Noname manuscript No.
(will be inserted by the editor)

Timing under Individual Evolutionary Learning in a Continuous Double Auction

Michiel van de Leur · Mikhail Anufriev

Received: date / Accepted: date

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 submit different offers at different submission moments whose distribution is uni-modal with a peak moving from late to early as the market size increases. This behaviour exacerbates efficiency loss from learning. If traders evaluate profitability of their strategies over longer history, orders are submitted later with the same effect of market size.

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.

Michiel van de Leur
CeNDEF, Amsterdam School of Economics, University of Amsterdam,
Valckenierstraat 65-67, NL-1018 XE Amsterdam, Netherlands
E-mail: M.C.W.vandeLeur@uva.nl

Mikhail Anufriev
University of Technology Sydney, Business School, Economics Discipline Group,
PO Box 123, Broadway, NSW 2007, Australia
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 traders submit their orders at random moments during a trading period.[1] This assumption simplifies the traders' behaviour, as they are then left with a one-dimensional decision, such as choosing a bid or ask price (as in LiCalzi and Pellizzari, 2006, 2007) or forecasting a future price and condition its offer on this forecast (as in Chiarella and Iori, 2002; Yamamoto and LeBaron, 2010). On the other hand, it is plausible that if traders are allowed to submit their orders at preferred moments during the trading session, their other decisions will change, as traders may condition their bids and 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 order 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 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 call markets and CDA, respectively.[2] The IEL algorithm assumes that traders select their active strategies from a pool

[1] Agent-based methodology avoids making an extreme assumption of full rationality and models traders as boundedly rational who learn from the past. The approach is appropriate, given complexity of the order-drive markets, see a review in Chakraborti et al. (2011).

[2] Anufriev et al. (2013) focussed 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).

of potential strategies. After a trading period the hypothetical payoffs of all strategies are calculated using the past information and some strategies are 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 length of the memory used to evaluate hypothetical payoffs.

The distribution of submission moments is studied in a benchmark symmetric environment under full information about trading history. We find that under the IEL-algorithm investors learn to submit their orders in various moments of the trading session. The distribution of submission moments is unimodal and its peak moves to the earlier moments as the market size increases. At the same time, traders arriving earlier make more aggressive offers (lower bids, higher asks) than later during the session. Extending the IEL-algorithm by allowing to evaluate strategies over multiple past periods affects the distribution of submission moments, shifting its peak to the later moments, which on average results in the higher profits of traders. We also discuss an effect of learning about timing on market allocative efficiency and price volatility comparing it 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 do not choose a submission moment, but these are modelled to follow a Poisson stochastic 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).

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 an environment where the buyers and sellers occupy a symmetric position in the market. The impact of the size of the market is considered in Section 5. In Section 6 the IEL-algorithm is extended to multi-period learning. Section 7 concludes. The Appendix is devoted to the analysis of an asymmetric environment.

2 Market setup

In this paper we study the market organised as a Continuous Double Auction (CDA) 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. Each trader has to make two decisions simultaneously: about the time of order submission and the submitted offer. To evaluate efficiency of the CDA markets, we will first introduce the Competitive Equilibrium outcome as a benchmark.

2.1 The environments and competitive equilibrium

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 valuation or cost, but not the values of the other agents. We denote the environments by vectors of valuations and costs.[4] When buyer b transacts with seller s at price p, the buyer's payoff is $V_b - p$, and the seller's payoff is $p - C_s$. The sum of all traders' payoffs defines the *trade surplus*.

Given the valuations and costs, the aggregate demand and supply functions can be determined in a standard way. The *competitive equilibrium* (CE) is a pair (q^*, p^*), where the quantity q^* is the largest possible quantity that clears the market and the price p^* is such that both demand and supply at this price are q^*. There might be an interval of equilibrium prices $[p^*_L, p^*_H]$, i.e., multiple

[3] See, in particular, Dawid (1999a), Bottazzi et al. (2005), Chiarella and Iori (2009), Anufriev and Panchenko (2009), Fano et al. (2013), Ladley and Pellizzari (2014), Chiarella et al. (2014) and Leal et al. (2015).

[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.

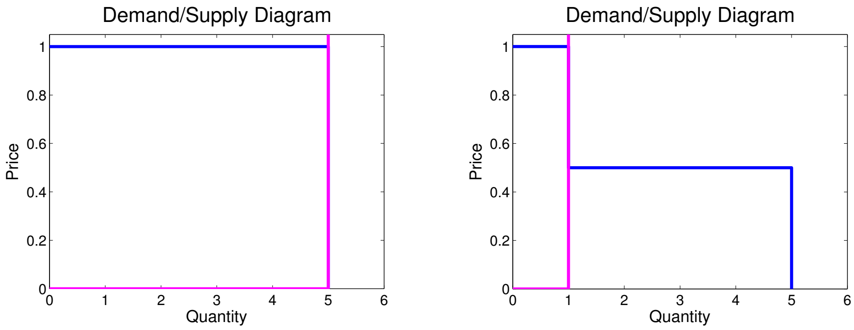

(a) The benchmark environment. (b) The main symmetric environment.

Fig. 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 $\beta = 0.5$

CE. The traders who trade in the CE are called *intramarginal* traders, and those who do not trade in the CE (and hence have zero payoff) are called *extramarginal* traders. It is easy to see that in all CEs the trade surplus is the same and equal to $\sum_b V_b - \sum_s C_s$, where both sums are over the intramarginal traders. The trade surplus under CE is the largest possible trade surplus under voluntary trade.

In this paper we will use the following two environments (see illustrations in Fig. 1). In the main body of the paper we will analyze a *symmetric* environment $\{[1, \ldots, 1], [0, \ldots, 0]\}$ with $B = S$ identical buyers and identical sellers. The CEs are given by $q^* = B = S$ and $p^* \in [0, 1]$ and all traders are intramarginal. This environment is simple, it does not have a competition with extramarginal traders, and for this reason is well suited to introduce timing into the model. In Appendix we investigate the affects of market asymmetry, when we discuss the so-called *GS*-environment[5] $\{[1, \beta, ..., \beta], [0]\}$, in which one buyer has valuation 1, $B - 1$ buyers have valuations $\beta < 1$, and one seller has cost 0. The CEs are given by $q^* = 1$ and $p^* \in [\beta, 1]$; in all CEs the only seller and the buyer with valuation 1 are intramarginal, whereas all other buyers are extramarginal.

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.

[5] 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.

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: the trader with the smallest submission moment arrives first, the next order will be from the trader with the second smallest submission moment, and so on. It might happen, however, that a group of two or more traders decides to submit their orders at the same submission moment. Moreover, there might be several such groups, each with an identical 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.[6] 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, to avoid a large random component in the sequence of submission.[7]

The submission moments of buyers and sellers in period t are denoted as $n_{b,t}$ and $n_{s,t}$, respectively. We assume the price-time priority rule for the order book. Thus, when 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. Finally, to focus on learning of traders about their timing of order submission, we assume that no order can be cancelled.

Given the set of submission moments and offers from all the traders, the rules of the CDA determine all transactions during period t. Traders payoffs, as defined in Section 2.1, depend on the own valuation or cost and also on the transaction price. As opposed to the benchmark case of the competitive equilibrium, under CDA different traders will generally transact at different prices. Let $p_{b,t}$ denote the transaction price of buyer b in period t. It follows that 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. Similarly the payoff of a seller equals $U_{s,t} = p_{s,t} - C_s$ if she traded at price $p_{s,t}$, and zero otherwise. Transaction prices are determined by the CDA mechanism and, obviously, depend on the trading sequence. Therefore also the trader's payoff depends not only on the own valuation or cost and the offer, but also on the moment of arrival.

Summing all trader's payoffs we can compute the trade surplus of period t. The *allocative efficiency* of CDA during period t is defined as the ratio of this trade surplus to the maximum possible trade surplus for the same environment, i.e., to the trade surplus under competitive equilibrium.

[6] 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.

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

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). As opposed to other similar models[8], 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 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.[9] Agents compute the hypothetical payoffs of the strategies using all information available to them. We assume that past offers and the submission moments of all traders are publicly available, which is the case nowadays for most of the stock exchanges. 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 by 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 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 distribu-

[8] 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).

[9] For the previously active strategy this payoff is the actual payoff of that strategy.

tion 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. 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, 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 in 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

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

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

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

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. Moreover, we assume that the submission moments of all traders in the last period are publicly available.[10] 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

[10] The case of publicly available order book is referred to as the OpenBook system in Anufriev et al. (2013).

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 $b > a_{s,t}$, the 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 price); $V_b - b$, in the case when $n < n_{s,t}$ (because then the transacted price is the buyer's early submission price); 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).

Selection of an active strategy from the pool
In the first period every strategy from the pool is equally likely to be chosen. In 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 multi-dimensional decision. Every trader is given an individual pool of strategies of size $K = 300$. A part of a strategy mutates with a probability of $\rho = 0.033$. 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.[11] 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

[11] Anufriev et al. (2013) used uniform mutation for the offers mutation.

and the result is rounded to the nearest integer. If the resulted strategy lies outside the strategy space, we repeat the procedure for that strategy.

First, we are interested in the outcome of learning regarding the submission moments. We observe that the traders face the *timing trade-off* between a higher probability to trade (which would imply earlier submission) and higher realised profit from a trade (which would imply later submission). We investigate the outcome for this trade-off that emerges from learning. Second, we study the *allocative efficiency* of CDA, which is defined in Section 2.2 as the ratio between the trade surplus in a given period and the maximal possible trade surplus. If all intramarginal buyers and sellers trade during a period and no extramarginal trader trades, the allocative value is maximal, and the allocative efficiency is equal to 1. Efficiency can be lower when an extramarginal agent trades, or when an intramarginal agent does not trade. Third, we look at 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 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 $D = 0.101$.

4 Symmetric environment

We consider simulations of the 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.[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

[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.

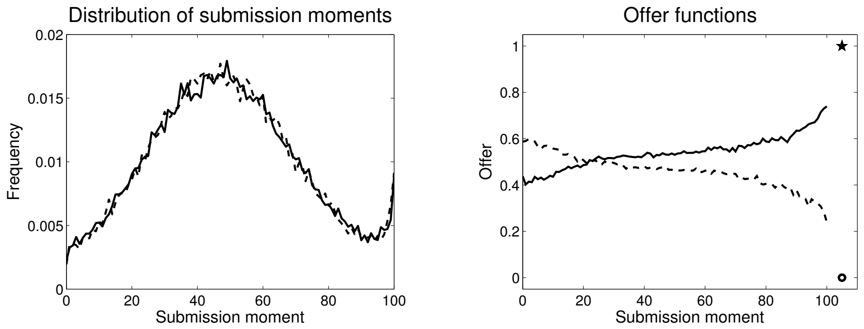

Fig. 2 Distribution of the submission moments (*left*) and offers (bids and asks) as functions of the submission moment (*right*) in the environment $\{[1, 1, 1, 1, 1], [0, 0, 0, 0, 0]\}$ under IEL. The average buyer's behaviour is shown by the solid line and the seller's behaviour by the dashed line

moments.[13] We observe that the resulting *offer functions* are symmetric and the offers are more conservative the later they are submitted.

This result can be considered in the following context. Any individual trader with a given offer to submit faces the *timing trade-off*. Submitting the same offer *earlier* increases the probability that a trade will occur, as otherwise an offer from a potential counter-party could have been matched with some other offer. On the other hand, conditional on the fact that trade occurs, trader would prefer to submit the same offer *later* to trade at a more favourable price of earlier submissions.[14]

The traders' emerging behaviour, illustrated in Fig. 2, reflects how this trade-off is resolved by the IEL. Agents desire to submit their orders in the middle part of the trading period, but they also condition the offer on the submission moment. Agents that submit early are rather aggressive, in the sense that their offers are relatively far from their valuation or cost, so that they would receive a high profit if they trade, but at the cost of a low probability to trade. Agents that submit late are very conservative, and get a very low profit if the trade occurs at their own submitted price. Their conservative submission behaviour compensates the low probability of trade existing at the end of the trading period.

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

[14] As an example, consider first the market with two traders, one buyer with valuation 1 and one seller with cost 0. Let their submitted offers be $b_1 = 0.6$ and $a_1 = 0.1$, respectively, so that transaction will occur. If the buyer submitted his order *before* the seller, the transaction price equals the buyer's bid, $p_{b,t} = p_{s,t} = 0.6$. If the buyer submitted his order *after* the seller, the transaction price equal to the ask of the seller, $p_{b,t} = p_{s,t} = 0.1$. The buyer has a higher profit (0.9 against 0.4) in the case of submitting late. Continue the example and now add another buyer and another seller with offers $b_2 = 0.4$ and $a_2 = 0.9$, respectively. The first buyer will not trade at all and get profit 0 by submitting too late (after the first seller and the second buyer).

We find a positive correlation between bids and submission moments and negative correlation between asks and submission moments. It implies that the bids' markdowns relative to the buyers' valuations and the asks' markups relative to the sellers' costs are changing during a trading period.[15] This finding is common across various environments and hence we formulate it in general.

Result 1 *Under IEL the intramarginal buyers (resp. sellers) learn to use the markdown (resp. markup) 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, in the sense that it has a lower markdown (resp. markup).*

Conditional on the submission moment within the period, the markdown strategies of buyers and markup strategies of sellers are identical, which is a feature of this symmetric environment. The offer functions are non-linear. Indeed, the traders are very aggressive in the beginning of the trading period, by bidding for a larger part of the total surplus. However, their markdowns and markups decrease fast and already after about one-fifth of the trading period, the traders are ready to accept less than half of 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 they are even ready to give up 75% of their surplus to the other party.[16]

The learning approach is often used to select one equilibrium from many, see, e.g., Conlisk (1996). In the context of CDA the issue of multiple equilibria is very relevant and becomes even stronger when the traders can choose submission moments.[17] It is therefore remarkable that in all our simulations of this market, we get an almost identical distribution of the submission moment and offer functions. Therefore, our simulations suggest that IEL has a unique and stable equilibrium, the one which is illustrated in Fig. 2.

The distribution of the submission moments is bell-shaped in this equilibrium with a small peak at the right tail. An interesting question is then how well this distribution consistent is with profit maximisation. We address this question by looking at how, in this equilibrium, various characteristics depend on the submission moment; see Fig. 3. As before, the curves for different traders were similar, and we display only 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

[15] Markdown is the amount by which bid is lower than the buyer's valuation or the amount by which ask is higher than the seller's cost. Various markdown and markup strategies have been studied recently in Zhan and Friedman (2007) and in Cervone et al. (2009).

[16] Note that this does not necessarily mean that such a low profit is achieved, as the offer will often be matched with an outstanding offer for a higher profit.

[17] For instance, in this environment none of the traders could improve the payoff under the strategy profile with *any* distribution of submission moments and where all buyers and sellers submit an arbitrary (but identical) offer $p \in [0, 1]$. All such profiles are equilibria.

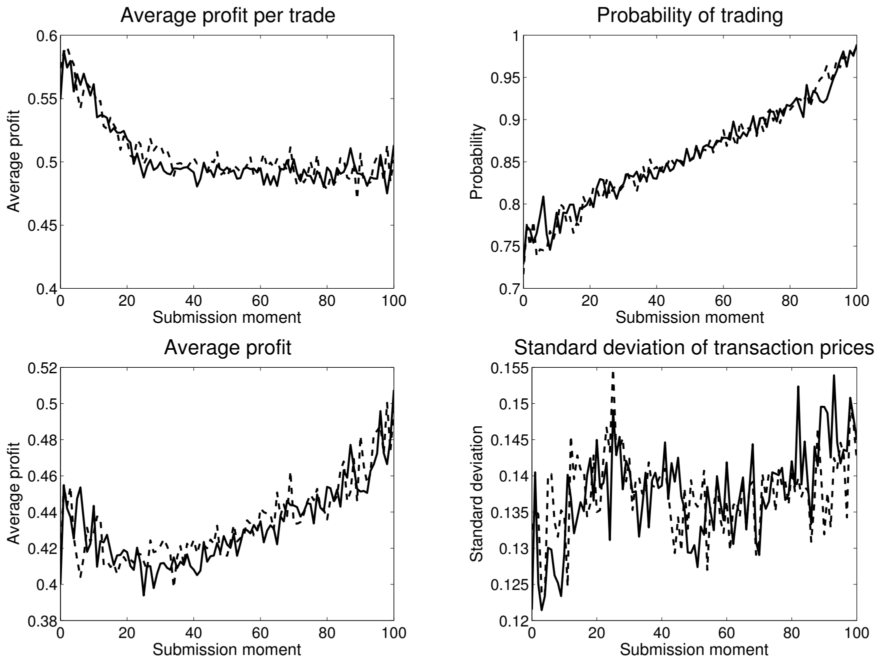

Fig. 3 Characteristics of the environment $\{[1,1,1,1,1],[0,0,0,0,0]\}$ per possible submission moment. The average buyer's behaviour is shown by the solid line and the seller's behaviour by the dashed line

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 aggressive offers, bidding below or asking 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. The average profit per trade is rather constant for all late submissions, because the trade often occurs at the price of the offer submitted previously and hence the own offer only rarely affects the profit. 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 (the bottom left panel of Fig. 3). 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 directly 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 right panel of Fig. 3).

Result 2 may sound counter-intuitive, but it is not uncommon to find that traders learning *individually* with no access to the profile of average profit do not end up as a group in the point with the highest expected payoff. In our case, the IEL algorithm selects those strategies that perform best *given the others' past behaviour* by filtering away the strategies that would bring lower profit. Hence strategies that more often yield the lowest profit are more likely to be removed from the strategy pool, even if they yield the highest average profit. In Section 6 we show the differences that would occur in the same model with a multi-period IEL model (i.e., with a longer past memory to evaluate performances of the strategies). Furthermore, in our model with multi-dimension learning, it might be that the strategy should mutate both in the submission moment and in the offer at the same time. The probability of such event is very small and in particular it occurs less often than strategies are replaced.[18]

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 (see Section 3.1) and the information efficiency (i.e., the price volatility) for our model. As in this environment there are no extramarginal traders, the surplus can be lower than the maximum only when some traders did not trade. Thus, 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).[19] We found that when traders start to learn the timing of their decisions, the average efficiency

[18] 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, which 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.

[19] For the sake of comparison we simulate the model from Anufriev et al. (2013) 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 (0.0171)	0.8930 (0.0046)
Price volatility	0.0208 (0.0101)	0.0191 (0.0080)
Number of transactions	4.2837 (0.0861)	4.4649 (0.0232)

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

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 *Under the IEL-algorithm, allowing traders to learn their preferred submission moment has a negative effect on the market allocation and informational efficiency.*

The extra freedom agents have in the model with timing, gives them actually more opportunities to seek a higher profit. For example, if there is a conservative buyer on the market, any seller can adjust the moment of submission to be most likely to trade against that buyer, and also submit an aggressive offer to take advantage and extract a higher profit. This strategy is certainly not foolproof, because both sellers and buyers learn constantly. Even if such a strategy can indeed yield a higher profit, it may come at a cost of having a lower probability of trading. Under IEL, such experimentation would pass the reproduction stage and could hurt the overall market efficiency.

5 Effects of size of the market

In this section we study robustness of our results to the number of traders in the symmetric environment. The effects of asymmetry of the environment is analyzed in Appendix. Our conclusion 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, for efficiency and the distribution of the submission moments.

In Fig. 4 we show the distribution of submission moments and the average offer per submission moment for different numbers of traders. Arranged from top to bottom, the results are given for the cases of 1, 3, 4, 10 and 50 traders on either side of the market. We compare the results with Fig. 2 for the benchmark with 5 traders on either side.

In light of the *timing trade-off* discussed before, the submission moment has opposite effects on the two components of total profit: the probability of

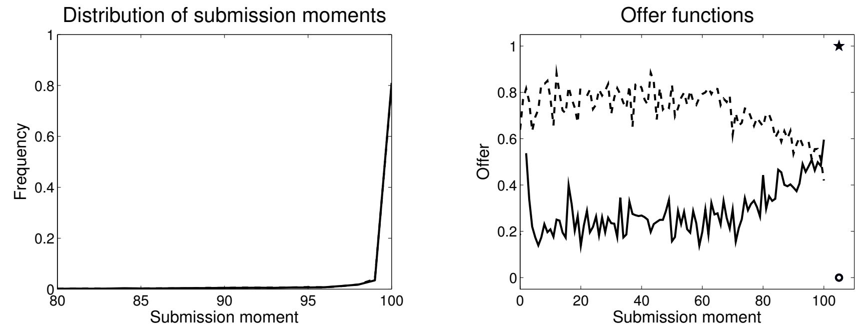

Equilibrium under IEL in the environment $\{[1],[0]\}$.

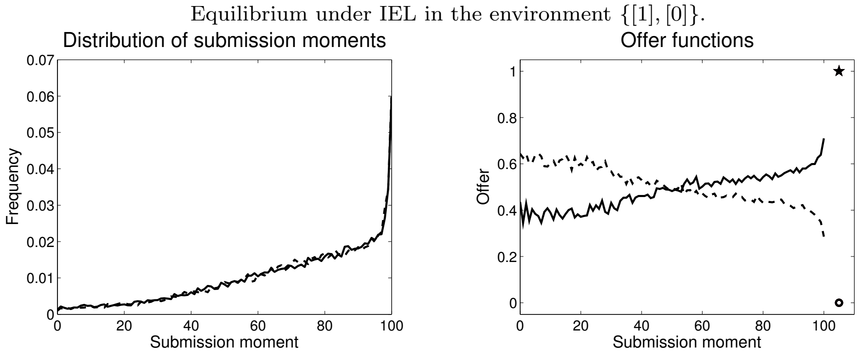

Equilibrium under IEL in the environment $\{[1,1,1],[0,0,0]\}$.

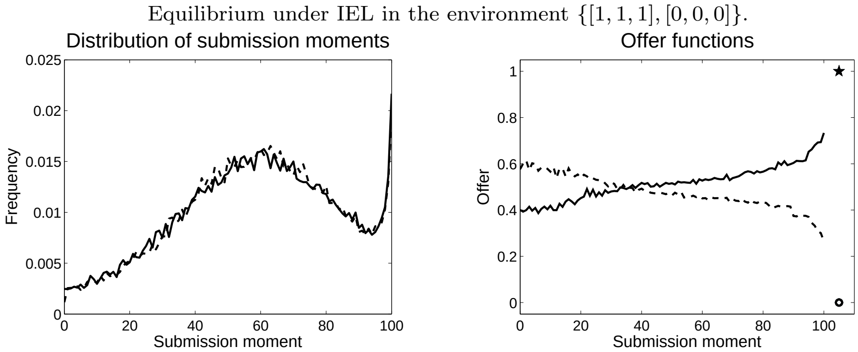

Equilibrium under IEL in the environment $\{[1,1,1,1],[0,0,0,0]\}$.

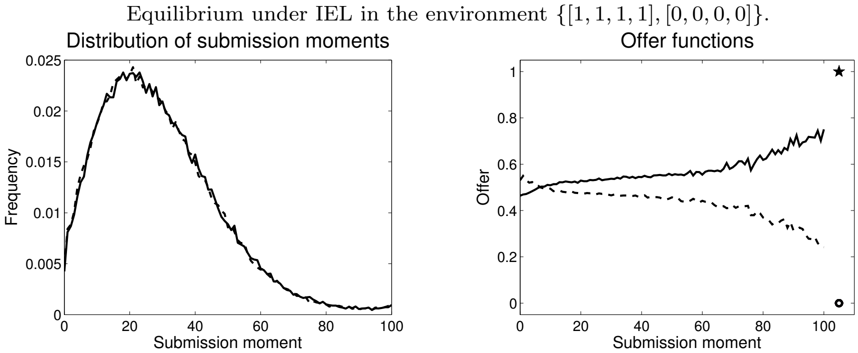

Equilibrium under IEL in the benchmark environment with 10 buyers and 10 sellers.

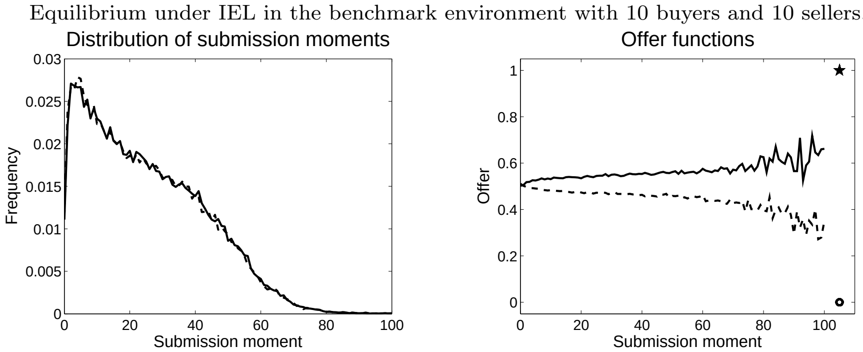

Equilibrium under IEL in the benchmark environment with 50 buyers and 50 sellers.

Fig. 4 Distribution of the submission moments (*left*) and offers (bids and asks) as functions of the submission moment (*right*), in the symmetric environment with 1, 3, 4, 10 and 50 traders on either side of the market.

trading and the expected profit from a transaction. It is intuitive that in thin markets the former effect has a larger impact than the latter, and it is the opposite in large markets. Indeed, in an extreme case of a market with only one buyer and one seller (see the top panel of Fig. 4), the submission moment does not affect the probability of trading. There, each trader would prefer to submit at the last possible moment, since submitting second results in a trade at the price of the other trader. On the other extreme, in an infinitely large market, the probability that the transaction price equals the own offer is one half. Therefore it is optimal to submit at the first possible moment when the probability of trading is maximised.

Simulations in Fig. 4 support the conclusion that the larger the size of the market is, the earlier 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. It turns out that the environment we discussed in Section 4 is an intermediate case, where both effects are comparably strong. Looking at the offer functions in Fig. 4 we conclude that Result 1 holds irrespectively of the size of the market.[20]

Table 2 shows the effect of the size of the symmetric market on allocative efficiency, price volatility and the number of transactions. Except for the environment with one buyer and one seller, the allocative efficiency and number of transactions does not change significantly with the market size at 1% significance level. However, we can still observe that the number of transactions (and therefore allocative efficiency) slowly increases. This is consistent with the shape of offer functions in the right panels of Fig. 4. When the market gets larger, the offers are, generally, less aggressive (e.g., the offer curves of buyers and sellers intersect earlier and earlier). Therefore, there will be more trades and less inefficiencies. The price volatility is decreasing, which is expected as the deviation in individual transaction prices is averaged out with a larger market size.

[20] As the size of the market increases, the offer becomes less aggressive for a given submission moment. The offers at the most frequently chosen submission moments are similar.

Traders per side	1	3	4	5	10	50
Efficiency	0.9838	0.8541	0.8569	0.8567	0.8609	0.8702
	(0.0126)	(0.0195)	(0.0182)	(0.0171)	(0.0139)	(0.0134)
Price volatility	0.0664	0.0351	0.0258	0.0208	0.0128	0.0070
	(0.0653)	(0.0179)	(0.0133)	(0.0101)	(0.0060)	(0.0036)
Number of trades	0.9838	2.5622	3.4275	4.2837	8.6086	43.5106
	(0.0126)	(0.0586)	(0.0729)	(0.0861)	(0.1391)	(0.7177)

Table 2 Averages of different characteristics over periods 41-50 and 3000 seeds for the benchmark environment with 1, 3, 4, 5, 10 and 50 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. All the comparisons are significant at a level of 0.01%

6 Multi-period IEL-algorithm

We have shown that the decisions about the timing matter in the CDA under the traditional IEL-algorithm. Despite a possibility that with a multi-dimensional decision agents may find difficulty to learn, Result 2 is puzzling. We found there that traders decide to submit in the middle of the period, even though this yields the lowest expected profit. In this section we extend the IEL-algorithm to allow traders to compare possible strategies on the basis of their hypothetical foregone payoffs over *multiple* previous periods. This might give a better chance to select the strategies with frequent zero but higher expected payoffs.

The submission moment distribution in the market with 5 and 10 traders on either side are shown in Fig. 5. The difference with respect to simulations shown in Fig. 2 is only in the calculations of hypothetical payoffs of IEL. Agents are assumed to remember 2 last periods, and the hypothetical payoff of every strategy is computed as the mean of the payoff that this strategy would generate in *both* the previous period and the second-last period. When calculating their payoffs, agents assume that the behaviour of other traders in a given period is exactly the same as it was in that period.

Multi-period IEL results in a slight alteration of the above results, in the sense that traders's submission moment distribution shifts to the right, i.e., traders submit later. At the same time, the effect of the size of the market is the same. These effects come indirectly from the increased efficiency in the market. Indeed, in the one-period learning model, traders' hypothetical calculation may often lead them to accept mutations with rather aggressive strategies and this behaviour will result in no trade. For instance, for a buyer who arrived before the trading partner whose ask is zero, the IEL gives the highest hypothetical foregone payoff to submitting a bid of zero next time. But if this strategy would be chosen, it will likely not result in a trade, as the trading partner could meanwhile increase her ask. When considering multiple periods such mutations from equilibrium strategies could occur less often, as the trader would look at the performance of the mutated startegy on several past periods with more variability in the behaviour of others. Thus, the probability of trading in multi-period IEL will increase and the pressure to submit earlier decreases. Of course, the traders that submit an aggressive offer do so at the very start of the period, as for them the probability of trading is still very important.

We have simulated the model with up to five periods memory, and find that traders submit later and later in the same market when they take more historical periods into account. On the other hand, the effect of market size still holds: when keeping the memory constant, the submission moments moves to the earlier moments as the size of the market increases. In relatively 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. Also, as the size of the market increases agents offers for a given moment are less aggressive.

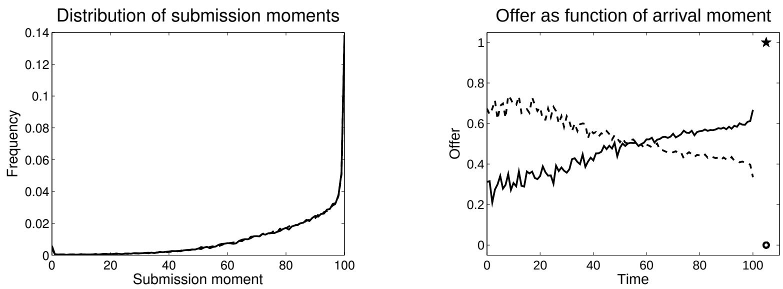

Equilibrium under IEL in the benchmark environment with 5 buyers and 5 sellers.

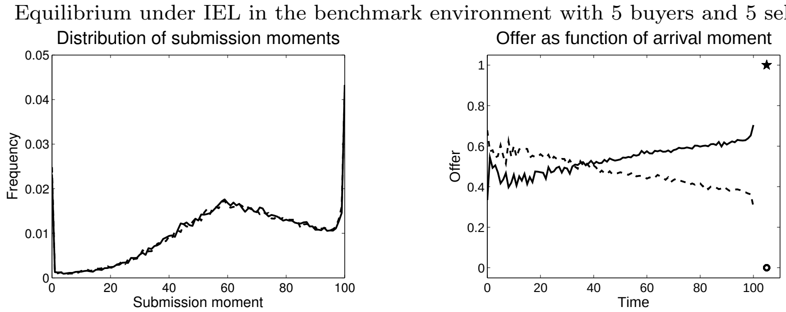

Equilibrium under IEL in the benchmark environment with 10 buyers and 10 sellers.

Fig. 5 Distribution of learned submission moments(*left*) and offers (bids and asks) as functions of the submission moment (*right*) in the benchmark environment with 5 resp. 10 traders on either side, 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 still converges to $n = 0$

7 Concluding Remarks

In the Continuous Double Auction the timing of order submission plays an important role. Keeping the offer constant, submitting it at the end of the trading period may yield a lower probability of trading but higher profit from trade, than when the offer is submitted in the beginning of the period. However, by and large, the issue of timing has been neglected in the literature.

In this paper we introduce learning about timing in the standard model of Individual Evolutionary Learning. Considering one of the most simple environments, a symmetric market with an equal number of identical buyers and identical sellers, all of whom would trade in competitive equilibrium, we find that the traders do not learn to submit at uniformly random times as many models just assume. At the same time their offers do not have constant markups and markdowns during the trading session. Instead, the equilibrium to which IEL converges is characterised by a decrease in aggressiveness of the offers (over the trading session) and a uni-modal distribution of the submission moments, whose peak moves from late to early as the market size increases.

We also find that the allocative and informational efficiency of the market changes as traders are allowed to learn about their timing. In fact, the efficiency would go down with respect to the case with uniformly random timing. This is a result of experimentation in extra dimension of agents whose individual learning is based on the past and does not take into account that other traders are learning as well.

An intriguing result of this paper is that the IEL selects an equilibrium which is hard to reconcile with individual profit maximisation. The expected profit at the moment of trading session when most of the traders submit is not the highest but rather lowest. We have generalised the IEL model to allow traders to take more past experience into account when making their decisions. The longer memory comes at a higher cognitive cost but cannot fully reconcile the discrepancy between aggregate learning outcome and individual profit.

The choice of our environment was motivated by its extreme symmetry in both sides of the market. If in the symmetric market the outcome of learning is not symmetric with non-uniform distribution of submission moments and non-constant offer functions, timing is indeed important. Further work should show how learning about timing would affect asymmetric markets. Appendix of this paper provides a first step in this direction.

It would also be interesting to compare the results that this paper generates through simulations with a behaviour of human subjects in the experimental environment. Arifovic and Ledyard (2011) show that IEL is quite successful in replicating experimental results for call markets. The CDA markets are much more complicated and traders' strategies are necessary multi-dimensional. We showed that introducing extra dimensionality in timing may already change the results in terms of behaviour and market efficiency. Further analysis may also look at other natural generalizations, such as possibility to submit multiple orders and cancel them.

Appendix

In this Appendix we consider the GS-environment (see the right panel of Fig. 1) to illustrate the outcome of IEL in an asymmetric market setup. We fix the number of traders (one seller and five buyers of which only one is intramarginal) and study how the valuation of the extramarginal traders, β, affects the outcome.[21]

As before we focus on the distribution of the submission moments and on the offer functions. Fig. 6 displays the results for different values of β. We show the seller's behaviour (dashed line), the intramarginal buyer's behaviour (solid line) and an average of the extramarginal buyers' behaviour (dotted line).

We find that the seller uses her market power only in the timing of the submission but not in the size of the offer. Indeed, the average ask is close to

[21] 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.

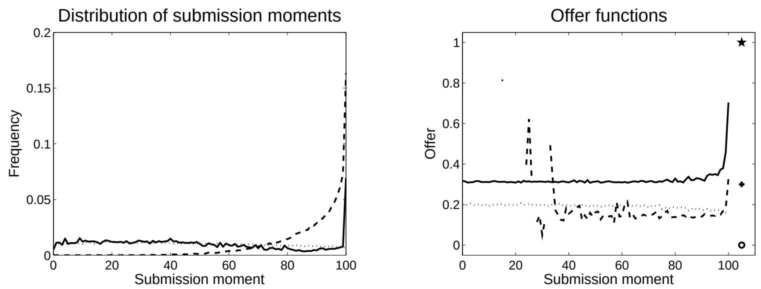

Dynamics in the GS-environment $\{[1, 0.3, 0.3, 0.3, 0.3], [0]\}$.

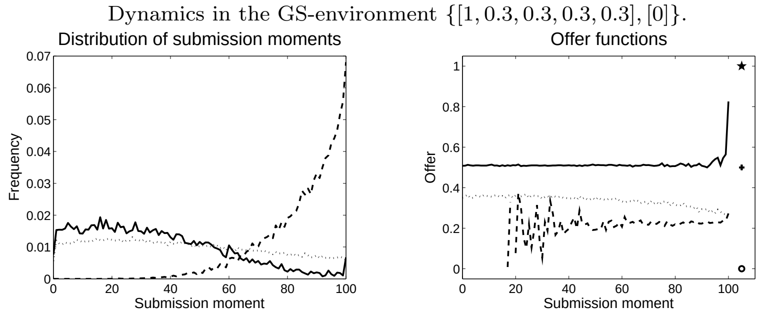

Dynamics in the GS-environment $\{[1, 0.5, 0.5, 0.5, 0.5], [0]\}$.

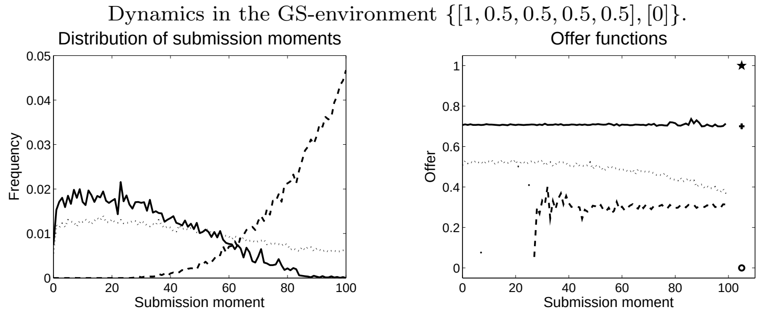

Dynamics in the GS-environment $\{[1, 0.7, 0.7, 0.7, 0.7], [0]\}$.

Fig. 6 Distribution of the submission moments (*left*) and offer functions (*right*), in the GS-environment with 1 intramarginal buyer (solid line), 4 extramarginal buyers with valuation β (dotted line for average behaviour) and 1 seller (dashed line).

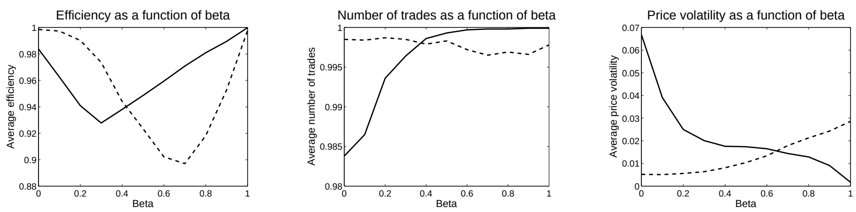

Fig. 7 Characteristics of the GS-environments with 4 extramarginal buyers with valuation β, with timing (*solid line*) and without timing (*dashed line*)

the mid point of the interval of equilibrium prices. At the same time, the seller submits her offer as late as possible to trade at the best bid submitted by the buyers. Under such behaviour the size of the ask does not matter, provided that it is sufficiently low to attract at least one buyer. The intramarginal buyer faces more competition but mostly through offers. With a late submitting seller, he is forced to trade earlier than the seller, to make sure he only has to submit a high enough order to outbid the extramarginal buyers. As β increases, the competition between the intramarginal and extramarginal buyers places a higher 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's and with the efficiency in the IEL model without timing of Anufriev et al. (2013). The three panels of Fig. 7 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 intramarginal buyer faces almost no competition and hence can be quite aggressive. Both intramarginals are now aggressive similar to the model without timing, but they now have two decisions in which they can be (too) aggressive, leading to a lower efficiency. When β is high, the intramarginal buyer faces a lot of competition, and is forced to arrive early and submit a conservative bid to outcompete the extramarginals. Hence the probability of trade increases, but efficiency can still be low if an extramarginal buyer is trading. Intermediate values of β correspond to the lowest possible efficiency. The number of trades is increasing and the volatility is decreasing in β. In the model without timing, the seller cannot optimise his profit solely by arriving last and therefore attempts this by submitting an aggressive offer, which leads to a lower efficiency for large values of β when both intramarginals are (too) aggressive. Comparing with the model without learning about timing, we find a lower efficiency, a lower number of trades and a higher price volatility for the case of low β and the opposite for larger values of β.

Compliance with Ethical Standards:

Michiel van de Leur has received a research grant from the European Commission in the framework of the European Doctorate in Economics-Erasmus Mundus (EDE-EM).
Mikhail Anufriev has received a research grant from the Australian Research Council through Discovery Project DP140103501.

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
Bottazzi G, Dosi G, Rebesco I (2005) Institutional architectures and behavioral ecologies in the dynamics of financial markets. J Math Econ 41:197-228
Camerer C, Ho T-H (1999) Experience-weighted attraction learning in normal form games. Econometrica 67:827-874
Cervone R, Galavotti S, LiCalzi M (2009) Symmetric equilibria in double auctions with markdown buyers and markup sellers. In: Hernandez C et al. (eds) Artificial Economics, Springer, pp 81-92
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
Conlisk, J (1996) Why bounded rationality? J Econ Lit 34(2): 669-700.
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
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

Zhan W and Friedman D (2007) Markups in double auction markets. J Econ Dyn Control 31:2984-3005