{b,t} = V_b - p_{b,t}$ if he traded, and zero otherwise. The payoff of a seller equals to $U_{s,t} = p_{s,t} - C_s$ if she traded, 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 two traders, one buyer and one seller, and their submitted offers are such that $b_{b,t} > a_{s,t}$, then the 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, $p_{b,t} = p_{s,t} = a_{s,t}$.

To focus on learning of traders about their timing of submitting the order and about the price of submission, we assume that no order can be cancelled and also 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.

2.3 Nash equilibria

In the model where traders are required to make a two-dimensional decision, *multiple* equilibrium outcomes may result. To illustrate this possibility with a simple example, we consider a one-period version of our model, where valuations and costs are common knowledge.[7] In this one-period model, traders are required to select a strategy consisting of an order and a submission moment. We focus on any competitive equilibrium in the market, and denote the corresponding equilibrium price as p^*.

Let us consider a generic case first, when no extramarginal trader demands or offers a good at this price. This is the case, for example, in any competitive equilibrium in the benchmark environment illustrated in Fig. 1. Consider the set of strategies, where all intramarginal traders submit p^*, but their submission moments are arbitrary. (Extramarginal traders submit any offer within their budget constraints at any submission moment, but they do not trade.) Such configuration is a Nash equilibrium. Indeed, under such strategy profile, all intramarginal traders trade and get a non-negative profit. On the other hand, submitting a different offer may result in a loss (either because of no trading due to a more aggressive offer or because of a worse transaction price due to a more conservative offer and unfortunate timing), though

[6] In particular, $M \gg B + S$. In benchmark simulations, when $B = S = 5$, we choose $M = 100$.

[7] This is a gross simplification. However, it would be hard to analyse the model with incomplete information even in the case of few participants. The point of this example is to describe possible equilibria and demonstrate the issue of their multiplicity even in the simplest case of complete information.

submitting at any other submission moment is payoff-equivalent. Moreover, all Nash equilibria would have this structure: if there are two intramarginal traders submitting different prices, there is always profitable deviation for at least one of them. We observe that adding submission moments to the strategy of traders significantly increases the number of equilibria.[8]

A non-generic case is also possible in which one side of the market has *extramarginal* traders willing to trade at the equilibrium price. An example is the environment $\{[1, 0.8], [0.4, 0.4, 0]\}$, where the unique equilibrium price is $p^* = 0.4$, but at this price one extramarginal seller can also supply a unit of the good. When one (long) side of the market has a number of extramarginal traders willing to trade at the equilibrium price, their competition changes the behaviour of intramarginal traders on the same side of the market (as they may lose an opportunity to trade) and, consequently, the behaviour of intramarginal traders at the other (short) side of the market. In the example above, note that the intramarginal sellers are no longer indifferent about the moment of order submission, as they may end up with no trading at all. It suggests that in this, non-generic case, the number of equilibria is smaller than in the generic case. However, this is not true, as now new types of equilibria are possible. For example, in one such equilibrium all three sellers submit at n, both buyers submit at $n + 1$, but the submitted orders are as follows: both buyers submit an offer of at least p^* and the two sellers with cost 0.4 submit $p^* = 0.4$, whereas the seller with cost 0 submits a lower price $0.4 - \varepsilon$ with $\varepsilon > 0$ to guarantee trading.

The Nash equilibria of this one-period model are natural candidates for the long run outcomes of the multi-period learning model used in this paper. Under the Individual Evolutionary Learning algorithm, we find in symmetric environments that the offers of traders converge towards the equilibrium price range, such as in the 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.

3 Individual Evolutionary Learning algorithm

In this paper we study the limiting outcome of repeated trading by the boundedly rational agents who learn (between the trading periods) both an offer to be submitted and also the moment when this offer should be submitted. Traders learn their strategies by the Individual Evolutionary Learning (IEL) algorithm. Under IEL, every agent has a finite memory allowing them to carry over only a limited number of potential strategies. The probabilistic choice of one active strategy from the set of potential strategies is based on the past hypothetical performances of strategies, i.e., performances that the strategy would have had in the past assuming the same market conditions (e.g., the same behaviour of other market participants).

[8]There are other sources of multiplicity of equilibria in this example: one is due to irrelevance of offers of the extramarginal traders, and another is due to a possibility of multiple competitive equilibria for given demand and supply schedules.

As opposed to other similar models[9], the IEL can be applied to environments with a large strategy space. In this respect, the IEL is similar to the Genetic Algorithms (see, e.g., Arifovic, 1994 and Dawid, 1999b) but has simpler and more intuitive interpretation.

We start with generalising the algorithm to handle multi-dimensional learning. Under the IEL, at each trading session t, every agent chooses an active strategy from the pool of potential strategies. In our case, every strategy is a pair of numbers specifying an offer to be submitted (bid for a buyer and ask for a seller) and a submission moment during the trading session when this submission should be made. At the beginning of the trading period the active strategy is selected from the pool with a probability increasing in its expected payoff, which is computed as the *hypothetical* payoff that this strategy would have had in the past.[10] Agents compute the hypothetical payoffs of the strategies using all information available to them. In the benchmark simulations we assume that past offers and the submission moments of all traders are publicly available. Therefore, each agent can determine exactly what his or her payoff would have been for each possible strategy in the previous periods, assuming no changes in the behaviour of other agents. Between the trading periods, the pool of strategies evolves. First, every strategy might mutate with a small probability allowing for some sort of experimentation. Second, some strategies in the pool are replaced with their positions filled the strategies with higher hypothetical payoffs.

Individual pool of strategies, mutation and replication

At period t every trader has an individual pool of strategies of size K. Buyer b has the set $B_{b,t}$ of K pairs of bids and submission moments $B_{b,t} = \{(b,n)_i\}_{i=1}^{K}$. Seller s has the set $A_{s,t}$ of K pairs of asks and submission moments $A_{s,t} = \{(a,n)_j\}_{j=1}^{K}$. To form the initial pools, we draw the offers (bids or asks) and the submission moments independently for every trader. The bids are drawn from a uniform distribution on $[0, V_b]$, the asks are drawn from a uniform distribution on $[C_s, 1]$, and the submission moments are drawn uniformly from the set $\{0, 1, \ldots, M\}$.

After every trading period the pool can be changed as a result of two operations. First, *mutation* takes place. For every strategy in the pool, each of the two parts mutates independently with a fixed small probability ρ. Given that a part of the strategy mutates, a normally distributed variable with mean zero is added to that part. The variance of the random variable depends on which part of the strategy mutates. The mutated submission moment is then rounded to the nearest integer. In any case, if the mutated strategy lies outside the strategy space, the previous mutation is discarded and a new normally distributed variable is drawn.

After some strategies in the old pool have possibly mutated, whereas others remained intact, the *replication* stage takes place during which the new pool is formed. For this purpose,

[9] Among closely related learning models we mention the model of reinforcement learning (Erev and Roth, 1998), the Experience-Weighted Attraction learning model (Camerer and Ho, 1999) and the Heuristic Switching Model (Anufriev and Hommes, 2012).

[10] For the previously active strategy this payoff is the real payoff of that strategy.

the foregone payoffs are calculated (as explained below) for all strategies that are currently in the pool, while taking the behaviour of all other traders from trading period t constant. Replication consists of a comparison of two strategies randomly selected from the old pool and placing the strategy with the highest hypothetical payoff to the new pool of strategies. For every agent this procedure is independently repeated K times to fill the entire pool. This new pool, denoted as $B_{b,t+1}$ for buyer b and $A_{s,t+1}$ for seller s, will be used in trading period $t+1$.

Hypothetical foregone payoff functions

The hypothetical foregone payoff functions correspond to the preference of traders imposed in the model. In general, for buyer b and strategy (b, n), it is given by

$$U_{b,t}(b,n) = \begin{cases} V_b - p_{b,t}(b,n), & \text{if strategy } (b,n) \text{ would result in a trade} \\ & \qquad\qquad \text{at price } p_{b,t}(b,n) \\ 0, & \text{otherwise.} \end{cases}$$

Similarly, for seller s and strategy (a, n), it is given by

$$U_{s,t}(a,n) = \begin{cases} p_{s,t}(a,n) - C_s, & \text{if strategy } (a,n) \text{ would result in a trade} \\ & \qquad\qquad \text{at price } p_{s,t}(a,n) \\ 0, & \text{otherwise.} \end{cases}$$

However, to apply these formulas, one should make specific assumptions about whether the given strategy resulted in a transaction and at which price. First of all, we will assume that the entire order book of the last period is publicly available, so that agents know the orders submitted by the others.[11] Moreover, we assume that the submission moments of all traders in the last period are publicly available. When the submission moments are available, every agent can simply simulate the whole trading session for every own (hypothetical) strategy, assuming that all other traders use the same strategies as in the previous period. Recall from Section 2.2 that in the cases when at least two traders have equal submission moments, the mechanism will choose one of several possible submission sequences randomly. We assume that the traders know it and, consequently, compute their foregone payoffs making corresponding expectations of the expressions given in the formula above.

To illustrate this with an example, let us compute the hypothetical foregone payoffs for the buyer in the market with only two traders, where the seller used strategy $(a_{s,t}, n_{s_t})$ in the previous period. Strategy (b, n) will result in zero payoff if $b < a_{s,t}$ for every n. On the other hand, when $b > a_{s,t}$, payoff of the strategy will be $V_b - a_{s,t}$, in the case when $n > n_{s,t}$ (because then the transacted price is the seller's early submission); $V_b - b$, in the case when $n < n_{s,t}$ (because then the transacted price is the buyer's early submission); and $\frac{1}{2}(V_b - a_{s,t}) + \frac{1}{2}(V_b - b)$, in the case when $n = n_{s,t}$ (because then both arriving orders of traders occur with equal probability).

[11] This case is referred to as the OpenBook system in Anufriev et al. (2013).

One may argue that knowledge of the submission moments is a rather strong assumption in our setup. Indeed, in reality the traders observe only the order of submission of other traders. Recall, however, that the submission moments were introduced in our setup as a tool to minimise the effect of randomness for the cases when, say, two traders are willing to arrive in the same order in the queue. In reality, in the small markets that we study, such simultaneous submissions may almost never take place, making the knowledge of past order of submission sufficient for traders to compute the hypothetical payoffs precisely.

As mentioned above, the hypothetical foregone payoffs are used to define the new pool through the pairwise comparison of the strategies. These payoffs are also used to determine the active strategy from the pool.

Selection of an active strategy from the pool

In the first period every strategy from the pool is equally likely to be chosen. At the later stages, when the hypothetical foregone payoffs are known, the active strategy is selected with probability proportional to this payoff. For example, for buyer b the probability of selecting strategy $(b,n)_i$ to be used in period $t+1$ is given by

$$\pi_{b,t+1}\big((b,n)_i\big) = \frac{U_{b,t}\big((b,n)_i\big)}{\sum_{k=1}^{K} U_{b,t}\big((b,n)_k\big)}\,.$$

The Individual Evolutionary Learning algorithm depends on some parameters, such as the size of the individual pools, the probability and the distribution of mutation and the replication rate. We will now discuss the values of these parameters used in the simulations, as well as the characteristics used to describe the overall outcome in a trading period.

3.1 Methodology of simulations

In the remaining part of the paper we simulate our learning model. Most of the parameters of the Individual Evolutionary Learning algorithm are the same as used in Arifovic and Ledyard (2007). Compared to their paper, we solely increase the size of the pool of strategies, since traders are now required to make a multidimensional decision. Every trader 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 submission moment mutates, a normally distributed term with mean 0 and a standard deviation of 10 is added to the submission moment. The mutated offer is truncated on the bounds of the interval $[0,1]$ and the mutated submission moment is rounded to the nearest integer in $\{0,1,\dots,100\}$. When the mutated strategy lies outside the strategy space, a new normally distributed variable is drawn. Note that the mutation we use differs from Anufriev et al. (2013), where the new strategy was selected from a uniform distribution. In the replication phase K pairs are compared.

We begin to simulate the IEL for the benchmark environment, when the market is symmetric and small ($B = S = 5$), as illustrated in the left panel of Fig. 1. This gives a flavour of what our IEL model produces. Next we simulate similar environment with a larger number of traders (keeping the number of buyers and sellers equal) to study the effect of the size in such symmetric market. Finally, we will simulate a asymmetric environment, namely the GS-environment illustrated in the right panel of Fig. 1.

We are especially interested in the outcome of learning regarding the submission moments. Recall that the traders face the timing trade-off between higher realised profit and higher probability to trade, and we investigate the outcome for this trade-off that emerges from learning. First, we look at the distribution of the submission moments, by evaluating *for each trader* the number of times this trader submitted at a given submission moment, divided by the total number of submissions by this trader (i.e., the number of trading periods). Next, we look at the dependency between the submission moments and the offer being submitted, when, again, for each trader we average all the offers made at a given 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. If all intramarginal buyers and sellers trade during a period, the allocative value is maximal, and the outcome of trading is fully efficient. Efficiency can be lower when an extramarginal agent trades, or when an intramarginal agent does not trade. Furthermore, we study the *price volatility* and the *number of transactions*. All these characteristics are considered per trading period as well as per possible submission moment and are compared with the one-dimensional model.

All the *averages* are calculated over $S = 3000$ random seeds and over ten trading periods $41 - 50$. Simulations in the benchmark environment with 5 buyers and 5 sellers show that already after about twenty initial 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*. There 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. The transitory periods that we skip reflect the randomness due to initialisation. But a general feature of the IEL algorithm is to evolve and converge very fast. We check that the distribution of submission moments is stable after 40 periods by conducting a two-sample Kolmogorov-Smirnov test on the submission distribution for the periods between 41 and 50. It turns out that during these periods the distribution solely fluctuates randomly, resulting in a insignificant change when comparing periods 41 and 50, with a test statistic equal to $D = 0.101$.

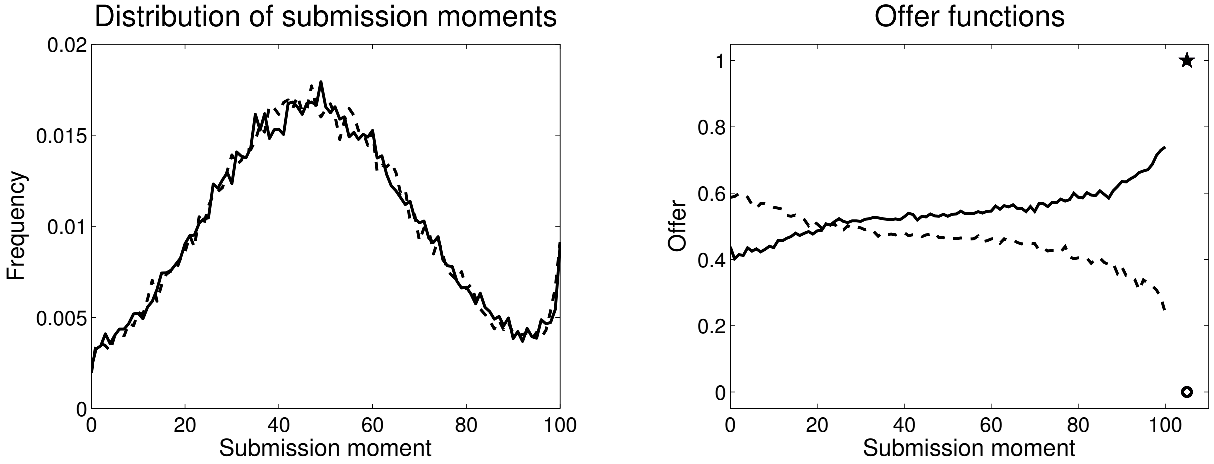

Figure 2: Distribution of the submission moments (*left*) and offers (bids and asks) as functions of the submission moment (*right*) in the environment $\{[1,1,1,1,1],[0,0,0,0,0]\}$ under IEL

4 Benchmark environment

We consider simulations of the basic environment with the following valuations and costs: $\{[1,1,1,1,1],[0,0,0,0,0]\}$. In this environment all buyers and all sellers are identical, and the market is symmetric. The left panel of Fig. 2 shows the distribution of the submission moments, where buyers are represented by the solid line and sellers by the dotted line.[12] Despite small variations, it is clear that there is no significant difference between the behaviour of buyers and sellers. The right panel of Fig. 2 shows how the offers depend on the submission moments.[13] We observe that offers do depend on the submission moment and that the resulting *offer functions* differ for buyers and sellers. In particular, the bids are positively correlated with the submission moment, resulting in an increasing offer function for the buyers. The asks are negatively correlated with the submission moment, resulting in a decreasing offer function for the sellers.

Any individual trader with an offer to submit faces the timing trade-off: submitting the offer earlier increases the probability that a trade will occur at the price of the own offer which results in a lower expected profit, but submitting later decreases the probability of trading. We observe in Fig. 2 that agents in our learning model resolve this trade-off by desiring to submit their orders in the middle part of the trading period. The concerns about trade are also reflected in the shape of the agents' offer functions. If an agent enters the market early, there will be relatively many trading possibilities and the agent hopes to get a higher profit by submitting a bid further from the valuation or submitting an ask further from the costs. If an agent submits at a late moment, then it might be that the trade will be at the price

[12] As it can be expected in the environment with identical buyers, the distributions for all five buyers turned out to be very similar. Thus, we show only the averaged distribution and not all five different distributions. The same remark holds for the sellers.

[13] In the right part of this and similar figures, we show the buyers' valuations as stars and the sellers' costs as circles.

stored in the book (and for this the offer should be less aggressive than the stored offer in the book) or the offer will be stored in the book, and then a higher possible profit will decrease the probability of trading. As a result, we find a positive correlation between bids and submitting time and negative correlation between asks and submitting time.

Generally, this last result is consistent with the mark-up strategies of traders found empirically and studied theoretically. We found that the mark-ups are changing during a trading period. Since this finding is common across various environments, we formulate it in general.

Result 1. *Under IEL the intramarginal buyers and sellers learn to apply the mark-up strategies that depend on the submission moment in a monotonic way. Namely, the later an offer is being submitted, the less aggressive this offer is.*

Conditional on the submission moment within the period, the mark-up strategies of buyers and sellers are identical, which is a feature of this symmetric environment. The offer functions are non-linear. Indeed, the traders are quite aggressive in the beginning of the trading period, by bidding for a larger part of the total surplus. However, the mark-up decreases fast and already after about one-fifth of the trading period, the traders are ready to accept less than half the surplus. Then the mark-up does not decrease as fast, except for the very end of the trading period, when the traders are so impatient to trade that are even ready to give up to 75% of the surplus to the other party.

The distribution of the submission moments has a small peak at the right tail. It is a consequence of the specific form of the offer functions emerged in the equilibrium. Traders who submit their offers late, submit offers close to their valuation or cost to have a high probability of trading. It makes sense for such a late offer then 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. This conclusion is also drawn in Fano and Pellizzari (2011).

Note that in this market environment, traders learn that some submission moments perform better than others, only when the offers of intramarginal traders are not identical, and, on the other hand, traders make different offers only when there is a possibility to gain by choosing a timing (as otherwise, the traders would be indifferent; recall also the discussion in Section 2.3). In this sense, the emerging properties depicted in the two panels of Fig. 2 support one another. Selection of this asymmetric equilibrium happens, at large, due to the mutation stage of the IEL allowing for some experimentation. Due to heterogeneity brought by mutation, the trade-off between the probability of trading and the expected profit is of importance even in the symmetric environment and the moment of submission starts to play a crucial role. It results in the distribution shown in the left panel of Fig. 2. An interesting question then is how well this distribution reflects some goal such as profit maximisation.

We address this question by looking at how various characteristics depend on the submission moment in Fig. 3. As before, the curves for different traders were similar, and we display only

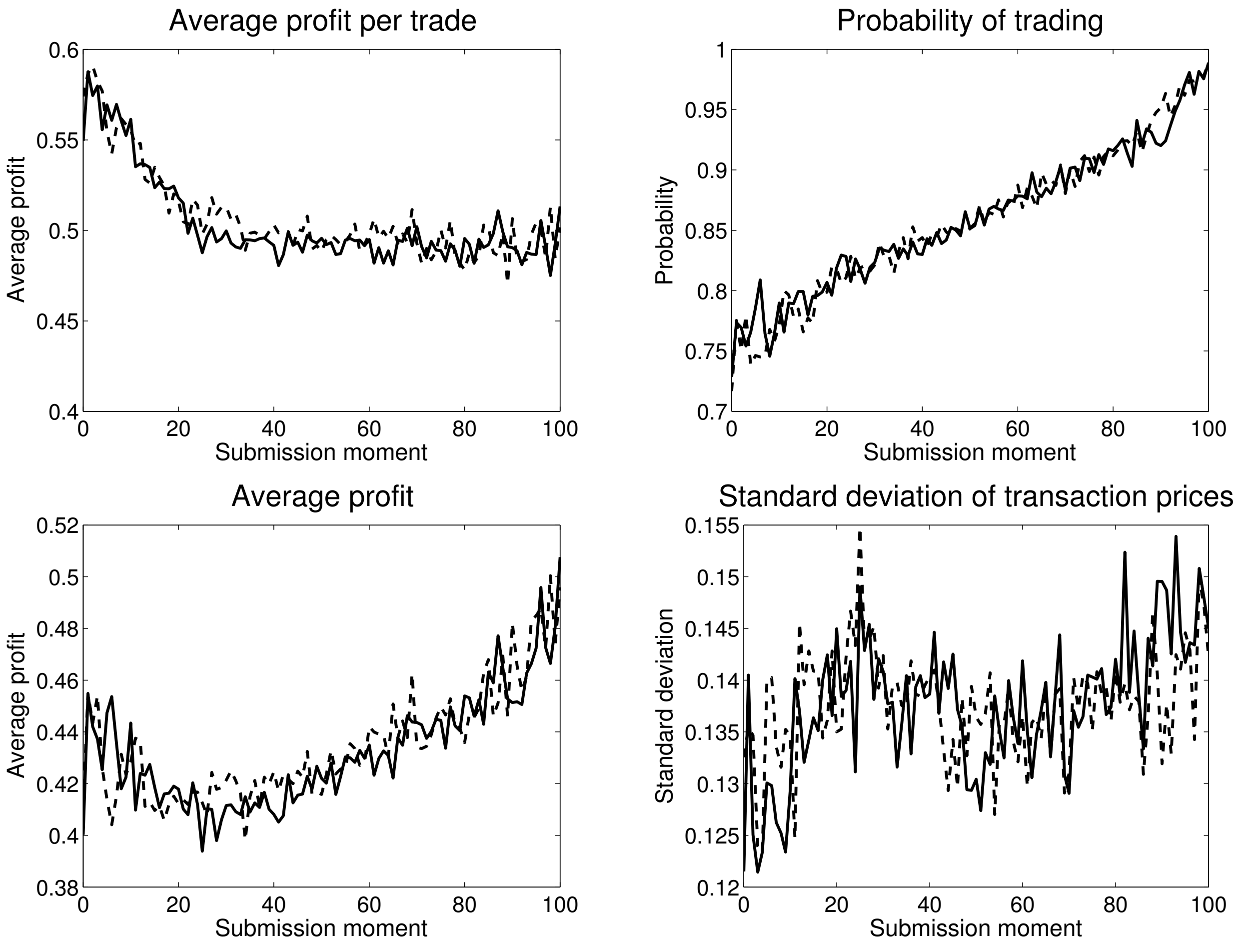

Figure 3: Characteristics of the environment $\{[1,1,1,1,1],[0,0,0,0,0]\}$ per possible submission moment

the averages over all the buyers (solid line) and over all the sellers (dashed line). The top left panel shows the average profit per transaction conditional on the fact that such transaction did take place (hence, the instances when a trader submitted an offer but did not trade are excluded). The top right panel shows the probability of trading, defined as the number of trades divided by the number of submissions. The bottom left panel shows the average profit over all the instances in which an offer was submitted at a given moment. Finally, the bottom right panel shows the standard deviation of transaction prices for every submitting moment.

The average profit per transaction is decreasing and the probability of trading is increasing with the submission period. Early submitting traders submit a very aggressive offer, such as a bid below or an ask above the middle of the interval of equilibrium prices (1/2). This behaviour will relatively often not result in a trade, but if a trade occurs this will yield a relatively high profit. Submitting late increases the probability of trading, but does not affect the average profit much, because the trade often occurs at the price of the offer submitted previously. The product of the first curve (the top left panel of Fig. 3), which is decreasing, with the second curve (the top right panel of Fig. 3), which is increasing, results in the curve for the average profit. Interestingly, this curve has a U-shape, indicating that submitting in the middle of the trading session would actually result in a low average profit. This reflects our next finding:

Result 2. *The Individual Evolutionary Learning algorithm is not consistent with maximisation of the profit in the sense, that the submission moments dominating in equilibrium do not bring the highest average profit.*

This result is especially remarkable, as it cannot be explained by concerns about risk. First, the traders do not care about risk in their motives. Second, there is no large difference in terms of uncertainty for the price of the transaction with the submission moment (see the bottom left panel of Fig. 3). Result 2 may sound counter-intuitive, but such features are common for agent-based models with traders learning *individually* with no access to the profile of average profit. The algorithm selects those strategies that perform best given the others' behaviour by filtering away the strategies with the lowest profit. Mutation that changes the submission moment should mutate in the right direction and, most probably, be accompanied by a correct mutation in the offer. The probability of such event is very small.[14]

The last question which we address in this environment is about efficiency of the equilibrium that the IEL finds. In Table 1 we compute the allocative efficiency (the percentage of the total surplus extracted in the market) and the information efficiency (the price volatility) for our model. In this environment, the surplus can be lower than the maximum only when some traders did not trade, and so we also report the total number of transactions. We compare the two-dimensional learning model analysed in this paper (the second column) with the one-dimensional model of Anufriev et al. (2013) where timing is random and traders learn only the offers (the third column).[15] We found that when traders start to learn the timing of their decisions, the average efficiency and the average number of trades decrease, whereas the average price volatility increases. All these comparisons are significant at a significance level of 1%.

Result 3. *Allowing traders to learn their preferred submission moment has a negative effect on the market allocation and informational efficiency.*

This result is a caveat to the inference about efficiency made on the basis of models with random timing.

[14]It is hard to analyse the multi-dimensional learning algorithm analytically. To have a feeling why the submissions at the moments when the average (and, hence, expected) profit is at minimum are frequent, consider the following calculations, performed by looking at the different panels of Fig. 3 and to the offer functions of Fig. 2. A trader submitting at $n \approx 25$ trades with probability around 0.8 and does not trade with probability around 0.2. In the former case, profit may increase only for those mutations that result in an earlier submission, whose probability is less than half of 0.033. But if the offer part of the strategy did not mutate (probability of such event is $1 - 0.033 = 0.967$), this offer will not be aggressive, will not lead to payoff-improvement, and, hence, will not be selected at the replication stage. In the latter case (of no trade), only mutations in submission moment leading to later submissions can be selected, but if the offer part is not mutated, the offer may turn out to be too aggressive to trade. Even after such events, not every offer mutation will be selected, but in any case the probability of profitable mutation is less than $0.8 \times 0.033 \times 0.5 \times 0.033 + 0.2 \times 0.033 \times 0.033 \approx 0.0007$. This shows that the IEL mechanism will result in mutation to a strategy that yields a higher average profit with very low probability.

[15]For the sake of comparison we simulate Anufriev et al. (2013) model with parameters of the IEL algorithm used in this paper. The results are not very different from those reported in the original paper.

	With timing	Without timing
Efficiency	0.8567	0.8930
	(0.1718)	(0.0464)
Price volatility	0.0208	0.0191
	(0.0101)	(0.0084)
Number of transactions	4.2837	4.4649
	(0.8591)	(0.2318)

Table 1: Averages of different characteristics over periods 41-50 and 3000 seeds for the environment $\{[1,1,1,1,1],[0,0,0,0,0]\}$ in the models with and without timing. The standard deviations are given in parentheses and are calculated using the average values per seed, since the individual values per seed are correlated

5 Effects of size and of market asymmetry

In this section we study robustness of our previous results to the number of traders in the symmetric environment and to asymmetry of the environment. Our general conclusion from this study will be that despite some qualitative changes in the shape of the offer functions, Result 1 holds for the environment with competition between intramarginal traders. Results 2 and 3 are also extended to the other environments. The size and composition of the market still matters, however. It affects efficiency and the distribution of the submission moments.

5.1 Size of the market

In Fig. 4 we show the distribution of submission moments and the average offer per submission moment for different number of traders. Arranged from top to bottom, the results are given for the cases of 1, 3, 10 and 15 traders on either side of the market. We compare the results with Fig. 2 for the benchmark with 5 traders on either side.

As we discussed before, two forces are important for the payoffs of traders and, hence, affect the distributions: the probability of trading and the expected profit from a transaction. It is intuitive that in thinner markets the former effect has a larger impact than in thicker markets. It turns out that the environment we discussed in Section 4 is an intermediate case, where both effects are equally strong. In an extreme case of a market with only one buyer and one seller, the submission moment does not affect the probability of trading. There, each trader would prefer to submit at the last possible moment, $n = 100$, 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. When the size of the market increases, the probability of trading does play a role. The effect of the moment of submission on the expected profit from trade decreases, since the probability that the transaction price equals the own offer tends towards one half for every submission moment. Thus, the larger the size of the market is, the earlier

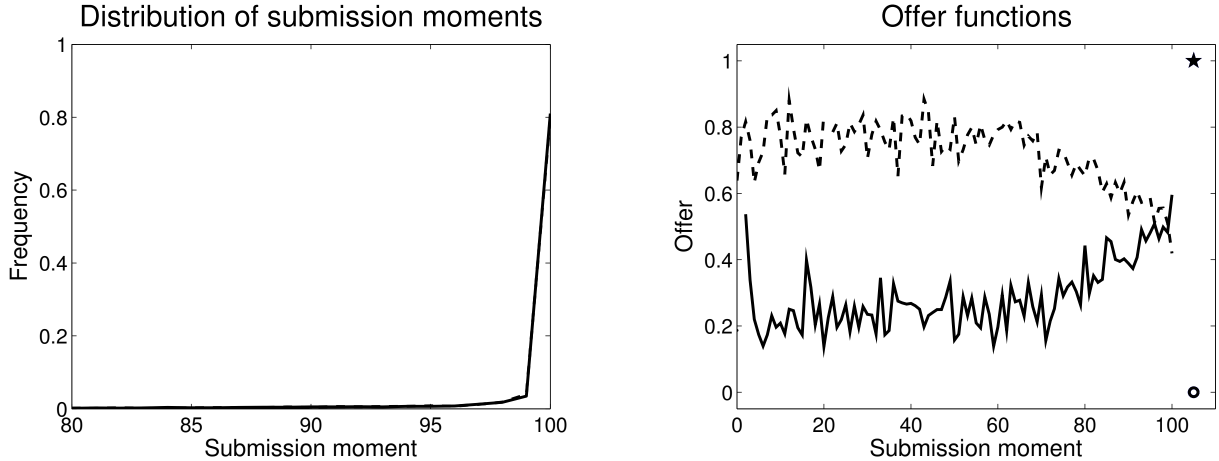

Equilibrium under IEL in the environment $\{[1],[0]\}$.

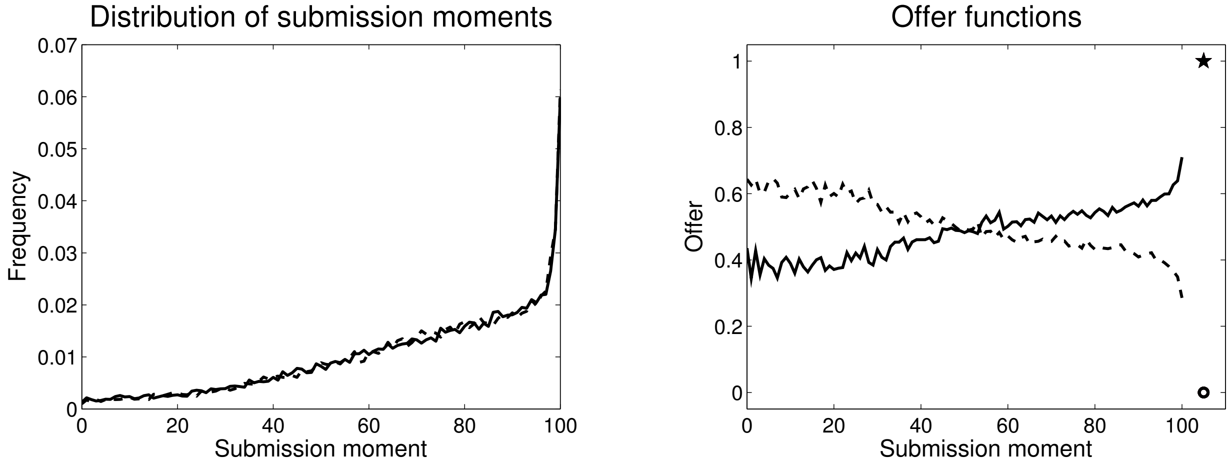

Equilibrium under IEL in the environment $\{[1,1,1],[0,0,0]\}$.

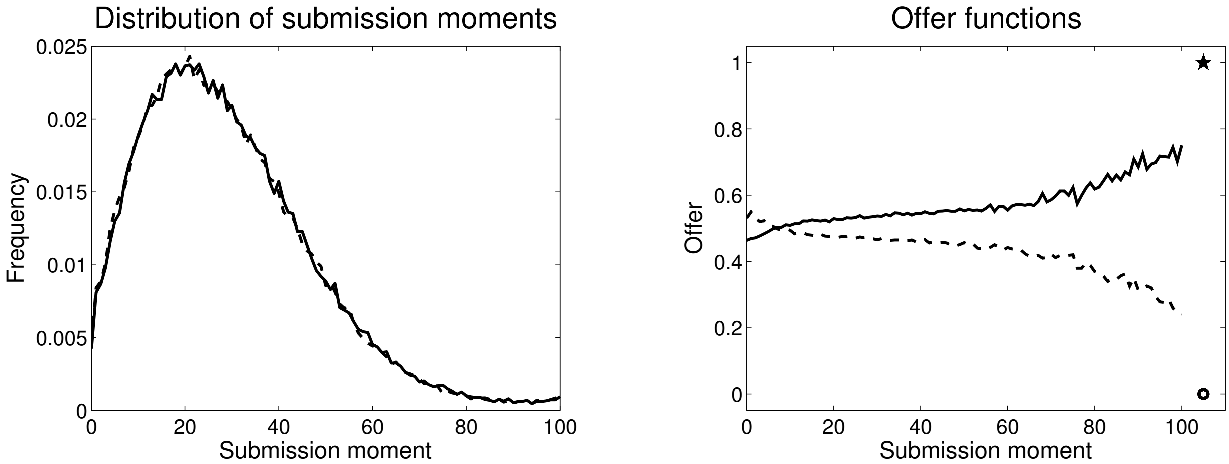

Equilibrium under IEL in the environment with 10 buyers with valuation 1 and 10 sellers with cost 0.

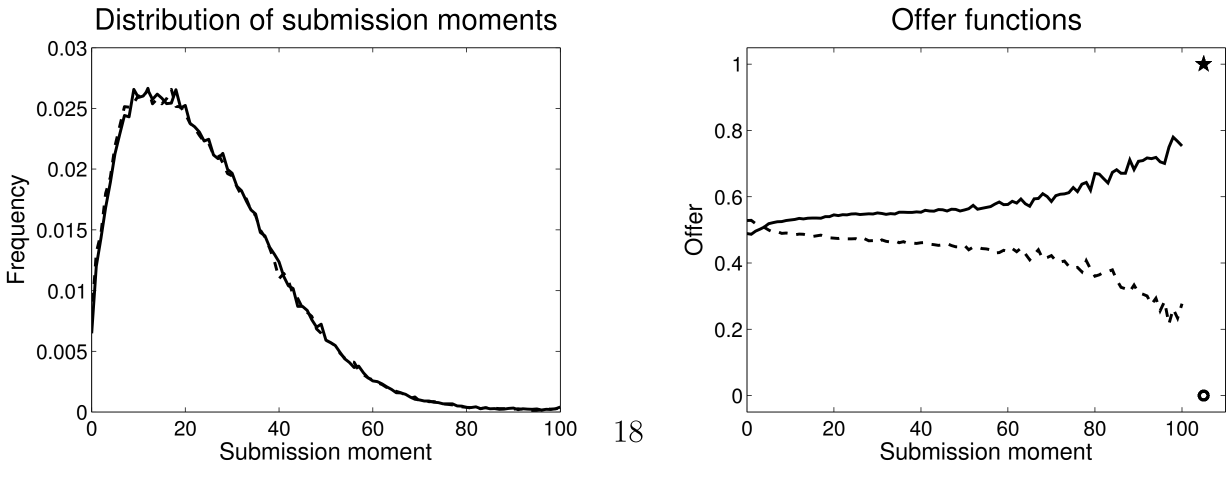

Equilibrium under IEL in the environment with 15 buyers with valuation 1 and 15 sellers with cost 0.

Number of traders per side	1	3	5	10	15
Efficiency	0.9838	0.8541	0.8567	0.8608	0.8656
	(0.0126)	(0.0195)	(0.1718)	(0.0151)	(0.0139)
Price volatility	0.0664	0.0351	0.0208	0.0128	0.0106
	(0.0653)	(0.0179)	(0.0101)	(0.0061)	(0.0053)
Number of transactions	0.9838	2.5622	4.2837	8.6076	12.9842
	(0.0126)	(0.0586)	(0.8591)	(0.1506)	(0.2084)

Table 2: Averages of different characteristics over periods 41-50 and 3000 seeds for the benchmark environment with 1, 3, 5, 10 and 15 traders on either side of the market. The standard deviations are in parentheses and are calculated using the average value per seed, since the individual values per seed are correlated

traders prefer to arrive. The simulations suggest that the peak of the distribution will converge to $n = 0$ as the size of the market converges to infinity, reflecting that in such case the effect of the expected profit from a transaction disappears. Table 2 shows the effect of the size of the market on characteristics as efficiency, volatility and the number of transactions. Interestingly, even though the moment of order submission shifts as the market size increases, the efficiency and percentage of transactions remain rather constant, except for the environment with one buyer and one seller. The sole effect is found in the decreasing price volatility, which is expected as the deviation in individual transaction prices is averaged out.

5.2 Gode-Sunder asymmetric environment

We now consider the GS environment (see the right panel of Fig. 1) to study the impact of an asymmetric market setup. We fix the number of traders (one seller and five buyers of which only one of the buyers is intramarginal) and study how the valuation of the extramarginal traders, β, affects the outcome. Anufriev et al. (2013) show that the efficiency under the IEL-algorithm (with random timing) is significantly larger than under Zero Intelligent behaviour studied in Gode and Sunder (1997) and is close to 1.

As usually we focus on the distribution of the submission moments and on the offer functions. Fig. 5 displays the results for different values of β, with $\beta = 0.3$, $\beta = 0.5$ and $\beta = 0.7$ (from top to bottom). We show the seller's behaviour (dashed line), the intramarginal buyer's behaviour (solid line) and an average of extramarginal buyers' behaviour (dotted line).

We find that the seller uses his market power only in the timing of the submission but not in the size of the offer. Indeed, the average ask is close to the mid point of the interval of equilibrium prices. At the same time, the seller submits his offer as late as possible to trade at the best bid submitted by the buyers. Under such behaviour the size of the offer does not matter, provided that it is sufficiently low. The intramarginal buyer faces more competition. She is forced both to trade earlier than the seller and also to submit a high

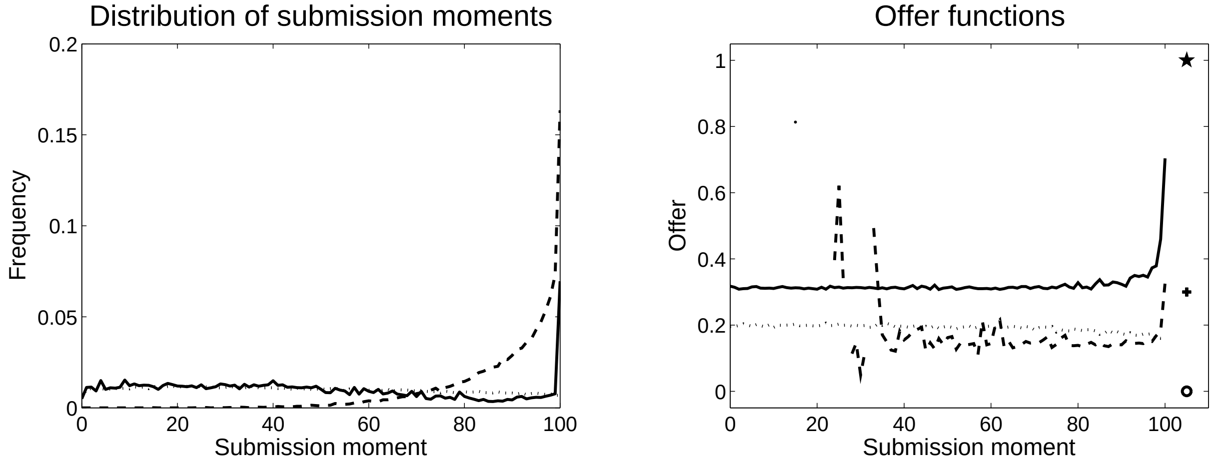

Dynamics in the GS-environment $\{[1, 0.3, 0.3, 0.3, 0.3], [0]\}$.

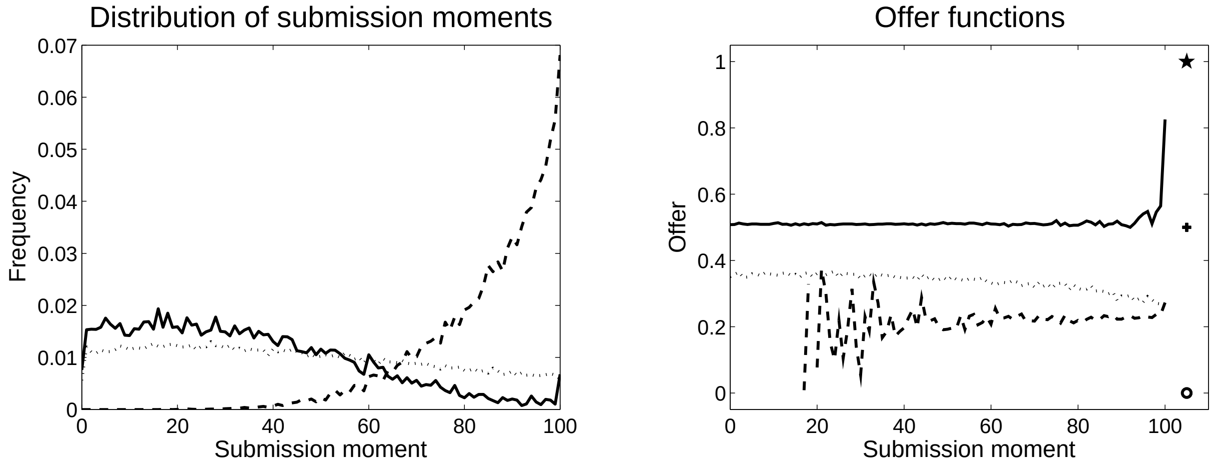

Dynamics in the GS-environment $\{[1, 0.5, 0.5, 0.5, 0.5], [0]\}$.

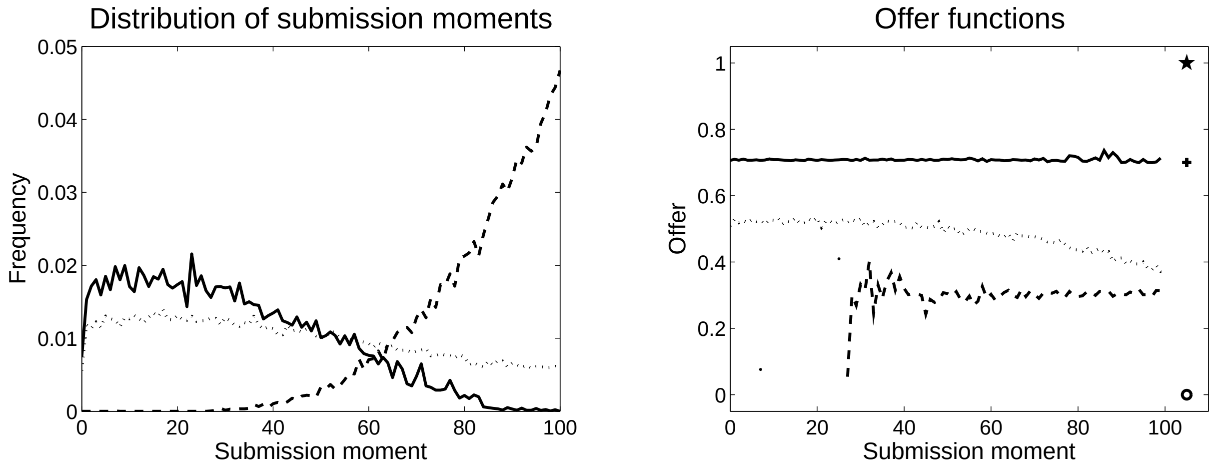

Dynamics in the GS-environment $\{[1, 0.7, 0.7, 0.7, 0.7], [0]\}$.

Figure 5: Distribution of the submission moments (*left*) and offers (bids and asks) as functions of the submission moment (*right*), in the GS-environment with 4 extramarginal buyers with valuation β. The valuation of extramarginal buyers is shown by a plussign and the average behaviour is shown by the dotted line, and as before the seller's behaviour by the dashed line and the intramarginal buyer's behaviour by the solid line

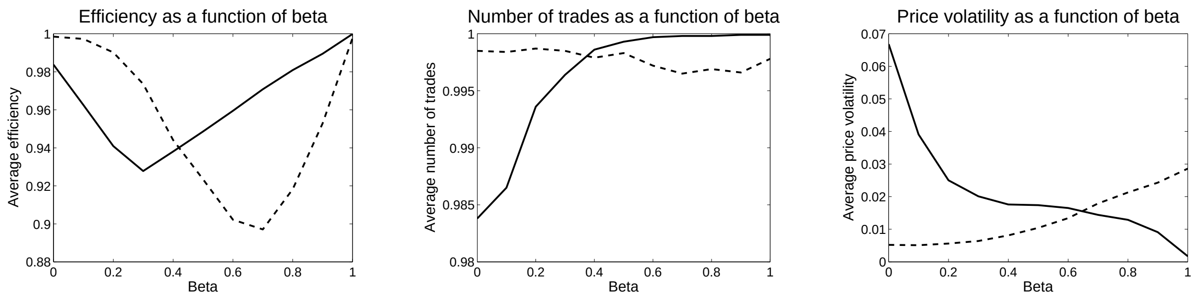

(a) Average efficiency. (b) Average number of trades. (c) Average volatility.

Figure 6: Characteristics of the GS-environments with 4 extramarginal buyers with valuation β, with timing (*solid line*) and without timing (*dashed line*)

enough order to outbid the extramarginal buyers. As β increases, the competition between the intramarginal and extramarginal buyers places more weight on the probability to trade. Thus, the intramarginal buyer prefers to come earlier than the extramarginals. This, in turn, weakens the seller's incentives to submit late.

Next, we compare efficiency of the markets for different β between themselves and with the efficiency in the IEL model without timing of Anufriev et al. (2013). The three panels of Fig. 6 show average efficiency, the number of trades and price volatility as a function of β. We note that the allocative efficiency function has a U-shape. When β is low, the probability that none of the buyers will match the seller's offer is high and the loss of efficiency is mostly due to absence of trade. When β is high, the probability of trade increases, but efficiency still can be low if an extramarginal buyer is trading. Intermediate values of β correspond to the lowest possible utility. Comparing with the model without learning about timing, we find a lower efficiency for the case of low β and a higher efficiency for larger values of β. The number of trades is increasing and the volatility decreasing in β.

6 Multiperiod IEL-algorithm

The previous sections have shown the distribution of submission moments under the traditional IEL-algorithm with a modification that allows traders to choose their submission moment. For a medium size market traders decide to submit in the middle of the period, even though this yields the lowest expected profit. This is caused by the many instances an early or late submission results in a low or zero profit. In this section we extend the IEL-algorithm even further, to allow traders to compare possible strategies on the basis of their unweighted average hypothetical foregone payoff over multiple previous periods.

The symmetric benchmark environments with 5 and 10 traders on either side of the market are studied in Fig. 7, when the unweighted average of the hypothetical payoff in the previous and the second-last periods of each strategy is taken in order to compare strategies. We find that

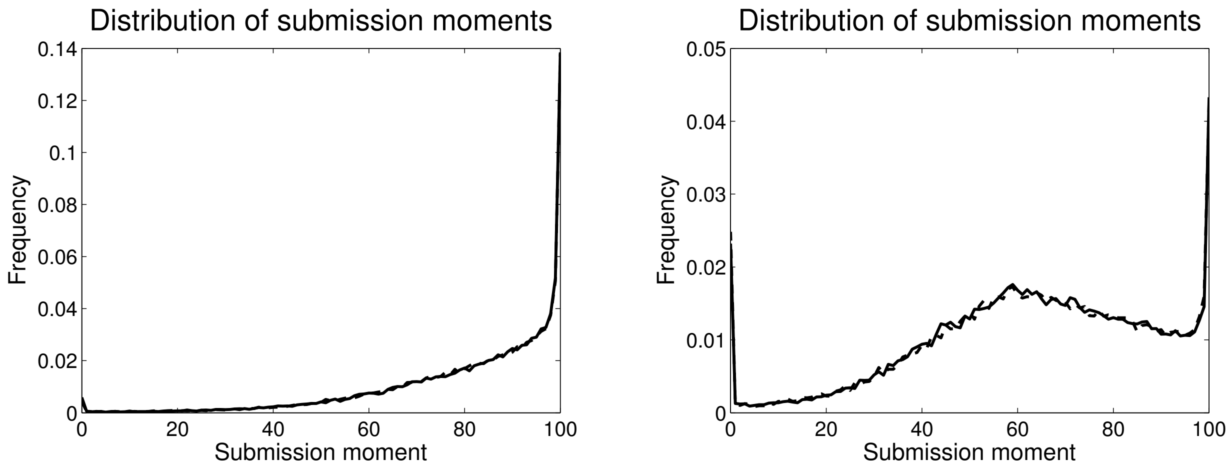

Figure 7: Distribution of learned submission moments in the environment with 5 buyers with valuation 1 and 5 sellers with cost 0 (*left*) and 10 buyers and 10 sellers (*right*), for the extended IEL-algorithm in which strategies are compared on the basis of their unweighted average hypothetical payoff over the previous two periods. Compared to the one-period learning algorithm traders submit their offers later, but as the size of the market increases the submission moment converges to $n = 0$

traders submit later than in the one-period learning model. However as the number of traders is increased the submission moments again converge towards $n = 0$. If even more previous periods are used in the IEL-algorithm this effect increases and traders submit later, even though the submission moments converge towards $n = 0$ as the size of the market increases. For small markets traders are attempting to obtain a high profit by submitting late and in larger markets traders submit early in order to increase the probability of trading. It is remarkable that even if the average profit over two periods is considered (and in fact also when the three-period average is taken), we find that in the benchmark environment with 10 traders the distribution of submission moments is inversely U-shaped, which means that the moments are favoured that result in a low average profit.

7 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 paper 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 influences this distribution.

The distribution of submission 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 submission results in a higher expected profit. However, an early submission increases the risk of not trading and a late submission 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.

If the IEL-algorithm is based on the average profit over multiple past periods we notice a shift in behaviour; orders are submitted as late as possible for medium size markets and as early as possible for large markets. It is however remarkable that there still exist sizes of the market in which the distribution of submission moments is inversely U-shaped and submission moments with a low expected profit are preferred.

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.

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. However the expected profit is the lowest when traders submit in the middle of the period, and hence in an extended IEL-algorithm in which strategies are evaluated over multiple periods traders submit as late as possible which ensures the highest expected 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 size of the market.

References

Anufriev M, Arifovic J, Ledyard J, Panchenko V (2013) Efficiency of continuous double auctions under individual evolutionary learning with full or limited information. J Evol Econ 23:539-573

Anufriev M, Hommes CH (2012) Evolutionary selection of individual expectations and aggregate outcomes. AEJ: Micro 4:35-64

Anufriev M, Panchenko V (2009) Asset prices, traders' behavior and market design. J Econ Dyn Control 33:1073-1090

Arifovic J (1994) Genetic algorithm learning and the cobweb model. J Econ Dyn Control 18:3-28

Arifovic J, Ledyard J (2011) A behavioral model for mechanism design: individual evolutionary learning. J Econ Behav Organ 78:374-395

Arifovic J, Ledyard J (2007) Call market book information and efficiency. J Econ Dyn Control 31:1971-2000

Biais B, Hillion P, Spatt C (1995) An empirical analysis of the limit order book and the order flow in the Paris Bourse. J Financ 50:1655-1689

Bottazzi G, Dosi G, Rebesco I (2005) Institutional architectures and behavioral ecologies in the dynamics of financial markets. J Math Econ 41:197-228

Brock WA, Hommes CH (1997) A rational route to randomness. Econometrica 65:1059-1096

Brock WA, Hommes CH (1998) Heterogeneous beliefs and routes to chaos in a simple asset pricing model. J Econ Dyn Control 22:1235-1274

Camerer C, Ho T-H (1999) Experience-weighted attraction learning in normal form games. Econometrica 67:827-874

Chakraborti A, Toke IM, Patriarca M, Abergel F (2011) Econophysics Review 2: Agent-based Models. Quant Financ 11:1013-1041

Chiarella C, He X-Z, Shi L, Wei L (2014) A behavioural model of investor sentiment in limit order markets. Research Paper Series, Sydney

Chiarella C, Iori G (2002) A simulation analysis of the microstructure of double auction markets. Quant Financ 2:346-353

Chiarella C, Iori G (2009) The impact of heterogeneous trading rules on the limit order book and order flows. J Econ Dyn Control 33:525-537

Dawid H (1999a) On the convergence of genetic learning in a double auction market. J Econ Dyn Control 23:1545-1567

Dawid H (1999b) Adaptive Learning by Genetic Algorithms, analytical results and applications to economic models. Springer, Berlin-Heidelberg

Erev I, Roth AE (1998) Predicting how people play games: reinforcement learning in experimental games with unique, mixed strategy equilibria. Am Econ Rev 88:848-881

Fano S, LiCalzi M, Pellizzari P (2013) Convergence of outcomes and evolution of strategic behavior in double auctions. J Evol Econ 23:513-538

Fano S, Pellizzari P (2011) Time-dependent trading strategies in a continuous double auction. In: Odinga S et al. (eds) Lecture Notes in Economics and Mathematical Systems 652, Springer, Berlin-Heidelberg, pp 165-176

Friedman D (1991) A simple testable model of double auction markets. J Econ Behav Organ 15:47-70

Gode D, Sunder S (1993) Allocative efficiency of markets with zero-intelligence traders: market as a partial substitute for individual rationality. J Polit Econ 101:119-137

Gode D, Sunder S (1997) What makes markets allocationally efficient? Q J Econ 112:603-630

Hommes C (2013) Behavioral rationality and heterogeneous agents in complex economic systems. Cambridge University Press.

Ladley D, Pellizzari P (2014) The simplicity of optimal trading in order book markets. In: Dieci R et al. (eds) Nonlinear Economic Dynamics and Financial Modelling, Springer, Switzerland, pp 183-199

Leal SJ, Napoletano M, Roventini A, Fagiolo G (2015) Rock around the clock: an agent-based model of low- and high-frequency trading. J Evol Econ, forthcoming, DOI 10.1007/s00191-015-0418-4.

LeBaron B (2001) A builder's guide to agent-based financial markets. Quant Financ 1:254-261

LiCalzi M, Pellizzari P (2006) The allocative effectiveness of market protocols under intelligent trading. In: Bruun C (ed) Advances in Artificial Economics, Springer, Berlin-Heidelberg, pp 17-29

LiCalzi M, Pellizzari P (2007) Simple market protocols for efficient risk-sharing. J Econ Dyn Control 31:3568-3590

Lux T (2009) Stochastic behavioral asset pricing models and the stylized facts. In: Hens T and Schenk-Hoppé K (eds) Handbook of Financial Markets: Dynamics and Evolution. Amsterdam, NorthHolland, pp 161-215

Parlour CA, Seppi DJ (2008) Limit order markets: a survey. In: Thakor AV, Boot AWA (eds) Handbook of financial intermediation and banking, Elsevier, North-Holland, pp 63-96

Roşu I (2009) A dynamic model of the limit order book. Rev Financ Stud 22:4601-4641

Yamamoto R, LeBaron B (2010) Order-splitting and long-memory in an order-driven market. Eur Phys J B 73:51-57