IEL-CDA Model: A More Accurate Theory of Behavior in Continuous Double Auctions*

Mikhail Anufriev† Jasmina Arifovic‡ Anil Donmez§
John Ledyard¶ Valentyn Panchenko‖

11 December 2023

Abstract

The continuous double auction (CDA) is a well-studied and widely used trading institution. However, there is no universally accepted theory regarding the dynamics of price formation, especially within the first period, that has endured experimental testing. In this paper, we introduce a behavioral model called IEL-CDA, which builds upon the Individual Evolutionary Learning (IEL) model of Arifovic et al. (2022). We enhance IEL by (a) incorporating Marshallian Selection, favoring traders with a higher expected surplus in making offers, and (b) allowing a trader's hypothetical reasoning to depend on the history of transactions. Using new experimental data, we test the hypothesis that efficiencies and average prices observed in the experiments follow the same distribution as those produced by simulations with various models. The hypothesis is rejected for both the Zero Intelligent theory of Gode and Sunder (1993) and IEL; however, it is not rejected for IEL-CDA. Therefore, IEL-CDA emerges as a more accurate theory of behavior in the continuous double auction.

Keywords: Continuous Double Auction; Learning; Evolutionary Dynamics; Zero Intelligence.

JEL Classification: C63; C92; D44; D82.

*We thank Steve Gjerstad, Megan Tjandrasuwita, Robert Sherman, students and faculty at 2022 Barcelona Workshop and participants of the Computational and Experimental Economic Workshop in memory of Jasmina Arifovic, held at Simon Fraser University, and two anonymous referees for their helpful comments and discussions. John Ledyard gratefully acknowledges partial financial support from the Australian Research Council (ARC) under grant DP190101718. Mikhail Anufriev gratefully acknowledges support from the ARC under grant DP200101438, and support from the Czech Science Foundation (GACR, project 22-28882S) and from the VSB-TU Ostrava under the SGS project SP2022/4.

†University of Technology Sydney, Department of Economics and Technical University of Ostrava, Faculty of Economics, Department of Finance.

‡Jasmina Arifovic passed away on 24 January 2022 while we were working on this project. Establishing IEL as a general theory of learning and, in particular, putting IEL to work in continuous double auctions, was a major goal of hers. She was a true scientist and a good friend and mentor.

§Simon Fraser University.

¶Corresponding author. California Institute of Technology, Email: jledyard@caltech.edu.

‖Economics, UNSW Business School, University of New South Wales

1 Introduction

A fundamental question about market economies is how prices are determined in centralized marketplaces. The process of price formation in continuous double auctions (CDAs) *at the very beginning* of trading is particularly intriguing. At this stage, all traders have just privately learned their valuations and costs, markets are not in equilibrium, and there is no common knowledge. We propose a novel model of trader behavior called IEL-CDA, which extends the Individual Evolutionary Learning model developed by Arifovic et al. (2022). Additionally, we present the results of a new experiment in which human subjects trade assets according to standard CDA rules. Our findings indicate that the IEL-CDA model aligns better with these and other experimental data compared to alternative models.

We are interested in the CDA due to its widespread use as one of the most common institutions for trading financial assets. In this market, buyers and sellers are free to make offers and accept the offers of others at any time. Unexecuted offers are stored in the 'order book' and can be either withdrawn or executed later. While specific trading rules may vary, traders typically have access to the current content of the order book. This 'open book CDA' is employed in trading on major stock exchanges, as well as popular online trading platforms. Our model and experiment take into account this information structure inherent in the CDA.

After the pioneering work of Smith (1962), numerous experiments have been conducted to examine the properties of CDA, see reviews in Holt (1995) and Plott (2008).[1] Almost without exception, these experiments indicate that the CDA serves as an efficient market mechanism, leading to prices and allocations that converge to their competitive equilibrium values after a sufficient number of replicated trading periods. However, there is currently no universally accepted theory regarding the initial, first-period dynamics of prices and allocations. This paper aims to address this gap.

Several theories have been proposed to explain price adjustment within one period in CDA, each based on different assumptions about the rationality of agents.[2] The three earlier theories are the Bayesian equilibrium of the waiting game of Wilson (1987), the Bayesian game against the nature of Friedman (1991), and the zero-intelligence (ZI) model of Gode and Sunder (1993). While all these theories are generally consistent with experimental findings of high efficiency in CDAs, differences

[1] An important advantage of the experimental method is its ability to facilitate direct control and manipulation of fundamentals, such as traders' valuations and costs, and market rules.

[2] Some models focusing on learning over repeated trading periods also incorporate within-period dynamics, e.g., Easley and Ledyard (1993); Gjerstad and Dickhaut (1998); Gjerstad (2007). While our focus is on first-period data, our approach shares common elements with these models.

emerge when they are closely matched to experimental data. These theories diverge in their predictions about the dynamics of order submission and transactions, price adjustment, and the sources of market inefficiency.[3]

Cason and Friedman (1996) employ these *discriminative* predictions to test the three theories on experimental data. While finding that none of the theories "adequately explains price formation in double auction markets," they highlight ZI as "the only natural source for a null hypothesis in assessing the performance of any more complicated mechanism." In a recent study, Lin et al. (2020) tested the same theories on over $2,000$ observations from classroom experiments[4] and found "much stronger support for zero-intelligence theory, compared to earlier evidence." Interestingly, they observe that "a non-negligible portion of our data falls outside of the simulated 95% confidence region" of the ZI model, suggesting a need for the development of alternative models.

Motivated by this finding, Arifovic et al. (2022) employ the Individual Evolutionary Learning (IEL) model in the CDA setting. Originally designed as a learning model for games with complex strategy spaces (Arifovic and Ledyard, 2001), IEL incorporates the principles of genetic algorithms, emphasizing experimentation and reinforcement of successful behaviors, within economic environments.[5] IEL agents are more sophisticated than ZI as their behavior is shaped by counterfactual analysis. However, the simulations with IEL, as reported in Arifovic et al. (2022), demonstrated mixed success when applied to the data from Lin et al. (2020): while the distribution of allocative efficiency closely resembled the data compared to ZI, the results for average prices were reversed.

[3]For instance, concerning transaction price changes, Wilson's model implies no predictable structure (zero autocorrelation), Friedman's model predicts positive autocorrelation, and the ZI model predicts negative autocorrelation. See Cason and Friedman (1996) for a detailed comparison.

[4]The data were collected through MobLab, an educational platform for conducting in-class experiments in economics courses. Full disclosure: Ledyard is on the Board of Directors of MobLab. The data are publicly available at `https://osf.io/9mfws/`.

[5]The effectiveness of genetic algorithms in optimizing individual and social learning in economics has been emphasized in Sargent (1993). For pioneering applications of GA to economic problems, see Arifovic (1994, 1995, 1996) and Dawid (1996). Other applications include Dawid and Kopel (1998), Dawid (1999), Vriend (2000), Arifovic et al. (2013), Fano et al. (2013), Hommes and Lux (2013), Anufriev et al. (2019). The IEL model has been previously used to investigate the impact of information availability from past trading sessions in both the call market (Arifovic and Ledyard, 2007) and CDA (Anufriev et al., 2013, 2022) environments. Additional applications of IEL encompass the principal-agent problem (Arifovic and Karaivanov, 2010), the Groves and Ledyard (1977) mechanism (Arifovic and Ledyard, 2001, 2004, 2011), the voluntary contribution mechanism (Arifovic and Ledyard, 2012), the timing of order submission (van de Leur and Anufriev, 2018), the repeated battle of the sexes game (Arifovic and Ledyard, 2018), correlated equilibrium behavior (Arifovic et al., 2019), and Keynesian beauty contest games (Anufriev et al., 2024).

Arifovic et al. (2022) argue that the classroom experiments, the source of data in Lin et al. (2020), deviated from the standard incentives and controls typically enforced in economic experiments. This resulted in numerous instances of individual rationality violations, such as selling for less than the cost or buying for more than the valuation. To address this issue, Arifovic et al. (2022) conducted controlled incentivized experiments at Simon Fraser University (SFU) with the same demand-supply configuration as in Lin et al. (2020) and did observe significant differences. However, once again, they found mixed support for the IEL model, which demonstrated improved fit over the ZI for efficiencies but not for average prices.

Motivated by these findings, this paper introduces a new model that aims to enhance the understanding of *first-period* behavior in CDAs, seeking improvement over the IEL model in the context of experimental data. This new model, termed IEL-CDA, incorporates two innovations into the established IEL framework. Firstly, it employs 'Marshallian' selection of an active trader (i.e., the trader chosen by the algorithm to submit an offer), prioritizing those with a higher expected surplus from trade. This diverges from IEL and ZI algorithms, where an active trader is selected uniformly at random. Secondly, the counterfactual analysis, used for making offers, is enhanced by introducing conditioning based on recent transactions within the period. The first modification is grounded in experimental evidence suggesting that buyers with higher valuations and sellers with lower costs tend to make offers earlier during the period. The second modification is grounded in behavioral considerations, where agents expect terms of trade as favorable as those observed in the recent past.[6] These modifications operate in opposing directions: Marshallian selection may improve efficiency by offering earlier trading opportunities to traders expected to trade in equilibrium, while conditioning offers on the recent past may inhibit trades as traders await better terms of trade. This paper demonstrates that both modifications are necessary for a more accurate fit to the data.

We evaluate the performance of IEL-CDA using two sets of data derived from controlled and incentivized experiments conducted at SFU. The first set employs the same data as reported in Arifovic et al. (2022), incorporating the demand and supply schedules from Lin et al. (2020), where each participant has three units to trade. Hence, we refer to this set as SFU3. However, experiments involving three units per subject pose strategic complexities. To provide an additional control, experiments with one unit per subject are often conducted in this literature.[7] Thus, we also

[6]Elements of Marshallian selection and history-dependent pricing strategies were considered in the Gjerstad and Dickhaut (1998) model.

[7]Moreover, as noted by Cason and Friedman (1996), in multiple-unit experiments, statistics such

utilize the results of a new experiment, labeled SFU1, where each participant has one unit to trade. The SFU1 design, characterized by greater symmetry compared to the SFU3 design, allows us to assess the robustness of IEL-CDA across different designs and cross-validate the parameters. Our simulations demonstrate that IEL-CDA is robust and exhibits closer alignment to both sets of experimental data than either ZI or IEL.

The paper is organized as follows. The next section reviews the ZI and IEL models, which have been the predominant theories concerning first-period dynamics in CDA. Section 3 introduces the IEL-CDA model of behavior. Section 4 details the experimental data (Section 4.1) and the simulation process (Section 4.2). Following that, Section 4.3 presents our main results, demonstrating that IEL-CDA is a more accurate model of human behavior in CDAs compared to IEL and ZI. Section 4.4 provides further insights through the robustness analysis, and Section 5 concludes. Supplementary results are included in Appendices A and B, while experiment instructions and screenshots are available in the Online Appendix.

2 A Review of ZI and IEL in the CDA

In this section, we first outline the key features of the CDA market and then review two existing algorithms, ZI and IEL.

Throughout the paper, we use the term 'offer' to refer to a limit order placed by any trader. Specifically, we use 'bid' to refer to an offer placed by a buyer and 'ask' to refer to an offer placed by a seller.

The CDA market operates based on the following four rules:

CDA-1. The market has a publicly visible book that lists all offers that have been made but not yet accepted.

CDA-2. Traders can submit offers at any time while the market is open. These offers are added to the book.

CDA-3. Traders can cancel any of their offers in the book at any time.

CDA-4. If a new bid is higher than the lowest ask in the book, a trade occurs at that ask. If a new ask is lower than the highest bid in the book, a trade occurs at that bid.

as buyers' and sellers' rank orders are influenced by the rule requiring individuals to trade their units in the order of their valuations/costs.

These market rules were implemented in our experiments, discussed later in Section 4.1, as they closely resemble those used in actual stock exchanges.[8]

Rule **CDA-4** refers to the highest bid and lowest ask in the book. Throughout the paper, we denote them as B_t and A_t, respectively. By construction, B_t represents the best bid available to sellers, and A_t represents the best ask available to buyers at the moment t. Rule **CDA-4** ensures that B_t is always less than A_t. The $[B_t, A_t]$ interval is referred to as the *spread*.

2.1 ZI (Zero Intelligence) traders

In Gode and Sunder (1993), it is argued that the high allocative efficiency of CDA markets may be attributed to the market rules, as it can be achieved even with seemingly random behavior of traders. Specifically, the ZI algorithm is employed, where a trader randomly determines when and what to offer. All offers fall within the interval of 'allowable' offers, denoted as $[s_\ell, s_u]$. The following algorithm runs for a fixed number of iterations.

ZI-1. (Random Selection) Offering occurs in a sequence of iterations. In each iteration, a trader is selected uniformly at random, with replacement, from the set of *active* traders, i.e., those traders who still have assets to trade.

ZI-2. (Random Bidding) The selected trader, if a buyer, places a bid for 1 unit at a price drawn uniformly from $[s_\ell, V]$, where V represents the marginal value to the buyer for the unit to be bought.

If a seller, the trader places an ask for 1 unit at a price drawn uniformly from $[C, s_u]$, where C denotes the marginal cost to the seller for the unit to be sold.

ZI-3. (Single Order in Book) The trader cancels their offer in the book, if they have one, whenever they are about to submit a new offer.

In markets where each trader has multiple units to trade, active buyers bid for the unit with the highest current valuation, while active sellers offer the unit with the lowest current cost.

Cason and Friedman (1996) and Lin et al. (2020) observed that the simulations generated by the ZI algorithm generally display more similarities to the earlier experimental data than more rational models. Consequently, it has been argued that

[8] Some literature on CDA markets includes additional rules. For instance, the 'spread reduction rule' requires new offers to improve the offers in the book. Another rule clears the book after each transaction. However, we did not incorporate these rules as they are not currently employed in practical applications of CDAs.

ZI should be considered as the null hypothesis for any theory of CDA. Therefore, we use ZI in this model as an important benchmark.

Although a ZI agent may appear random, it possesses a certain, non-zero degree of intelligence. Rule **ZI-2** ensures that the ZI trader behaves rationally on an individual level, avoiding losses. This is encapsulated in the following assumption:

IR. An artificial trader behaves *individually rational* if, as a buyer, they never bid above the valuation of the unit they intend to purchase, and as a seller, they never ask below the cost of the unit they wish to sell.

As pointed out by Gjerstad and Shachat (2021), this assumption is crucial not only for achieving high allocative efficiency and price convergence but also is not inherently part of the CDA market rules in reality.[9] It is worth noting that our experiments also do not enforce individual rationality. However, deviations from individual rationality were relatively rare in the experiments, with less than one case per trading period.

2.2 IEL (Individual Evolutionary Learning) traders

Arifovic et al. (2022) applied the IEL algorithm to the CDA. An IEL agent closely resembles a ZI agent in most aspects, with similarities in the selection of an agent to make an offer from the set of active agents and canceling the offer (as per **ZI-1** and **ZI-3**). Additionally, the **IR** assumption that traders avoid offers resulting in losses is retained. The submitted bid or ask corresponds to the agent's "strategy" and is chosen in a manner similar to other IEL studies.

The IEL trader maintains a pool of J remembered offers, which are all numbers within the range of allowable offers $[s_\ell, s_u]$. These offers are evaluated by a utility function U based on current market conditions.[10] At the start of each iteration, the pool undergoes experimentation and replication processes. During *experimentation*, each element in the pool is randomly replaced with probability μ by an element drawn from a normal distribution centered around the element being replaced.[11] *Replication* involves filling the new version of the pool by comparing utilities of members in the

[9]Hurwicz et al. (1975) provide a very similar model for Edgeworth box environments with random bidding from a trader's upper contour set. They prove that, with recontracting, the random process converges to a Pareto-optimal allocation. The simple demand-supply environment with quasi-linear utility functions widely used in CDA experiments is a special case of the Edgeworth box environment, where all Pareto-optimal allocations have the same commodity allocation. With recontracting, ZI traders would be expected to converge to the unique allocation associated with Pareto-optimal allocations, resulting in 100% efficiency.

[10]In IEL applications with simultaneous play, the utility function was referred to as *foregone* utility since it was based on what would have happened during the previous move. Here play is sequential.

[11]One drawback of this *local* experimentation is its reliance on a variance parameter. In our new

current pool through a pairwise tournament. After the pool is updated, the trader randomly chooses an offer from the new pool based on proportional utility. The key to the algorithm's success is how the utility function is defined.

At iteration t, when the current best bid and ask in the book are respectively B_t and A_t, the buyer's utility for the potential bid b is

$$U(b) = \begin{cases} V - A_t & \text{if } b \in (A_t, V] \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

Similarly, a seller's utility for the potential ask a is

$$U(a) = \begin{cases} B_t - C & \text{if } a \in [C, B_t) \\ 0 & \text{otherwise.} \end{cases} \tag{2}$$

IEL traders derive positive utility only from offers enabling an immediate profit. This condition necessitates that trading at the best price on the opposite side of the book adheres to their **IR** constraint. Buyers are inclined to make an immediate purchase if the best asking price is lower than their valuation, while sellers are eager to sell immediately as long as the best bid exceeds their cost. In all other scenarios, such as when the other side of the book is empty, traders have zero utility.

To initiate the IEL algorithm, at the start of a trading period, the pools of all traders undergo initialization by randomly selecting J items from the sets $[s_\ell, V_i^1]$, where V_i^1 denotes the valuation of the most valued unit for buyer i, and from the set $[C_j^1, s_u]$, where C_j^1 denotes the cost of the less expensive (i.e., most efficient) unit for the seller j. Following this initialization, the IEL algorithm for CDA is defined as follows.

IEL-1. (Random Selection) Offering occurs in a sequence of iterations. In each iteration, a trader is selected uniformly at random, with replacement, from the set of *active* traders.

IEL-2. (IEL Bidding) In each iteration, all traders update their pools of remembered

algorithm, IEL-CDA, introduced in Section 3, we utilize *global* experimentation, where the new element is drawn uniformly at random. Additionally, after experimentation, in the IEL algorithm of Arifovic et al. (2022), *all* strategies in the pool were checked for **IR**. Strategies failing this check were removed and replaced with random strategies satisfying **IR**. This check was especially necessary for multi-unit environments, where (for instance) a buyer who has just purchased their most valuable unit could have bids in the pool exceeding the valuation of the next unit, thus failing to satisfy **IR**. Our design of the IEL-CDA algorithm eliminates the need for this check.

offers by applying experimentation, **IR** validation, and replication. Buyers use the utility function (1), while sellers use the utility function (2).

The selected trader randomly chooses a strategy from its pool in proportion to the utility value. This selected strategy becomes the trader's offer.

IEL-3. (Single Order in Book) The trader cancels their offer in the book, if they have one, whenever they are about to submit a new offer.

To understand the behavior of IEL, recall that after an agent (e.g., a buyer) is selected to trade according to Rule **IEL-1**, the choice of its offer follows a proportionality rule based on utility function (1). This utility function yields at most two possible values: zero and $V - A_t$. If the best ask in the book, A_t, is affordable for the buyer (i.e., $V - A_t > 0$), the IEL agent is highly likely to accept it.[12] In contrast, a ZI agent in such a situation accepts the offer A_t with a probability of $(V - A_t)/(V - s_\ell) < 1$. On the other hand, when $V - A_t \leq 0$, all bids have the same value of 0, and thus a random offer from the IEL agent's remembered set is submitted. It can not be traded immediately and thus is placed into the order book. Similarly, in this case, a ZI agent places in the order book an offer uniformly selected from the $[s_\ell, V]$ interval.

In summary, IEL traders exhibit less patience than ZI traders, as they accept some offers with a higher probability than ZI agents. Consequently, trading progresses more rapidly with IEL agents than with ZI agents.

Simulations conducted in Arifovic et al. (2022) indicate that IEL provides a more accurate representation of subjects' behavior in CDA experiments compared to ZI. While neither IEL nor ZI precisely replicates the experimental data, the allocative efficiencies predicted by IEL are much closer to the MobLab data from (Lin et al., 2020) than those generated by ZI. IEL also outperforms ZI in explaining the efficiencies of the SFU3 incentivized experiment. However, ZI outperforms IEL in explaining average prices for both data sets. Thus, there is a need for a more effective model, which we explore in the following sections.

3 IEL-CDA

Two observations shed light on why IEL struggles to explain experimental data in CDA. Firstly, experiments indicate that traders with a larger potential surplus (i.e., buyers with higher valuations and sellers with lower costs) are more likely to

[12] The only scenario where the IEL agent is not accepting A_t is when their remembered set contains no offers within the $(A_t, V]$ interval.

make offers earlier. This is supported by rank correlation coefficients between valuations/costs and the timing of offers in our experimental data.[13] In contrast, the IEL procedure selects traders to bid with equal probability. Secondly, when aware of past transaction prices, traders will likely avoid reckless trading by refraining from immediately accepting any ask less than their valuation or any bid greater than their cost. Buyers may hesitate to accept an ask that seems too expensive based on previous trade prices. Even when the best ask in the book is affordable, they may submit a bid lower than this best ask, hoping for better terms of trade in the future. Similarly, sellers may avoid accepting a bid that they consider too cheap, even when it would not violate the **IR** assumption, anticipating higher bids in the future. This strategic behavior involves a forward-looking approach beyond the current bid or ask in the book, a nuance not captured by the myopic IEL algorithm. In contrast, IEL only uses information about the best bid or ask in the book and is not strategic and history-dependent.

Motivated by these observations, we introduce two modifications to the IEL model of Arifovic et al. (2022). First, we replace the random selection process **IEL-1** with what we term *Marshallian Selection*. Instead of randomly choosing the next bidder, we select bidders with a probability proportional to their expected surplus. Second, we alter the bidding rule **IEL-2** by modifying the utility function. Instead of using utilities (1) and (2), the utility functions now condition on the history of trades. These new functions are defined below in (8) and (9). Rule **IEL-3** remains unchanged.

Additionally, we streamline the strategy formulation process by defining them not in terms of bids and asks, but in terms of the *relative markdowns* and *relative markups*, defined for buyers and sellers, respectively, as

$$\varepsilon_b = \frac{V - b}{V - s_\ell} \quad \text{and} \quad \varepsilon_a = \frac{a - C}{s_u - C}, \tag{3}$$

where V is the valuation of the item to be bought, C is the cost of the item to be sold, b is the bid submitted by the buyer, and a is the ask submitted by the seller. When $\varepsilon_b = 0$, the buyer bids its valuation, and when $\varepsilon_b = 1$, the buyer submits the

[13] For the SFU1 experimental data (see Section 4.1), the corresponding rank correlations are -0.15 and 0.29 for buyers and sellers, respectively. Both are significantly different from 0 at the 5% level (one-sided p-values are 0.0311 and 0.0003, respectively). Note that the signs of these coefficients precisely indicate that earlier offers come from traders with more favorable valuations/costs. Cason and Friedman (1996) and Lin et al. (2020) reported the rank correlations between valuations/costs and the order of *trades*. In accordance with the literature, in what follows, we also study the rank-order correlations for trades, reporting them in Tables 1 for experimental data and in Table 2 for models.

lowest allowable bid. Similarly, for sellers.

Running the learning algorithm in terms of markdowns and markups, rather than bids and asks, eliminates the need for a separate **IR** validation step. By imposing that strategies $\varepsilon_b, \varepsilon_a \in [0,1]$, the **IR** assumption is automatically satisfied. Furthermore, this approach ensures that the strategies learned by traders for buying or selling their first unit can be used for subsequent units without any changes. When the algorithm selects a strategy, the offers submitted to the market are determined as:

$$b = V - \varepsilon_b(V - s_\ell), \tag{4}$$

and

$$a = C + \varepsilon_a(s_u - C), \tag{5}$$

which are equivalent to (3).

The algorithm for this new IEL-CDA model is defined as follows.

IEL-CDA-1. (Marshallian Selection) At each iteration, a trader is randomly drawn from the set of active traders in proportion to their perceived expected surplus. In other words, the selection of a trader to participate from the set of active traders, denoted as $\mathcal{A}_t$ (those who still have units to trade), follows the probabilities

$$\text{Prob}(i \text{ is selected}) = \begin{cases} \dfrac{PES_{i,t}}{\sum_{j\in\mathcal{A}_t} PES_{j,t}} & \text{if } \sum_{j\in\mathcal{A}_t} PES_{j,t} \neq 0 \\ \dfrac{1}{\#\mathcal{A}_t} & \text{otherwise}\,. \end{cases} \tag{6}$$

Here $\#\mathcal{A}_t$ is the number of active traders and the perceived expected surplus is calculated as

$$PES_{i,t} = \sum_{\varepsilon\in P_{i,t}} \tilde{U}(\varepsilon)\pi(\varepsilon) \tag{7}$$

where $P_{i,t}$ is the pool of agent i, $\tilde{U}(\varepsilon)$ is the utility function for strategy ε, and $\pi(\varepsilon)$ is the probability of selecting strategy ε from $P_{i,t}$.

IEL-CDA-2. (IEL-CDA Bidding) At each iteration, all traders update their pools of remembered strategies (relative markups and markdowns) by applying experimentation and replication based on the foregone utility functions.

Utilities are evaluated based on:

- the current best bid and ask in the book, B_t and A_t, and

- the largest and the smallest of the last L transaction prices, denoted by $\bar{p}_t$ and $\underline{p}_t$, respectively.[14]

Buyers use the utility function

$$\tilde{U}(\varepsilon_b) = \begin{cases} V - A_t & \text{if } b \in \big(A_t, \min\{V, \bar{p}_t\}\big] \\ 0 & \text{otherwise}, \end{cases} \tag{8}$$

where b is determined according to (4).

Sellers use the utility function

$$\tilde{U}(\varepsilon_a) = \begin{cases} B_t - C & \text{if } a \in \big[\max\{C, \underline{p}_t\}, B_t\big) \\ 0 & \text{otherwise}, \end{cases} \tag{9}$$

where a is determined according to (5).

The selected trader chooses a strategy from their pool at random proportionally to the utility $\tilde{U}$. The buyer's bid b is derived from the chosen strategy according to (4). The seller's ask a is derived from the chosen strategy according to (5).

IEL-CDA-3. (Single Order in Book) The trader cancels their offer in the book, if they have one, whenever they are about to submit a new offer.

We now discuss these modifications starting with the bidding rule **IEL-CDA-2**.

An IEL-CDA trader places bids similar to an IEL trader but with a distinct utility function. In **IEL-CDA-2**, the trader's utility function incorporates not only the current best ask and bid in the book but also depends on the largest and smallest of the past L transaction prices. This modification enables traders to use some of the past transaction history to anticipate their future options.[15] Traders act under the assumption that they can expect to encounter an offer in the future that is at least as favorable as the worst among the last L transactions. With this adjustment, traders become less inclined to quickly accept not-so-good offers. L is a new parameter of the model. Note that when $L = 0$, the new utility function coincides with the old one, i.e., $\tilde{U} = U$.

[14] Before the first transaction, we set $\bar{p}_t = s_u$ and $\underline{p}_t = s_\ell$.

[15] The conditioning on prices $\bar{p}_t$ and $\underline{p}_t$ draws parallels with the model of Easley and Ledyard (1993) for minimal expectations *between periods*. Additionally, there are resemblances with the model of Gjerstad and Dickhaut (1998), where traders form beliefs about the acceptance of their offers based on empirical frequencies observed in the order book. However, our approach is behaviorally simpler.

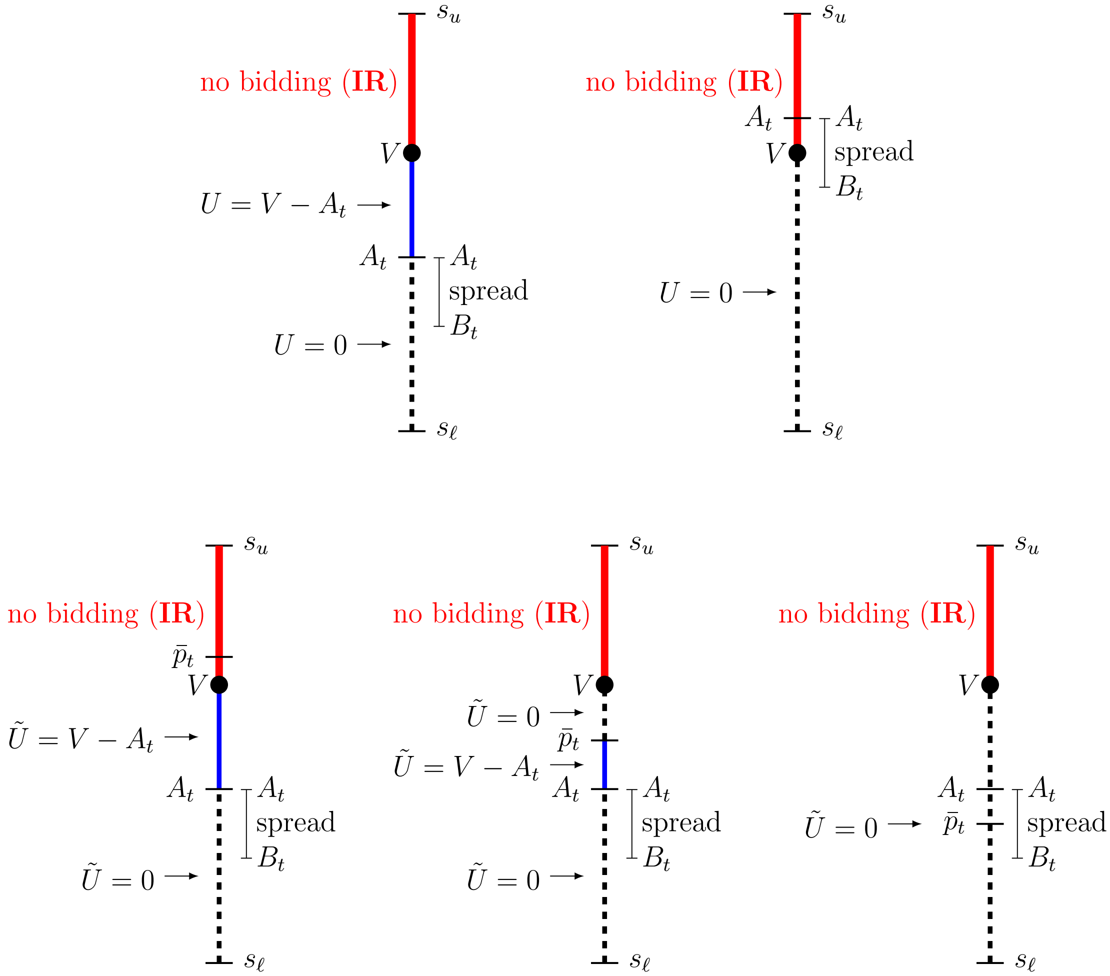

Figure 1: Illustration of the difference between utilities U (two upper panels) and $\tilde{U}$ (the right upper and all lower panels). The buyer with valuation of the unit V never bids in red thick intervals due to **IR** assumption. Their utility is positive in blue intervals and zero in the intervals shown by the dashed line.

Fig. 1 illustrates how this idea is realized in utility (8) for a buyer. (The case for a seller can be illustrated similarly.) Every panel shows the vertical interval of admissible offers, $[s_\ell, s_u]$. A buyer bids for a unit with valuation V, shown by the black dot. Due to assumption **IR**, the buyer never bids in the interval $[V, s_u]$.

The two upper panels illustrate the IEL utility function U defined in (1) and used in **IEL-2**. The utility depends on the current best ask in the book, A_t. There are two possibilities.[16] If $V > A_t$ (upper left panel), the buyer assigns a positive value $V - A_t$ to any possible bid above A_t, and 0 utility for bids below A_t. If $V \leq A_t$ (upper right panel), the buyer assigns zero utility for all possible bids.

[16]The spread $[B_t, A_t]$ is shown for illustrative purposes. The exact value of B_t does not matter.

IEL-CDA-2 does not differ from **IEL-2** when $V \leq A_t$; the upper right panel stays the same. However, when $V > A_t$, the case shown in the upper left panel, there are changes illustrated in the bottom three panels. In these cases, all bids above the maximum price of the previous L transactions are considered as 'not-so-good' and have zero utility. We have three subcases, $\bar{p}_t > V$, $\bar{p}_t \in (A_t, V)$ and $\bar{p}_t < A_t$. The lower panels illustrate these sub-cases for the utility function $\tilde{U}$ defined in (8). When $\bar{p}_t > V$ (left lower panel), there is no change in the buyer's utility calculations.

The changes occur, and **IEL-CDA-2** differs from **IEL-2**, when $\bar{p}_t < V$, that is, the maximum price of the last L transactions is within the interval of admissible bids of a buyer. In **IEL-CDA-2**, the buyer assigns positive utility only to those strategies that result in bids inside the interval $(A_t, \bar{p}_t)$. This interval is non-empty when $\bar{p}_t > A_t$ (middle lower panel) and empty otherwise (right lower panel). Thus, an IEL-CDA trader may not accept the best offer in the book in these cases when the IEL trader would accept it immediately. Therefore, rule **IEL-CDA-2** concentrates prices and, due to slower trading, lowers the efficiency of the final allocation.

However, we have also changed the rule for choosing the trader to make an offer, replacing random selection with Marshallian selection. **IEL-CDA-1** is a 'reduced-form rule' in the sense that its application requires knowledge of PES for all active traders, as these are all used in the proportionality rule (6). This raises a natural question of how this rule can be realized in a decentralized market. One way to justify it is to apply the idea from the models of Gjerstad and Dickhaut (1998) and Gjerstad (2007) about the timing of bidding by individual traders in the CDA. These models run in continuous time, and agents make offers at random instances. For each trader i, their individual waiting time before making an offer follows the exponential distribution with a rate proportional to their current expected surplus, ES_i. When traders make offers according to these processes, the probability that i will make the next offer is proportional to ES_i. That is, this decentralization of bidding implies **IEL-CDA-1**.

The Marshallian selection rule **IEL-CDA-1** speeds up trading of infra-marginal units relative to random selection. This tends to increase the efficiency of the final allocation. On the other hand, **IEL-CDA-2** tends to lower the efficiency of the final allocation. The combined effect of the two changes is to bring the distributions of both efficiency and average prices generated by IEL-CDA closer to the distributions from the experiments. This will be demonstrated in the next section.

4 Testing IEL-CDA

Our methodology for evaluating the suitability of IEL-CDA as a model for human behavior in the CDA comprises both experiments and simulations. We conduct experiments with human subjects and perform IEL-CDA simulations under identical market schedules. Furthermore, we compare these simulations to each other, as well as to simulations of the ZI and IEL models, to assess the performance of each model.

4.1 Experiments

Two experiments, SFU3 and SFU1, were conducted through Simon Fraser University. The data from SFU3 were previously described and used in Arifovic et al. (2022), while the data from SFU1 are new. The procedures for each experiment were similar.

For each experimental session, we recruited 10 participants with no prior experience in this type of experiment. Due to COVID-19 restrictions, the experiments were conducted online. Participants received a show-up fee of $7 and profits earned during the non-practice trading rounds. They were paid for each trading round, and the total earnings ranged from $7 to $25, with an average of $18. Each session lasted approximately 70 minutes.

Participants traded a non-divisible commodity in the CDA market for several rounds, using the Flex-E-Markets CDA software.[17] Each session included two practice rounds and 10 regular trading rounds (used in this paper), with an additional two human-to-robot trading rounds (not analyzed in this paper). Each trade round lasted for $T = 90$ seconds in the SFU3 experiment and for $T = 75$ seconds in the SFU1 experiment, with a 60-second break between rounds during which participants could view a screen displaying their earnings from the previous round. The roles of five buyers and five sellers were fixed for each participant in each round, with no contracting. Participants maintained the same role during the first five and last five regular rounds and switched roles in the middle. These details, along with the value of T, were publicly known.

Given our focus on the initial period of trade, we aimed to present each experimental round to participants as involving a distinct demand-supply configuration. One approach to achieve this would be to randomly draw new valuations and costs for each round, as conducted in Cason and Friedman (1993). However, we also sought to gather as much data as possible for the same configuration. To strike a balance, we adopted a compromise. We selected a benchmark demand and supply configura-

[17] See `http://flexemarkets.com`.

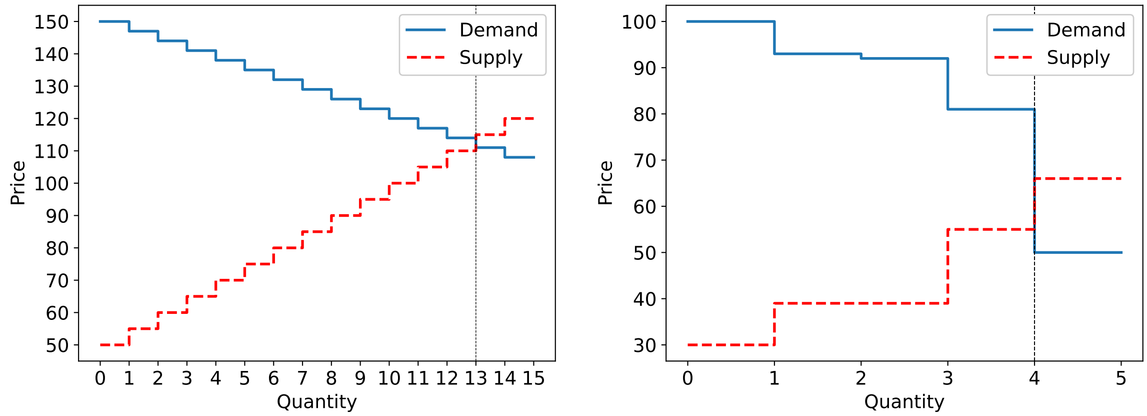

(a) SFU3 configuration

(b) SFU1 configuration

Figure 2: Benchmark demand and supply configurations used in the experiments. The dashed line indicates the equilibrium quantity.

tion, and for each round, we drew a random number from a distribution with a zero mean and added it to each valuation and cost, effectively shifting the demand and supply schedules up or down. Additionally, we randomly reallocated buyers along the demand curve and sellers along the supply curve, essentially shuffling valuations and costs. We believe that this approach created an environment for each round that participants perceived as new while being equivalent under quasi-linearity. Consequently, each experimental round is treated as an independent observation in this paper for the first period of trading.

For the SFU3 experiment, we adopted the benchmark demand-supply configuration from Lin et al. (2020). Each buyer could purchase up to 3 units, and their corresponding valuations were

$$\Big\{(150, 135, 120), (147, 132, 117), (144, 129, 114), (141, 126, 111), (138, 123, 108)\Big\}.$$

Sellers, on the other hand, could sell up to 3 units, and their costs were

$$\Big\{(50, 75, 100), (55, 80, 105), (60, 85, 110), (65, 90, 115), (70, 95, 120)\Big\}.$$

These valuations and costs resulted in the demand-supply configuration depicted in Fig. 2a. For this schedule, and any variations obtained by a random shift, the competitive equilibrium quantity is 13, comprising 13 infra-marginal and 2 extra-marginal units. The interval of competitive equilibrium prices is $[110, 114]$ for the benchmark schedule.

In the SFU1 experiments, the benchmark demand-supply configuration is taken from Arifovic and Ledyard (2007) and Anufriev et al. (2022). Each buyer is allowed to purchase up to 1 unit, and each seller can sell up to 1 unit. The valuations for buyers are $\{100, 93, 92, 81, 50\}$, while the costs for sellers are $\{30, 39, 39, 55, 66\}$. These valuations and costs led to the demand-supply configuration illustrated in Fig. 2b. For this schedule, and any variations obtained by a random shift, the competitive equilibrium quantity is 4, consisting of 4 infra-marginal and 1 extra-marginal unit. The interval of competitive equilibrium prices is $[55, 66]$ for the benchmark schedule.

We conducted three sessions each for the SFU3 and SFU1 experiments.[18] At the beginning of each session, we informed all traders about the range of admissible offers, $[s_\ell, s_u]$. We set this range to $[0, 300]$ in SFU3 and $[0, 150]$ in SFU1.[19] Before the start of each trading period, each trader saw their valuation(s) or cost(s) for that period. Traders did not know the valuations and costs of others, and neither did they know the whole schedule. The instructions and screenshots are provided in the Online Appendix.[20]

Table 1 provides some summary statistics from the two experiments. When we report prices, we adjust the actual experimental prices to the shifts in the schedules, allowing us to report all results for the benchmark schedules. We report the averages of all available observations (rounds) for the following statistics:

- Efficiency, E: the percentage of the maximum possible surplus realized in a round;
- Average price, P: the average price of all transactions during a round;
- Quantity, Q: the number of transactions in a round;

[18]The SFU3 sessions were conducted on April 20 and 21, 2021, while the SFU1 sessions were conducted on October 6, 8, and 11, 2021. For each experiment, we have 30 regular trading rounds, treated as independent observations. However, the software does not control for the maximum allowable quantity to buy. Although the instructions clearly indicated that participants should not exceed their allowable quantity, as their payoff would become negative in a round, some violated this. In SFU3, there were 7 rounds, and in SFU1, there were 6 rounds where these violations occurred, and we eliminated these rounds as the efficiency is not well-defined in these cases. The average price statistics did not change much after elimination.

[19]Note that the valuations and costs in the benchmark schedule for SFU3 belong to the interval $[50, 150]$. We achieved the shifting up and down for each round by adding an integer number drawn from the uniform distribution on the interval $[-50, 50]$ to this benchmark schedule. Similarly, the valuations and costs in the benchmark schedule for SFU1 are in $[30, 100]$, and we added integer numbers drawn from the uniform distribution on the interval $[-25, 25]$. This procedure ensured that all valuations and costs belonged to the range of admissible offers.

[20]The units we report in the paper are cents. The software presented these units to the subjects as fractions of dollars. For example, if a subject bids 1.73, we report it as 173.

		SFU3	**SFU1**
	Equilibrium quantity	13	4
	Equilibrium price range	$[110, 114]$	$[55, 66]$
E	Efficiency	70.2	81.1
P	Average price	106.2	65.0
Q	Quantity	7.7	3.4
ρ_B	Buyers' rank-order correlation	-0.44	-0.06
ρ_S	Seller's rank-order correlation	0.50	0.21
I	Inefficiency ratio	0.06	0.29
α	Smith's alpha	0.17	0.23
ρ_p	Price change autocorrelation	-0.30	-0.81
O	Offers	40.0	24.0
	Number of observations (rounds)	24	23

Table 1: The equilibrium quantity and price range for the schedules used in the SFU3 and SFU1 experiments, the averages of various statistics derived from the experimental data (per round), and the number of observations included in the analysis.

- Buyer's rank-order correlation coefficient, ρ_B, and seller's rank-order correlation, ρ_S: correlation between the rank of the valuation/cost (in ascending order) of an item and the order in which it traded;

- Inefficiency ratio, I: the percentage of the surplus lost due to traded extra-marginal units (as opposed to under-trading of infra-marginal units). If I is close to 0, then most of the efficiency loss is due to untraded infra-marginal units;

- Smith's alpha, α: a measure of transaction price deviations from the mid-point of the equilibrium price range, as first defined in Smith (1962);

- Price change autocorrelation, ρ_p: autocorrelation between consecutive transaction price changes in a round.

We also provide the average number of offers made in a round, denoted as O. This count includes updates of existing orders. This statistic is essential for our simulations as it is used for calibration, as further discussed in Section 4.2.

There are some differences in the experimental results for SFU3 and SFU1. Firstly, the average transaction price is within the equilibrium range for SFU1 but not for

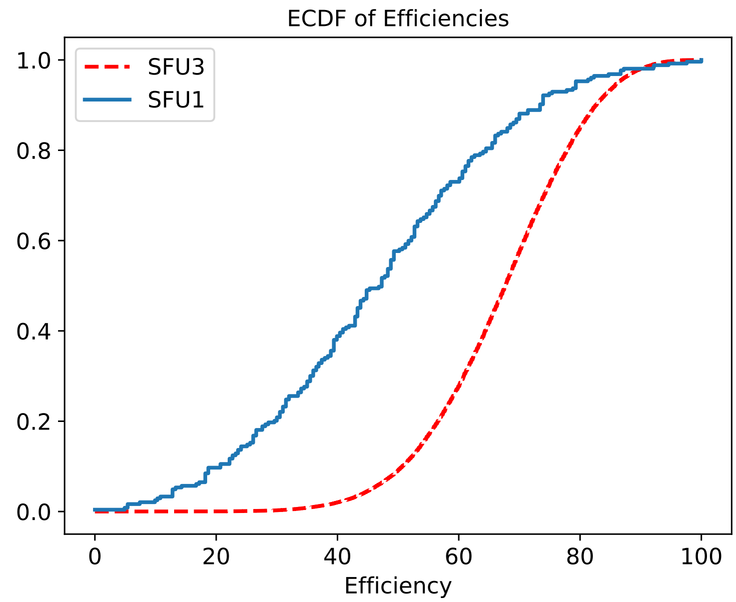

Figure 3: Empirical CDFs of efficiencies for random allocation of units for both experimental schedules. The SFU1 efficiencies are stochastically dominated by the SFU3 efficiencies.

SFU3. One reason for this difference may be the asymmetry in the SFU3 demand-supply configuration, which is much stronger than in the SFU1 configuration. In SFU3, at the mid-point of the competitive equilibrium price range, 112, the sellers' surplus is 422, significantly larger than the buyers' surplus of 253. In comparison, for SFU1, at the mid-point of the competitive equilibrium price range, 60.5, the buyers' surplus of 104 and the seller's surplus of 99 are much more similar. If units were to trade strictly in their rank order (an extreme Marshallian hypothesis), then in the SFU3 design, prices would start at 100 and climb as trades occur. With asymmetry favoring sellers and Marshallian trading, the average prices will be below competitive equilibrium prices. In the SFU1 design, prices would start at 65 and stay near that. With symmetry and Marshallian trading, the average price will be a competitive equilibrium price. Note that this story is supported by the price change autocorrelation coefficient, ρ_p. While it is negative in both experiments, reflecting the tendency of oscillations that are typical for the CDA, it is lower (in absolute value) for SFU3, indicating that the price indeed may have a trend there.

Secondly, while the average quantity traded per round is below the equilibrium level in both experiments, it is close to the competitive equilibrium for SFU1 but not for SFU3. Additionally, the average efficiency is higher in SFU1 than in SFU3. One might conjecture that this difference is caused by the demand-supply configurations, suggesting that it could be easier to attain high efficiencies in SFU1. To test this conjecture, Fig. 3 displays the cumulative distribution functions (CDFs) of the efficiencies for a random allocation.[21] However, we found that, for this exercise, SFU1 is a more challenging environment in which to achieve high efficiencies, contrary to the conjecture. The SFU1 efficiencies are stochastically dominated by the SFU3 efficiencies. Hence, the schedules do not help to explain the differences in efficiencies.

Another potential reason for the lower efficiency in SFU3 could be the higher likelihood of an aggressive extra-marginal unit replacing an infra-marginal unit, given the larger number of units per trader to trade. However, this conjecture is not supported by the inefficiency ratio I, which is lower in SFU3, indicating that the loss of efficiency due to the trading of extra-marginal units is less likely there than in SFU1. Yet another factor may be that in the multi-unit environment of SFU3, there is a higher chance for traders to be satisfied with their initial trades and not trade all their units. This would result in an inefficiently low quantity traded and lower efficiencies in SFU3, aligning with the experimental observations.

Finally, note that the rank-order correlations, ρ_B and ρ_S, are higher in absolute values for SFU3. This inherent bias is caused by the requirement for each trader to trade their units in the order of decreasing valuations and increasing costs.

4.2 Simulations

We conducted a series of simulations to evaluate the performance of our new IEL-CDA model and to compare it with the performances of the ZI and IEL models. Each simulation run corresponds to the initial period of a single experimental CDA market, with $100,000$ runs performed for each simulation.

While we aimed to replicate the experimental setting as closely as possible, some unavoidable differences exist. In the experiment, time is continuous, and trade is asynchronous, allowing for numerous offers (or updates to existing offers) within a trading round of known length T. In the simulations, time is discrete, indexed by t, and each iteration of an algorithm randomly selects a trader and their offer. The offer is either executed or stored in the book. The time is then incremented by one,

[21] We simulated the following process for each schedule: we took all the possible units for sale from the sellers, i.e., 5 units for SFU1 and 15 units for SFU3. We then randomly allocated those items across both buyers and sellers up to the maximum each could hold (one unit for SFU1 and 3 units for SFU3). We then calculated the surplus for that allocation by adding up the valuations of the items now held by buyers and subtracting the costs of the items not held by the sellers. Dividing that total by the maximum possible surplus across all allocations and multiplying by 100 yields the efficiency of that allocation.

and another active trader is selected with replacement.

The algorithm is set to terminate after D iterations, with D referred to as 'draws,' a critical parameter. If D is too small, there may be missed trading opportunities, resulting in low allocative efficiency. Conversely, if D is too large, the pool of active traders may become empty before all D iterations are completed, resulting in no market activity for the remaining iterations. To avoid optimizing the number of draws *ex-post*, we fixed the number of draws so that the average number of offers per period in the simulations closely matches the average number of offers in the experiments, as reported in the last row of Table 1.[22]

IEL-CDA involves three parameters: the size of the pool of remembered strategies J, the probability of experimentation μ, and the length of the history of transactions, L, along with a probability distribution used for experimentation. The J and μ parameters were originally introduced in the context of IEL and have been consistently used in various applications, some of which are listed in footnote 5. Our goal, aligning with previous research, is to identify a version of IEL that generalizes across different environments and mechanisms. Consequently, we have maintained the tradition of using the same 'off-the-shelf' parameter values for IEL-CDA as those employed in diverse applications of IEL (Arifovic and Ledyard, 2007, 2012; Anufriev et al., 2013, 2024). Specifically, the size of the pool of remembered strategies is set at $J = 100$, and the experimentation probability is $\rho = 0.033$. Our results demonstrate robustness to variations in these parameters, as we discuss in Section 4.4.

During experimentation in IEL-CDA, a new strategy (relative markup or markdown) is generated using a probability distribution function. In contrast to IEL's use of a truncated normal distribution centered around the previous strategy, IEL-CDA draws the new strategy from the uniform distribution on $[0, 1]$. This adjustment has minimal impact on the outcomes but makes a distribution parameter-free.

The new and only free parameter in IEL-CDA is L, representing the length of the transaction history considered in the foregone utility functions $\tilde{U}$ in (8) and (9). We explored the model's fit across different values of L (refer to Table 5 in Section 4.4). Our analysis suggests that intermediate values of L (specifically $L = 2$ and 3) offer the best fit for both SFU3 and SFU1 experiments. Applying parsimony, we use $L = 2$ in all IEL-CDA simulations reported below.

[22]See Arifovic et al. (2022) for a comprehensive discussion of this modeling choice. Note that, it is possible for the number of offers to be less than the number of draws. This happens when a buyer who has already bought the maximum number of units is chosen to bid or when a seller who has sold all their units is chosen to ask. This is more likely to occur when the number of units allowable per trader is smaller, that is, in the SFU1 schedule.

	D	O	E	P	Q	ρ_B	ρ_S	I	α	ρ_p	T^2	p-val
SFU3 schedule												
Experiment		40.0	70.23	106.23	7.70	−0.44	0.50	0.06	0.17	−0.30		
			(2.28)	(4.03)	(0.34)	(0.06)	(0.08)	(0.02)	(0.03)	(0.11)		
ZI	40	40.0	40.4	109.0	3.5	−0.30	0.30	0.00	0.17	−0.40	85.72	0.00
IEL	40	40.0	76.0	105.0	8.1	−0.48	0.48	0.06	0.19	−0.43	15.49	0.05
IEL-CDA	40	40.0	67.8	106.6	6.8	−0.41	0.46	0.02	0.17	−0.48	8.01	0.43
SFU1 schedule												
Experiment		24.0	81.04	65.01	3.42	−0.06	0.21	0.29	0.23	−0.81		
			(5.50)	(2.25)	(0.22)	(0.15)	(0.15)	(0.09)	(0.03)	(0.15)		
ZI	24	24.0	50.4	67.0	2.0	−0.07	0.08	0.10	0.24	−0.40	19.77	0.01
IEL	25	23.7	86.3	65.8	3.9	−0.11	0.10	0.75	0.27	−0.45	15.55	0.05
IEL-CDA	25	24.4	85.5	65.8	3.7	−0.16	0.24	0.58	0.25	−0.50	7.83	0.45

Table 2: A comparison of average measures in the experimental data and ZI, IEL, and IEL-CDA algorithms. Standard errors for the averages in the experimental data are indicated in parentheses. The sampling error for computing averages in simulations with $100,000$ runs is negligible compared to the experiments. The last two columns provide the test statistics and p-values for a joint Hotelling's T^2 test for the eight measures.

4.3 Results

In each experimental period and for each simulation run, a sequence of trades is generated, and all eight measures introduced in Section 4.1 can then be computed. The averages of these measures over all periods for each experiment are initially presented in Table 1. To facilitate comparison with models, these experiment averages are included in the first row of the corresponding block in Table 2. Standard errors for these averages are provided in parentheses below. The following three rows in each block display the values of the same measures obtained for each of the three models (ZI, IEL, and IEL-CDA) by averaging results from $100,000$ independent simulation runs. The sampling errors for the averages in the simulations are negligible compared to those in the experiments and not reported.

The second and third columns of Table 2 display the number of draws for the algorithms and the average number of offers made in the experiments and simulations, respectively. Recall from Section 4.2 that the number of draws is controlled in the simulations to match the average number of offers observed in the experiment.

The average values for most of the measures, including allocation efficiency, E, and average transaction prices, P, indicate that the IEL-CDA model closely approximates the experimental data, differing strongly from the ZI and closely aligning with the IEL model. To formalize this observation, we conduct a joint Hotelling's T^2 test for the eight considered measures. The test accommodates potential dependencies between the measures by utilizing the full variance-covariance matrix in the computation of the test statistics. The test statistics and corresponding p-values are reported in the last two columns of Table 2. Our conclusion is that, considering all the measures, the null hypothesis – that the averages generated by the IEL-CDA model are not significantly different from the averages of the experimental data – is not rejected for both SFU3 and SFU1 experiments. In contrast, the same hypothesis formulated for either ZI or IEL model is rejected at least at the 10% nominal significance level.[23]

However, the closeness in averages does not necessarily imply closeness in distribution. To enhance our confidence in asserting that IEL-CDA outperforms ZI and IEL, we compare the distributions of efficiency and average price, the two key measures, from both experiments and simulations. In Fig. 4, we present a comparison of the empirical CDFs of efficiency and average prices in the SFU3 experiment with simulations from three models. Similarly, Fig. 5 provides the same comparison for the SFU1 experiment.[24]

Regarding the efficiency measure, the ZI simulations substantially deviate from the data in both SFU3 and SFU1 cases. IEL and IEL-CDA closely track the efficiency measure, but a noticeable difference exists between them, particularly in SFU3. In general, it seems that IEL achieves higher efficiency compared to IEL-CDA, as also reflected in the average of this measure reported in the E column in Table 2. However, the IEL-CDA model provides a closer fit to the experimental data, improving upon the IEL model.

Concerning average prices, all simulations show relatively similar results compared to the data. By examining the tails of the empirical CDF, we observe that ZI generates the largest variance, while IEL generates the smallest. IEL-CDA falls in between, indicating an improvement over IEL, with this difference being particularly visible for the SFU3 data. This improvement is also reflected in the average of this measure reported in the P column in Table 2.

Based on these observations, we formulate the hypothesis that the efficiency and average prices observed in the experimental data come from the same distribution as those in the simulation data. To test this hypothesis, we perform two sets of two-sample Kolmogorov-Smirnov (KS) tests: one for the differences in the distributions of efficiency and another for the differences in the distribution of average price. Ad-

[23]We conduct multiple comparisons using two datasets, SFU3 and SFU1. To account for the increased risk of making a Type I error, we apply the Bonferroni correction by dividing the desired nominal significance level by the number of comparisons, which is 2.

[24]For a more comprehensive comparison across various dimensions, Appendix A includes additional figures of empirical CDFs for other measures.

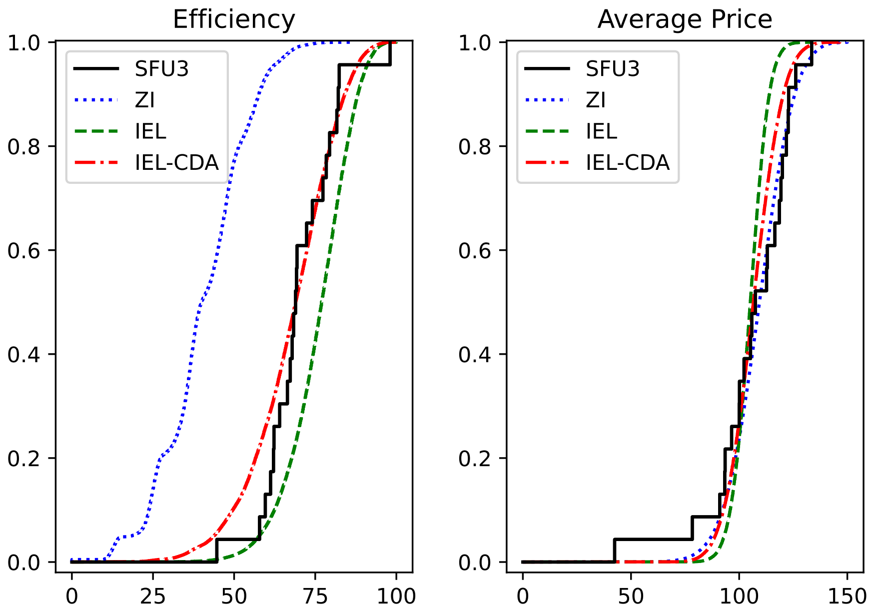

Figure 4: Empirical CDF of Efficiency (*left panel*) and Average Price (*right panel*) for the SFU3 experiment (black solid curve) and three models: ZI, IEL, and IEL-CDA.

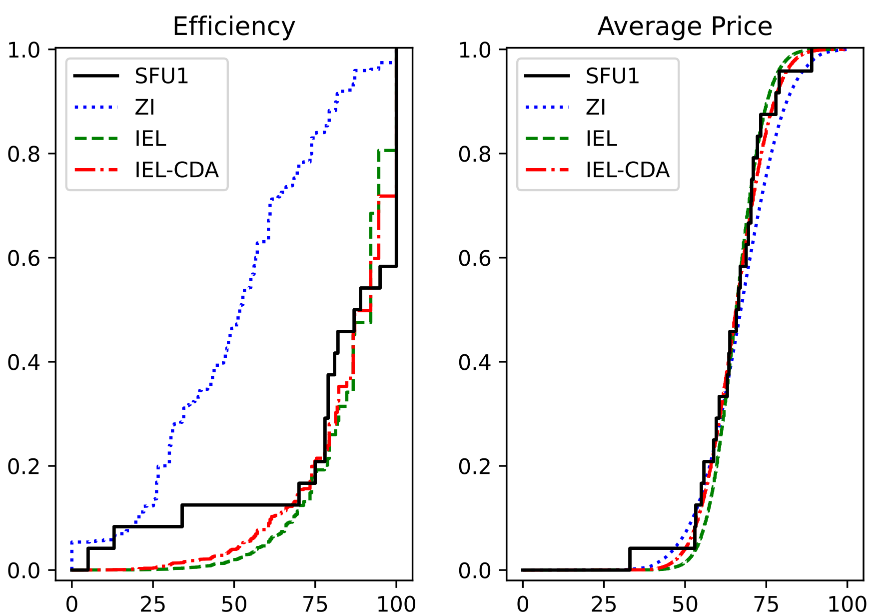

Figure 5: Empirical CDF of Efficiency (*left panel*) and Average Price (*right panel*) for the SFU1 experiment (black solid curve) and three models: ZI, IEL and IEL-CDA.

	Efficiency, E		Average Price, P		Joint E and P	
	KS stat	p-value	KS stat	p-value	2D KS stat	p-value
SFU3 schedule						
ZI	0.86	0.00	0.11	0.93	0.78	0.00
IEL	0.37	0.00	0.34	0.01	0.45	0.00
IEL-CDA	0.18	0.41	0.21	0.22	0.22	0.29
SFU1 schedule						
ZI	0.67	0.00	0.21	0.23	0.65	0.00
IEL	0.26	0.06	0.13	0.80	0.33	0.03
IEL-CDA	0.18	0.40	0.08	0.99	0.20	0.38

Table 3: Test statistics and p-values for the KS tests for equality of the distributions in the experimental data and simulations.

ditionally, a 2-dimensional KS test assesses the difference in the joint distribution of efficiency and average price. Table 3 reports the KS test statistic and p-values for these tests.

For efficiency, the hypothesis is rejected at the 5% level for both ZI and IEL models in the SFU3 data, and only for ZI in the SFU1 data. However, the hypothesis is not rejected for the IEL-CDA model in both experiments. Regarding average price, the hypothesis is rejected for IEL but not for ZI and IEL-CDA models in the SFU3 data, while it is not rejected for any model in the SFU1 data. The 2-dimensional KS test rejects the hypothesis for both ZI and IEL in both SFU3 and SFU1 data. In both cases, the hypothesis is not rejected for the IEL-CDA model. Consequently, we can conclude that neither ZI nor IEL survive the KS test, and instead, the test supports the IEL-CDA algorithm.

In summary, we have established that the IEL-CDA model, featuring Marshallian selection and the incorporation of the largest and smallest of the last L transaction prices in the utility functions (8) and (9), is not rejected by the data from SFU3 and SFU1 experiments. This contrasts with both the ZI model (Gode and Sunder, 1993) and the IEL model (Arifovic et al., 2022) that are rejected by the data.

4.4 Robustness and cross-validation

The IEL-CDA model utilized so far employs $L = 2$ as the length of the history, along with off-the-shelf values for the size of the pool, $J = 100$, and the experimentation probability, $\mu = 0.033$. In this section, we explore the robustness of IEL-CDA simulations to variations in its parameter values.

We begin by exploring variations in J and μ, keeping $L = 2$. We consider the grid with values of $J \in \{50, 100, 150, 200, 300\}$ and values of $\mu \in \{0.0033, 0.033, 0.1, 0.25\}$. In Table 4, we present the averages of efficiency (E) and average price (P) over 100,000 simulations of the IEL-CDA model for each pair of parameter values, separately for SFU3 and SFU1 experiments.

	$J = 50$	$J = 100$	$J = 150$	$J = 200$	$J = 300$
SFU3 schedule					
Efficiency, E					
$\mu = 0.0033$	*63.0*	67.5	69.4	70.7	71.5
$\mu = 0.033$	*63.2*	67.8	69.7	70.2	71.2
$\mu = 0.1$	*63.7*	67.6	70.1	70.1	71.2
$\mu = 0.25$	*63.1*	67.1	69.3	70.9	71.8
Average Price, P					
$\mu = 0.0033$	106.6	107.0	106.7	106.1	106.9
$\mu = 0.033$	107.0	106.6	106.7	106.5	106.2
$\mu = 0.1$	107.0	106.9	106.3	106.9	105.8
$\mu = 0.25$	106.7	107.1	107.8	106.9	106.2
SFU1 schedule					
Efficiency, E					
$\mu = 0.0033$	81.9	85.0	86.2	86.7	87.8
$\mu = 0.033$	80.3	85.5	86.4	87.6	88.4
$\mu = 0.1$	81.2	86.3	86.3	87.6	88.3
$\mu = 0.25$	84.4	88.0	88.0	87.9	88.3
Average Price, P					
$\mu = 0.0033$	66.2	65.8	65.2	65.9	65.8
$\mu = 0.033$	65.4	65.8	65.6	65.6	65.4
$\mu = 0.1$	65.9	66.3	65.8	65.5	65.7
$\mu = 0.25$	65.9	65.5	66.1	65.1	65.4

Table 4: Robustness: The size of the pool J and the probability of experimentation μ in IEL-CDA simulations. Values in *italics* are rejected by the data.

We observe that increases in either J or μ result in somewhat higher efficiency values in both the SFU3 and SFU1 schedules. The KS test results (not reported for brevity) indicate that the null hypothesis of equal efficiency distribution is only rejected for $J = 50$ and all considered μ values (for this J) in the SFU3 schedule simulations. The efficiency values generated by all other variations are not rejected by the experimental data. The average price shows even smaller sensitivity to the selected variations of μ and J and does not exhibit any clear pattern. The KS test suggests that the average price values generated by all variations are not rejected by the experimental data. In conclusion, our main results remain robust for values of $J \geq 100$ and all μ within the chosen grid.

Now, let us examine the impact of the length of the history of transactions, L. Given the common robustness of the IEL algorithm to J (as long as it is sufficiently large) and μ across applications, we can consider L as the sole free parameter in the IEL-CDA model. In Table 5, we present the results of simulations of IEL-CDA for values of $L \in 0, 1, 2, 3, 4, 5$. As before, we calculate the averages of efficiency and average price measures over 100,000 simulations for both SFU3 and SFU1 schedules. Additionally, for each value of L, we perform two-sample Kolmogorov-Smirnov (KS) tests (for differences in the distributions of efficiency and the distribution of average price) and report their p-values in the last two columns. The data for baseline simulations with $L = 2$ are in bold.

L	Efficiency, E	Average Price, P	p-value E	p-value P
SFU3 schedule				
0	92.2	105.1	0.00	0.00
1	61.8	106.8	0.01	0.24
2	**67.8**	**106.6**	**0.41**	**0.22**
3	71.2	106.5	0.24	0.21
4	73.2	106.3	0.07	0.18
5	74.5	106.3	0.03	0.19
SFU1 schedule				
0	94.0	65.3	0.00	0.77
1	82.1	65.7	0.14	1.00
2	**85.5**	**65.8**	**0.40**	**0.99**
3	86.3	65.6	0.48	1.00
4	86.2	65.7	0.47	0.99
5	86.2	65.7	0.47	0.99

Table 5: Sensitivity of the IEL-CDA model to the length of the history of transactions, L, in SFU3 and SFU1 simulations. The last two columns show the p-values of the KS-tests.

In SFU3 schedule simulations, only values of $L = 2$ and $L = 3$ are supported by the KS statistics (when applied to the efficiency measure, E). In contrast, in SFU1 schedule simulations, the model is rejected by the KS test only for $L = 0$ (when applied to the efficiency measure). Note that for this schedule, only 3 or 4 units per period are often traded in the experiment and simulations. Therefore, any value of $L > 3$ yields nearly the same outcomes, keeping in mind that L is always truncated to the available past history. Overall, this indicates that the choice of L is important for the IEL-CDA model.

Recall that the SFU3 and SFU1 schedules are not only different in the number of units agents are allowed to trade but also in the (a)symmetry between buyers and sellers' equilibrium surplus. These differences give a way to cross-validate our choice of L across the schedules. For the SFU3 schedule, $L = 2$ gives the best overall fit to the SFU3 experimental data (suggested by the highest p-values). Checking this choice of L against the IEL-CDA model from the SFU1 schedule shows that $L = 2$ is not rejected by the SFU1 experimental data. Conversely, $L = 3$ is the best choice for the IEL-CDA model when applied to the SFU1 schedule. Checking this choice of L against the model from the SFU3 schedule shows that $L = 3$ is not rejected by the SFU3 data. Therefore, either $L = 2$ or $L = 3$ is a reasonable choice for L in the simulations. Using the principle of parsimony, we selected $L = 2$.

Appendix B provides more detailed simulations and a discussion about the interactions of the two introduced features, Marshallian selection, and the length of the history of transactions, L. The analysis re-confirms that both these features are necessary for a good fit of the model with the experimental data.

5 Concluding Remarks

Although the Continuous Double Auction has been extensively studied for over 60 years, there is still no universally accepted theory regarding the dynamics of price formation within the crucial first period that has withstood experimental testing. In this paper, we propose a behavioral model, IEL-CDA, that addresses this gap. Our new model of behavior, IEL-CDA, builds upon the previous stochastic individual evolutionary learning (IEL) model, introducing two key modifications. First, the probability of traders being selected by the IEL-CDA algorithm to make an offer is proportional, at each time, to their expected surplus. Second, when selected to make an offer, an IEL-CDA trader generally accepts the best offer in the book only if both their valuation/cost allows it and the best offer does not seem unfavorable compared to the previous L transaction prices.

We compared simulations of three behavioral models – ZI by Gode and Sunder (1993), IEL by Arifovic et al. (2022), and our new IEL-CDA – with the new data from incentivized experiments. Our primary focus was on matching the *entire distributions* of allocative efficiency and average trading prices. Additionally, we considered various other features of the data identified in the earlier literature. The econometric analysis in Section 4.3 led us to conclude that the IEL-CDA model is not rejected by the data, in contrast to the ZI and IEL models.

There are several questions that future research can consider. First, it is of interest to study whether the IEL-CDA model passes the experimental tests that we explored in this paper, when analyzing the data with repeated trading over multiple periods and in richer environments, such as those involving multiple assets and the possibility of resale. It is known that repeated trading often converges to a competitive equilibrium in experiments via the so-called saw-tooth pattern (Plott, 2008). The asset pricing models of multiple assets include cross-commodity effects (such as whether assets are substitutes and complements) that influence trade patterns (Asparouhova et al., 2020). Some additional modifications to the IEL-CDA may be necessary to account for these empirical regularities.

Second, one might be interested in studying the implications of violations of the individual rationality assumption, **IR**, in IEL-CDA. The models we considered all adhere to this assumption, prohibiting agents from trading at a loss. However, even in controlled and well-incentivized experiments, occasional violations of individual rationality occur. It may well be that incorporating occasional errors in offers, which would eventually be mitigated by the IEL mechanism, improves the fit of IEL-CDA to the experimental data, similar to how Quantal Response Equilibrium improves over Nash predictions in explaining game-theoretical experiments.

Third, as mentioned before, our model can be seen as a behaviorally simpler, stochastic version of the Gjerstad and Dickhaut (1998) model. In that model, agents use historical frequencies to form their beliefs about the acceptance of different offers and then maximize utility. One related possibility is to consider IEL over strategies based on various evolving and competing belief schemes, such as in the Heuristic Switching Model of Anufriev and Hommes (2012). Some data are better explained by the models with learning over heuristics than learning over direct actions (Anufriev et al., 2019), and it is of interest to check if this holds in the CDA environment. While the experimental data we use are too short and the markets are too thin for such an extension of the model, future research may bring more suitable data. Examining how the Gjerstad and Dickhaut (1998) model of behavior in a CDA compares with IEL-CDA would further test both models on new data.

References

Anufriev, Mikhail, Jasmina Arifovic, John Ledyard, and Valentyn Panchenko (2013) "Efficiency of Continuous Double Auctions under Individual Evolutionary Learning with full or limited information," *Journal of Evolutionary Economics*, 23, 539–573.

——— (2022) "The role of information in a continuous double auction: An experiment and learning model," *Journal of Economic Dynamics & Control*, 104–387.

Anufriev, Mikhail, John Duffy, and Panchenko Valentyn (2024) "Individual Evolutionary Learning in Repeated Beauty Contest Games," *Journal of Economic Behavior & Organization*, forthcoming.

Anufriev, Mikhail and Cars Hommes (2012) "Evolutionary selection of individual expectations and aggregate outcomes in asset pricing experiments," *American Economic Journal: Microeconomics*, 4 (4), 35–64.

Anufriev, Mikhail, Cars Hommes, and Tomasz Makarewicz (2019) "Simple forecasting heuristics that make us smart: Evidence from different market experiments," *Journal of the European Economic Association*, 17 (5), 1538–1584.

Arifovic, Jasmina (1994) "Genetic algorithm learning and the cobweb model," *Journal of Economic Dynamics & Control*, 18 (1), 3–28.

——— (1995) "Genetic algorithms and inflationary economies," *Journal of Monetary Economics*, 36 (1), 219–243.

——— (1996) "The behavior of the exchange rate in the genetic algorithm and experimental economies," *Journal of Political Economy*, 104 (3), 510–541.

Arifovic, Jasmina, Joshua F Boitnott, and John Duffy (2019) "Learning correlated equilibria: An evolutionary approach," *Journal of Economic Behavior & Organization*, 157, 171–190.

Arifovic, Jasmina, James Bullard, and Olena Kostyshyna (2013) "Social learning and monetary policy rules," *The Economic Journal*, 123 (567), 38–76.

Arifovic, Jasmina, Anil Donmez, John Ledyard, and Megan Tjandrasuwita (2022) "Individual evolutionary learning and zero-intelligence in the continuous double auction," in Füllbrunn, Sacha and Ernan Haruvy eds. *Handbook of Experimental Finance*, Chap. 19, 225–249: Edward Elgar Publishing.

Arifovic, Jasmina and Alexander Karaivanov (2010) "Learning by doing vs. learning from others in a principal-agent model," *Journal of Economic Dynamics & Control*, 34 (10), 1967–1992.

Arifovic, Jasmina and John Ledyard (2001) "Computer testbeds and mechanism design: application to the class of Groves-Ledyard mechanisms for provision of public goods," Working Paper, Available at https://www.researchgate.net/publication/2381600.

——— (2004) "Scaling up learning models in public good games," *Journal of Public Economic Theory*, 6 (2), 203–238.

——— (2007) "Call market book information and efficiency," *Journal of Economic Dynamics & Control*, 31 (6), 1971–2000.

——— (2011) "A behavioral model for mechanism design: Individual evolutionary learning," *Journal of Economic Behavior & Organization*, 78 (3), 374–395.

——— (2012) "Individual evolutionary learning, other-regarding preferences, and the voluntary contributions mechanism," *Journal of Public Economics*, 96 (9-10), 808–823.

——— (2018) "Learning to alternate," *Experimental Economics*, 21 (3), 692–721.

Asparouhova, Elena N, Peter Bossaerts, and John O Ledyard (2020) "Price Formation in Multiple, Simultaneous Continuous Double Auctions, with Implications for Asset Pricing," Working Paper, Available at SSRN: https://ssrn.com/abstract=2774301.

Cason, Timothy and Daniel Friedman (1993) "An empirical analysis of price formation in double auction markets," in Friedman, Daniel and John Rust eds. *The Double Auction Market: Institutions, Theories, and Evidence*, 253–284: Routledge.

——— (1996) "Price formation in double auction markets," *Journal of Economic Dynamics & Control*, 20, 1307–1337.

Dawid, Herbert (1996) *Adaptive learning by genetic algorithms: Analytical results and applications to economic models*: Springer-Verlag New York, Inc.

——— (1999) "On the convergence of genetic learning in a double auction market," *Journal of Economic Dynamics & Control*, 23, 1545–1567.

Dawid, Herbert and Michael Kopel (1998) "On economic applications of the genetic algorithm: a model of the cobweb type," *Journal of Evolutionary Economics*, 8, 297–315.

Easley, David and John Ledyard (1993) "Theories of price formation and exchange in double oral auctions," in *The Double Auction Market: Institutions, Theories, and Evidence*, 63–98: Routledge.

Fano, Shira, Marco LiCalzi, and Paolo Pellizzari (2013) "Convergence of outcomes and evolution of strategic behavior in double auctions," *Journal of Evolutionary Economics*, 23 (3), 513–538.

Friedman, Daniel (1991) "A simple testable model of double auction markets," *Journal of Economic Behavior & Organization*, 15 (1), 47–70.

Gjerstad, Steven (2007) "The competitive market paradox," *Journal of Economic Dynamics and Control*, 31 (5), 1753–1780.

Gjerstad, Steven and John Dickhaut (1998) "Price formation in double auctions," *Games and Economic Behavior*, 22 (1), 1–29.

Gjerstad, Steven and James Shachat (2021) "Individual rationality and market efficiency," *Nonlinear Dynamics, Psychology, and Life Sciences*, 25 (4), 395–406.

Gode, Dhananjay and Shyam Sunder (1993) "Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality," *Journal of Political Economy*, 119–137.

Groves, Theodore and John Ledyard (1977) "Optimal allocation of public goods: A solution to the "free rider" problem," *Econometrica*, 783–809.

Holt, Charles A (1995) "Industrial organization: A survey of laboratory research," in Kagel, J. and A. Roth eds. *The Handbook of Experimental Economics*, 349, 402–03: Princeton University Press Princeton, NJ.

Hommes, Cars and Thomas Lux (2013) "Individual expectations and aggregate behavior in learning-to-forecast experiments," *Macroeconomic Dynamics*, 17 (2), 373–401.

Hurwicz, Leonid, Roy Radner, and Stanley Reiter (1975) "A stochastic decentralized resource allocation process: Part I," *Econometrica*, 43 (2), 187–221.

van de Leur, Michiel and Mikhail Anufriev (2018) "Timing under Individual Evolutionary Learning in a Continuous Double Auction," *Journal of Evolutionary Economics*, 28 (3), 609–631.

Lin, Po-Hsuan, Alexander L Brown, Taisuke Imai, Joseph Tao-yi Wang, Stephanie W Wang, and Colin F Camerer (2020) "Evidence of general economic principles of bargaining and trade from 2,000 classroom experiments," *Nature Human Behaviour*, 4 (9), 917–927.

Plott, Charles R (2008) "Properties of disequilibrium adjustment in double auction markets," in Plott, Charles R and Vernon L Smith eds. *Handbook of Experimental Economics Results, Volume I*, Chap. 1, 16–21: Elsevier.

Sargent, Thomas J (1993) *Bounded Rationality in Macroeconomics*: Oxford University Press.

Smith, Vernon L (1962) "An experimental study of competitive market behavior," *Journal of Political Economy*, 70 (2), 111–137.

Vriend, Nicolaas J (2000) "An illustration of the essential difference between individual and social learning, and its consequences for computational analyses," *Journal of Economic Dynamics & Control*, 24 (1), 1–19.

Wilson, Robert B (1987) "On equilibria of bid-ask markets," in Feiwel, G ed. *Arrow and the ascent of modern economic theory*, 375–414: Palgrave Macmillan.

APPENDIX

A Empirical CDFs comparison

Fig. 6 complements the empirical CDF illustrations for efficiency and average prices presented in Fig. 4 in the main text for the SFU3 experiment. Its six panels display the empirical CDFs for the data in comparison with the ZI, IEL and IEL-CDA models for the remaining measures reported in Table 1, refer to the caption for details. Similarly, Fig. 7 complements Fig. 5 in the main text for the SFU1 experiment. The averages of these distributions are reported in Table 2.

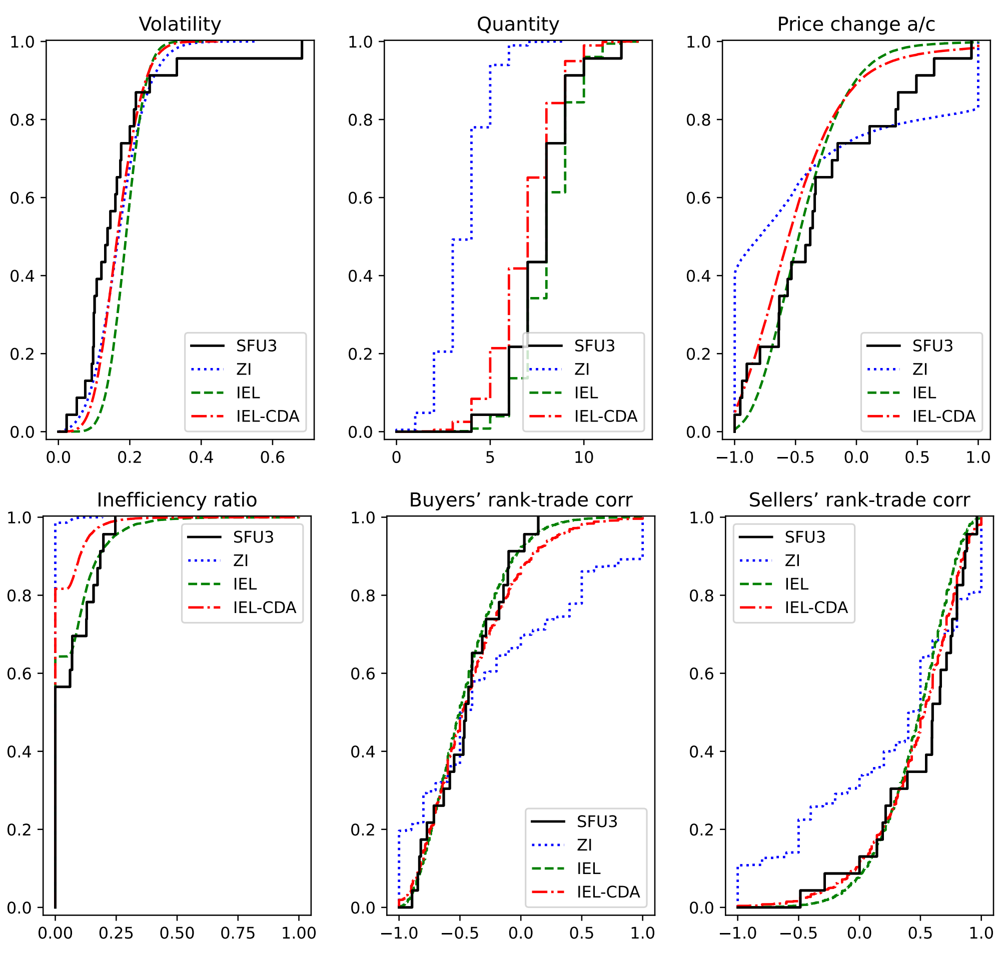

Figure 6: Empirical CDF of volatility (Smith's alpha, α), quantity (Q), price change autocorrelation (ρ_P), inefficiency ratio (I), buyers' and sellers' rank-order correlation for the transactions (ρ_B and ρ_S) for the SFU3 experiment (black solid curve) and three models: ZI, IEL, and IEL-CDA.

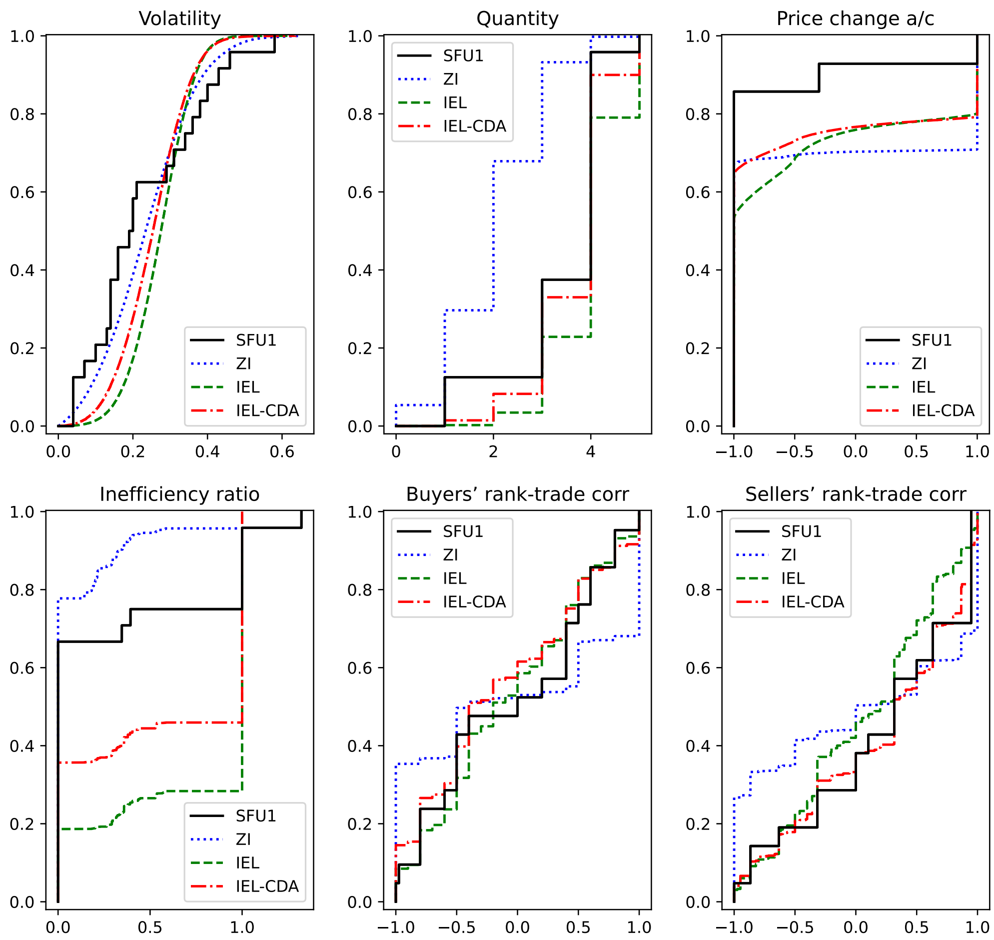

Figure 7: Empirical CDF of volatility (Smith's alpha, α), quantity (Q), price change autocorrelation (ρ_P), inefficiency ratio (I), buyers' and sellers' rank-order correlation for the transactions (ρ_B and ρ_S) for the SFU3 experiment (black solid curve) and three models: ZI, IEL, and IEL-CDA.

B Supplementary simulation results

We study the effect of Marshallian selection (MS) **IEL-CDA-1** and different values of the memory parameter L, which is part of the new utility function in **IEL-CDA-2**, on different measures and compare them with experimental data.

Table 6 reports the average values of statistics introduced in Table 1 for uniformly random (column MS is 'no') and Marshallian (column MS is 'yes') selection, and for six different values of L, ranging from 0 to 5. Note that no Marshallian selection with $L = 0$ is the specification of IEL-CDA that is closest to the IEL model. The only difference is in the strategy specification (IEL-CDA is in markups and markdowns; IEL is in offers), which affects the mutation and subsequent selection. In all simulations, the number of draws is chosen so that the average number of offers is close to this number in the experimental data. See columns D and O for the values of these variables that we use.

We note that for any value of L, Marshallian selection results in higher efficiency, a slightly higher average price, and a larger traded amount than uniform selection. Moreover, just from construction, it leads to higher (in absolute values) rank-order correlations for both buyers and sellers. Also, the values of the inefficiency ratio (proportion of loss of efficiency due to extra marginal trades) and Smith's alpha (measuring deviation from the equilibrium price) seem to be close or sometimes higher in the case of Marshallian selection. There is no clear effect of Marshallian selection on the price correlation.

For a fixed selection (either uniform or Marshallian), we observe that as the algorithm starts conditioning utilities on past transaction prices (so that L increases from 0 to 1), efficiency, quantity traded, and the absolute values of rank-order correlations drop significantly. Then, as L increases from 1 onwards, all these statistics start to increase. Quantitatively, the case of Marshallian selection and $L = 2$ that we use in the body of the paper matches both the SFU3 and SFU1 data reasonably well. Looking at the p-values, we can see that the model with no Marshallian selection is rejected by all our data, as well as the IEL-CDA model with $L = 0$, $L = 1$, and $L = 5$. Only the IEL-CDA model with $L = 2$ and $L = 3$ is strongly supported for both the SFU3 and SFU1 data.

MS	L	D	O	E	P	Q	ρ_B	ρ_s	I	α	ρ_P	p-val E	p-val P
SFU3 experiment			40.0	70.2	106.2	7.7	-0.44	0.50	0.06	0.17	-0.30		
no	0	40	40.0	79.8	105.2	8.9	-0.50	0.50	0.12	0.18	-0.43	0.00	0.01
no	1	40	40.0	53.5	107.7	5.1	-0.38	0.38	0.01	0.17	-0.40	0.00	0.37
no	2	40	40.0	58.1	107.7	5.7	-0.41	0.41	0.02	0.17	-0.46	0.00	0.35
no	3	40	40.0	60.5	107.6	6.1	-0.43	0.42	0.02	0.17	-0.46	0.00	0.35
no	4	40	40.0	62.0	107.5	6.3	-0.44	0.43	0.03	0.16	-0.45	0.02	0.31
no	5	40	40.0	62.8	107.5	6.4	-0.45	0.44	0.03	0.16	-0.44	0.03	0.32
yes	0	40	40.0	92.2	105.1	11.0	-0.58	0.67	0.28	0.18	-0.46	0.00	0.00
yes	1	40	40.0	61.8	106.8	6.0	-0.42	0.46	0.01	0.18	-0.38	0.01	0.24
yes	**2**	**40**	**40.0**	**67.8**	**106.6**	**6.8**	**-0.45**	**0.50**	**0.02**	**0.17**	**-0.48**	**0.41**	**0.22**
yes	3	40	40.0	71.2	106.5	7.3	-0.47	0.53	0.03	0.17	-0.47	0.24	0.21
yes	4	40	40.0	73.2	106.3	7.6	-0.49	0.54	0.04	0.17	-0.46	0.07	0.18
yes	5	40	40.0	74.5	106.3	7.8	-0.50	0.55	0.05	0.17	-0.46	0.03	0.19
SFU1 experiment			24.0	81.0	65.0	3.4	-0.06	0.21	0.29	0.23	-0.81		
no	0	26	24.4	87.6	65.3	4.0	-0.12	0.11	0.79	0.27	-0.44	0.07	0.86
no	1	25	24.8	70.3	66.2	3.0	-0.05	0.07	0.34	0.25	-0.39	0.00	0.93
no	2	25	24.8	73.3	66.1	3.2	-0.06	0.09	0.40	0.25	-0.45	0.01	0.94
no	3	25	24.7	74.1	66.2	3.2	-0.07	0.09	0.42	0.25	-0.45	0.01	0.93
no	4	25	24.7	74.1	66.2	3.2	-0.07	0.09	0.42	0.25	-0.42	0.01	0.93
no	5	25	24.7	74.1	66.2	3.2	-0.07	0.09	0.42	0.25	-0.42	0.01	0.93
yes	0	26	23.0	94.0	65.3	4.3	-0.23	0.31	0.96	0.28	-0.43	0.00	0.77
yes	1	25	24.6	82.1	65.7	3.5	-0.14	0.23	0.48	0.26	-0.40	0.14	1.00
yes	**2**	**25**	**24.4**	**85.5**	**65.8**	**3.7**	**-0.16**	**0.24**	**0.58**	**0.25**	**-0.50**	**0.40**	**0.99**
yes	3	25	24.3	86.3	65.6	3.7	-0.17	0.25	0.61	0.25	-0.48	0.48	1.00
yes	4	25	24.3	86.2	65.7	3.7	-0.17	0.25	0.61	0.25	-0.48	0.47	0.99
yes	5	25	24.3	86.2	65.7	3.7	-0.17	0.25	0.61	0.25	-0.47	0.47	0.99

Table 6: Comparison of the measures in the experimental data and the IEL-CDA algorithm.

C Online Appendix: Instructions and Screenshots

Below, we provide the instructions for the SFU1 experiment. The instructions for the SFU3 experiment are analogous.

Experiment Instructions

Trading on Flexemarkets

Welcome to this experiment! This is an experiment in trading. You have a very simple role to play. You are either a buyer or a seller. You will be buying or selling through an online trading market, the details of which will be provided soon. You will trade against other human subjects and trading robots. Before the experiment starts, please listen to the experimental instructions carefully. If you follow the experimental instructions and make the right decisions, you will receive additional income. The extra income is determined by your decisions and that of other experimental participants. Items and trades are valued in *Experimental Currency Unit* (ECU). Earnings will be accumulated and paid at the end of the experiment. In this experiment, 1 ECU worth 5 Canadian Dollar.

After listening to the experiment instructions, you will take a short quiz that assesses your understanding on the market structure. The quiz link will sent to you after the instructions video. In this quiz, you are also asked to write your participant number, which will be assigned to you at the beginning of the experiment, and your e-Transfer email for payment.

Please remain quiet during the experiment. Please do not talk to other participants or exchange information in other ways during the experiment. Please do not use communication tools such as mobile phones. If you have any questions or need any help, please use the private chat.

The Marketplace

This experiment will be conducted at a market on the Flexemarkets website. It can be found at https://flexemarkets.com/signin. You will need an account name, an email address, and a password to join the marketplace. These will be provided to you before the experiment starts.

The market is organized as a Continuous Double Auction market with an electronic book. Trading will occur in rounds. In each round you will be either a buyer or seller. Buyers submit bids to buy. This is an amount they are willing to pay for 1 unit of the commodity. Sellers submit asks to sell. This is an amount they are willing to receive for 1 unit of the commodity. **You will only be able to bid/ask for 1 unit of the commodity at a time.**

The following figures display the Flexemarket interface you will use in today's experiment. The upper (lower) figure displays the interface for buyers (sellers).

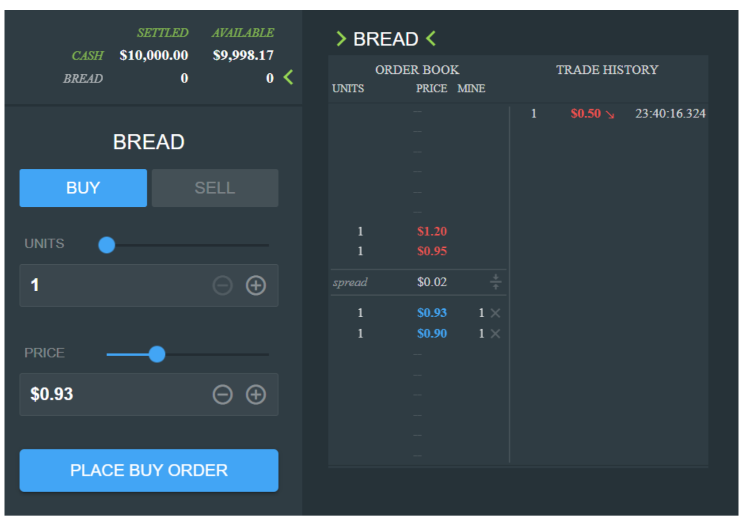

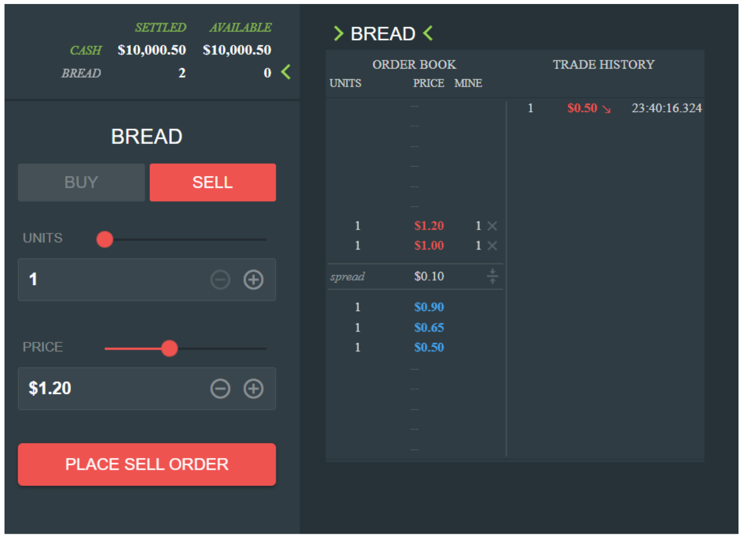

BREAD is being traded. CASH and BREAD holdings are displayed on the upper left part of the screen.

Bids and asks are submitted using the left side of the screen. To submit a new bid/ask, set the price and then click on PLACE BUY ORDER/SELL ORDER button. *Note: In these instructions, "buy order" = "bid" and "sell order" = "ask".* You can set the price, that you want to bid or ask by, using either the slider or the +/− buttons. In this experiment, you will not be allowed to submit orders for multiple units, so the units will always be set to 1.

As bids and asks are submitted, they are stored in the ORDER BOOK section which is located in the middle. Bids will be listed in blue from high at the top of the list to low at the bottom, and will be located under the asks. Asks will be listed in orange from low at the bottom of the list to high at the top of the list, and will be located above the bids. Your bids/asks are identified there with a 1. The SPREAD indicates the difference between the lowest ask and highest bid. The spread is located in the middle of the bids and asks.

A TRADE occurs when either a buyer submits a bid that is *higher than the lowest ask* in the book or a seller submits an ask which is *lower than the highest bid* in the book. The trade price is the price of the order in the book. In the first image above, if the buyer submits a bid of $1.50, it is larger than the lowest ask of $0.95 and a trade occurs at the price of $0.95. That is the seller gives 1 unit to the buyer in return for $0.95. In the second image, if the seller submits an ask of $0.50, then a trade occurs at the price of $0.90. The bid and the ask of the trade are removed from the book.

Past transactions and their prices are located to the right under TRADE HISTORY section. Your trades there are identified with a 1 at the left of that trade.

If you have a bid/ask in the ORDER BOOK that you want to CANCEL, you should click on the "x" button placed to the right of that order.

How you make money

The Buyers

Each buyer starts the experiment with no units. Each buyer will be able to buy only 1 unit each round. Each buyer will be given a valuation that specifies how much the experimenter will pay them for each unit they hold at the end of the round. The buyer will keep the difference between the value of the unit and the price paid for it.

Example #1: Suppose your valuation is 2.00. This implies that you will receive 2.00 for the next unit you buy. Suppose that during a round, you bought a unit for 1.25. Then, you will receive 2.00-1.25 = 0.75 ECUs for that round.

You can bid and pay more than the value of an item to you, but you will lose money if your bid is accepted.

Your valuation will change from round to round.

Example #2: Suppose your valuation is 0.50. This implies that you will receive 0.50 for the next unit you buy. Suppose that during a round, you bought a unit for 1.75. *In this scenario, you are paying higher than your valuation implying that your resulting payoff will be* ***negative***. Particularly, you will receive 0.50-1.75 = -1.25 ECUs for that round.

The Sellers

Each seller starts with 1 unit each round. Each seller will be given a cost value that specifies how much the experimenter will charge them for each unit they sell during the round. The seller will keep the difference between the price received for a unit and its cost.

Example #3: Suppose your cost is 2.00. This implies that it will cost you 2.00 for the next unit you sell. Suppose that during a round, you sold 1 unit for 5.00. You will receive 5.00-2.00 = 3.00 ECUs for that round.

You can ask and receive less than the cost of an item to you, but you will lose money if your ask is accepted.

Your cost value will change from round to round.

Example #4: Suppose your cost is 1.50. This implies that it will cost you 1.50 for the next unit you sell. Suppose that during a round, you sold 1 unit for 1.25. *In this scenario, you are receiving less than your cost implying that* *your resulting payoff will be* ***negative***. Particularly, you will receive 1.25-1.50 = -0.25 ECUs for that round.

The Experiment

Each experiment will consist of 14 rounds. The first two rounds are warm-up rounds. Your trades in warm-up rounds will not contribute to your earnings. After completing warm-up rounds, the actual experiment will start. The warm-up rounds will last 100 seconds. The remaining 12 rounds will last 75 seconds. The experimenter will announce the time left at various points in each round. At certain points during the experiment, you will switch to a new trading environment, the details of which will be provided in your **personal information sheet**. The experimenter will make announcements prior to switching to a new marketplace.

You can display your earnings (in ECU) from each round through an external website. The guideline for using this website will be provided in the next section.

You can submit bids/asks from [0, 1.50] interval. That is, it's not possible to bid/ask less than 0 or greater than 1.50.

You will receive an error message if
-you try to submit an order for multiple units
-you try to sell when you have 0 units of holdings
-your cash holdings are not enough to satisfy your order.

In each round, you are allowed to buy or sell only 1 unit. If you buy or sell more than 1 unit in a given round, you will be punished and your payoff will be deducted by the sum of trade prices of additional trades that you make after the first trade.

Example #5: Suppose in a given round, you traded 1 unit and your profit from this trade is 2.50 ECU. Suppose before the round ends, you make two more trades at prices of 1.00 and 1.50 and the total number of trades you make becomes 3. Since you violated the rules by trading more than a single unit, your payoff will be deducted by 1.00+1.50=2.50 ECU and the profit you make in this round will be 2.50-2.50=0 ECU.

Earnings Displaying

The Flexemarkets website does not allow subjects to see their earnings. In order to display per round earnings, a separate website will be used. This website can be accessed at https://flexemarkets-earnings-display.netlify.app/. In order to login, you will use the same login information that you are provided for the Flexemarkets login. The left and the right sides of the below figure show the login page and the main that you will be redirected upon a successful login, respectively.

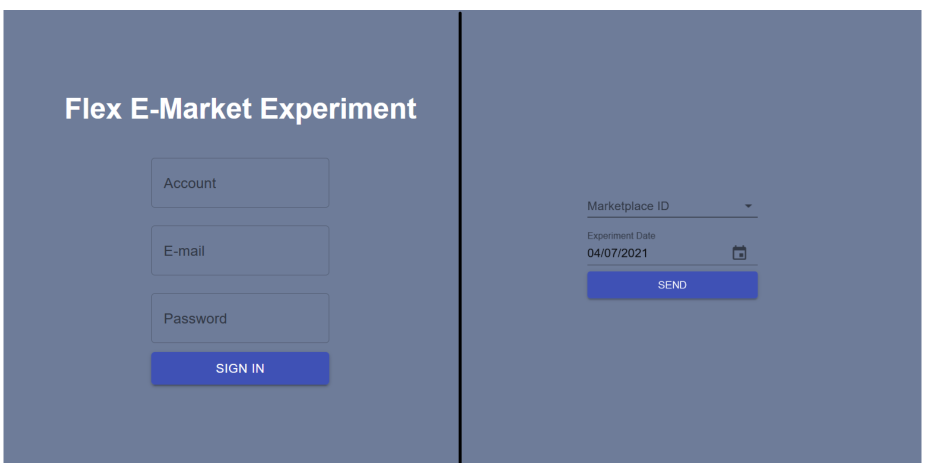

To display your profit in a given round, you need to input two things: the experiment date and the marketplace ID that you are currently trading in. Then, you should click on "SEND" and this will redirect you to profit displaying page. If you want to display your profit before any trading round is held in a given marketplace, then you will see the left hand side of the below figure. It tells you that the marketplace has not opened yet in this marketplace. If you display your earnings after the marketplace is open, then you will see the right hand side of the below figure. If you click on "Profit" button, your current earnings will be displayed.

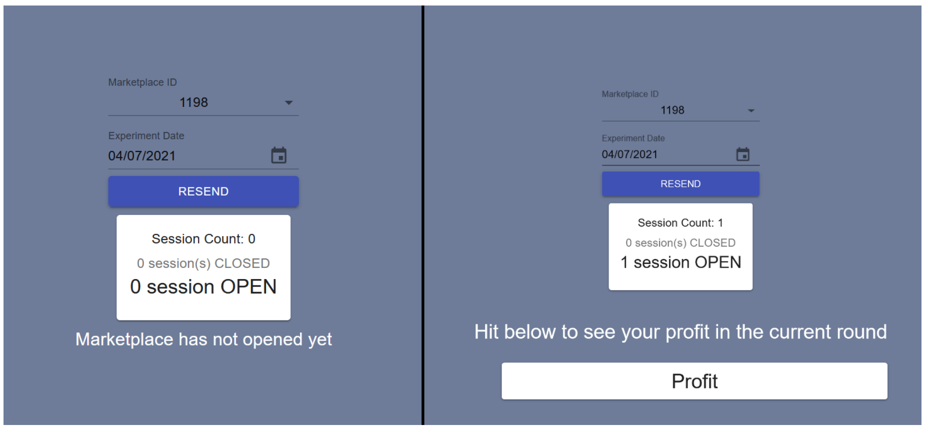

Your profit will start from 0 at the beginning of a trading round (left side of the below figure). As you make trades, your profit will be updated (right side of the below figure).

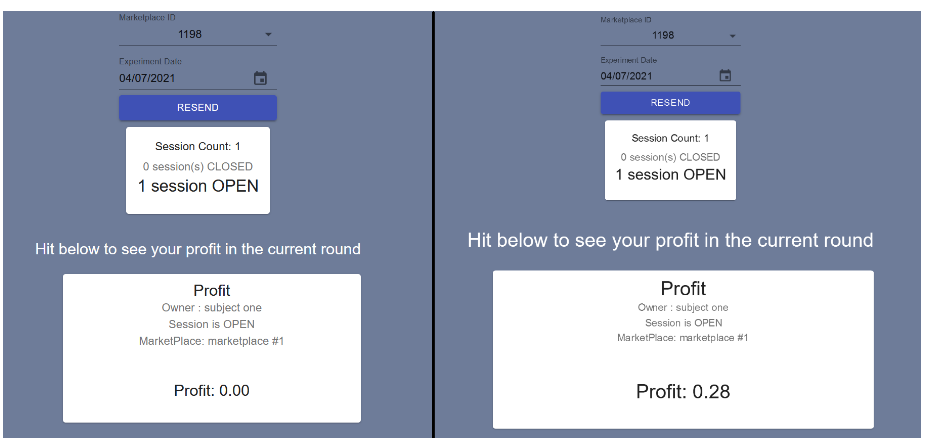

When you complete trading rounds in a given marketplace, the experimenter will make an announcement to remind subjects that they should switch to a different marketplace. When this happens, you should select the ID of this new marketplace from the "Marketplace ID" list to display the correct earnings.

You can use this external website to track your earnings in two possible ways. First, you can use it simultaneously with your trading activity. If you choose to do that, then you need to switch between trade page and earnings displaying page as you make a trade. Second, you can use it at the end of each trading round. If you choose to do that, then the displayed earnings will be from the last trading round. Recall that there will be 60 seconds break between each trading round. During the break time, the marketplace will close and there will be no trading activity. This will give you enough time to check your earnings from the last trading round. For example, suppose you completed the first trading round and now you want to see your earnings from this trading round during the 60 seconds break until the second trading round starts. If you choose to do that, you will be able to see your total profit from the last trading round, or the first trading round in this example. **We suggest you to follow the second way as this would minimize the distraction that could be caused by earnings displaying while a trading round still continues.**

Recall that the displayed earnings are in Experimental Currency Unit. In order to calculate your actual profit, you should use the conversion rate provided at the beginning of experimental instructions. The experimenter will also inform you regarding your overall profit in CAD at the end of the experiment.

An Extended Example

There is only 1 buyer and 1 seller who are trading BREAD. After the period starts, the seller submits an ask of 1.00.

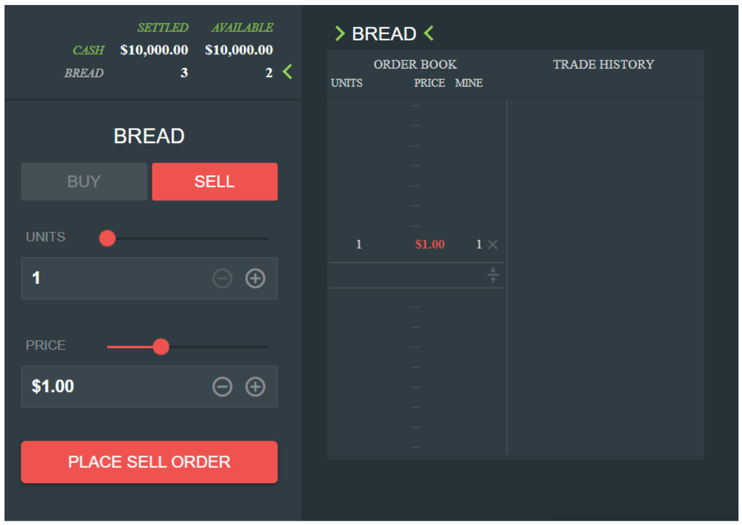

Next the buyer submits a bid of 0.50. Since this is smaller than the lowest ask available in the book, 1.00, there is no trade.

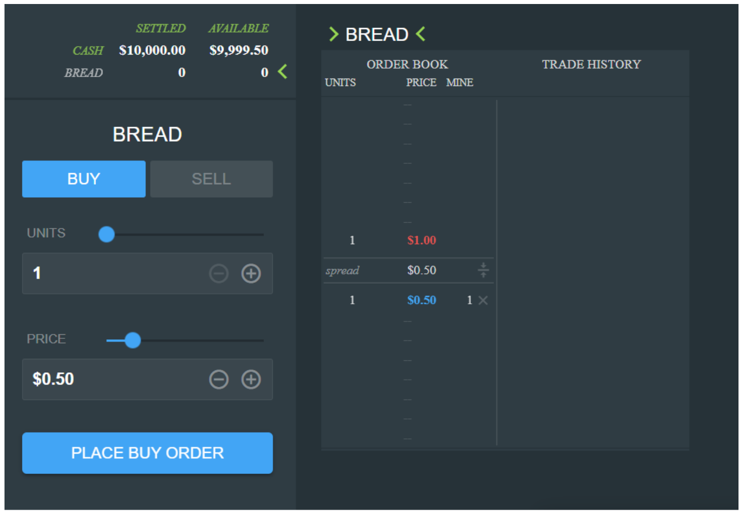

Next the seller submits a new ask of 0.40. This will be the new lowest ask in the book since it is less than the initially submitted ask price of 1.00. Furthermore, this ask is smaller than the currently highest bid of 0.50 in the book. Therefore, a trade occurs at the bid price of 0.50.

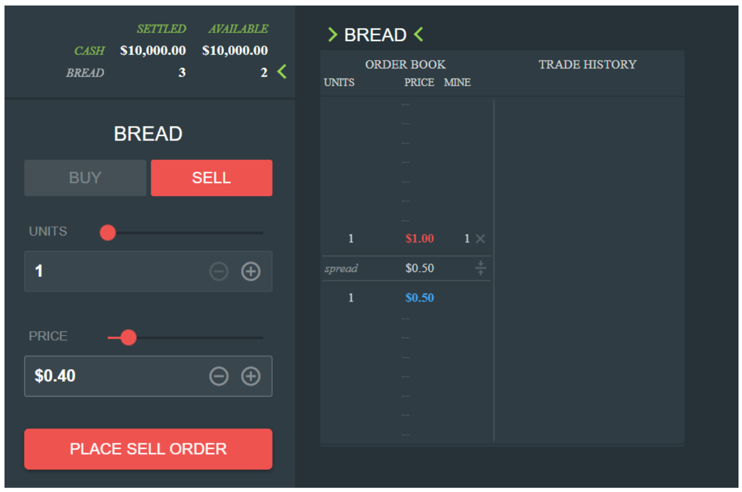

Note that the seller's initial ask remains in the book after the trade.

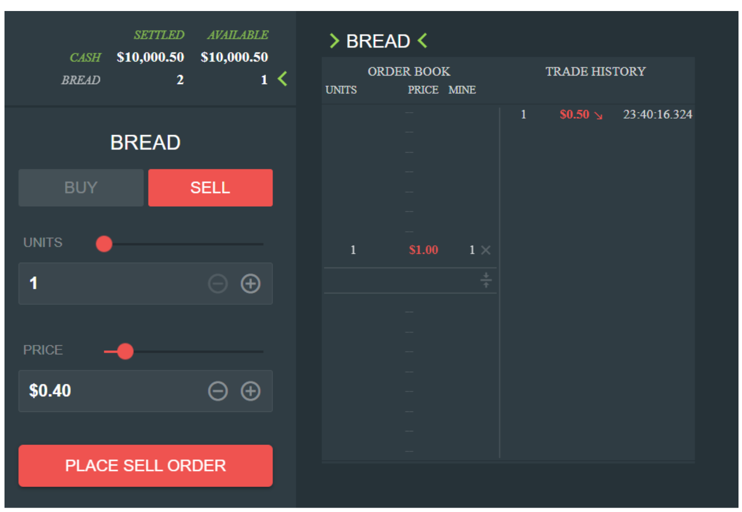

Over the next few seconds, the buyer and seller submit new buys and asks, none of which generate a trade as the lowest ask price (1.00) is still higher than the highest bid price (0.90).

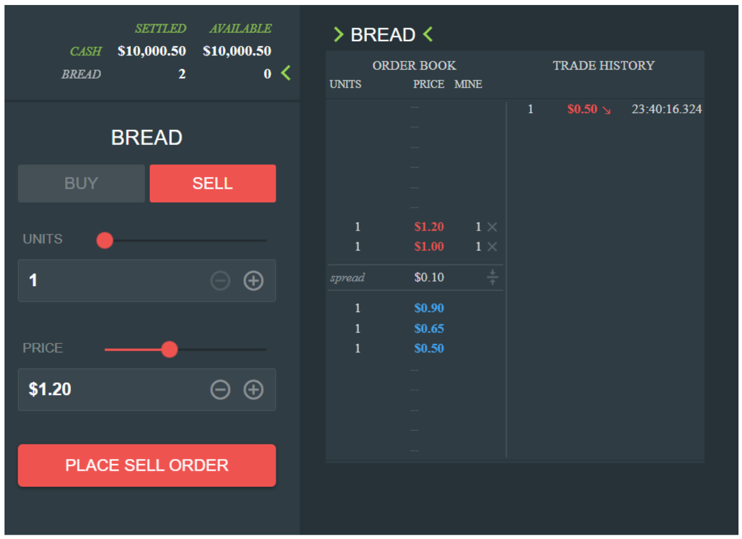

Next the seller cancels the initial ask at 1.00 and submits new ask at 0.95. The buyer then cancels two of her bids (0.50 and 0.65), and submits a new one at 0.93. Again, since the lowest ask price of 0.95 is greater than the highest bid price of 0.93, there are no trades.

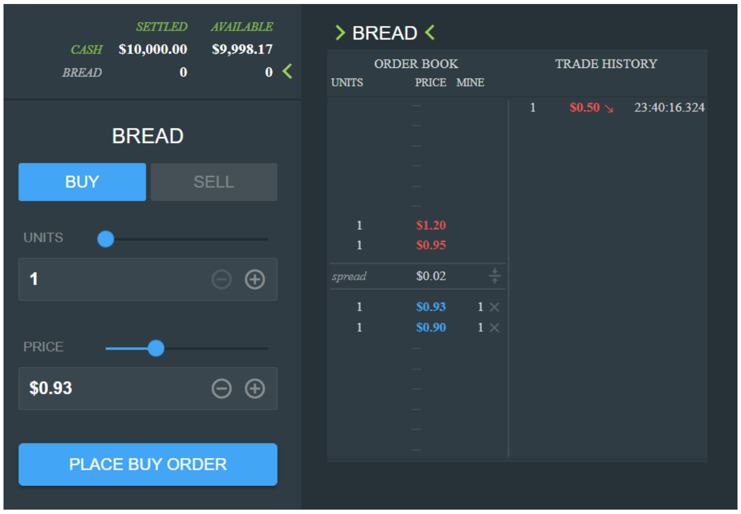

Near the end of the period, the buyer submits a new bid of 1.30.

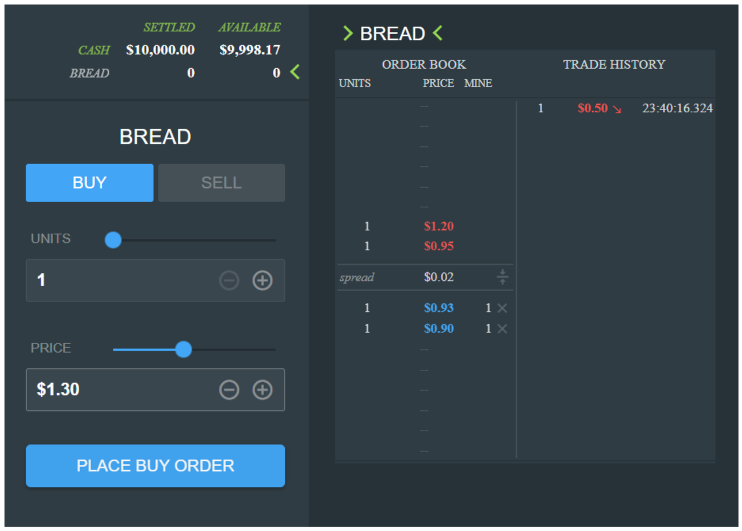

Since this is greater than the currently lowest ask of 0.95, a trade occurs at the price 0.95.

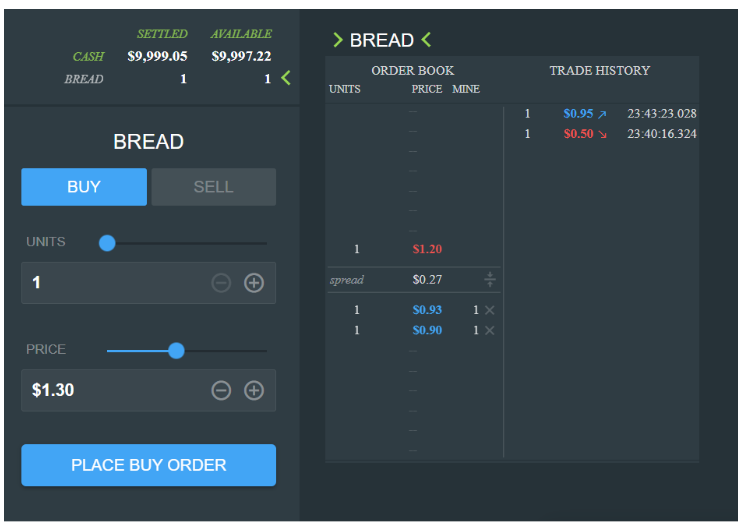

When time runs out, the market will be close and you'll see the following message on your screen.

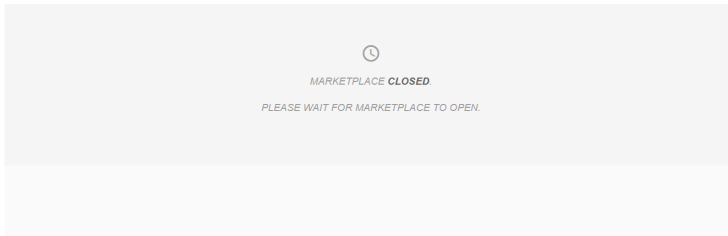

You will then need to wait until the next round starts.