Journal of Economic Dynamics & Control 141 (2022) 104387

Contents lists available at ScienceDirect

Journal of Economic Dynamics & Control

journal homepage: www.elsevier.com/locate/jedc

The role of information in a continuous double auction: An experiment and learning model☆

Mikhail Anufriev[a,*], Jasmina Arifovic[b,1], John Ledyard[c], Valentyn Panchenko[d]

[a] Economics Discipline Group, Business School, University of Technology Sydney, NSW 2007, Australia
[b] Department of Economics, Simon Fraser University, Burnaby, B.C. V5A 1S6, Canada
[c] Division of the Humanities and Social Sciences, California Institute of Technology, MC 228-77, Pasadena, CA 91125, USA
[d] School of Economics, UNSW Business School, University of New South Wales, Sydney, NSW 2052, Australia

ARTICLE INFO

Article history:
Available online 14 April 2022

JEL classification:
D83
C63
D44

Keywords:
Continuous double auction
Experiments
Individual evolutionary learning

ABSTRACT

We analyze trading in a modified continuous double auction market. We study how more or less information about trading in a prior round affects allocative and informational efficiency. We find that more information reduces allocative efficiency in early rounds relative to less information but that the difference disappears in later rounds. Informational efficiency is not affected by the information differences. We complement the experiment with simulations of the Individual Evolutionary Learning model which, after modifications to account for the CDA, seems to fit the data reasonably well.

1. Introduction

Is more detailed past information important for the performance of financial markets? In particular, does access to past trading information affect the allocative efficiency of the market? Will more information lead to higher or lower market volatility? These questions are receiving growing attention in the literature, especially as access to information gets easier and algorithmic trading spreads. In the experimental literature, one of the earliest contributions on the importance of information in trading is by Smith (1980). In the context of the double action, he compares the environment with incomplete information, when each trader knows only its marginal value or cost, and complete information, when all traders know values and costs of all traders. He finds that convergence to the competitive equilibrium price occurs slower under complete information. Arifovic and Ledyard (2007) study the effect of the availability of past information for allocative efficiency in double auctions organized as a call market. In this paper, we address these questions for the markets that have an order book, such as the continuous double auction (CDA). For this purpose and to disentangle several confounding effects that the

☆ We thank the organizers and participants of the 2021 Conference on Markets and Economies with Information Frictions, and especially the discussant, Te Bao, for many suggestions that helped to improve the paper. We also thank participants of the ESA meeting in Tucson, and the seminar at Simon Fraser University for their comments on the earliest results of this research. We are grateful to Michiel van de Leur for providing research assistance during the earlier stages of this research. Mikhail Anufriev acknowledges financial support from the Australian Research Council through Discovery Project DP200101438.

* Corresponding author.
E-mail address: Mikhail.Anufriev@uts.edu.au (M. Anufriev).

[1] Sadly, Jasmina Arifovic passed away on January 24, 2022. We are endlessly thankful to Jasmina for developing individual evolutionary learning, getting us excited about the topic, and bringing us together as a research team. She was a true scientist and had a great personality.

https://doi.org/10.1016/j.jedc.2022.104387

use of the order book brings to the behavior of traders, we propose a new trading mechanism and run an experiment that matches this mechanism with treatments corresponding to different access to past information. Further, we complement the experimental results with an evolutionary learning model and show that the model is able to capture the experimental outcomes, particularly during the initial, learning stage of the market.

The double auction, a centralized market with many buyers and sellers, comes in many flavors; see, e.g., Friedman (2018) and Easley and Ledyard (2018). Two of the most well known implementations are the call market (CM) and the continuous double auction (CDA). The key difference between these two mechanisms is that in the call market, orders are submitted and cleared simultaneously, while in the CDA market orders are submitted asynchronously and are cleared at different times during the trading session. In the call market, after all orders are submitted, the market is cleared at a single equilibrium price. Buyers with bids at or above that price and sellers with offers below or at that price trade. In the CDA, the order book keeps all unsatisfied orders (bids and offers). When a new order arrives, if possible it is matched with the closest order at the opposite side of the book, or, otherwise, it is stored in the book. There are specific prices for every transaction.[2]

In this paper, we are interested in a hybrid double auction, called the “Continuous Double Auction with Synchronous Decisions” (CDA-SD). In the CDA-SD, the orders are formulated simultaneously, as in the call market, but they are cleared asynchronously, as in the CDA.[3] We focus on this market mechanism for two reasons. First, due to legal aspects as well as considerations of efficiency, many participants in financial markets delegate trading to brokers. This leads to a separation between investors’ decisions (to formulate the order, i.e., the highest price to buy or the lowest price to sell) and the activity of brokers (who have access to the market and execute trades). We focus on investors who naturally make their decisions at some specified time, such as at the beginning of a day, and from whose perspective the trading mechanism (which is CDA, as in most of the markets) is not as important. Thus, we think of the CDA-SD – that combines synchronous decisions with realistic trading protocol – as the mechanism useful to study the behavior of investors.

Second, we address the question of the impact of past information availability for market efficiency by building on the earlier contribution of Arifovic and Ledyard (2007), AL henceforth.[4] The CDA-SD mechanism allows us to disentangle the effect of access to information about the trading history from the effect of an influence of the contemporaneous state of the order book on the trading strategies. We focus on the former effect, keeping our analysis closer to the AL.[5]

AL found that more information is harmful for call markets: in their experiment, allocative efficiency is smaller in the treatment where participants could see all offers submitted in the previous session, in comparison with the treatment where such detailed information is not available. AL complement their experiments with simulations of the computational model where artificial agents use the Individual Evolutionary Learning (IEL) algorithm.[6] Our previous work, Anufriev et al. (2013), AALP henceforth, theoretically addresses the efficiency of CDA under two different information feedbacks. Assuming that the artificial agents use the same IEL as in AL, AALP focus on the bidding strategies that the algorithm selects in the stationary state.[7] AALP find that the past available information affects the set of these selected strategies, though this difference does not translate into different allocative efficiency. Some other market characteristics are affected, however. For instance, price volatility decreases when the amount of past information increases.

The experiment designed in this paper not only tests these conjectures but also allows us to go beyond the analysis based on the stationary state. Indeed, the initial, learning phase is as important as the stationary dynamics, because in the

[2] As a bid arrives, it is matched with the lowest offer from the book among those that are below or equal to the bid. If there are several offers in the book satisfying this criteria, the bid is matched with the offer that arrived the earliest (i.e., according to the time/price priority). As an offer arrives, it is matched with the highest bid in the book among those that are above or equal to the offer. Again, time/price priority is applied. In both cases, the transaction price is the price of the order that was in the book, i.e., arrived earlier. During the trading session, brokers see the dynamically changing book but agents do not. As explained later, in this paper, agents cannot condition their behavior on the book’s current state, but may condition it on the historical data from the book.

[3] In this way, we abstract from optimal timing considerations in order placement. Optimal order timing in the CDA is investigated by van de Leur and Anufriev (2018) who that the optimal arrival has a skewed hump-shaped distribution which depends on the environment and the market size.

[4] The quest for the sources of market efficiency attracts a lot of attention in the economic literature. The results of experiments with human subjects starting with Smith (1962) show quick convergence towards competitive equilibrium, resulting in high allocative efficiency of the continuous double auction (CDA). Market efficiency is determined by both market rules and traders’ behaviors, and disentangling one from another is challenging (Bottazzi et al., 2005). Gode and Sunder (1993, 1997) show that the “Zero-Intelligent” (ZI) agents submitting orders randomly are able to achieve high level of efficiency. However, the results strongly depend on the behavioral assumption that the ZI are trading within the budget constraints, see Gjerstad and Shachat (2007) and LiCalzi and Pellizzari (2008).

[5] Even when traders are observing the order book closely, their orders arrive to the market with delays, resulting in some randomness in the outcome of trading. Furthermore, many financial exchanges deliberately introduce “speed-bumps”, i.e., (random) delays in the execution of orders to avoid market instabilities, see Table 2 in He et al. (2020), who model uncertainty in order execution timing. Thus, whereas the sort of decisions that our subjects made in our CDA-SD experiments differ from the decisions in the standard CDA experiments, our construction is, in this respect, closer to the real markets.

[6] This learning model is an appropriate modeling tool for repeated complex environments, see Arifovic and Ledyard (2004, 2011). For recent examples of IEL applications, see Arifovic and Ledyard (2018), van de Leur and Anufriev (2018) and Arifovic et al. (2019). The agents in IEL are boundedly rational, as they learn without taking into account that other agents are learning as well. IEL can be thought of as a simplified version of genetic algorithm learning, introduced in the economic literature in Arifovic (1994) and used recently in Anufriev et al. (2019). The word “individual” in IEL stresses that the agents are not involved in social learning, as in the models based on imitation behavior like in Dawid (1999); see Vriend (2000) for an example stressing the difference between the two types of learning.

[7] The AALP prove, for selected demand/supply schedules, that certain bidding strategy profiles are evolutionary stable under the IEL. They run simulations for other, more complicated, demand/supply schedules, but study these simulations after 100 transitory periods, where initial learning takes place.

real markets the environment (sets of valuations/cost of traders) may often change from period to period, and even experienced traders may find themselves in a situation that is similar to the initial few periods of our experiment. We find that in the experiments there is a significant effect of information on allocative efficiency in the first periods with *more information leading to lower efficiency*. This is the same effect as AL found for call markets. However, learning in our experiment is quicker, when there is more past information available, so that eventually allocative efficiency in both information treatments becomes comparable, as in the stationary state studied in AALP. By looking at several other measures, including the price volatility, we also find that the exact configuration of the demand/supply schedules affects the outcome at least as much as the information structure. Finally, we ask whether the IEL model, as introduced in AALP, can capture both short and long-term features of the dynamics. The availability of experimental data allows us to find a simple but intuitive improvement in the IEL algorithm that captures the data relatively well in both experimental treatments. This is an example of how experiments are crucial in complementing theoretical analysis, especially when that analysis is based on a computational model with bounded rationality assumptions.[8]

The rest of the paper is organized as follows. Section 2 introduces the information differences that we study, discusses related literature, and formulates the experimental hypotheses on the basis of the previous theoretical study of the AALP, Anufriev et al. (2013). Section 3 explains the experimental design and provides the details of the experiment that we run in two locations. Section 4 presents the experimental results. In Section 5 we revisit the Individual Evolutionary Learning model with the experimental data and demonstrate that a slight variation in the model can fit the data well. Section 6 concludes. The Appendix contains the experiment instructions, screenshots, and some estimation results.

2. Role of information and IEL

In this section, we introduce two different information settings used in the literature. We also introduce the Individual Evolutionary Learning (IEL) model that Anufriev et al. (2013) (AALP) studied theoretically.

2.1. Two information feedback scenarios

Financial markets have witnessed an increasing access of the traders to past detailed information, see, for instance, Boehmer et al. (2005) and Easley et al. (2016). Arifovic and Ledyard (2007) pose the question: will the allocative efficiency of markets increase if agents use richer past information? AL do this in the context of a call market, focus on learning under repeated trading, and disentangle the two information scenarios. In the *closed book* scenario, agents have no access to the individual level data from the past market session; they only know the past market clearing price. In contrast, in the *open book* scenario, the information about past individual orders is available to traders. The latter scenario provides traders with strictly more information than the former. It turns out, perhaps surprisingly, that access to more information results in a lower market efficiency. AL find this by conducting experiments with human subjects and via simulations of the Individual Evolutionary Learning (IEL) model with artificial traders.

The difference between the two information scenarios above, when extended to the CDA markets, is in the access to the orders from the previous session. To separate any strategic effects that information from the current order book may produce when decisions are made *during* the trading session, we introduce the CDA with Synchronous Decisions (CDA-SD) mechanism. This is the market with an order book matching the orders that *arrive* asynchronously, but where the traders *formulate* their orders simultaneously. Thus, in the CDA-SD, the orders are decided before the trading session and arrive at random times during the session. The session is organized as the standard CDA with the order book. This construction allows us to keep the distinction between the same two information scenarios that AL studied, which we call *Aggregate-level* (market) *feedback* (**AF**) and *Individual-level* (orders) *feedback* (**IF**). In the **Aggregate Feedback** case, all traders have access only to the average transaction price of the past session. In the **Individual Feedback** case, all traders have a detailed access to the order book of the previous session. The **AF** corresponds to the closed book scenario in AL and the **IF** corresponds to the open book scenario in AL.[9]

Several related studies vary information availability and compare market performance. Apart from allocative efficiency, this literature focuses on the market characteristics related to "information efficiency". Information efficiency refers to the market's ability to aggregate individual information (for example, individual valuations and cost) into the price. When markets are informationally efficient, there are no systematic deviations of the price from its equilibrium values. Low price volatility (for repeated trade in a fixed environment) is another measure reflecting information efficiency. An empirical study of Boehmer et al. (2005) found behavioral and market changes that followed the New York Stock Exchange decision to open past order books to traders to increase transparency. In particular, this decision resulted in lowering the price volatility and increasing market liquidity, that is in higher information efficiency. The theoretical study of Baruch (2005), where the call market mechanism is assumed, compares the case of "open book" information environment that occurred as the result of an NYSE decision (our **IF** case) with the "closed book" information environment that existed before (our **AF** case). The study

[8] Computational agent-based models have gained popularity in Economics and Finance, for recent advances see Hommes and LeBaron (2018), Chen et al. (2018) and Dieci and He (2018).

[9] We do not use the open/closed book terminology that AL and AALP used to avoid a confusion with the cases when agents may access the *current* session order book.

finds that larger transparency favors traders that demand liquidity (e.g., those who submit market orders) at the expense of the traders who provide liquidity (e.g., specialists and limit order traders).

Theoretical studies of the CDA typically rely on simulations and abstract from many factors that are in place in real markets (order size, market orders, possibility of cancellations, and so on), and focus on the impact of information only. Ladley and Pellizzari (2014) investigate trading strategies in a market organized as a CDA with order book and find that the amount of information about the book (beyond the best quotes) has little effect on the performance. Fano et al. (2013) compare call markets with CDAs. The strategies of traders evolve over time following a genetic algorithm that favours strategies of better performing agents. It turns out that, in the call market, traders become price-takers, offering their valuation or cost, and in the CDA, they become price-makers, bidding the equilibrium market price. Anufriev et al. (2013) use IEL and compare the **AF** and **IF** information environment under the CDA market. They find that traders behave more like price-takers in the **AF** environment and tend to be more like price-makers in the **IF** environment, as more information becomes available.

All of the studies discussed above are focused on outcomes after some learning stage, that is on some "equilibrium" outcome. In this paper, instead, we are interested in the experimental evidence of immediate learning of human subjects under the two feedback environments. We will use the equilibrium predictions of the IEL algorithm from AALP to formulate hypotheses and organize the results. However, we will extend their IEL algorithm in Section 5 to better match our short-run experimental data.

2.2. The IEL model

The Individual Evolutionary Learning (IEL) model (Arifovic and Ledyard, 2003; 2007) is an appropriate computational test-bed to study the market design questions discussed above. Indeed, in a complex environment with a large strategy space, it is fairly impossible to get analytic solutions to equilibrium models, and it is also unlikely that traders will behave fully rationally from the outset. The IEL model defines a computational algorithm that simulates the process of learning. Arifovic and Ledyard (2004) show that this algorithm performs better than other learning rules, and the outcomes of the IEL model are very similar to experimental outcomes in many situations where subjects have continuous or large strategy spaces.

Given the three building blocks – (i) a specific environment, consisting of traders' endowments and valuations and costs, (ii) a trading protocol, defining the outcome of the trading session given the strategies of traders, and (iii) information feedback, i.e., what information is available to traders between consecutive trading sessions – the IEL algorithm defines a multi-dimensional stochastic process. The state variables of this process are the individual bidding strategies that agents use and the aggregate market variables, such as prices. IEL can be used to produce theoretical predictions. For example, the two information feedback cases, **AF** and **IF**, as defined above, create a contrasting set of IEL-simulations and corresponding statistics for allocative efficiency, average price, price volatility, and so on.

IEL, in its simplest form, has only two parameters, the size of the strategy space, J, that reflects the cognitive capacity of traders, and the probability of experimentation, ρ, that models the rate at which new strategies are experimented with at each stage. IEL also depends on the specification of "hypothetical" utility; that is, the utility that traders would have received from playing a strategy in the past. Traders are boundedly rational and compute hypothetical utility without taking into account the learning processes of others. Hypothetical utilities depend, generally, on the environment, the trading protocol, and the information feedback.

2.3. IEL for the CDA-SD markets

The IEL model that AALP study can be used on the CDA-SD markets to confront the **AF** and **IF** settings. We present their results in order to formulate our experimental hypotheses.

The trading sessions are in discrete time, with periods indexed by t. Several buyers and sellers, whose valuations and cost are exogenously given and fixed over all periods, trade repeatedly. Each trader can buy or sell at most one unit of a good every period. Traders know their own valuations and costs, but not the valuations and costs of others, neither do they know the distributions. Utilities are linear; that is, they are valuation minus transaction price for buyers and transaction price minus cost for sellers. Agents who do not trade, get utility 0. Let V_b and C_s denote the valuation of buyer b and cost of seller s, respectively. The set of valuations and costs define an environment.

Trade on the CDA-SD market is organized as follows. At the beginning of each period, each trader submits one order (bid for a buyer, offer for a seller). These orders arrive at the market organized as the CDA in random order. The CDA defines a (possibly empty) set of trades and corresponding transaction prices, according to the standard rules (as described in the Introduction).

After the period ends, all traders receive the same between-period feedback. Two feedback scenarios are considered. The richer **IF** scenario provides each trader with a detailed information about the order book from the previous period. Specifically, each trader can see all individual bids and offers, as well as how the order book evolved. The **AF** scenario provides each trader with the average transaction price only.

Under the IEL algorithm, every artificial agent is endowed with an individual pool of strategies, evolving in time. The pools are denoted as $B_{b,t}$ and $A_{s,t}$ for buyer b and for seller s, respectively, and are composed of J real numbers belonging to the *admissible* intervals, which are $[0, V_b]$ for buyer b and $[C_s, 100]$ for seller s.[10] When simulations start, at period $t = 1$, the initial pools are formed by J uniform draws from the corresponding admissible interval, independently for all agents. In this period, each agent takes one of the strategies from the pool with equal probability and submits it. The trading mechanism matches orders and defines prices. After the period, each trader receives an information according to the feedback. Before the next period starts, the IEL model plays a role in (i) updating each agent's pool and (ii) selecting a new order from that pool. The same process then is repeated and so on.

At the beginning of every period, the pools of all traders are updated independently, in two consecutive stages. At the *experimentation* stage, every element of the pool is either removed with probability ρ, or remains with probability $1 - \rho$. If an element is removed, it is replaced by an element drawn from some distribution truncated to the admissible interval of the trader.[11] After this procedure is repeated independently for each of the J positions in the previous period pool, an intermediate pool is formed. At the *replication stage*, the final pool is obtained, element-by-element, repeating the following process J times. Two randomly chosen strategies from the intermediate pool are compared with each other (with replications), with the best of them occupying a place in a new pool. The comparison between strategies is made according to a performance measure, called hypothetical utility, denoted as U^I or U^A since this utility depends on the information feedback. After the new pools are formed, the strategy is selected for each trader randomly from their pool, with probabilities proportional to the hypothetical utilities. For instance, the probability that buyer b selects bid b_i in period t under **IF** is given by

$$\pi_{b,t}(b_i) = \frac{U^I_{b,t}(b_i)}{\sum_{k=1}^J U^I_{b,t}(b_k)} ,$$

where $U^I_{b,t}(b_k)$ is the hypothetical utility of buyer b at time t from bidding b_k.

AALP specified the hypothetical utilities as follows. For the **AF** setting ("closed" book in AL), where only the average price of all transactions from the previous period, $\bar{p}_{t-1}$, is reported back to each trader, the hypothetical utilities are

$$U^A_{b,t}(b_i) = \begin{cases} V_b - \bar{p}_{t-1} & \text{if } b_i \geq \bar{p}_{t-1} , \\ 0 & \text{otherwise} , \end{cases} \tag{1}$$

for buyer b. Analogously, the hypothetical utility of seller s is

$$U^A_{s,t}(a_i) = \begin{cases} \bar{p}_{t-1} - C_s & \text{if } a_i \leq \bar{p}_{t-1} , \\ 0 & \text{otherwise} . \end{cases} \tag{2}$$

For the **IF** setting ("open" book in AL), where traders can see the whole book of the previous session, the agents substitute their orders to the last session book and find their corresponding utility. The hypothetical utility of buyer b from bid b_i is

$$U^I_{b,t}(b_i) = \begin{cases} V_b - p^*_{b,t-1}(b_i) & \text{if bid } b_i \text{ would lead to a trade at } p^*_{b,t-1}(b_i) , \\ 0 & \text{otherwise} . \end{cases}$$

Analogously, seller s computes the hypothetical utility of offer a_j as

$$U^I_{s,t}(a_i) = \begin{cases} p^*_{s,t-1}(a_i) - C_s & \text{if offer } a_j \text{ would lead to a trade at } p^*_{s,t-1}(a_i) , \\ 0 & \text{otherwise} . \end{cases}$$

In other words, the hypothetical utility is what a trader would get last period with the order, if all other traders submitted the same orders they did and the sequence of the orders in the book would be the same.[12]

2.4. AALP results and experimental hypotheses

AALP study the stationary state of the IEL algorithm after a long transitory period. They found that, in the long-run, the strategies employed by the artificial traders depend on the information feedback.[13] In the **IF** case, traders learn to become "price-makers", as they tend to submit similar orders belonging to the range of the equilibrium prices. This leads to relatively stable prices, with low volatility and high information efficiency. The occasional experimentation of IEL, leads however to the possible loss of allocative efficiency due to missing transactions. The strategies and dynamics are different under the **AF**

[10] The size of the pool is kept constant, but note that it is possible (and in fact common) that the pool has repeated strategies. In all our environments, the valuations and costs are between 0 and 100. We impose individual rationality constraints, not permitting traders to submit offers that could result in a negative profit.

[11] Simulations in AALP use the uniform distribution. However, we found that our experimental data are matched better when the experimentation occurs from the truncated normal distribution, with the mean given by the element that is replaced. In other words, local experimentation describes the data better.

[12] The assumption of the same sequence is a reasonable behavioral assumption. Alternatively, traders could "simulate" all possible sequences of orders' arrival and evaluate expected hypothetical utility. However, the number of computations for this is very large.

[13] More precisely, the major differences are in the strategies of the infra-marginal traders, i.e., traders who should trade in the demand/supply imposed equilibrium model. The other, extra-marginal traders submit random admissible strategies but trade very rarely. We define the infra- and extra-marginal traders in Section 3.

Table 1
Treatments of the experiment and the number of observations.

Treatments		UNSW	Caltech	Total
AF-1	aggregate feedback; schedule **S1**	4	5	9
AF-2	aggregate feedback; schedule **S2**	4	5	9
IF-1	individual feedback; schedule **S1**	4	5	9
IF-2	individual feedback; schedule **S2**	4	5	9

case, where traders submit the orders outside of the equilibrium price range, close to their valuations and cost (to make sure that they transact), and behave thus similar to the "price-takers". As they do so, the price gets volatile, sometimes leaving the equilibrium price range, and lowering informational efficiency. The loss of allocative efficiency occurs due to possible trading of the extra-marginal traders (who should not trade in the equilibrium). In both cases, the allocative efficiency is in the range of 88% – 95% for most of the IEL parameters.

One important caveat to these results is that the AALP describe only some but not necessarily all stationary states. Moreover, the results hold precisely only for the specific schedules taken from Gode and Sunder (1997). Simulations for more sophisticated environments, similar to those that we use in the experiment reported in this paper, suggest that the result about low information efficiency in the **AF** setting can be extended, but the impact of the feedback on allocative efficiency is somewhat uncertain due to a strong interaction with the environment.[14]

This discussion leads to the following two hypotheses that our experiment will test.

Hypothesis 1. Full allocative efficiency of 100% is not achieved under the CDA-SD both in the **AF** ('closed' book) and in the **IF** ('open' book) information feedback scenario. The ranking of the allocative efficiency for the two information feedback scenarios depends on the schedule.

Hypothesis 2. Price volatility is significantly higher in the **AF** ('closed' book) than in the **IF** ('open' book) information feedback scenario. The average transaction price stays in the equilibrium price range in the **IF**, but may leave it in the **AF**.

Both parts of Hypothesis 2 suggest that the **AF** will have lower information efficiency than the **IF**.

3. Experiment

The identical experiment sessions were run in two locations. 8 sessions were conducted in October 2010 at the Business Experimental Research Laboratory (BizLab) at the UNSW, Sydney; and 10 sessions were conducted in October 2012 and February 2013 at Caltech's Laboratory for Experimental Economics and Political Science (EEPS). The experiments were computerized using the zTree software (Fischbacher, 2007). In both locations, the subjects were mostly undergraduate and some postgraduate students majoring in different areas. The participants were recruited from the large pools through the ORSEE system (Greiner, 2015).

There were 10 participants in each session. The session incorporated two blocks, with 20 trading rounds each. Blocks corresponded to two different environments, i.e., sets of valuations and cost. The valuation and cost of every participant were kept the same during all trading rounds of each block. In each round, 5 buyers and 5 sellers traded according to the CDA-SD protocol. Each buyer demanded one unit of a commodity whose valuation was privately known, and each seller could sell one unit of that commodity whose cost was privately known. Every participant was able to submit an order with up to two decimal digits. These individual (limit) orders, bids and offers, were collected before the trading period, and then the order book was simulated with random arrival of these orders. The calculation of the payoff per trading round is standard for trading experiments: the buyer's payoff is given by valuation minus transaction price, if the buyer traded, and 0, otherwise. The seller's payoff is given by the transaction price minus cost, if the seller traded, and 0, otherwise.[15] The procedure and incentives were explained to the participants before the experiment.

We ran two treatments of the experiment that differed in the feedback that participants received between trading rounds. Those corresponded to the *Aggregate-level* (market) *feedback* (**AF**) and *Individual-level* (orders) *feedback* (**IF**) scenarios explained in Section 2.1. Before every trading round, the participants could see information from the previous round and analyze it for 20 seconds. Then the information window was supplemented on the screen with a decision window and subjects had additional 60 seconds to submit the offer. In the treatments with the **AF** scenario, participants only knew the previous average price of all transactions, their previous period earning (from which they could infer whether they traded, and if yes, then their transaction price) and their cumulative earnings. In the treatments with the **IF** scenario, in addition to this information, participants were faced with the whole order book of the previous period, i.e., with all 10 submitted bids and asks (without identities of traders) in the order of their arrival. If no transaction was recorded during the previous

[14] For most of the IEL parameters in the AALP, the allocative efficiency is slightly higher in the **AF** case, for our **S1** schedule in Section 3. Instead, it is substantially higher in the **IF** case, for the schedule that is very similar to our **S2** schedule in Section 3.

[15] To prevent negative payoffs, the bids could not exceed the buyer's valuations and the offers could not exceed the seller's cost.

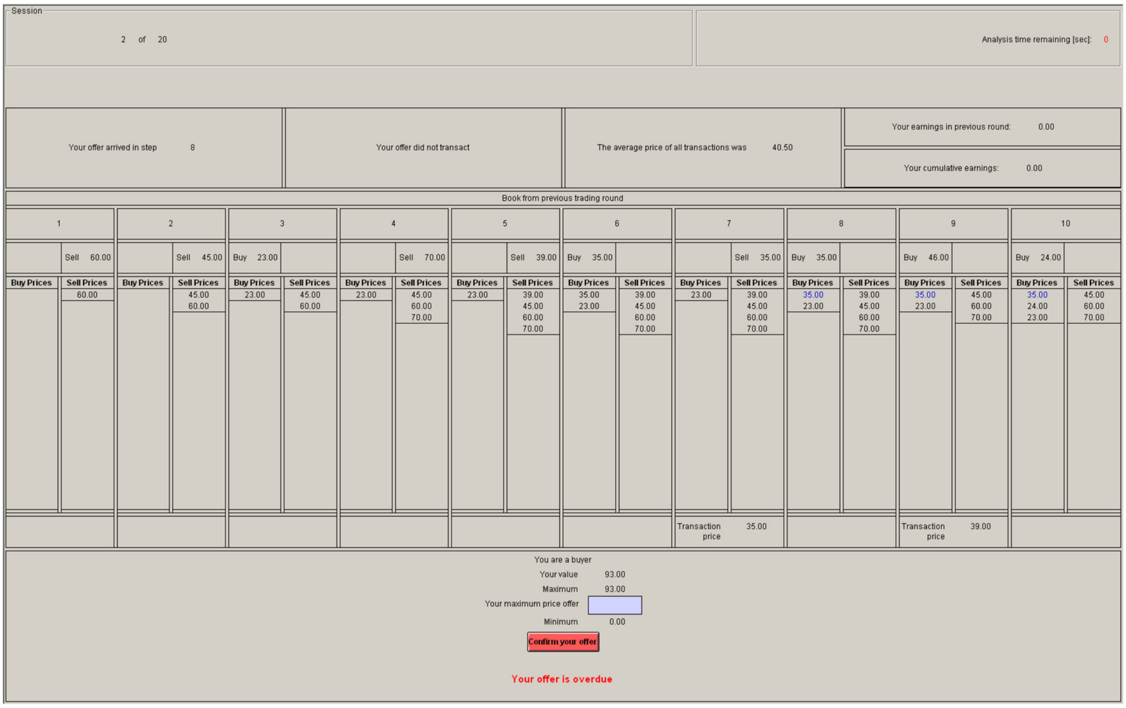

Fig. 1. Information and decision windows shown in the treatment with individual level information feedback, **IF**. The decision part contained the counter. When the offer was overdue, the overdue message appeared.

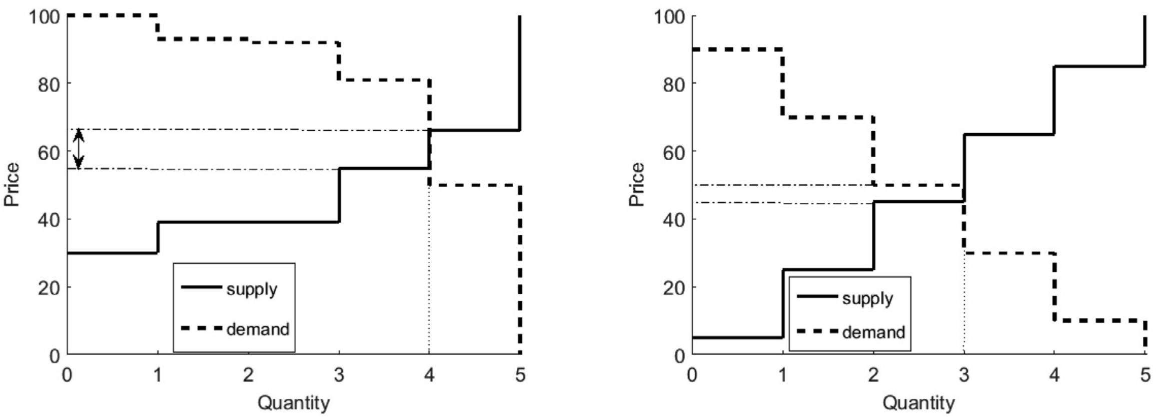

Fig. 2. Demand/Supply diagrams for the market configurations considered in the paper. *Left:* **S1**-market. *Right:* **S2**-market.

round, participants in both treatments were informed about this, and the participants in the **IF** treatment could see the evolution of the order book.[16]

Fig. 1 shows the screen from the **IF** treatments with combined information and decision windows. The decision window (identical in both **AF** and **IF** treatments) is in the lower part of the screen, below the table. It shows the role of the participant, the valuation (or cost), contains the window to type the offer, and displays the time counter. The information window for the **IF** treatment contains, in the upper part, the table with four columns showing (from left to right) the step when the order of the participant arrived, whether it resulted in a transaction, the average price of all transactions, and the last period and cumulative earnings. Then, in the middle part, it shows how the order book evolved in the previous session. The information screen in the **AF** treatment had only the last two columns of the upper table and no middle part. The examples of the screens from all treatments can be found in Appendix B.

[16] There were only 3 periods with no transactions. This is less than 0.5% of the total number of 720 trading periods ($9 \times 4 = 36$ trading blocks with 20 periods each).

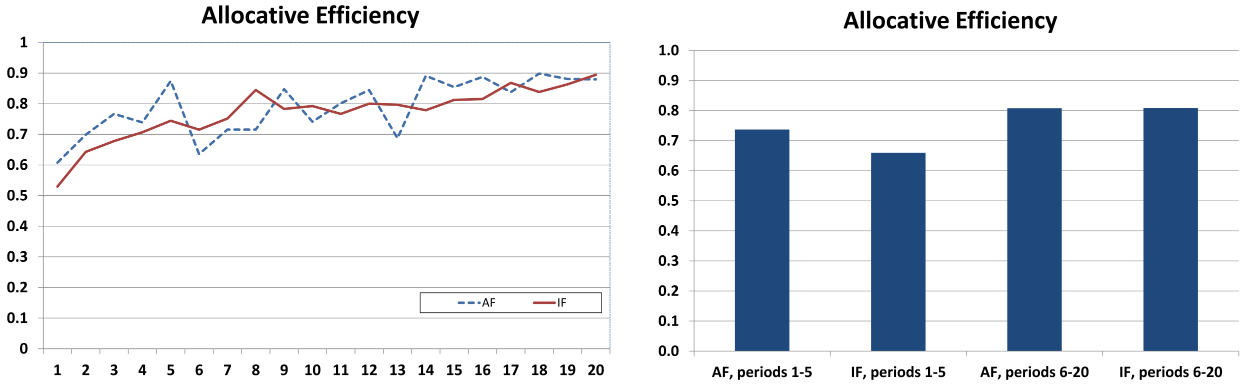

Fig. 3. Allocative efficiency in the experiment: evolution over 20 periods (left) and split for each treatment into the initial periods, 1-5, and subsequent periods, 6-20 (right).

In the experiment, we used two different market schedules, both with 5 buyers and sellers, see Fig. 2. Running an experiment with different schedules is necessary to check robustness of the differences between the information scenarios. We picked two schedules that differ in the equilibrium quantity and price ranges and have been used in the previous studies.[17] Schedule **S1** (the left panel) has 4 trades in equilibrium with the range of equilibrium prices [55,66). In schedule **S2** (the right panel), the equilibrium quantity is 3 and equilibrium price range is [45,50]. Thus, **S1** allows more trades and has a larger range of equilibrium prices, leading to potentially higher efficiency than **S2**.[18] The traders that would trade in equilibrium are called infra-marginal and those who would not trade in equilibrium are called extra-marginal. The total surplus of trade (regardless of the mechanism) is maximized when the infra-marginal buyers (IMB) trade with infra-marginal sellers (IMS). There are 4 IMB and 4 IMS in **S1** with the largest surplus equal to 203. There are 3 IMB and 3 IMS in **S2** with the largest surplus equal to 135. The *allocative efficiency* is defined as the fraction of this surplus extracted during the trading session.

After reading the instructions (see Appendix A) and completing a test designed to check the subjects' understanding of the instructions, the first block of the experiment started. At the beginning of this block each participant was randomly assigned to be a buyer or a seller; every buyer saw their valuation and every seller saw their cost. After 20 trading rounds, a new block began. Every buyer changed their role to become a seller, every seller changed their role to become a buyer, and the valuations and costs were assigned according to a new schedule.[19] Then the next 20 trading rounds started. When those were over, subjects filled in a questionnaire and collected earnings equal to a show-up fee of 5 dollars and all their accumulated payoff over 40 trading rounds.

The experiment has a 2×2 design, as we varied information feedback and the schedule.[20] In total we have 4 treatments (2 information feedback scenarios and 2 different schedules), see Table 1. Taking into account sessions in both locations, we have 9 observations for each treatment.

4. Experimental results

We present the results in the same order as the two hypotheses formulated in Section 2.4. We start with the allocative efficiency in Section 4.1 and then discuss the measures of informational efficiency in Section 4.2. Our findings are supported by a regression analysis, the details of which can be found in Appendix C.

Before presenting the results, we provide an overview of how we analyzed the data. We illustrate the time evolution of average allocative efficiency in Fig. 3 (left) for the two information treatments (averaging the results from all sessions for both schedules). Allocation efficiency initially goes up quickly but exhibits variability in time over the whole course of the experiment. Variability is observed for other characteristics we are interested in, as well, and is present in almost any

[17] For schedule **S1**, the valuations/costs are $V_1 = 100$, $V_2 = 93$, $V_3 = 92$, $V_4 = 81$, $V_5 = 50$, $C_1 = 30$, $C_2 = C_3 = 39$, $C_4 = 55$ and $C_5 = 66$. This is schedule 1 in Arifovic and Ledyard (2007) and it was referred as 'AL' in simulations in Anufriev et al. (2013). For schedule **S2**, the valuations/costs are $V_1 = 90$, $V_2 = 70$, $V_3 = 50$, $V_4 = 30$, $V_5 = 10$, $C_1 = 5$, $C_2 = 25$, $C_3 = 45$, $C_4 = 65$ and $C_5 = 85$. This schedule is very similar to schedule 2 in Arifovic and Ledyard (2007) and the symmetric 'S5' market in Anufriev et al. (2013).

[18] We verified it by running a simple simulation where agents from a schedule are matched randomly and the total surplus is computed. For **S1** we achieve an average efficiency of 46% in comparison to an efficiency of 10% for **S2**.

[19] The subjects knew about the change of the role between blocks. To balance expected payoffs for the participants, we re-assigned the buyers with higher valuations in the first block to become the sellers with higher cost in the second block, and the same for the sellers from the first block who became buyers. The subjects did not know this.

[20] In each session we used both schedules, running them in different orders. We found no significant impact of the order in which the schedule appeared. We also found no evidence that the experience of participants with one schedule would affect their behavior with another schedule, which is not surprising given that participants changed their roles between blocks.

Table 2
Means (with standard errors in parentheses) of different characteristics for experimental data. All characteristics are averaged over all trading periods, except for the price volatility that is averaged over sessions. The last two columns report the *t*-test statistics and *p*-values for differences in the means between treatments for the data pooled across the two schedules. When computing the *t*-statistics, we subtract from the upper mean the lower mean.

		Schedule 1	Schedule 2	Pooled	*t*-stat	*p*-values
Allocative efficiency						
periods 1–5	**AF**	0.745 (0.021)	0.729 (0.036)	0.737 (0.202)	2.219	**0.013**
	IF	0.666 (0.030)	0.654 (0.050)	0.660 (0.028)		
periods 6–20	**AF**	0.822 (0.014)	0.794 (0.028)	0.808 (0.016)	−0.007	0.497
	IF	0.830 (0.021)	0.786 (0.023)	0.808 (0.016)		
Number of "missed" transactions						
periods 1–5	**AF**	0.911 (0.111)	1.089 (0.111)	1.000 (0.080)	−2.505	**0.006**
	IF	1.1778 (0.091)	1.333 (0.089)	1.256 (0.064)		
periods 6–20	**AF**	0.489 (0.062)	0.785 (0.097)	0.637 (0.067)	−0.447	0.328
	IF	0.563 (0.082)	0.793 (0.081)	0.678 (0.062)		
Indicator of whether the average price is outside of the equilibrium range						
periods 1–5	**AF**	0.356 (0.104)	0.889 (0.048)	0.622 (0.085)	0.096	0.462
	IF	0.467 (0.111)	0.756 (0.093)	0.611 (0.078)		
periods 6–20	**AF**	0.422 (0.101)	0.748 (0.065)	0.585 (0.070)	0.667	0.252
	IF	0.370 (0.117)	0.659 (0.086)	0.515 (0.079)		
Distance from the average price to the equilibrium range						
periods 1–5	**AF**	1.357 (0.446)	8.037 (1.408)	4.697 (1.081)	0.410	0.341
	IF	2.305 (0.795)	5.985 (1.133)	4.145 (0.806)		
periods 6–20	**AF**	1.534 (0.580)	3.074 (0.778)	2.304 (0.506)	−0.676	0.249
	IF	1.284 (0.634)	4.475 (0.975)	2.880 (0.684)		
Price volatility						
periods 1–5	**AF**	4.278 (0.876)	6.664 (1.054)	5.471 (0.725)	0.718	0.236
	IF	4.767 (0.902)	4.703 (1.192)	4.735 (0.725)		
periods 6–20	**AF**	3.273 (0.606)	3.976 (0.651)	3.625 (0.440)	−1.001	0.158
	IF	3.043 (0.750)	5.850 (0.992)	4.446 (0.693)		

session of the experiment. Therefore, we start by averaging the characteristic of interest (such as allocative efficiency) in each session over a specified time period and focus on the corresponding mean.[21] We distinguish, in particular, the initial periods (periods 1–5) and the subsequent periods (periods 6–20). Distinguishing the initial five periods from the rest of the treatment allows us to focus on the learning phase of the experiment. We can also separately look at the later periods that may correspond to a stationary state. Recall that the hypotheses in Section 2.4 are based on the theoretical IEL results from AALP. Those results, in turn, reflected one of the *stationary states* of the model.

After the data of a specific characteristic, such as allocative efficiency, are averaged over time for each session, the session-specific statistics are averaged over all sessions of a given treatment. As the results are independent between different sessions, we compute the mean values and the standard errors of such averaging. We use those means and standard errors to perform the statistical tests. Specifically, for the same characteristic we make the pair-wise comparisons between treatments, compute the *t*-statistics for the difference in mean test, and report in the text whether the difference is significant at the 5% level.[22] Table 2 reports the *p*-values and other data behind the figures and results.

To be more precise in our findings, we run regressions of the following type for various characteristics, using the trading periods as the observation units:[23]

$$\text{Allocative Efficiency} = \beta_0 + \beta_1\text{DInfo} + \beta_2\text{DInit} + \beta_3(\text{DInfo} \times \text{DInit}) + \beta_4\text{DLoc} + \beta_5\text{DSch} + \beta_6\text{DExp} + \text{error}. \quad (3)$$

In this regression the dummy variables are set as follows: DInfo is 1 for **AF** and 0 for **IF**; DInit is 1 for periods 1–5 and 0 for periods 6–20; DLoc is 1 for Caltech and 0 for UNSW; DSch is 1 for schedule **S1** and 0 for **S2**; DExp is 1 for the first block and 0 for the second block of the experiment. Appendix C collects the corresponding estimates for different characteristics.

Running an experiment in two different locations (UNSW and Caltech) is useful to make sure that the differences in treatments are not affected by the participant pool. However, we often find a significant effect of location, with average efficiencies being higher in the Caltech sessions than in the UNSW sessions, when compared for the same treatments.[24] On

[21] To limit the effect of possible outliers, we have done the same analysis based on the medians (not on the means) of data over time. The results are similar and are available upon the request.

[22] The appropriate *t*-statistics is the difference of two means divided by the squared root of the sum of the squares of the two standard errors. The *p*-values are reported for the one-sided tests and are based on 18 independent observations (as we pool the data over locations and schedules).

[23] The other characteristics (defined later) are: the number of "missed" transactions, an indicator of whether the average price is outside of the equilibrium range, and price volatility.

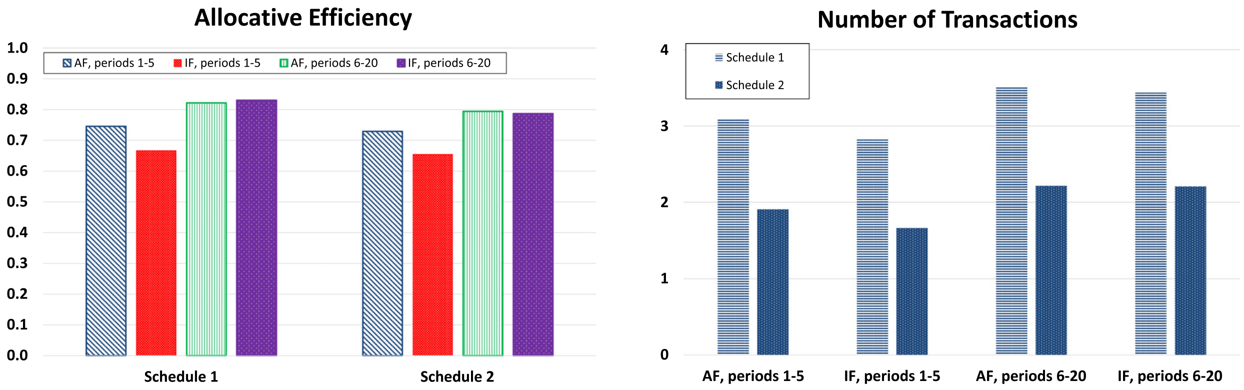

Fig. 4. Comparison of the allocative efficiency across information treatments for the two schedules separately (left) and comparison of the number of transactions across information treatments for the two schedules (right).

the other hand, we found that the effect of information feedback (**AF** vs. **IF**) is in the same direction for both locations.[25] Thus, we pool data along location when we present them visually and make the pairwise tests for the mean.

4.1. Allocative efficiency

We begin with an illustration in Fig. 3 (right) of one of the key results of the experiment. This figure compares the allocative efficiency across two information treatments, **AF** and **IF**, for the first five and remaining 15 periods of the experiment. We make three observations from this figure. First, in the initial periods, allocative efficiency is larger for the **AF** than for the **IF** treatment. Thus, similar to the call markets analyzed in AL, an increase in an amount of information available to the traders results in a lower allocative efficiency. Second, in the remaining periods of the experiment, allocative efficiencies under the two information treatments are very close to each other. Thus under the CDA, the negative effect of more information is only temporary. The 'Allocative efficiency' part of Table 2 confirms that the difference between treatments is significant in the initial periods but not in the remaining periods. Finally, we notice that allocative efficiency reaches 80%. On the one hand, this is lower than that achieved in the IEL simulations in the AALP, where average efficiency is above 88% in both schedules, for most of the parameters. On the other hand, 80% is much larger than the efficiencies we would get under the random allocation process (that adjusts for the difficulty of the schedule) described in footnote 18.[26]

In Fig. 3, the data are pooled across schedules. To justify that this is appropriate, we show the allocative efficiency disaggregated across schedules in Fig. 4 (left). (The data behind this figure are also in the 'Allocative efficiency' part of Table 2). We can see that the effect of the information treatments is the same for both schedules for the first five periods of the experiment and for the remaining periods.[27]

To find the reason for low allocative efficiency, we investigate in Fig. 4 (right) the average number of transactions. For schedule **S1**, the equilibrium number of transactions is 4. We can see that, with time, the number of transactions increased, but it is never close to the equilibrium 4 transactions. For schedule **S2**, the equilibrium number of transactions is 3. Again, the actual average number increases in time but it is smaller than 3. The 'Missed transactions' part of Table 2 reports the statistics for the difference between the equilibrium and experimental number of transactions. Testing across treatments, we find that the difference is significant but only for the initial periods, consistently with the allocative efficiency results. From Fig. 4 (right) and Table 2, we can see that the gap between the equilibrium and experimental number of transactions is larger for the **S2**. Apparently, this contributes to a smaller allocative efficiency for schedule **S2**.

The regression analysis in Appendix C confirms and expands these findings. We find that the allocative efficiency is significantly lower in the initial periods and significantly lower for schedule **S2**. Both these effects are also significant in the

[24] We find the same effect using the coefficient estimate on the DLoc dummy in the regressions reported in Appendix C. At the 5% level of significance, the location effect is significant for the allocative efficiency (that is higher in the Caltech sessions), and the out-of-equilibrium index and volatility (both are lower in the Caltech sessions).

[25] Our findings are similar to those in Snowberg and Yariv (2021) who compare performances of different subject pools in a number of elicitation tasks. They find that the Caltech students' behavior is closer to rational and that the direction of comparative statics between treatments is consistent across locations.

[26] The allocative efficiency in periods 6–20 of the experiment, in both treatments and for both schedules, is larger than the allocative efficiency for 97% of simulations for the random allocation. This can be compared with similarly computed numbers of 93% and 97% in the **AF** treatment, and 85% and 82% in the **IF** treatment, for schedules 1 and 2, correspondingly. This suggests that, once learning is done after first five periods, in both treatments the efficiency is pretty high. It was also high in the first 5 periods for the **AF** treatment, which is significantly better than for the **IF** treatment.

[27] For periods 1–5, the difference in the allocative efficiency between the **AF-S1** and **IF-S1** is significant with p-value 0.023; this difference between the **AF-S2** and **IF-S2** is marginally significant with p-value 0.120. For periods 6–20 the differences are not significant in both cases with p-values 0.622 and 0.414, respectively.

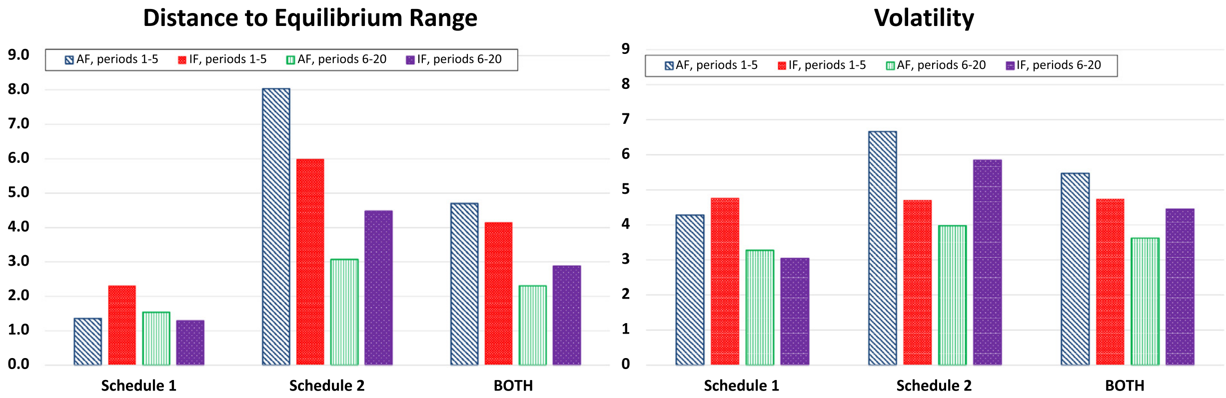

Fig. 5. Comparison of the distance to the range of equilibrium prices (*left*) and price volatility (*right*) across information treatments for the two schedules separately.

regression for the missed number of transactions. Interestingly, the effect of the schedule on the allocative efficiency is in the same direction as expected from simulations with random agents (cf. footnote [17]).

We summarize the above discussion in the next result supporting Hypothesis 1.

Result 1. Over all sessions, allocative efficiency is about 77 – 80%. Allocative efficiency increases over time, indicating learning. Allocative efficiency depends on the schedule and is significantly higher under the **AF** than under the **IF** but only for the initial stage of 5 periods, after which there is no significant difference between the two information treatments.

4.2. Information efficiency

Measures of "information efficiency" reflect how market is able to aggregate individual information into the price. In our repeated setting, where no new fundamental information is revealed between periods, information efficiency is largely about the ability of the mechanism to keep the price within the equilibrium range. Still, information efficiency is interesting to study, in light of the empirical evidence discussed in Section 2.1. Information efficiency is strongly affected by the individual bidding strategies, and it is also related, but not identical to allocative efficiency.[28]

We analyze two measures of information efficiency: an indicator of whether the average transaction price during the session falls outside of the equilibrium range (in which case it is set to 1 for the session, and otherwise it is 0), and, the distance of the average transaction price to the equilibrium range (which is 0 when the average transaction price belongs to the equilibrium price range). Note that the latter measure is more refined than the former. The results are very similar and we illustrate them only for this latter, more refined measure in the left panel of Fig. 5. What stands out from this figure is a very strong effect of the schedule for this measure. The distance is especially high for schedule **S2**, and especially for the initial five periods. This is also confirmed in the regression analysis in Appendix C. The effect of the information treatment is mixed, and because of the interaction between these two effects, we find no significant difference between the **AF** and **IF** treatments, when we pool observations over the schedules in Table 2.

We also look at the variability of the average price over time, i.e., price volatility. We measure that, in a given session, during the corresponding time interval. The right panel of Fig. 5 compares the volatility across treatments, see also Table 2. The directions of the effects of schedule, information, and time periods are the same for the price volatility as for the distance to the equilibrium range. The effect of schedule seems to dominate the other effects with no significant effect of information treatment.

Taken together, our results indicate very low information efficiency in the experiment. For example, the indicator in Table 2 shows that the average price of transactions is very often outside of the equilibrium range. For schedule **S2**, this occurs for almost 89% of the initial periods for **AF-S2**, dropping to 75% for the remaining periods, and for 76% of the initial periods for **IF-S2**, dropping to 66% for the remaining periods. The index is lower for schedule **S1**, where the equilibrium price range is larger, but it is still around 40% of the cases over the experiment.

By and large, the experiments do not support Hypothesis 2, suggesting that the IEL generates much less volatile simulations in the stationary state than those that we observe in the experiment. We summarize all of this in the next result.

Result 2. Information efficiency, as measured by the fraction of periods when the average transaction price is outside the equilibrium range, the distance to the range, or price volatility is small in the experiment and is especially small for schedule **S2**. There is no significant difference in price volatility across information treatments.

[28] Indeed, when an extra marginal trader transacts and lowers allocative efficiency, the *average* price of transactions in the trading session may fall within the equilibrium range with no impact on information efficiency. On the other hand, infra marginal traders can transact at prices outside of the range, lowering information but not allocative efficiency.

5. IEL and experimental data

To verify that the IEL algorithm is able to replicate the key data patterns observed in the experiment, we simulated the CDA-SD environment using the IEL specification of AALP, described in Section 2.3, for each of the four treatments, **AF-S1, AF-S2, IF-S1** and **IF-S2**. We conducted 10,000 simulations with different random seeds. In contrast to AALP who looked at the stationary IEL dynamics over long time and after 100 transitory periods, in the simulations presented below, we closely follow the experimental setting and focus on the first 20 periods of the IEL simulations, without any burn-in periods. More specifically, we separately consider periods 1–5 and periods 6–20. We do this to reflect learning in the experiment, evaluate an ability of the IEL to replicate this learning phase, and see the effect of the IEL estimated parameters when it is fitted to the short-term data.

We optimized the key parameters of the IEL, that is, the size of the pool, J, the experimentation probability, ρ, and the assumption about the distribution that is used at the experimentation stage. Specifically, we minimize the Mahalanobis distance between the experimental data and simulations, defining this distance jointly for all five statistics listed in Table 2 and for all treatments and for both blocks of periods.[29] In terms of distribution used at the experimentation stage, we found that the *local experimentation* under which an element from the last pool is replaced by a random variable drawn from the normal distribution, whose mean is equal to the removed element, truncated at the admissible interval of the trader, fits data better than the uniform distribution. Therefore, we also optimized the standard deviation σ of this normal distribution representing local experimentation. A grid-search over a wide range of parameter values led to the following optimized values: $J = 10$, $\rho = 0.1$ and $\sigma = 5$.

When performing the search for parameters, we discovered that the allocation efficiency is mainly driven by the combination of J and ρ. Using a larger experimentation pool and lower experimentation probability results in much higher allocative efficiencies than observed in the experimental data. The information efficiency measures such as price volatility are, instead, more sensitive to the distribution used at the experimentation stage of the IEL. Specifying experimentation from the uniform distribution on all range of admissible values, that was used in AALP, resulted in much lower information efficiency, and higher price volatility, than we observed in the experimental data.

Table 3 reports the means of the **AF-IEL** and **IF-IEL** simulations including a new modified AF-IEL algorithm, **AF-IEL mod**, which we explain later in this section. The three columns with "Means" (for schedule 1, schedule 2, and pooled data over the schedules) can be directly compared with the corresponding statistics for the experimental data reported in Table 2. In the last three columns of Table 3 we report the p-values of the two-sided t-test for the null hypothesis that the IEL-based means are the same as the experiment-based means. A sufficiently high p-value indicates a good match with the experimental data. Comparing Table 3 with Table 2, we observe that the optimized IEL matches the experimental data relatively well in terms of the allocative efficiency but less so for the information efficiency measures, even if they are in reasonable ranges.

We also notice that the **AF-IEL** algorithm strongly overpasses the allocative efficiency in periods 6–20 for both schedules. That may happen for the following reason. IEL, as described in Section 2.3, was initially designed in AL for the agents participating in a *call* market. In the call market, bids and offers are cleared simultaneously and everyone who trades does so at the same price. In the **AF** treatments, where the feedback is limited to the past price, it is then reasonable to assume that a trader's own impact on the clearing price is negligible and that the best strategy to maximize the probability of transaction is to submit a bid at or just below the trader's own valuation or an offer at or just above the trader's own cost. In the CDA, however, orders are cleared at different prices, that is, at the bid or offer prices that are in the book at the time of the match. In contrast to the call market, it often happens that a buyer will pay and a seller will receive exactly what they bid or offer. This provides traders incentives to reduce a bid or increase an offer to realize a higher profit. This "hawkish" (and more price-making) bidding behavior may, in turn, reduce the number of transactions and lower the allocative efficiency.

This observation from our experiments suggests that, when adjusting the IEL algorithm to the CDA, and designing the hypothetical utility functions for the **AF** treatments, AALP did not fully reflect the conditioning of traders on their own bids or offers. Indeed, Eq. (1) for buyers' hypothetical utility assumes that as long as a bid is above the average past price, the order will be cleared at the average past price, regardless of the exact bid; and similarly for sellers and their hypothetical utility (2) for sellers. Under this setting, the IEL produces volatile price and pushes individual strategies towards the valuations/cost of traders, as was shown in AALP. We did not observe this in the experiment. That suggests that the behavioral assumption of the hypothetical utility in the IEL algorithm calls for a change.

5.1. IEL for the aggregate feedback information treatment

Armed with these observations, we now introduce the **AF-IEL mod** algorithm, which is the modified IEL algorithm for the **AF** treatment.

[29] The Mahalanobis distance is a generalization of the Euclidean distance for heteroskedastic data. It is defined as $\sqrt{(x-\mu)'S^{-1}(x-\mu)}$, where x and μ are column vectors and S is the covariance matrix. Since we conduct 10,000 simulations, the standard errors of the IEL-based means are negligible relative to the standard errors of the experimental means. Therefore, in computing the Mahalanobis distance as well as subsequent t-tests, IEL-based means were treated as population parameters, μ.

Table 3
Means of different characteristics for the IEL-simulated data for the best fitted parameters over 10,000 simulations. The last three columns report the *p*-value of the *t*-test for the null hypothesis that the means for the IEL simulated data are the same as the means of the experimental data for each corresponding schedule and period. Combining all 5 measures and 2 phases (periods 1–5 and 6–20), this null hypothesis cannot be rejected at the 5% significance level in 4 out of 10 instances for **AF-IEL**, 6 out of 10 for **AF-IEL mod**, and 6 out of 10 for **IF-IEL**.

		Means			*p*-values (Ho: experim=IEL)		
		Sched 1	Sched 2	Pooled	Sched 1	Sched 2	Pooled
Allocative Efficiency							
periods 1–5	**AF-IEL**	0.750	0.684	0.717	0.825	0.213	0.921
	AF-IEL mod	0.732	0.669	0.700	0.532	0.095	0.856
	IF-IEL	0.712	0.650	0.681	0.124	0.892	0.674
periods 6–20	**AF-IEL**	0.907	0.870	0.888	0.000	0.007	0.000
	AF-IEL mod	0.846	0.806	0.826	0.082	0.662	0.254
	IF-IEL	0.806	0.820	0.813	0.262	0.138	0.741
Number of “missed” transactions							
periods 1–5	**AF-IEL**	0.884	1.222	1.053	0.810	0.232	0.508
	AF-IEL mod	1.004	1.330	1.167	0.401	0.030	0.037
	IF-IEL	1.080	1.250	1.165	0.280	0.349	0.153
periods 6–20	**AF-IEL**	0.109	0.553	0.331	0.000	0.016	0.000
	AF-IEL mod	0.582	1.032	0.807	0.135	0.011	0.011
	IF-IEL	0.722	0.964	0.843	0.053	0.035	0.008
Indicator of whether the average price is outside of the equilibrium range							
periods 1–5	**AF-IEL**	0.616	0.778	0.697	0.012	0.020	0.380
	AF-IEL mod	0.566	0.752	0.659	0.044	0.004	0.664
	IF-IEL	0.570	0.778	0.674	0.356	0.814	0.422
periods 6–20	**AF-IEL**	0.656	0.855	0.756	0.020	0.100	0.015
	AF-IEL mod	0.475	0.747	0.611	0.602	0.982	0.715
	IF-IEL	0.497	0.772	0.634	0.279	0.191	0.132
Distance from the average price to the equilibrium range							
periods 1–5	**AF-IEL**	4.459	9.018	6.738	0.000	0.486	0.059
	AF-IEL mod	3.747	7.958	5.852	0.000	0.955	0.285
	IF-IEL	3.681	8.024	5.853	0.084	0.072	0.034
periods 6–20	**AF-IEL**	4.516	8.058	6.287	0.000	0.000	0.000
	AF-IEL mod	2.089	4.985	3.537	0.339	0.014	0.015
	IF-IEL	2.371	5.314	3.842	0.086	0.390	0.160
Price volatility							
periods 1–5	**AF-IEL**	9.580	12.357	10.969	0.000	0.000	0.000
	AF-IEL mod	6.970	9.133	8.051	0.002	0.019	0.000
	IF-IEL	8.106	11.399	9.752	0.000	0.000	0.000
periods 6–20	**AF-IEL**	10.025	11.631	10.828	0.000	0.000	0.000
	AF-IEL mod	3.616	5.201	4.409	0.571	0.060	0.075
	IF-IEL	4.833	6.905	5.869	0.017	0.288	0.040

We made two modifications. First, we take into account the most recent experience of a trader. Note that in the experiment subjects could see their profits from the last trading session. This allows them to deduce the price of their transaction. A buyer who traded at a price that was lower than the average price optimistically assumes, when deriving hypothetical, that all bids equal to or above their last transaction price will result in a trade. On the other hand, if their transaction price was higher than the average price or if the buyer did not trade, they expect that successful bids should be equal to or above the last average price. Thus, buyer b sets the *reference* price as

$$P^{\text{ref}}_{b,t} = \begin{cases} \min\left(p_{b,t-1}, \bar{p}_{t-1}\right) & \text{if buyer } b \text{ traded at } t-1 \text{ at price } p_{b,t-1} \\ 0 & \text{otherwise}. \end{cases}$$

Analogously, seller s sets the reference price

$$P^{\text{ref}}_{s,t} = \begin{cases} \max\left(p_{s,t-1}, \bar{p}_{t-1}\right) & \text{if seller } s \text{ traded at } t-1 \text{ at price } p_{s,t-1} \\ 0 & \text{otherwise}, \end{cases}$$

as the price below which the transaction is expected to occur.

Second, we reflect on the fact that a transaction may occur either at the bid/offer submitted by the trader or at the price of the counter-party. We, therefore, modify the hypothetical utilities for buyers and sellers as follows

$$U^A_{b,t}(b_i) = \begin{cases} w(V_b - P^{\text{ref}}_{b,t}) + (1-w)(V_b - b_i) & \text{if } b_i \geq P^{\text{ref}}_{b,t} \\ 0 & \text{otherwise} \end{cases} \tag{4}$$

$$U^A_{s,t}(a_i) = \begin{cases} w(P^{\text{ref}}_{s,t} - C_s) + (1-w)(a_i - C_s) & \text{if } a_i \leq P^{ref}_{s,t} \\ 0 & \text{otherwise,} \end{cases} \tag{5}$$

where $w \in [0, 1]$ is a weighting parameter that determines the probability of the trade being executed at the submitted bid/offer vs. at the counter-party price (proxied by the past average price). The rest of the algorithm is identical to the IEL in Section 2.3.

We re-optimized the parameters of this **AF-IEL mod** algorithm, including the new parameter w. The optimal parameters turned out to be as follows: pool size $J = 10$, experimentation probability $\rho = 0.1$, the standard deviation of the truncated normal (local) experimentation is $\sigma = 5$ (same as before) and weight on the own bid/offer $w = 0.5$. Since the order of arrival in the experiment is random, it is reasonable to see the optimal value $w = 0.5$ indicating an equal chance of the order being executed at the submitted bid/offer or the counterparty order.

Returning back to the IEL-based statistics in Table 3 and comparing these to experimental data in Table 2, we observe that the **AF-IEL mod** algorithm matches the experimental data on the allocative efficiency and most of the other statistics better than the initial **AF-IEL** algorithm.

6. Summary

In this paper, we study how information about the previous trading session impacts the allocative and informational efficiency of a market organized as a continuous double auction (CDA). In this CDA, buyers and sellers submit orders simultaneously, as in a call market, but the market clears as in the standard CDA with orders arriving sequentially. This corresponds to the retail financial market. Guided by previous contributions in the literature, we formulate two hypotheses and conduct an experimental study to verify them.

In the analysis, we split the data into early periods (1–5) and later periods (6–20). We find that efficiency increases over time - there is learning by the participants. We also find that in the early periods, more detailed information about the last period's prices and transactions actually reduces efficiency. That difference disappears in the later periods when efficiency under both information treatments becomes comparable and relatively high. With respect to informational efficiency, we find it to be low and not significantly different between both information treatments.

The experimental data reject the earlier IEL model of AALP (adapted from the AL model for a call market), especially in the information condition where only aggregate feedback (AF) about the average transaction price of the previous period is available. The IEL model predicts allocative efficiencies that are too high and informational efficiencies that are too low. We propose an improvement to the way the AF treatment is modeled in IEL to better fit the details of the CDA market and then show that the modified IEL matches the experimental data relatively well.

Appendix A. Experimental instructions

A1. Common part

Introduction You are about to participate in an economic experiment where your earnings depend on your decisions and decisions of others. You will be paid in cash at the end of the experiment. Please read the following instructions carefully. You may make annotations or write with a pen, if it helps. You may also take notes or make calculations during the experiment. At the end of the instructions you will be asked a number of questions to check your understanding. If you have any questions while reading the instructions or during the experiment, please raise your hand and the experimenter will come to you.

This is an experiment in trading. You are either a buyer or a seller. All other participants are also buyers or sellers with the total amount of 5 buyers and 5 sellers. Depending on your role you will be buying or selling an item through an online trading market, the details of which will be provided below. Trading will occur in *rounds*. There are 40 rounds in total; your role (buyer or seller) will change after the first 20 rounds. Total earnings for all rounds will be paid in the end of the experiment. All monetary values and prices in the experiment are given in the Experimental Currency Units (ECU). They will be converted to US dollars (US) at the rate of US 0.03 (or 3 cents) per one ECU. A show-up fee of US 10 will be added to your earnings at the end of the experiment. All payments will be rounded to US 1.

Earnings of buyers and sellers

A *buyer* may buy one item in each round. Each buyer will be given a valuation for the item. You will not know the valuations of other participants, which may be different. When you are a buyer, your valuation will remain the same for all

20 rounds. If you successfully buy one item in that round, you will earn your valuation minus the price of the item. This price will depend on your decision and decisions of others. As a buyer you will submit a buy order with a *bid price*, that is, the *maximum* price you wish to pay for the item. This cannot exceed your valuation. You will find more details about how to submit a bid at the end of these instructions. The price that you actually pay for an item (the transaction price) will be computed by the market. This is described below. If you did not trade in a round, you will receive 0 ECU.

$$\text{Buyer's earnings in the round} = \begin{cases} \text{valuation} - \text{transaction price}, & \text{if your order transacted} \\ 0, & \text{if your order did not transact} \end{cases}$$

A *seller* may sell one item in each round. Each seller will be given a cost for the item. You will not know the costs of other participants, which may be different. When you are a seller, your cost will remain the same for all 20 rounds. If you successfully sell one item in that round, you will earn the price of the item minus your cost. As a seller you will submit a sell order with an *ask price*, that is, the *minimum* price you wish to receive for the item. This cannot be below your cost. The price you actually receive for an item (the transaction price) will be computed by the market. This is described below. If you did not trade in the round, you will receive 0 ECU.

$$\text{Seller's earnings in the round} = \begin{cases} \text{transaction price} - \text{cost}, & \text{if your order transacted} \\ 0, & \text{if your order did not transact} \end{cases}$$

Market matching

The market will determine whether transactions occur and compute the prices of the transactions. This will be done by matching the orders from buyers and sellers, including your own order. You can submit your order only once per round and you cannot change it during the round. All orders will be collected before the matching process starts. You will not know the bids and asks of the other participants before the end of the round.

The occurrence and price of a transaction will depend not only on the submitted orders, but also on a *sequence of their arrival* to the market. The sequence of arrival is *random* in this experiment and you cannot influence it. It changes from round to round. The sequence is unknown to you and to other participants. It does not depend on the time when you submit your order.

The market uses the following rules to match the orders of buyers and sellers as they arrive to the market. If no match is possible an order (bid or ask price) is stored in an *order book*. Therefore, the first order to arrive is always stored in the book. There are two separate sides in the book: buy (for bids) and sell (for asks). A transaction occurs if a newly arrived order is matched with one of the opposite type orders stored earlier in the book. The *transaction price* is equal to the price of the order which was stored in the book and satisfied the match. This order is removed from the book after the transaction occurs. For example, if a new order with bid P_B comes from buyer B, and there is an old order with ask P_S in the book submitted by seller S, such that $P_B \geq P_S$, that is, the bid is higher than the ask price, there is a match. Transaction between B and S occurs at the price P_S. After the transaction, the orders from buyer B and seller S are removed from the book.

If there are many orders in the book which can satisfy a newly arrived order of the opposite type, the matching happens with the best priced order available in the book from the perspective of the newly arrived order (the highest bid in the book or the lowest ask in the book). The matching process continues until all orders arrive to the market. The orders for which no match is found will not transact in the current round. All unmatched orders will be removed from the book at the end of the round.

Here is a detailed example of one round with 2 buyers, B1 and B2, and 2 sellers, S1 and S2. Notice that all numbers in this example are for illustration only and differ from the experiment. Suppose that the B1's valuation is $V_{B1} = 180$ and the B2's valuation is $V_{B2} = 150$, and that the S1's cost is $C_{S1} = 30$ and the S2's cost is $C_{S2} = 10$. Recall that all buyers and sellers know only their own valuations or costs. Before the matching starts buyer B1 submits bid $P_{B1} = 100$ and buyer B2 submits bid $P_{B2} = 50$. Also seller S1 submits ask $P_{S1} = 35$ and seller S2 submits ask $P_{S2} = 60$. After these orders are submitted the computer randomly chooses the sequence of their arrival as follows: B2, B1, S1, and S2.

Thus, the B2's buy order with bid $P_{B2} = 50$ is the first to arrive to the market. Since there is no available sell order, no transaction is possible and the buy order is stored in the order book (on the buy side). The next order from B1 is also a buy order with bid $P_{B1} = 100$. Because again there is no sell order available, bid $P_{B1} = 100$ will be saved in the book (on the buy side). However, it will receive a higher priority in the book, because $P_{B1} > P_{B2}$, that is, buyer B1 is ready to pay a higher price. Next, the S1's sell order with ask $P_{S1} = 35$ arrives. It can be matched with the existing buy orders. Namely, it is matched with the highest price order in the buy side of the book, that it, $P_{B1} = 100$. Seller S1 sells an item to buyer B1. The transaction price is $P_T = P_{B1} = 100$, which is the price of the order that arrived earlier and was stored in the book. The transacted buy order is removed from the book, so that the book now contains only one buy order with bid $P_{B2} = 50$. Finally, the S2's sell order with ask $P_{S2} = 60$ arrives. Because the buy side of the book contains only one buy order with $P_{B2} = 50$ and $P_{B2} < P_{S2}$, the match does not occur and the sell order is saved in the order book. Since in this example no new orders arrive, the book is cleared at the end of the round, and no transaction between buyer B2 and seller S2 occurs.

The following table displays graphically this example. The table should be read from left to right, which corresponds to the development of this round in time.

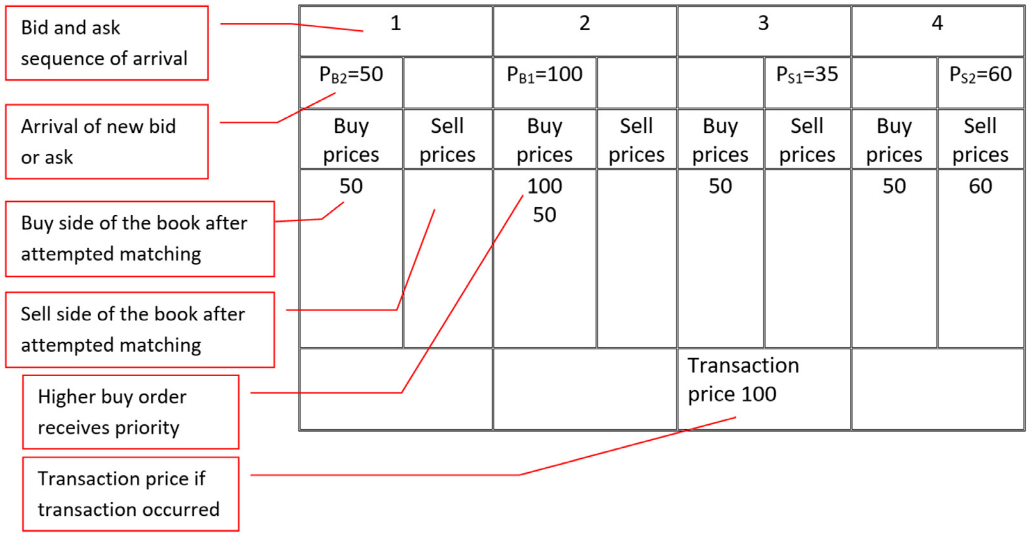

A2. Information part

Individual feedback treatment

Screen information and order submission

In order to help you to decide on your bid or ask in any round, the following information will be provided (*after* the first round):

1. The sequence of arrival of your order (called step) in previous round
2. Step in which your order transacted and transaction price if you order transacted
3. The above table (book) for the previous trading round for all 10 arriving orders (without IDs).
4. If your order did not transact when it arrived, you will see it in *blue color* in the book.
5. Your earnings from the previous round
6. Your cumulative earnings for all previous rounds.

For 20 seconds, only this information about the previous round will be on the screen. This is a good time to *analyze* the information. You can use pen and paper and take records or make calculations.

After the first 20 seconds, a *decision* window will appear. You will have an additional 1 minute for the analysis and decision. The decision window will show if you are buyer or seller, your valuation/cost and allowed range for your bid or ask:

- for bid price min is 0 and max is your valuation;
- for ask price min is your cost and max is 105.

You should *type* your bid or ask price with up to two decimal digits (for example 10.21 or 12.3 or 11). If you make a mistake, you can use Backspace key or arrows and Del to edit your order. Once you are satisfied, press the "Submit Order" button. After this your decision cannot be changed in this round.

Note: the chances of your order to transact and the transaction price of your order may depend on your valuation or cost; your valuation or costs will remain unchanged for the first 20 rounds; after your buyer/seller role changes at round 21 you will receive a new valuation or cost and this will remain the same for the last 20 rounds.

Any questions? Now or any time during the experiment raise your hand.

Scratch space (you may write here):

Aggregate feedback treatment

Screen information and order submission

In order to help you to decide on your bid or ask in any round, the following information will be provided (*after* the first round):

1. An average price of all transactions occurred in the previous round. For example, if 3 transactions occurred $\text{Ave}(P_T) = (P_{T1} + P_{T2} + P_{T3})/3$.

2. Your earnings from the previous round.
3. Your cumulative earnings for all previous rounds.

For 20 seconds, only this information about the previous round will be on the screen. This is a good time to *analyze* the information. You can use pen and paper and take records or make calculations.

After the first 20 seconds, a *decision* window will appear. You will have an additional 1 minute for the analysis and decision. The decision window will show if you are buyer or seller, your valuation/cost and allowed range for your bid or ask:

- for bid price min is 0 and max is your valuation;
- for ask price min is your cost and max is 105.

You should *type* your bid or ask price with up to two decimal digits (for example 10.21 or 12.3 or 11). If you make a mistake, you can use Backspace key or arrows and Del to edit your order. Once you are satisfied, press the "Submit Order" button. After this your decision cannot be changed in this round.

Note: the chances of your order to transact and the transaction price of your order may depend on your valuation or cost; your valuation or costs will remain unchanged for the first 20 rounds; after your buyer/seller role changes at round 21 you will receive a new valuation or cost and this will remain the same for the last 20 rounds.

Any questions? Now or any time during the experiment raise your hand.

Scratch space (you may write here):

Appendix B. Screenshots

Individual feedback treatment

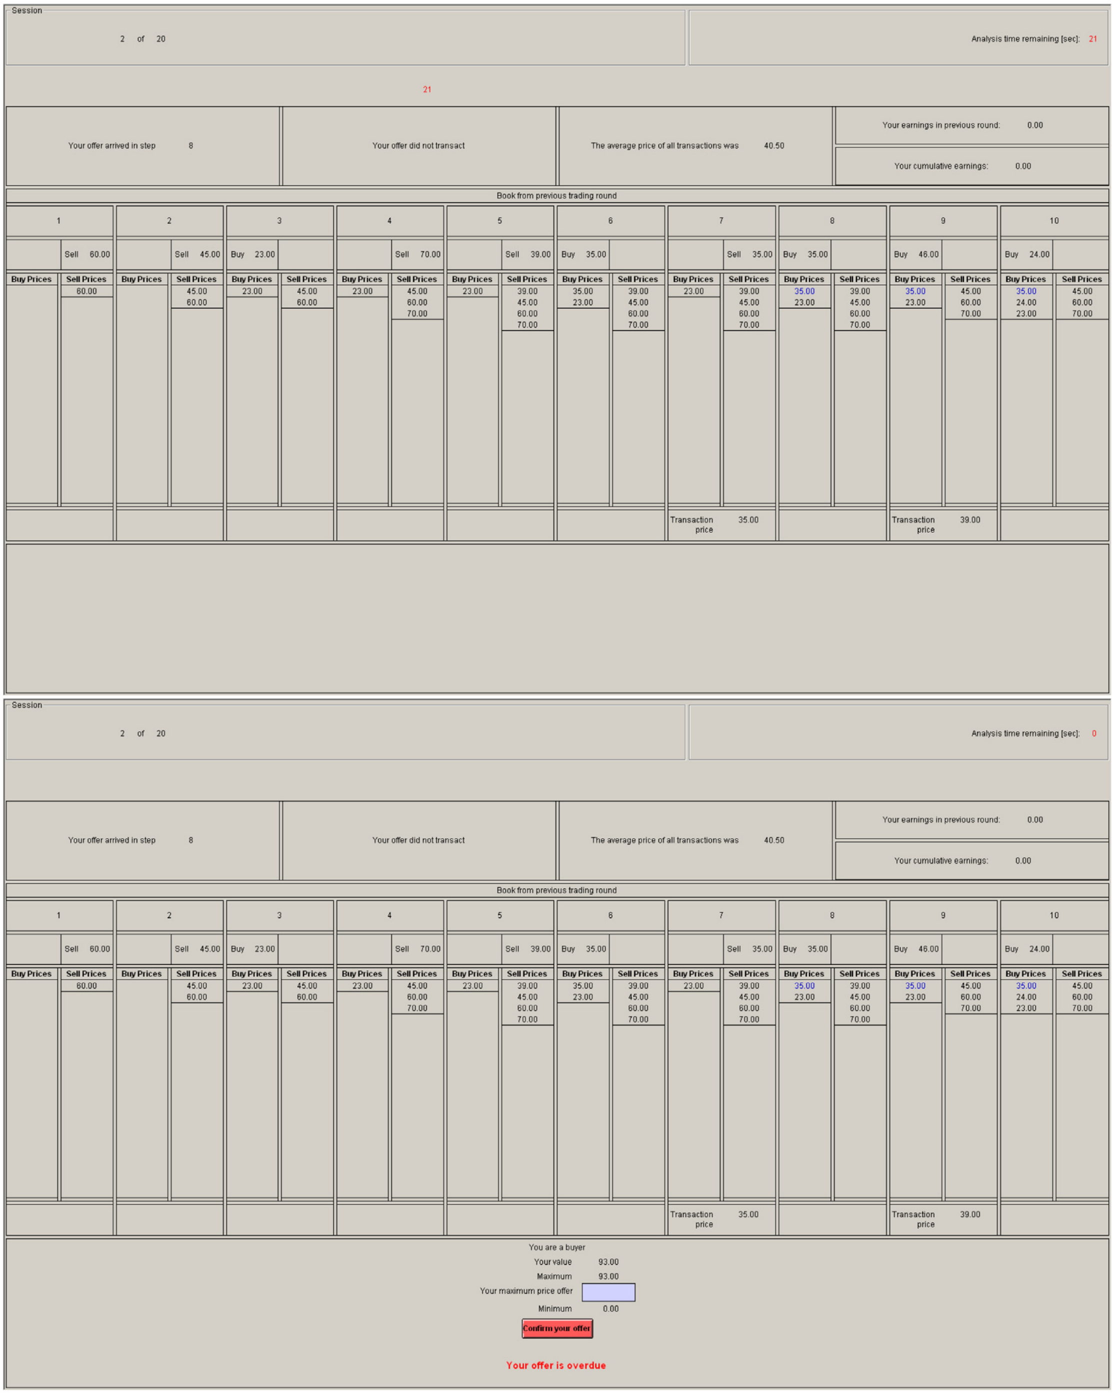

Fig. 6. Analysis screen (upper part) and analysis and decision screen (lower part) for the Individual feedback (**IF**) sessions.

Aggregate feedback treatments

Fig. 7. Analysis screen (upper part) and analysis and decision screen (lower part) for the Aggregate feedback (**AF**) sessions.

Table 4
Regression of allocative efficiency.

Allocative efficiency	Estimate	NW-robust st errors			Cluster-robust st errors		
		SE	*t*-stat	*p*-value	SE	*t*-stat	*p*-value
(intercept)	0.769	0.019	40.245	**0.000**	0.020	39.403	**0.000**
DInfo	0.000	0.018	−0.009	0.993	0.021	−0.008	0.994
DInit	−0.148	0.031	−4.811	**0.000**	0.038	−3.851	**0.000**
DInfo x DInit	0.077	0.043	1.816	0.069	0.042	1.823	0.068
DLoc	0.060	0.017	3.592	**0.000**	0.020	3.072	**0.002**
DSch	0.035	0.015	2.360	**0.018**	0.013	2.657	**0.008**
DExp	−0.023	0.015	−1.539	0.124	0.013	−1.734	0.083

Table 5
Regression of the number of missed transactions.

Missed transactions	Estimate	NW-robust st errors			Cluster-robust st errors		
		SE	*t*-stat	*p*-value	SE	*t*-stat	*p*-value
(intercept)	0.841	0.060	14.053	**0.000**	0.097	8.649	**0.000**
DInfo	−0.041	0.063	−0.648	0.517	0.094	−0.434	0.664
DInit	0.578	0.092	6.270	**0.000**	0.079	7.315	**0.000**
DInfo x DInit	−0.215	0.134	−1.603	0.109	0.096	−2.249	**0.025**
DLoc	−0.073	0.055	−1.336	0.182	0.091	−0.807	0.420
DSch	−0.237	0.052	−4.541	**0.000**	0.058	−4.108	**0.000**
DExp	−0.008	0.052	−0.162	0.872	0.058	−0.146	0.884

Table 6
Regression of the out-of-equilibrium index.

Out-of-equilibrium index	Estimate	NW-robust st errors			Cluster-robust st errors		
		SE	*t*-stat	*p*-value	SE	*t*-stat	*p*-value
(intercept)	0.773	0.045	17.262	**0.000**	0.085	9.120	**0.000**
DInfo	0.070	0.044	1.588	0.112	0.084	0.840	0.401
DInit	0.096	0.060	1.597	0.110	0.055	1.767	0.077
DInfo x DInit	−0.059	0.076	−0.783	0.434	0.069	−0.858	0.391
DLoc	−0.167	0.037	−4.525	**0.000**	0.077	−2.180	**0.029**
DSch	−0.334	0.041	−8.166	**0.000**	0.092	−3.641	**0.000**
DExp	0.002	0.041	0.048	0.962	0.092	0.021	0.983

Appendix C. Results of regression analysis

As explained at the beginning of Section 4, we regress different characteristics (allocative efficiency, measures of information efficiency, price volatility) to the following set of dummy variables:

- DInfo: set to 1 for **AF** and to 0 for **IF**;
- DInit: set to 1 for periods 1–5 and to 0 for periods 6–20;
- DInfo × DInit: measures the mixed effect of information and initial periods;
- DLoc: set to 1 for Caltech and to 0 for UNSW;
- DSch: set to 1 for schedule **S1** and to 0 for schedule **S2**;
- DExp: set to 1 for the first block and to 0 for the second block.

This regression allows us to control the results for differences between the two information treatments, **AF** and **IF**, for the location of the experiment, the schedules, phases, and so on.

Table 4 reports the results for the regression as presented in Eq. (3) of the main text. The second column contains parameter estimates, so that we have the following result

$$\begin{aligned}\text{Allocative Efficiency} = \mathbf{0.769} + 0.000\,\text{DInfo} - \mathbf{0.148}\,\text{DInit} + 0.077\,(\text{DInfo} \times \text{DInit}) + \mathbf{0.060}\,\text{DLoc} \\ + \mathbf{0.035}\,\text{DSch} - 0.023\,\text{DExp} + \text{error}.\end{aligned}$$

In the next columns of Table 4, we report standard errors, *t*-statistics and *p*-values. We use two methods (the “cluster-robust” and “Newey-West-robust”) of computing the standard errors. The bold *p*-values (and coefficients in the regression above) indicate the instances of significant effect at the 5% level.

We estimate the same regression for other characteristics, see Table 5 for the number of missed transactions (that is the number of trades in equilibrium minus the number of trades in the experiment), Table 6 for the indicator of when the average transaction price is outside of the equilibrium range, and Table 7 for the price volatility.

Table 7
Regression of the volatility.

Volatility	Estimate	NW-robust st errors			Cluster-robust st errors		
		SE	*t*-stat	*p*-value	SE	*t*-stat	*p*-value
(intercept)	6.429	0.853	7.538	**0.000**	0.936	6.869	**0.000**
DInfo	−0.821	0.739	−1.111	0.266	0.774	−1.062	0.288
DInit	0.288	0.927	0.311	0.756	0.810	0.356	0.722
DInfo x DInit	1.558	1.118	1.394	0.163	0.919	1.695	0.090
DLoc	−2.227	0.599	−3.718	**0.000**	0.685	−3.251	**0.001**
DSch	−1.449	0.663	−2.185	**0.029**	0.693	−2.091	**0.037**
DExp	−0.042	0.663	−0.063	0.949	0.693	−0.061	0.952

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.jedc.2022.104387.

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, C., Makarewicz, T., 2019. Simple forecasting heuristics that make us smart: evidence from different market experiments. J. Eur. Econ. Assoc. 17 (5), 1538–1584.
Arifovic, J., 1994. Genetic algorithm learning and the cobweb model. J. Econ. Dyn. Control 18 (1), 3–28.
Arifovic, J., Boitnott, J.F., Duffy, J., 2019. Learning correlated equilibria: an evolutionary approach. J. Econ. Behav. Organ. 157, 171–190.
Arifovic, J., Ledyard, J., 2003. Computer testbeds and mechanism design: application to the class of Groves–Ledyard mechanisms for the provision of public goods. manuscript.
Arifovic, J., Ledyard, J., 2004. Scaling up learning models in public good games. J. Public Econ. Theory 6 (2), 203–238.
Arifovic, J., Ledyard, J., 2007. Call market book information and efficiency. J. Econ. Dyn. Control 31 (6), 1971–2000.
Arifovic, J., Ledyard, J., 2011. A behavioral model for mechanism design: individual evolutionary learning. J. Econ. Behav. Organ. 78 (3), 374–395.
Arifovic, J., Ledyard, J., 2018. Learning to alternate. Exp. Econ. 21 (3), 692–721.
Baruch, S., 2005. Who benefits from an open limit-order book? J. Bus. 78 (4), 1267–1306.
Boehmer, E., Saar, G., Yu, L., 2005. Lifting the veil: an analysis of pre-trade transparency at the NYSE. J. Finance 60 (2), 783–815.
Bottazzi, G., Dosi, G., Rebesco, I., 2005. Institutional architectures and behavioral ecologies in the dynamics of financial markets: a preliminary investigation. J. Math. Econ. 41, 197–228.
, 2018. In: Chen, S.-H., Kaboudan, M., Du, Y.-R. (Eds.), The Oxford Handbook of Computational Economics and Finance. Oxford University Press.
Dawid, H., 1999. On the convergence of genetic learning in a double auction market. J. Econ. Dyn. Control 23, 1545–1567.
Dieci, R., He, X.-Z., 2018. Heterogeneous Agent Models in Finance. In: Handbook of Computational Economics, vol. 4. Elsevier, pp. 257–328.
Easley, D., Ledyard, J.O., 2018. Theories of price formation and exchange in double oral auctions. In: The Double Auction Market Institutions, Theories, and Evidence. Routledge, pp. 63–98.
Easley, D., OHara, M., Yang, L., 2016. Differential access to price information in financial markets. J. Financ. Quant. Anal. 51 (4), 1071–1110.
Fano, S., LiCalzi, M., Pellizzari, P., 2013. Convergence of outcomes and evolution of strategic behavior in double auctions. J. Evol. Econ. 23 (3), 513–538.
Fischbacher, U., 2007. z-tree: Zurich toolbox for ready-made economic experiments. Exp. Econ. 10 (2), 171–178.
Friedman, D., 2018. The Double Auction Market: Institutions, Theories, And Evidence. Routledge.
Gjerstad, S., Shachat, J.M., 2007. Individual Rationality and Market Efficiency. Working Papers. Institute for Research in the Behavioral, Economic, and Management sciences, Purdue University.
Gode, D.K., Sunder, S., 1993. Allocative efficiency of markets with zero-intelligence traders: market as a partial substitute for individual rationality. J. Polit. Econ. 119–137.
Gode, D.K., Sunder, S., 1997. What makes markets allocationally efficient? Q. J. Econ. 112 (2), 603–630.
Greiner, B., 2015. Subject pool recruitment procedures: organizing experiments with ORSEE. J. Econ. Sci. Assoc. 1 (1), 114–125.
He, X., Kang, J., Zhou, X., 2020. The fast and the furious: exchange latency and ever-fast trading. Available at SSRN 3681157.
Handbook of Computational Economics, 2018. In: Hommes, C., LeBaron, B. (Eds.), vol. 4. Elsevier.
Ladley, D., Pellizzari, P., 2014. The simplicity of optimal trading in order book markets. In: Dieci, R., He, X.-Z., Hommes, C. (Eds.), Economic Dynamics and Financial Modelling. Springer, pp. 183–199.
van de Leur, M., Anufriev, M., 2018. Timing under individual evolutionary learning in a continuous double auction. J. Evol. Econ. 28 (3), 609–631.
LiCalzi, M., Pellizzari, P., 2008. Zero-intelligence trading without resampling. In: Schredelseker, K., Hauser, F. (Eds.), Complexity and Artificial Markets. Springer, pp. 3–14.
Smith, V.L., 1962. An experimental study of competitive market behavior. J. Polit. Econ. 70 (2), 111–137.
Smith, V.L., 1980. Relevance of laboratory experiments to testing resource allocation theory. In: Evaluation of Econometric Models. Elsevier, pp. 345–377.
Snowberg, E., Yariv, L., 2021. Testing the waters: behavior across participant pools. Am. Econ. Rev. 111 (2), 687–719.
Vriend, N.J., 2000. An illustration of the essential difference between individual and social learning, and its consequences for computational analyses. J. Econ. Dyn. Control 24 (1), 1–19.