Opening the Book: Price Information's Impact on Market Efficiency in the Lab

Brett Williams*

Abstract

Human trader behavior and market convergence are studied in a general equilibrium two good setting through the use of the continuous double auction. The set of active bids and asks shown to traders in a session is either (1) only the best bid and ask in the market or (2) the full set of active orders; similarly, the set of visible transactions spans the full history of the trading period to only a trader's own transactions. Laboratory markets reveal no improvement in allocative efficiency and price convergence when improving price/order accessibility symmetrically across the set of bids and asks, as well as the transaction history. Improvements in accessibility in only one factor leads to reductions in efficiency. Effects on price discovery related outcomes appear to be asymmetric across transaction history and orderbook improvements, with the order of increased visibility impacting market outcomes. Results are compared against standard simulated benchmarks.

Keywords: Continuous Double Auction, Experiment, General Equilibrium, Exchange Economy
JEL Classifications: C92, D44, D47, D51, D80

*Ph.D. Candidate, Department of Economics at UC Santa Cruz, 1156 High Street, Santa Cruz, CA, USA 95064; bwillia4@ucsc.edu
Many thanks to my advisor Daniel Friedman for guidance and support. I am also grateful for comments and advice from Kristian Lopez Vargas, Natalia Lazzati, Duncan James, Shuchen Zhao, members of the UCSC experimental workshop, members of the IFREE 2020 Graduate Student Workshop cohort, and attendees of my talks at ESA Global 2021, ESA North America 2021, and Yale's ZIMI2 conference.
Funding: Thanks to IFREE for funding the creation of the novel UI/platform used in this paper via their small grants program. The experimental process was funded by UCSC department of economics' graduate research grant.

1 Introduction

Markets provide a service to their agents by creating, receiving, aggregating, and disseminating information. The structure and rules that define a market determine the type of information and how and when this information is transferred or received. One such piece of information, potentially the most integral to a market, is the price associated with an order. If one considers the set of orders in a market, their associated prices, and the structure the market places on how these orders are presented to the traders, one can order a market based on the amount of price/order information is given to the traders. Naturally, this means there is some minimum and maximum amount of price information accessible to traders. An understanding of the performance of structures that yield differing levels of price accessibility is thus crucial, as new markets are rapidly appearing and old markets are ripe for improvement. This paper aims to investigate prominent information structures in one of the most popular market institutions, the continuous double auction (CDA), and document the efficiency and trader behavior associated with each.

Given the question of *which level of information accessibility is best*, a common and natural response might be *maximum accessibility*. Why not give all of the price and order information to traders in the market? Shouldn't this make the market as efficient as possible? In response to the former, information inclusion may not be cost-less, either financially to the central body, or behaviorally to the traders. It may be the case that a structure which yields less accessibility may be just as effective as the maximum.

To the latter, while potentially true, this is not fully known. The efficiency associated with each potential structure has not been fully mapped, and the relationship between accessibility and efficiency may not be strictly positive. Take, for example, the introduction of the Openbook subscription software by the New York Stock Exchange (NYSE) in early 2002. The platform released order prices and quantities for the full book, as well as transactions, to traders away from the trading floor (information which previously had not been accessible in such a state). Boehmer et al. (2005) analyzed trading data from NYSE before and after the introduction of Openbook, finding a significant increase in cancellation rate and reduction in order size. A similar platform adjustment in the Toronto Stock Exchange a few years earlier, though with less of an increase in information accessibility (only the best bid and offer, as well as the depth, were made transparent), revealed wider post-platform spreads and increased volatility (Madhavan et al. (2005)).

Evidence of changing market outcomes due to adjustments in price accessibility is also seen in markets outside of finance; one example is the Kerala fishing market. Prior to a massive mobile phone rollout, studied by Robert Jensen (2007), the market experienced high levels of price volatility, likely due to the lack of knowledge about the prices on land and the amount of fish being supplied by the fisherman on the lake or being bought at the market. The access to phone service allowed for easier access to on-shore prices across multiple small fishing markets, rolled out in three waves to three separate lake regions. In each case, the rollout of the phone service provided (nearly) immediate, drastic reductions in price volatility. Despite the difference in trader types (specialized versus two-way), good type (durable or non-durable), or good unit(single, divisible, or multiple), adjustments in price accessibility lead to meaningful (though not always beneficial) adjustments in market performance.

This project explores the impacts of different levels of accessibility on market outcomes in a controlled environment through a series of laboratory CDA markets. Two-way traders induced with utility preferences through a novel interface (first implemented in Crockett et al. (2021)) trade in a two-good Edgeworth box general equilibrium setting, with markets varying in their level of orderbook and transaction history price accessibility. Most markets reveal moderate levels of convergence in price and allocation, with symmetry in accessibility (low in the book and the history or high in both) being important for outcomes such as allocative efficiency. Asymmetrically accessible markets hamper efficiency levels in exchange for more intense price discovery behaviors, including much higher order and trader frequencies. This experiment adds to a vast and well-known experimental market literature.

A long history of tests of market equilibrium behavior in the lab exists, starting with Chamberlin (1948) and Smith's (1962) seminal oral outcry markets. The two experiments differed in their outcomes, with the latter finding convergence to equilibrium predictions and the former not. Aside from a difference in formalism, the main reason appears to be the better access to price info in Smith's version.

The laboratory markets literature, particularly on simple versions of the CDA or call market, has flourished since then, with hundreds of market experiments being run over the past 50 years. The bulk of the literature resides in the partial equilibrium (PE) space, with simple one-way (specialized) agents trading single units of a single good based on cost and redemption value schedules. Among the expanse of PE papers, this paper relates to those investigating adjustments in market information and its relation to efficiency and price formation. Smith (1980) tests the existence of complete information (for value and cost schedules among traders) in a series of experiments with supply and demand schedules yielding extreme asymmetry in potential gains from trade. Results suggest the increase in information leads to inconsistent occurrences of convergence (contrary to the consistent convergence of earlier markets with incomplete information). Inspired by these findings, Kimbrough and Smyth (2018) provides a replication of a market similar to that from Smith and Williams (2000), testing complete and incomplete information. The papers finds the existence of complete information is not enough to cause deviation from competitive equilibrium, but adding in symmetric market power along with complete information is enough.

Even closer to the sentiment of the paper I present now, a couple papers test the information present in the orderbook. Kirchsteiger et al. (2005) endogenize the accessibility of the markets in their experiments, allowing subjects to choose which traders on the same side and opposite side of the market have access to their orders. Ikica et al. (2018) tests numerous market formats across hundreds of experimental markets, a subset of which test the difference in efficiency between full orderbook and transaction history accessibility and a black-box setting. Both papers suggest the accessibility of order or trade prices may have substantive impacts on market outcomes. Arifovic and Ledyard (2007) test call markets with either a closed and open orderbook using simulated traders, finding closed book markets outperform open book markets in both efficiency and price volatility.

Studies centered around market transparency or information accessibility are also prominent in a more complex market format, namely dealer markets. These markets are comprised of standard one-way traders, as well as market makers who set orders on both sides of the market via a spread and help provide liquidity to the market. Laboratory experiments in this literature (Bloomfield and Libby (1996), Pagano and Röell (1996), Flood et al. (1999), Bloomfield and O'Hara (1999)) generally agree in their findings (transparency provides more desirable market outcomes), while empirical studies, including the ones discussed above (Madhavan et al. (2005) (TSE), Boehmer et al. (2005) (NYSE), Board and Sutcliffe (1995) (LSE)) tend to have conflicting findings.

Another notable, though much smaller, sub-thread of experimental market papers that this project contributes to is the general equilibrium (GE) sub-literature. Early works naturally followed suit with the PE experiments, providing extensions close in sentiment to Smith (1962). Williams et al. (2000) induce buyers with constant elasticity of substitution (CES) preferences across two markets with two batches of sellers (driven by cost schedules), while Plott provides a GE replication of Smith's original experiments in Plott (2000) and follows up with a multi-market study of his own in Plott (2001) (experimentally applying the setting of Gale (1963)). Early theoretical contributions in the space received GE experimental attention from other projects as well. Anderson et al. (2004) and Goeree and Lindsay (2016) experimentally test the unique setting presented in Scarf (1960), and Crockett et al. (2011), much like Plott (2001), pay respects to Gale (1963) via a series of experimental tests. Much like many papers in this literature, the laboratory markets I run are situated in a two good Edgeworth box economy, providing the first test of price accessibility adjustment in this simple GE setting.

A new GE expansion of classic CDA trader behavior is also modelled in this paper. The first wave of such models appeared in the late 80's and early 90's. Wilson (1987) began the influx with likely the most complex model of the bunch, modelling the continuous double auction game theoretically through the use of bilateral bargaining dynamics. Friedman (1991) and Easley and Ledyard (1993) follow Wilson with simpler models, both equipping traders with reservation prices,. Friedman positions traders as playing a Bayesian game against nature and Easley and Ledyard assume traders' reservations adjust over the course of a period to their true valuations. Gjerstad and Dickhaut (1998) follow up with a similarly non-strategic model, with players playing in essence against nature, though with traders updating their beliefs on trade success in a frequentist manner as opposed to Bayesian (as in Friedman (1993)). Newer simulation based models have appeared in the last couple decades, including the individual evolutionary learning model (IEL) of Arifovic and Ledyard (2011) and Anufriev et al. (2013), Crockett and Oprea (2012) reference dependence model, and the timing-focused expansion of IEL in van de Leur and Anufriev (2018).

A closely related group of papers containing minimal intelligence agent-based models became popularized after Gode and Sunder's (1993) zero intelligence (ZI) model was introduced. The model provided agents with entirely random order choice processes (in the commonly used ZI-C version, these were given slightly more guidance via a budget constraint which restricts price submissions to weakly surplus increasing options), asserting that the efficiencies found in simulations were thus driven entirely by the structure of the CDA. Several papers proposed adjustments to the model, either slightly increasing the intelligence of the traders (e.g. profit margin targeting in the ZI-P model of Cliff and Bruten (1997)), or adjusting an attribute of the market format (e.g. the addition of an orderbook in Bollerslev and Domowitz (1993)). Models providing slightly more intelligence have thus been deemed as having traders with minimal intelligence. General equilibrium extensions of ZI have also been proposed, including Gode et al. (2004) and Crockett et al. (2008), with Hurwicz et al. (1975) possibly being an early predecessor. I use the model of Williams (2021), another extension of the sort, though adjusted for more complex rules within the standard two-good Edgeworth box, as a benchmark for this paper's human markets to compare against.

The rest of the paper continues as follows. Section 2 lays out the environment as well as a newly adjusted general equilibrium agent-based model build on Gjerstad and Dickhaut

(1998). Section 3 articulates the methodology for the human laboratory experiments. Section 4 presents price and allocation adjustments, efficiencies, and an adjusted version of Gode and Sunder's (1993) ZI agent-based model. Section 5 concludes the paper.

2 Environment

This section provides the market design and microstructure employed in this paper, as well as an agent based algorithm for general equilibrium trader behavior in a continuous double auction.

2.1 Two-Good Edgeworth Box Economy

This paper highlights a market containing a set N of traders who partake in the buying and selling of two non-durable goods, x and y. Each trader is endowed with some non-negative amounts of both x and y to begin each period of trading. Traders each have their own utility function, which is monotonically increasing and twice differentiable in both goods.

Each trader is allowed to act as both a buyer and a seller in the market within a trading period. The set N is partitioned between two subtypes of these two-way traders, namely the set of "natural" buyers B and the set of "natural" sellers A. "Natural" buyers, in this sense, are traders with a marginal rates of substitution at their endowment point that is higher than the competitive equilibrium price. An analogous definition holds for "natural" sellers.

Each trader i's objective is to maximize her utility given her budget constraint, for some prices p_x and p_y over single units of x and y and endowment m. The budget constraint can be simplified assuming y is a numeraire, yielding the constraint $px + y = \hat{m}$ where p is p_x/p_y (the price of a unit of x in terms of units of y) and $\hat{m} \equiv m/p_y$.

$$\max_{(x_i,y_i)} u_i(x_i, y_i) \quad s.t. \quad px_i + y_i = \hat{m} \tag{1}$$

For the purpose of this paper, model and the accompanying experiments, I allow the traders to have constant elasticity of substitution preferences $u_i(x_i, y_i) = c_i((a_i x_i)^{r_i} + (b_i y_i)^{r_i})^{\frac{1}{r_i}}$. At the beginning of trade, as well as after each trade in the market, trader i's excess demand can be defined as

$$Z_i^X(p|(x_{i,o}, y_{i,o})) = \frac{a^{\gamma_i}(y_{i,o} + p x_{i,o})}{p(a^{\gamma_i} + p^{\gamma_i} b_i^{\gamma_i})} - x_{i,o} \tag{2}$$

where $(x_{i,o}, y_{i,o})$ is the initial bundle of trader i and $\gamma_i = \frac{r_i}{1-r_i}$. Solving

$$Z^X(p|(x_o, y_o)) = \sum_{i=1}^{N} Z_i^X(p|(x_{i,o}, y_{i,o})) = 0 \tag{3}$$

yields p^*, a competitive equilibrium price. Plugging this back into each trader's Z_i^X gives their desired change in x, determining the net trades in competitive equilibrium.

2.2 Continuous Double Auction

The market type of choice in this paper is likely the most prolific both academically and in practice, the continuous double auction. Traders in this institution may actively submit or accept orders at any time, so long as their allocations can accommodate the trade(s). An order in this setting consists of price, quantity and time fields; in this paper, I delegate the time choice to be the length of the market's trading time (unless the trader wishes to cancel or replace it). As mentioned above, each trader can participate on both sides of the market, submitting both bids and asks at their leisure. Orders placed in the book in this market that do not immediately cross with an existing order are analogous to a limit order with no expiration time. Additionally, market orders can be mimicked in this market through the ability to instantly accept orders in the book.

Contrary to the majority of the previous lab experiments and many CDA trader behavior models, orders may be non-unitary in both senses of the word: (1) the quantity field accepts values larger than 1, and (2) non-integer values (i.e. partial units) are acceptable. A second order characteristic present in this market that is not overly common is the existence of retrade. Both X and Y can be “retraded” without limit, unimpeded by traditional unit ordering restrictions.

2.3 Agent Algorithm

The model described in the following subsection is set in the environment laid out above. The algorithm and beliefs which drive the behavior in this model are based heavily on those found in Gjerstad and Dickhaut (1998), with adjustments made to fit the model to the desired environment.

The general environment matches that characterized above. A set of two-way traders, N, are partitioned into natural buyers, B, and natural sellers, A. Each trader has some endowment (x_i, y_i) of goods X and Y. Traders place orders over the length, T, of the period of trade, where orders are defined as 3-tuples containing a price, quantity and time message, $\{p, q, t\}$. The price p and quantity q elements of an order are chosen by the trader. The time t field represents the time in the trading period at which the order was placed. Orders essentially are infinitely lived, however can be removed from the exchange by either directly cancelling the order or by replacing it. As trader's in this algorithm do not cancel an order as a stand-alone action, this can only occur through order replacement. An order's p is bounded naturally below by 0 and artificially above by some real number M.[1] Similarly, the q field of an order is bound above by a trader's current allocation of goods if the order is an ask, or by the trader's allocation of goods divided by the p element of the order.

The set of orders that have been posted to the orderbook or transacted over the duration of the market (up until the current is Ω. The elements of Ω are indexed in terms of submission to the orderbook, with o_k being the k^{th} order placed in the orderbook.

Traders are guided by an algorithmic trading behavior which can be defined in four main steps: (1) entry, (2) belief updating, (3) order selection, and (4) wait time selection for the

[1] This is common in the agent-based literature for CDA trader behavior. The value of M can be relatively low, though the gap between equilibrium price and M is usually larger than the lower half of the price domain.

next entry.

Entry

Traders enter the market one at a time. At the inception of the market, all traders make an uninformed random draw of price and quantity. Price is drawn from $[0, MRS_i]$ if trader i is a natural buyer and $[MRS_i, M]$ if he is a natural seller.[2] Quantity is similarly drawn uniformly randomly.

The surplus (gained utility) associated with each (price, quantity) choice is multiplied by the probability it will be accepted. As no price information is available yet in the market, this is just p/M for bids and $(M - p)/M$ for asks. (Later, I will define these probabilities as functions of the relative acceptability of potential order prices, written as $p_a(a)$ for sell prices and $p_b(b)$ for buy prices.) After finding the expected surplus of each trader's randomly drawn potential order, traders make a decision on how long they will wait to make their first move. The trader with the lowest wait time enters. Once the entrant is scheduled to enter, the history (re)observed to establish a base for the entrants belief updating process.

Beliefs

Let Ω_H be the set of orders, and P be the set of submitted prices, contained in the trader's limited history of the market. For each price $\rho \in P$, trader i can check the total number of orders o in Ω_H whose p is equal to ρ. In Gjerstad and Dickhaut (1998), this count is denoted $TA(\rho)$. The set of orders satisfying such a constraint shall be called $\tau_A(\rho)$. This measure of traded orders equally weights all orders that have been filled or accepted in some manner. For the original model, this is appropriate, as the setting only allowed single unit orders and trader; however, the setting of this paper is much more general. To accommodate the idea of partially filled orders, I propose a weighted version of this count $TA(\rho)$. Each order, instead of receiving a guaranteed count of 1, receives a count of $\sqrt{q_k}\frac{q_{k,traded}}{q_k}$, where q_k is the original

[2] MRS_i refers to the marginal rate of substitution of trader i.

quantity of order o_k and $q_{k,traded}$ is the number of units accepted in the trade. $TA(\rho)$ can then be written as

$$TA(\rho) \equiv \sum_{k=1}^{|\tau_A(\rho)|} \sqrt{q_k} \frac{q_{k,traded}}{q_k} \cdot 1 \tag{4}$$

An analogous definitions is presented for bids in the remembered history at price ρ in the original model. $\tau_B(\rho)$ is the set of bids in the remembered history at price ρ, and $TB(\rho)$ is the weighted count of accepted bids at ρ. The weighted number of rejected asks (bids) at price ρ, termed $RA(\rho)$ ($RB(\rho)$), is defined similarly. Each rejected order is weighted again by the portion of the trade that was cancelled (and scaled by the square root of the size of the order).

With weighted versions of TA, TB, RA, and RB now defined, the beliefs of a trader on the acceptability of an order can be defined. Traders consider the set of orders in Ω_H which yield information on the previous performance of a certain price. Consider an ask at price a. The success of all asks at worse prices than a, the number of bids filled at prices above a (i.e. the bids of buyers revealing a willingness to pay higher than the price in question), and the failure of all asks at more competitive prices than a all reveal information to the trader. Thus, the function $p_a(a)$ can be defined (in the same form as Gjerstad and Dickhaut's original model) as the probability of success for an ask at price a:

$$p_a(a) = \frac{\sum_{\rho \geq a} TA(\rho) + \sum_{\rho \geq a} TB(\rho)}{\sum_{\rho \geq a} TA(\rho) + \sum_{\rho \geq a} TB(\rho) + \sum_{\rho \leq a} RA(\rho)} \tag{5}$$

The analogous function $p_b(b)$ for bids at price b is defined as

$$p_b(b) = \frac{\sum_{\rho \leq b} TB(\rho) + \sum_{\rho \leq b} TA(\rho)}{\sum_{\rho \leq b} TB(\rho) + \sum_{\rho \leq b} TA(\rho) + \sum_{\rho \geq b} RB(\rho)} \tag{6}$$

A trader who has entered the market with the intent to sell (buy) solves for $p_a(a)$ ($p_b(b)$) for each unique price contained in the orders of τ_A (τ_B). Since $p_a(\tau_A) \equiv \{p_a(a) : a \in \tau_A\}$ is a discrete monotonically increasing set[3], a piece-wise linear interpolation is used to complete the trader's beliefs on acceptability. These beliefs are over the domain [0,M], with $p_a(0) = 1$ and $p_a(M) = 0$.

Order Choice

Once the entrant has updated his beliefs over the acceptability of each of the prices in his remembered history on the side of the market he entered, he can go through the process of choosing an order to submit. As with the original model, the entrant's main goal is to choose an order that maximizes his expected surplus. Unlike the original model, this task is now more tedious, as utility is used instead of cost/redemption schedules and orders are not restricted to have single unit quantity.

The maximized expected surplus of the entrant, trader i, is written simply as

$$S_i^k = \max\{\max_{\rho \in P}(u_i(x_{i,k}, y_{i,k}) - u_i(x_{i,k-1}, y_{i,k-1})) \cdot p(\rho), 0\} \tag{7}$$

where k refers to the k^{th} (potential) transaction taken in the market (in some future time in the period, t_k). Here, the entrant's current utility is that which is associated with his (x, y) bundle in time t_{k-1} (the time of the last transaction in the market). Given the desire to improve utility, it is natural to restrict the entrant's considered prices to $[0, MRS_{i,k-1}]$, or the lower bound to the his current marginal rate of substitution, if entering with the intent to buy. If the entrant is selling, the domain becomes $[MRS_{i,k-1}, M]$.[4] The entrant considers a fine grid of prices in this domain.

For each price considered, the entrant must determine an appropriate quantity, before determining the expected surplus. To do so, he considers a fine grid of quantities from zero to his current holdings of x (or his current holdings of y adjusted by the price being considered if buying), call this quantity $\bar{q}$. He solves for the utility gained from having the order fully

[3] See Gjerstad and Dickhaut (1998) for discussion on characteristics of p(a), including monotonicity.
[4] The entrant's marginal rate of substitution in this sense can be considered a reservation price.

accepted at each point on the grid. This set of quantities can naturally be partitioned into quantities which yield utility gains and utility losses. Given the concavity of the trader's utility function, this partition occurs at a single price, call it $\tilde{q}$, where $\tilde{q} \in [0, \bar{q}]$. The set of utility weakly improving quantities is thus $[0, \tilde{q}]$. Within this set, a quantity, $\hat{q}$ exists which yields a maximum utility improvement, meaning after this quantity, such a choice would have decreasing marginal gains. In other words, any $q > \hat{q}$ would be considered over trading. As such, the final reduced set of quantities considered by the entrant is contained in $[0, \hat{q}]$. To make his decision, the trader chooses randomly from this final set of quantities, with the weight associated with each q being that q's relative utility gain (q's utility gain divided by the sum of all utility improvements for the quantities in the reduced set).

Now that the entrant knows the set of ordered pairs (p, q) to maximize over, he selects the order which yields S_i^k when fully accepted. Once selected, the order is submitted to the orderbook. If it crosses with an order(s) currently in the market, it will fill the order(s) until either all crossing orders are filled or the order the entrant posted has fully filled. This finishes the actions of the entrant and begins the process of determining the next entrant.

Timing

To determine the next entrant (and the associated elapsed wait time), each trader needs to re-evaluate the current landscape. All traders update their remembered histories of the market to account for the most recent action (taken in time t_k), and perform the belief updating and order choice process described above for both sides of the market. After doing so, each trader i has two ordered pairs in consideration, an ask which yields some maximum expected surplus $S_{a,i}^{k+1}$ and a bid yield a maximum expected surplus $S_{b,i}^{k+1}$. To determine the side of entry, each trader flips a weighted coin. Trader i's chance of entering on the sell side of the market is $\frac{S_{a,i}^{k+1}}{S_{b,i}^{k+1}+S_{a,i}^{k+1}}$, and $\frac{S_{b,i}^{k+1}}{S_{b,i}^{k+1}+S_{a,i}^{k+1}}$ for the buy side.

Elapsed wait times are determined via draws from trader specific exponential distributions. The distribution parameter for trader i is a function of the expected surplus for their proposed next order, defined as $\alpha_{a,i} = S_{a,i}^{k+1} \cdot \frac{T}{T-t_k}$ for the sell side and $\alpha_{b,i} = S_{b,i}^{k+1} \cdot \frac{T}{T-t_k}$. Traders are then ranked by their elapsed wait time draws, with the lowest draw determining the next entrant.[5]

3 Design

The experimental design of this paper rests upon the continuous double auction, straddling various bundles of price accessibility.

3.1 Information Treatments

I impose variation in the price information presented in the open portion of the book as well as the transaction history. Within the bids/asks columns of the book, I test the upper bound (full), as well as the most externally relevant intermediate case: best bid and offer (BBO). The transaction history is partitioned in a similar manner. The higher accessibility level is a full transaction history, while the lower level of this factor is a common piece in financial markets from a decade or two ago: ticker tape (the most recent trade in the market, updated and replaced with each new trade).

3.2 Session Setup

I employ a between-subject full factorial design with two factors: transaction history and order book accessibility. Transaction history and orderbook accessibility each have two levels, full history and ticker-tape and full book and best-bid-and-offer, respectively. I run eight laboratory sessions, two in each of the four level combinations: Full-Full (FF), Full - Ticker-Tape (FT), BBO-Full (BF), and BBO - Ticker-Tape (BT).

Each session has 6-8 subjects who participated as two-way traders in 12-14 three-minute periods. Between periods, subjects can see an interim screen for 30 seconds. The traders are

[5] If all traders draw times outside the duration of the trading period, the current period ends.

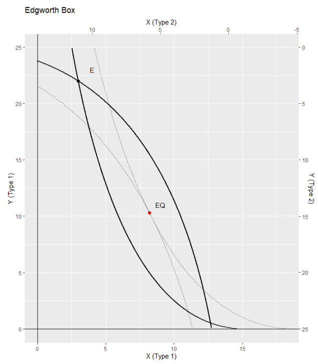

Figure 1: Edgeworth box displaying natural buyer and seller preferences. Type 1 refers to buyers, 2 to sellers.

	Buyers	Sellers
c	0.113	0.099
a	0.825	0.6875
b	0.175	0.3125
r	0.5	0.5
x_o	3	11
y_o	23	3
x_{eq}	8.2	5.8
y_{eq}	10.31	15.69

Table 1: CES parameters, starting endowments and equilibrium allocations.

split evenly into natural buyers and sellers. All natural buyers have the same endowment and heatmap (i.e. CES parameters) to begin each period; similarly, sellers match at the beginning of each period. Traders keep their role for all periods in the session. The utility parameters and endowments, as well as the equilibrium allocations, for each trader type are displayed in Figure 1 and Table 1.

3.3 Laboratory Realization

I use an updated version of the novel user interface first displayed in Crockett et al. (2021). Traders are induced with preferences through a large, continuous heatmap, as seen in Figure 3.3. Higher utility-yielding bundles are associated with warmer colors on the map. The traders are made aware of the indifference curve associated with their current endowment, and can see the indifference curve associated with any bundle they hover over. The map can be clicked to prompt an order placement in the orderbook, which takes up the the remainder of the user interface (aside from an error box which flags attempted market orders that are not feasible). A bids column, asks column and trades column make up the orderbook in the laboratory interface. Own orders and trades are highlighted red if the trader is buying x and green if selling x.

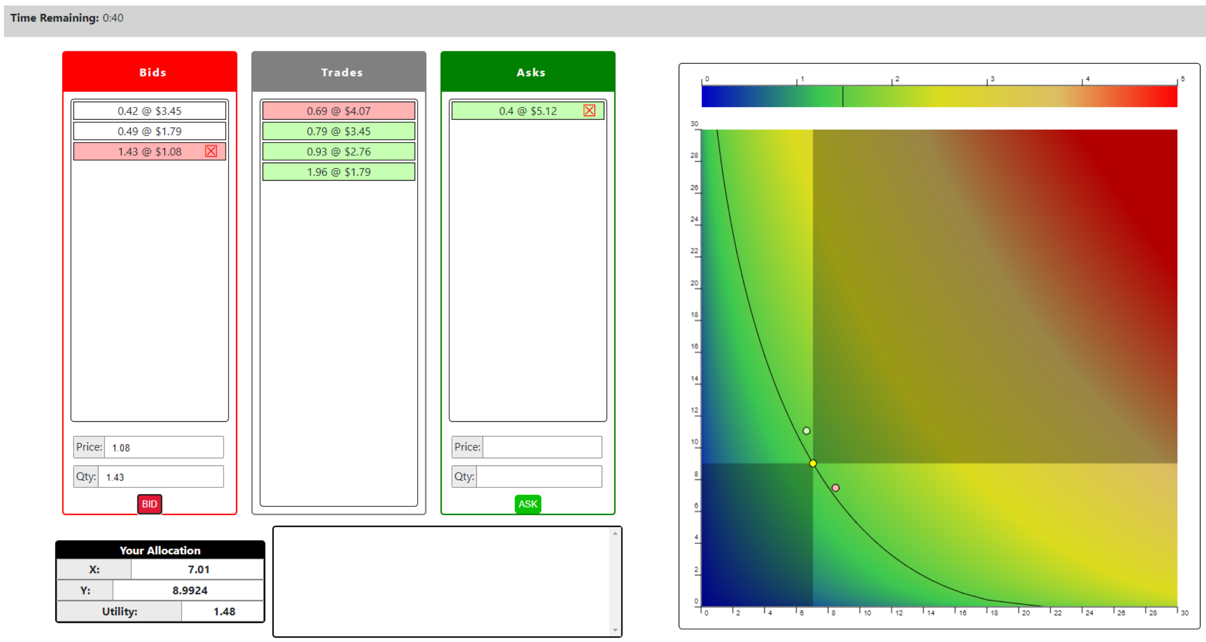

Figure 2: User interface for laboratory market sessions. Contains an orderbook, preferences heatmap, allocation box and error log box.

3.4 Implementation

Subjects were recruited via Orsee (Greiner, 2015), the overwhelming majority of whom were students of UC Santa Cruz[6]. Each session was comprised of eight subjects, with the exception of two sessions which each featured six person markets due to participation complications[7].

All sessions were run virtually, with subjects joining a zoom call for the duration (roughly 90 minutes) of the session. The average payment per subject was \$19.38, the maximum payment being \$37.68. All payments were made via venmo, with payment amounts determined via the following equation: $Pay = showup\ fee + \sum_{i=1}^{N}(\alpha * gained\ utility + \beta * initial\ utility)$, where $(\alpha = 2,\ \beta = 0.4)$.[8]

[6]The vast majority of subjects come from majors in buildings close to the economics department, e.g. computer science, biology and engineering. One subject in the inexperienced subject pool was from UC Berkeley; this subject was not included in the experienced subject pool

[7]These occur in mirrored treatments(BF and FT), so each level in the main 2x2 square was impacted equally.

[8]Concern over the single vs multiple round payment discussion can be felled as losses in utility and thus payment within trading rounds was achievable. While round payments could be negative, this was relatively rare in inexperienced rounds and very rare in experienced. Additionally, the sum of round payments was floored at the show-up fee.

4 Results

The following set of experimental results will discuss the equilibrating tendencies of the laboratory markets in both price and allocation space, as well as the differential impacts price accessibility has on price discovery.

4.1 Prices

Price trends within and across trading periods are a key (and aside from efficiency, likely those most studied) class of indicators for market performance. In this section, I interpret the trade price dynamics and qualitative characteristics, first at an individual trade level and then at a round-average level.

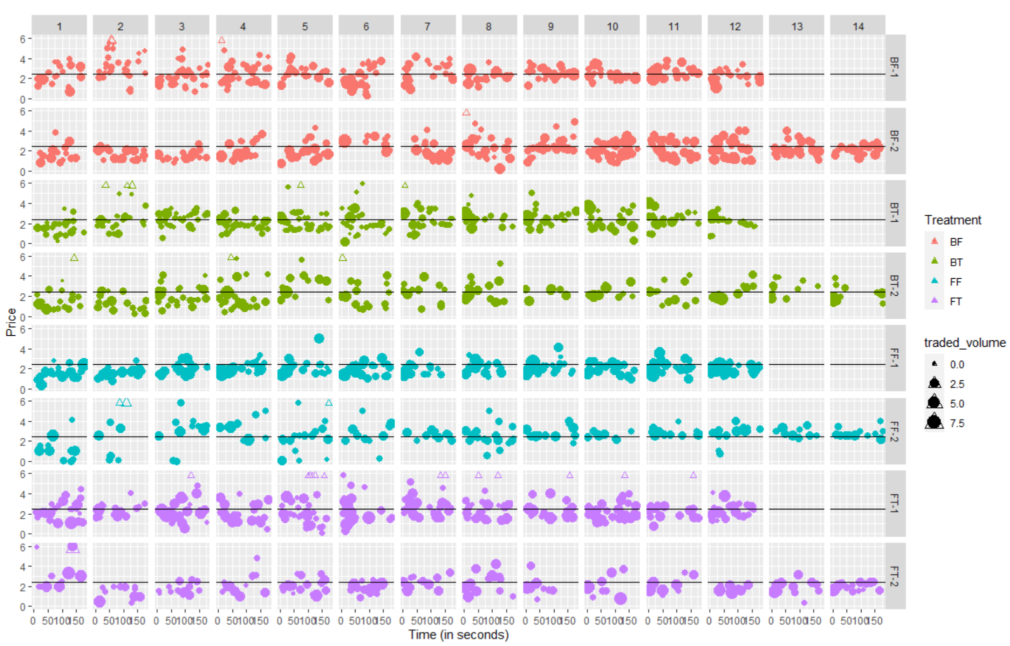

Figure 3: Individual transaction prices. Trades with prices above 5 units of y per unit of x are plotted as outliers (triangles along the $p = 5$ line). The black line plots the round-start competitive equilibrium price.

Figure 3 plots all trades[9] in each session against the time the exchange marked the

[9]The top 2% of trades, all over twice the equilibrium price, are excluded from this graph and the subsequent analysis.

transaction. A qualitative inspection of the price trends presents evidence towards greater stability and lower trade sizes in high accessibility sessions. Outlier trades over triple the equilibrium price appear far more often in sessions with Ticker-Tape transaction histories. As commonly seen in other market CDA's, buyers exhibit a bargaining advantage, with a majority of prices appearing below the equilibrium price line.

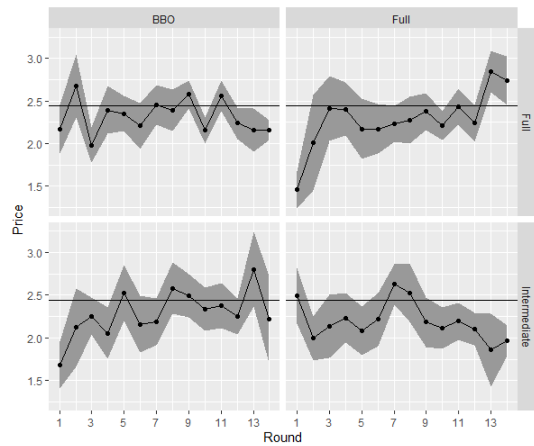

Figure 4: Round-average transaction prices. Shaded region shows 95% confidence interval. Columns show levels for orderbook factor. Rows show levels for transaction history level.

A round-average depiction of the prices seen in Figure 3 can be found in Figure 4. Congruent with the individual price findings, Full-Full markets converge quickly; however, not much is gained over the BBO-TT markets. Where performance diverges is in the markets with asymmetric levels of accessibility. BBO-Full markets converge, though not without substantial oscillatory behavior. Full-TT sessions perform much poorer, demonstrating divergent trends in later periods. Two main qualitative results can be summarized from these figures:

Result 1a: *Prices that largely deviate from equilibrium p^* are more prevalent in markets with lower transaction history. This holds for both levels of orderbook accessibility, and increases in severity when in BBO.*

Result 1b: *Introducing full transaction history accessibility without a full orderbook yields divergent behavior. All other treatments converge (or at least oscillate).*

		BBO		Full	
		All Rounds	Second Half	All Rounds	Second Half
	Price	2.29(0.39)	2.35(0.24)	2.35(0.51)	2.44(0.33)
	$\lvert Price - CE \rvert$	0.71(0.56)	0.69(0.55)	0.72(0.95)	0.71(0.72)
	SD	0.77(0.28)	0.71(0.25)	0.79(0.56)	0.56(0.20)
	RMSE	0.86(0.34)	0.75(0.12)	0.91(0.13)	0.64(0.05)
Full	# Orders	139.12(29.60)	141.71(21.32)	108.27(14.41)	113.21(12.87)
	Order Size	2.10(0.50)	2.20(0.38)	2.11(0.37)	1.94(0.34)
	# Trades	35.33(6.68)	35.17(6.11)	23.58(3.90)	22.83(4.17)
	Trade Size	1.28(0.41)	1.39(0.41)	1.78(0.23)	1.79(0.22)
	Seller MRS	2.19(0.33)	2.35(0.32)	2.15(0.29)	2.21(0.33)
	Buyer MRS	2.82(0.34)	2.63(0.29)	2.81(0.40)	2.74(0.45)
	Price	2.29(0.37)	2.39(0.19)	2.19(0.34)	2.19(0.28)
	$\lvert Price - CE \rvert$	0.87(0.96)	0.92(1.00)	1.04(2.40)	1.12(2.72)
	SD	0.97(0.37)	0.80(0.16)	0.82(0.30)	0.74(0.25)
	RMSE	1.02(0.36)	0.79(0.16)	0.91(0.25)	0.82(0.19)
TT	# Orders	88.77(20.57)	88.07(19.34)	107.35(47.54)	122.43(53.44)
	Order Size	1.62(0.43)	1.81(0.40)	2.25(0.36)	2.24(0.29)
	# Trades	22.42(10.96)	17.36(8.33)	19.35(8.35)	18.29(9.22)
	Trade Size	1.23(0.48)	1.46(0.44)	1.92(0.53)	2.13(0.53)
	Seller MRS	1.98(0.30)	2.09(0.25)	1.97(0.34)	1.99(0.28)
	Buyer MRS	3.07(6.68)	2.88(0.33)	3.04(0.53)	3.01(0.40)

Table 2: Descriptive Statistics at the round-level. Estimates are shown for all rounds and the rounds in the second half of sessions. The four quadrants relate to data from the four treatments, with the vertical panels denoting levels in the orderbook factor and horizontal panels representing levels of the transaction history factor. Seller MRS and Buyer MRS are using round-end estimates, while the rest of the outcomes are in round averages or averaged round totals.

Using BBO-TT markets as a reference point, a descriptive analysis of the adjustment from low to high accessibility in one or both dimensions in these outcomes is outlined. Table 2 can be segmented into three clear outcome types: price convergence as represented by the first four rows, price discovery (the next three rows), and allocative convergence in the final two rows. First, prices show no improvement when transitioning to BBO-Full, but drop significantly when moving to Full-TT. Moving to Full-Full shows a slight improvement from BBO-Full, while the move from Full-TT to Full-Full alleviates the initial reduction from BBO-TT and surpasses the original price by a moderate, but significant margin. The massive decline in price from BBO-TT to Full-TT may be driven by the sudden realization of how much competition as in the book. Buyers are naturally more aggressive early in trading periods[10]; the urgency induced by a realization of more competition, as well as the reinforcement of low prices (since traders would only see the most recent trade price, along with their own) may be the driving mechanism behind this mitigation in price level and convergence.

The next three rows provide a grouping of outcomes indicative of price variation or volatility. All three measures of volatility show improvement in performance as accessibility is increased, with the sole exception of average price deviation in Full-TT (which is quickly explained by the divergent behavior shown in Figure 4).

Order frequencies exhibit behavior similar to that found in the empirical literature of the early 2000's (e.g. Boehmer et al. (2005), Madhavan et al. (2005)). All treatments with full accessibility in at least one dimension see higher order frequency, though the more interesting differences appear when order of improvement is considered. Consider, first, right and then up in Table 2 (thus from BBO-TT to Full-TT to Full-Full) as a potential improvement path, markets show a monotonic progression in order count. However, following the other path from BBO-TT to Full-Full creates a non-monotonic adjustment in order frequency. Order and trade frequencies are indicative of price discovery and the efforts of the traders to aggregate information on their own. The massive increase in order and trade frequency in BBO-Full is likely induced by traders attempting to post a spread reducing order. The increased visibility of trades makes the importance of having a BBO order more apparent, as these are (most often) the trades appearing in the transaction history. The uptick in orders (and especially spread-reducing orders) essentially mechanically induces the increase in trades in these BBO-Full markets. Once the orderbook is fully accessible (Full-Full markets), and traders realize they can place non-spread-reducing orders that may either become BBO orders later in the round, or may be taken up in larger trades by aggressive traders on the opposite side.

[10]Common mechanisms for this include the existence of a natural lower bound for prices at 0, or the general fact that humans participate in society far more as buyers which leads to a better understanding of how to bargain on the buy side in environments like a laboratory market.

Result 2. *Order frequency increases in all three treatments with higher accessibility, relative to BBO-TT markets. Price discovery behavior is associated with the order in which accessibility is improved; increasing orderbook accessibility first leads to a smoother transition in order frequency, though at the cost of price convergence.*

The final two rows present the average marginal rate of substitution for aggregated agents who represent the set of four natural buyers and sellers in the market, respectively. Each buyer's (seller's) action is scaled by the number of similar traders (in this case each trade change in y and x for a given buyer is quartered and taken as the change in allocation for the aggregated buyer). Transactions between traders of the same natural side yield a null movement for the aggregated agent for that side. As explained in a much deeper sense later in section 4.4, ideally the MRS of these aggregated agents (and each individual trader) will be equal to the competitive equilibrium price. Table 2 suggest that an increase in transaction history visibility yields a large reduction in the final spread between the marginal rates of substitution of the aggregated natural buyer and seller. An increase in orderbook accessibility paired with no change in the accessibility of the the transaction history, however, yields no apparent improvement in MRS convergence to CE.

For further discussion on pricing tendencies in these laboratory markets, including within- and across-period dynamics, see Appendix B.

4.2 Allocations

The reallocation of goods throughout a markets existence, with the simultaneous movement of multiple goods and both sides of the market, is a defining factor of general equilibrium; and the final reallocation being just as indicative of the market's convergent behavior as its prices.

Figure 5 presents Edgeworth box depictions of each market's (trading period) final allocation, categorized by treatment. The box maps the average movement of natural buyers in their x and y holdings after each trade, with the lower horizontal axis and left vertical axis marking each respectively. The average movement of natural sellers is mapped similarly, as the average seller allocation is the average amount of x and y left in the market.

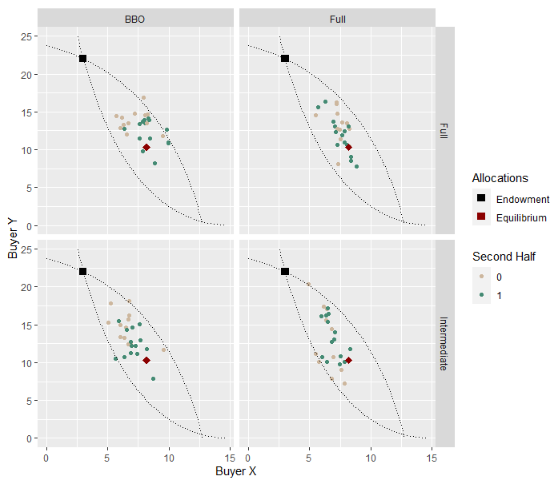

Figure 5: Round-end allocations.

Three criteria for performance in Figure 5 are: tightness of cluster, deviation from endowment-to-equilibrium path, and existence in the space of preferable points (between the two initial indifference curves). When considering the set of all periods, clustering is similar across all treatments. All treatments with at least one factor with less than full accessibility show at least one period outside the set of mutually preferred allocations. Perhaps more interesting, markets with either high accessibility in both dimensions or low accessibility in both exhibit re-allocations close to the the endowment-to-equilibrium path. Markets with asymmetry in their accessibility, however, appear to favor one player type: natural buyers in BBO-Full and natural sellers in Full-TT. The phenomenon likely ties in with the differences in price discovery behavior discussed in section 5.1. Higher order and trade frequencies in BBO-Full[11], driven by aggressive traders (often buyers) trying to make large gains (at low prices) early in rounds and spread-reducing orders throughout, are conducive to the better yields for buyers. Such behavior can be summarized as

[11] Lower price information accessibility benefiting one side of the market disproportionately is not unsurprising, as Ikica et al. (2018) saw buyer advantages appearing in "black box" settings.

Result 3. *Markets with symmetric levels of accessibility reallocate on average near the endowment-equilibrium path. Asymmetric accessibility treatments favor one trader type over the other. All markets reallocate, on average, to a similar distance from the equilibrium, with late rounds finishing near the contract curve relatively often.*

The information portrayed in Figure 5 is compressed and re-imagined in terms of distance from the equilibrium allocation in Figure 6 and Table 3. Distances are measured using the average Euclidean distance for each trader after normalizing the y dimension by the equilibrium price. Across-round final distance dynamics are estimated through a log-linearized regression representing the exponential decay function $d_t = d_1 e^{\gamma t}$. Log(distance) is thus regressed on log(round) and presented numerically in the upper panel of Table 3 and via a best fit line in Figure 6.

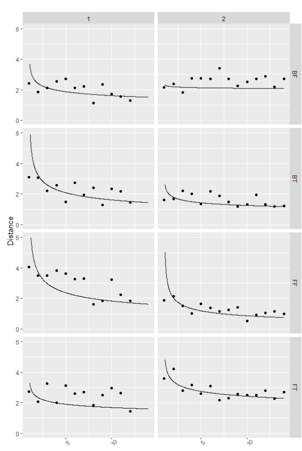

Figure 6: Edgeworth box displaying natural buyer and seller preferences. Type 1 refers to buyers, 2 to sellers.

		Distance			
		BBO		**Full**	
		Sess1	Sess2	Sess1	Sess2
Round-End	**Full**	-0.01	-0.17	**0.22**	-0.09
		(0.89)	(0.06)	(0.04)	(0.75)
	TT	-0.04	**-0.11**	-0.09	-0.04
		(0.78)	(0.08)	(0.33)	(0.84))
Timing	**Full**	-0.05	0.08	0.04	-0.08
		(0.36)	(0.40)	(0.64)	(0.77)
	TT	-0.02	-0.03	-0.09	-0.08
		(0.89)	(0.58)	(0.23)	(0.68)

Table 3: Estimates from regressing log(outcome) on log(round), where outcomes where final distance from equilibrium allocation (upper panel) and timing of shortest distance in each round (lower panel). () denotes p-values. Bold estimates are statistically significant at at least the 0.1 level.

The plots show linear decay in the majority of sessions, with session FF-2 being the closest to being exponential. FF sessions also displayed the steepest decay. Table 3's upper panel corroborates these claims, with estimates for γ being low and generally insignificant. What if the time at which each market reaches its minimum distance from equilibrium is not the final moment of the market's life? Given trader's propensity to learn in these laboratory markets, this is not unlikely. The lower panel of Table 3 shows estimates for a regression of the same form as before, but with time of shortest distance instead of distance itself. Markets show weak decay in this outcome as well, showing trader's are learning to converge allocations faster across periods in all treatments. Decays is slightly stronger in markets with full orderbook accessibility, though insignificant.

4.3 Efficiency

Given the mission of laboratory market experiments is often to test theoretical predictions over competitive equilibria, tests and measures of efficiency are crucial. Traditionally, allocative efficiency is measured by comparing the surplus gained by the set of traders to the the market's gained surplus theoretically maximum. In partial equilibrium, an equivalent depiction of the numerator is the aggregate of each trader's sum over the prices they traded at and the costs($c_{j,s}$)/values($v_{i,b}$) for each of the traded units

$$\sum_{b=1}^{B}\sum_{i=1}^{P_b}(p_{i,b} - v_{i,b}) + \sum_{s=1}^{S}\sum_{j=1}^{P_s}(c_{j,s} - p_{i,s})$$

where B and S are the cardinalities of the buyer and seller sets, and P_b and P_s are the number of buyer and seller units at the inception of a market.

For general equilibrium with induced utility functions, this definition can be adjusted in a natural way. Instead of summing over the surplus gains, I consider the sum of utility gained by all market traders divided by the theoretical utility gain of the market in competitive equilibrium. The numerator can be written as follows:

$$\sum_{n=1}^{N}(u_n(x_{n,Final}, y_{n,Final}) - u_n(x_{n,Endow}, y_{n,Endow}))$$

where N=B+S.

Before presenting the experimental results, I establish a simple theoretical benchmark. The next subsection provides a short synopsis of the well-known agent-based zero intelligence (ZI) model (Gode and Sunder (1993), adjusted for the general equilibrium setting.

4.3.1 Benchmark: Zero-Intelligence

First introduced in 1993 (Gode and Sunder), and then adjusted to accommodate over-simplified versions of general equilibrium, the zero intelligence (ZI) model provides a solid theoretical floor for human behavior in a continuous double auction. In the original partial equilibrium version, traders randomly choose prices uniformly over a specified range. ZI-Constrained (ZI-C) amends this decision space by raising the floor for sellers from 0 to their current unit-cost, and lowering the ceiling for buyers from the max m to their resale value. The model can aptly be summarized (along with the above) by the following set of rules:

- Each trader is either a one-way buyer or seller, endowed n_b and n_s units respectively. Buyers have resale value schedules $\{v_1, .., v_{n_b}\}$, and sellers have cost schedules $\{c_1, .., c_{n_s}\}$.
- Units are ordered by price, such that $v_i > v_k$ and $c_j < c_k$ for $i, j < k$. Each unit must be traded in this order (i.e. a seller must buy their highest cost unit first).
- Orders are single-unit only.
- Spread reduction: Only the best bid and ask are kept in the book, with new orders only being posted if they improve the best bid-ask spread.

Gode and Sunder, along with Spear (2004), provide an amended version of this model, adapted to suit a two-good pure exchange economy. Many of the above conditions still hold in Gode et al. (2004), with orders still being having an artificial step-size setand a spread reduction rule still existing. The resale and cost schedules are naturally replaced with utility functions. The major change, aside from the setting itself, is the method by which traders select their price (the number of units of the numeraire the trader is willing to send/receive for one unit of the commodity). Traders uniformly randomly select prices for both sides of the market. Prices are chosen as an angle in radians. On the sell side, the angle is bounded between $\pi/2$ and the angle of the step-size length vector that begins at the trader's current allocation and is secant to the trader's indifference curve. On the buy side, bounds of 0 and the angle of a similarly defined secant vector are used.

Williams (2021) provides a more generalized version of the 2004 model, in which the quantity entry of the normal order tuplet is not restricted to be a uniform step-size unit. Traders instead select a (p, q) ordered pair uniformly randomly from a fine lattice over the space of orders on the side of the market entered. A constrained version of ZI in this context is thus choosing over a constrained subset of this lattice where all ordered pairs are weakly utility improving.[12] Each of the remaining rules of ZI are also generalized or relaxed in some way:

- Traders are two-way traders, with natural dispositions towards one side of the market.
- Units are not ordered, and may be retraded within a trading period.
- Spread reduction is not enforced, and all orders (up to one per trader on each side) are stored in the book.

Along with the goal of generalization and relaxation of simplifying restrictions, the above rules were chosen to match those enforced on the laboratory traders.

4.3.2 Estimates

Table 4 presents the average allocative efficiency (as defined in section 5.1) and distance efficiency, or one minus the percentage of the original distance between the endowment and equilibrium left to be travelled at the end of a period, for all rounds (first column in block)

[12]The angle choice method used in Gode et al. (2004), when adjusted for variable quantities, produces more intelligence than desired. Uniform random choice of quantity after an angle is chosen does not create a uniform distribution over the feasible set of ordered pairs (instead orders closer to the current endowment are much more likely to occur).

and rounds in the second half of the sessions (second column). Zero Intelligence simulation (1000 runs) outcomes are also reported and used as a baseline.

		BBO		Full	
		All Rounds	Second Half	All Rounds	Second Half
Full	Alloc Eff	0.77(0.12)	0.79(0.14)	0.74(0.24)	0.83(0.16)
	Distance Eff	0.56(0.10)	0.57(0.12)	0.60(0.21)	0.69(0.16)
TT	Alloc Eff	0.80(0.09)	0.84(0.07)	0.71(0.15)	0.73(0.10)
	Distance Eff	0.63(0.11)	0.68(0.09)	0.49(0.11)	0.53(0.08)
ZI	Alloc Eff	0.83	0.83	0.83	0.83
	Distance Eff	0.65	0.65	0.65	0.65

Table 4: Round-level efficiencies. Allocative efficiency is the sum of utility gained in the market divided by the total utility gain if equilibrium is achieved. Distance efficiency is calculated as $1 - \frac{||(x_{eq},y_{eq})-(x_T,y_T)||}{||(x_{eq},y_{eq})-(x_o,y_o)||}$, where the norm is average Euclidean distance from final allocation to equilibrium for each trader (normalized in the y dimension by price).

BBO-TT markets show respectable levels of efficiency (in both allocation and distance), with allocative efficiencies landing in the realm of other studies in the literature. These low accessibility markets yield estimates remarkably close to those of the ZI-C markets. This is not to say that human traders behave similarly to ZI agents in these markets, but that BBO-TT levels of accessibility do not hamper the equilibrating powers of the market institution any more than minimally intelligent ZI-C agents.[13] Improving both factors to full accessibility provides no significant improvement in either measure of efficiency. Possibly even more interesting, both market types with asymmetric levels of accessibility exhibit estimates noticeably worse than markets with symmetric accessibility and markets with ZI-C agents. Traders in these markets are substituting efficiency in an effort to accommodate more aggressive price discovery behaviors.

Much like with final allocation distances (Table 3), end-of-round estimates may not be telling the whole story. Table 5 displays exponential decay coefficients for the gap in allocative efficiency at the round level, as well as decay for time of least inefficiency across

[13] ZI market performance is generally taken to be driven by the market institution. In versions of ZI with a budget or no-loss constraint, this is less convincing, though still provides some lower bound for performance (driven by the institution and minimally intelligent traders).

rounds. All treatments show some signs of significant decay for both outcomes. In efficiency gap, Full-Full markets show the strongest decay coefficients. In timing estimates, sessions with full orderbook accessibility show marginally stronger decay than those without.

Efficiency Gap							
BBO-TT		**BBO-Full**		**Full-TT**		**Full-Full**	
Sess1	Sess2	Sess1	Sess2	Sess1	Sess2	Sess1	Sess2
-0.16	**-0.32**	-0.02	**-0.18**	**-0.15**	-0.15	**-0.39**	**-0.32**
(0.02)	(0.05)	(0.78)	(0.04)	(0.02)	(0.14)	(0.01)	(0.01)

Timing of Smallest Gap							
BBO-TT		**BBO-Full**		**Full-TT**		**Full-Full**	
Sess1	Sess2	Sess1	Sess2	Sess1	Sess2	Sess1	Sess2
-0.12	**-0.72**	-0.24	**-0.58**	**-0.21**	**-0.50**	**-0.65**	**-0.75**
(0.52)	(0.04)	(0.26)	(0.03)	(0.02)	(0.08)	(0.06)	(0.02)

Table 5: Estimates from regressing log(outcome) on log(round), where outcomes include 1 minus final allocative (surplus) efficiency (upper panel), as well as d timing of highest efficiency in each round (lower panel). () denotes p-values. Bolded estimates are statistically significant at at least the 0.1 level.

Cumulative density functions for the round-end trader-level difference in utility gained and expected utility gain in equilibrium are reported in Figure A.1. Relative gains are close on average for all four treatments, however, symmetric accessibility treatments (BBO-TT and Full-Full) appear to second order stochastically dominate asymmetric treatments.

Result 4. *Performance in efficiency standards correlates with the symmetry of accessibility of a market. Symmetric treatments (BBO-TT and Full-Full) outperform asymmetric treatments (BBO-Full and Full-TT). Individual estimates of relative utility gain follow a similar trend.*

4.4 Inefficiency in Two-way Trading

In the more classical partial equilibrium setting, there are two driving forces that can lead to inefficiencies in the market: (1) extramarginal traders (units) being involved in trades, and (2) intramarginal traders (units) not all trading. Given the standard value and cost schedule set up of partial equilibrium theory and experiments, checks for either force is straight forward. The same cannot be said for a general equilibrium setting, as there are no simple cost or value schedules assigned to traders for their one or handful of single-unit trades over one good. A natural analog to this PE idea of intra- and extra-marginal units or traders can come from the analysis of marginal rates of substitution. Traders begin with MRS's away from the equilibrium MRS shared by all traders, namely that equivalent to the competitive equilibrium price. As trader's adjust their bundles through trade, their MRS's adjust, moving closer to the equilibrium MRS assuming they are making good trades.

At any point in a trading period, if a trader has not reached the equilibrium MRS, the trader is considered to be an intra-marginal trader as he still has an incentive to trade and room to provide competitive prices. Once the trader reaches or crosses the equilibrium MRS, he is deemed extra-marginal as he has essentially over-traded in his desired direction and can no longer provide competitive prices in his natural side of the market. In a setting with traditional one-way traders, this trader would be considered extra-marginal for the remainder of the trading period; however, with two-way traders, this trader would transition to become intra-marginal on the opposite side of the market compared to their natural preference.

If, by the end of a trading period, a trader's final MRS has not reached its equilibrium value, the trader can be classified to have under traded, or "left trades on the table". If instead, the trader has surpassed the desired endpoint, and become and extra-marginal trader on their natural side, then we can say they have over traded.

Figure 7 displays the MRS for the aggregated natural buyer and natural seller, described earlier in section 5.1, across rounds averaged between sessions for each treatment. As discussed in section 5.1 and summarized in Table 2, markets with higher transaction history accessibility display significantly lower MRS spreads than those with lower accessibility (while holding orderbook accessibility constant). Symmetry in accessibility provides more stable improvement across periods, as BBO-TT and Full-Full show decreasing trends in spread, while BBO-Full shows an over-trading inefficiency and Full-TT spreads diverge in

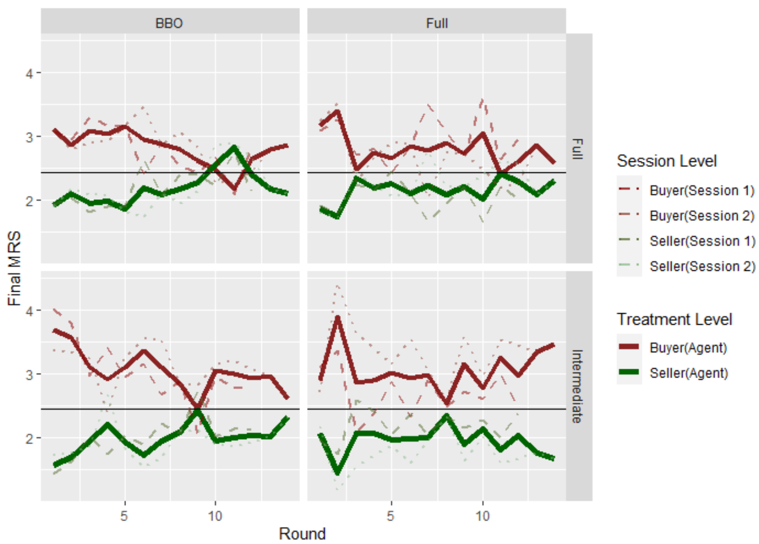

Figure 7: Marginal rate of substitution at the end of each round, averaged across session within treatment. Estimates are for aggregated buyer and seller traders who, after each transaction occurs, adjusts their allocations by the average individual adjustment made by the natural buyers and sellers, respectively. Dashed lines show the session level round-end estimates.

later rounds.

Result 5. *Consistent improvement in MRS spread is associated with symmetric accessibility, while MRS spread magnitudes overall are lower in markets with full transaction history accessibility.*

4.5 Treatment Effects

Table 6 provides regression analysis for the main outcomes discussed in sections 5.1-5.4. Main effects in the price regressions corroborate the story told in section 5.1, with variation slightly increasing, though insignificantly, as accessibility increased. The large negative estimate on $FullOB$ in column 2 matches the markedly low prices shown for Full-TT in Table 2, with the large positive increase of 0.160 matching the jump from Full-TT to Full-Full.

Price discovery impacts prove to be the strongest, with impressive main effects matching the surge in aggressive price discovery behavior in asymmetric accessibility markets and the equally large interaction effect showing the reversal in this behavior once both factors have high levels of accessibility and symmetry is restored. Similarly main effects in the final column support reductions in efficiency in asymmetric accessibility markets and an interaction effect which returns estimates to those of BBO-TT levels in Full-Full markets by the end of the session, though now without significance. Estimates for the time control *Round* are significant and in the direction associated with improvement/convergence in all all of the reported regressions.

	Dependent variable:							
	$\lvert Price - CE \rvert$	Price	Variance	Orders	Trades	RMSE	Distance	Alloc Eff
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	
FullOB	0.053	−0.101	−0.310***	18.577	−3.077	−0.109**	0.761***	−0.090
	(0.108)	(0.090)	(0.109)	(31.836)	(8.037)	(0.046)	(0.255)	(0.087)
FullT	0.046	0.005	−0.402***	50.346*	7.423	−0.160***	0.406	−0.032
	(0.072)	(0.155)	(0.078)	(28.986)	(7.067)	(0.058)	(0.321)	(0.020)
Round	−0.029***	0.018	−0.089**	1.794**	−0.341	−0.051***	−0.073***	0.014**
	(0.007)	(0.013)	(0.041)	(0.851)	(0.282)	(0.013)	(0.025)	(0.006)
FullOB:FullT	0.043	0.160	0.577	−49.423	−9.538	0.165	−0.997	0.058
	(0.126)	(0.307)	(0.485)	(41.104)	(9.906)	(0.165)	(0.716)	(0.141)
Constant	0.491***	2.161***	1.688***	76.140***	24.822***	1.380***	2.418***	0.701***
	(0.065)	(0.090)	(0.294)	(13.766)	(6.093)	(0.108)	(0.298)	(0.042)
Observations	104	104	104	104	104	104	104	103
R^2	0.221	0.051	0.137	0.264	0.256	0.314	0.274	0.159
Adjusted R^2	0.190	0.012	0.102	0.234	0.226	0.286	0.245	0.124
Note:						*p<0.1; **p<0.05; ***p<0.01		

Table 6: Regression estimates using round-level data. BBO-TT is the control in this setup. () denote standard errors, which are clustered at the session level.

5 Conclusions

The rapid progression of technology and increasing interconnectedness of economic agents has provided a constantly changing landscape for markets, making the investigation of market formats and their attributes' impacts on market outcomes increasingly valuable. An interesting and likely impactful class of attributes are the levels of price accessibility the market format offers. While the option of full accessibility seems like the obvious choice for a format, giving full access to all price information in a market is not always achievable or helpful. Providing such information is not necessarily costless, as realizations of larger financial markets may be mentally taxing to traders, or information dissemination may be excessively costly to the central agent or market itself in developing markets and countries.

In this respect, an understanding of the benefits and costs of adjusting price information accessibility in two major aspects of a market, namely its orderbook and history of transactions, is crucial. This paper presents a laboratory market experiment testing popular levels of price accessibility in the orderbook and history in a continuous double auction. A much more generalized, less restrictive (or guided) environment is implemented, with traders being induced with CES utility functions over the two goods of a simple pure exchange economy. Additionally, order quantities are highly flexible and optional market rules that generally provide (potentially too much) structure, such as the common spread reduction rule, are relaxed. A new general-equilibrium adjusted trader behavior algorithm is also presented to provide insights into market and trader responses to changes in accessibility.

Laboratory markets reveal a few of insightful impacts price accessibility adjustment has on CDA market outcomes. First, prices are more likely to have large deviations in markets with lower transaction history accessibility, though increasing accessibility here without full orderbook accessibility leads to an inability to converge in prices. Second, more aggressive price discovery tendencies appear in markets with asymmetric levels of accessibility between the book and history, coming at the cost of allocative efficiency. Third, improvement in accessibility between symmetrically accessible markets leads to improvements in volatility but no perceivable gains in efficiency.

This project reveals non-monotonic gains in efficiency and other main market outcomes when improving price accessibility. As such, there are clear implications for markets, their choice of format and what options more advanced markets could give to their traders to limit their own accessibility (and reduce the mental load). More information bundles within this framework and market format attributes outside of it should be studied to provide a more clear picture of how efficiency maps from across formats or bundles. Additionally, more trader behavior and agent-based models could be brought to more complex (yet still tractable) settings such as the two-good Edgeworth box, as the vast majority still reside in a partial equilibrium framework.

References

Christopher M Anderson, Charles R Plott, Ken-Ichi Shimomura, and Sander Granat. Global instability in experimental general equilibrium: the scarf example. *Journal of Economic Theory*, 115(2):209–249, 2004.

Mikhail Anufriev, Jasmina Arifovic, John Ledyard, and Valentyn Panchenko. Efficiency of continuous double auctions under individual evolutionary learning with full or limited information. *Journal of Evolutionary Economics*, 23(3):539–573, 2013.

Jasmina Arifovic and John Ledyard. Call market book information and efficiency. *Journal of Economic Dynamics and Control*, 31(6):1971–2000, 2007.

Jasmina Arifovic and John Ledyard. A behavioral model for mechanism design: Individual evolutionary learning. *Journal of Economic Behavior & Organization*, 78(3):374–395, 2011.

Robert Bloomfield and Robert Libby. Market reactions to differentially available information in the laboratory. *Journal of Accounting Research*, 34(2):183–207, 1996.

Robert Bloomfield and Maureen O'Hara. Market transparency: who wins and who loses? *The Review of Financial Studies*, 12(1):5–35, 1999.

John Board and Charles Sutcliffe. The effects of trade transparency in the london stock exchange: A summary. *Available at SSRN 6135*, 1995.

Ekkehart Boehmer, Gideon Saar, and Lei Yu. Lifting the veil: An analysis of pre-trade transparency at the nyse. *The Journal of Finance*, 60(2):783–815, 2005.

Tim Bollerslev and Ian Domowitz. Some effects of restricting the electronic order book in an automated trade execution system. In *The Double Auction Market Institutions, Theories, and Evidence*, pages 221–252. Routledge, 1993.

Edward H Chamberlin. An experimental imperfect market. *Journal of Political Economy*, 56(2):95–108, 1948.

Dave Cliff and Janet Bruten. Zero is not enough: On the lower limit of agent intelligence for continuous double auction markets. 1997.

Sean Crockett and Ryan Oprea. In the long run we all trade: Reference dependence in dynamic economies. *University of California Santa Barbara Working Paper*, 2012.

Sean Crockett, Stephen Spear, and Shyam Sunder. Learning competitive equilibrium. *Journal of Mathematical Economics*, 44(7-8):651–671, 2008.

Sean Crockett, Ryan Oprea, and Charles Plott. Extreme walrasian dynamics: The gale example in the lab. *American Economic Review*, 101(7):3196–3220, 2011.

Sean Crockett, Daniel Friedman, and Ryan Oprea. Naturally occurring preferences and general equilibrium: A laboratory study. *International Economic Review*, 2021.

David Easley and John O Ledyard. Theories of price formation and exchange in oral auctions. *The Double Auction Market: Institutions, Theory and Evidence, SFI Studies in Sciences of Complexity, Redwood City, Calif., Addison-Wesley*, 1993.

Mark D Flood, Ronald Huisman, Kees G Koedijk, and Ronald J Mahieu. Quote disclosure and price discovery in multiple-dealer financial markets. *The Review of Financial Studies*, 12(1):37–59, 1999.

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

David Gale. A note on global instability of competitive equilibrium. *Naval Research Logistics Quarterly*, 10(1):81–87, 1963.

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

Dhananjay Dan K Gode, Shyam Sunder, and Stephen Spear. Convergence of double auctions to pareto optimal allocations in the edgeworth box. 2004.

Jacob K Goeree and Luke Lindsay. Market design and the stability of general equilibrium. *Journal of Economic Theory*, 165:37–68, 2016.

Leonid Hurwicz, Roy Radner, and Stanley Reiter. A stochastic decentralized resource allocation process: Part i. *Econometrica: Journal of the Econometric Society*, pages 187–221, 1975.

Barbara Ikica, Simon Jantschgi, Heinrich H Nax, Diego Nunez Duran, and Bary Pradelski. Trading in a black box: Zero intelligence & lack of knowledge. *Available at SSRN 3131004*, 2018.

Robert Jensen. The digital provide: Information (technology), market performance, and welfare in the south indian fisheries sector. *The Quarterly Journal of Economics*, 122(3): 879–924, 2007.

Erik O Kimbrough and Andrew Smyth. Testing the boundaries of the double auction: The effects of complete information and market power. *Journal of Economic Behavior & Organization*, 150:372–396, 2018.

Georg Kirchsteiger, Muriel Niederle, and Jan Potters. Endogenizing market institutions: An experimental approach. *European Economic Review*, 49(7):1827–1853, 2005.

Ananth Madhavan, David Porter, and Daniel Weaver. Should securities markets be transparent? *Journal of Financial Markets*, 8(3):265–287, 2005.

Marco Pagano and Ailsa Röell. Transparency and liquidity: a comparison of auction and dealer markets with informed trading. *The Journal of Finance*, 51(2):579–611, 1996.

Charles R Plott. Market stability: backward-bending supply in a laboratory experimental market. 1999 presidential address western economic association. *Economic Inquiry*, 38(1): 1–18, 2000.

Charles R Plott. Equilibrium, equilibration, information and multiple markets: from basic science to institutional design. 2001.

Herbert Scarf. Some examples of global instability of the competitive equilibrium. *International Economic Review*, 1(3):157–172, 1960.

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

Vernon L Smith. Relevance of laboratory experiments to testing resource allocation theory. In *Evaluation of Econometric Models*, pages 345–377. Elsevier, 1980.

Vernon L Smith and Arlington W Williams. The boundaries of competitive price theory: Convergence, expectations, and transaction costs. In *Bargaining and market behavior: essays in experimental economics*, pages 286–319. 2000.

Michiel van de Leur and Mikhail Anufriev. Timing under individual evolutionary learning in a continuous double auction. *Journal of Evolutionary Economics*, 28(3):609–631, 2018.

Arlington W Williams, Vernon L Smith, John O Ledyard, and Steven Gjerstad. Concurrent trading in two experimental markets with demand interdependence. *Economic Theory*, 16 (3):511–528, 2000.

Brett Williams. Zero intelligence in an edgeworth box. 2021.

Robert B Wilson. On equilibria of bid-ask markets. In *Arrow and the ascent of modern economic theory*, pages 375–414. Springer, 1987.

Appendix A Distributional Tests

A.1 Tests for Descriptive Statistics

	BT→ BF	BT→ FT	BT→ FF	BF→ FT	BF→ FF	FT→ FF
Price	0.54(∼)	**0.04**(-)	0.73(+)	0.13(-)	0.54(+)	**0.08**(+)
$\|Price - CE\|$	0.54(-)	**0.04**(+)	0.73(-)	0.13(+)	0.54(∼)	**0.08**(-)
SD	0.23(-)	0.54(-)	**0.00**(-)	0.60(∼)	0.11(∼/-)	**0.04**(-)
RMSE	0.43(-)	0.73(-/∼)	**0.03**(-)	0.31(+)	0.38(+/-)	**0.02**(∼/-)
# Orders	**0.00**(+)	0.18(+)	**0.00**(+)	0.16(-)	0.21(-)	0.60(∼/-)
Order Size	**0.00**(+)	**0.00**(+)	0.21(+)	0.70(∼)	**0.04**(∼/-)	**0.01**(-)
# Trades	**0.00**(+)	0.91(-/∼)	0.82(∼/+)	**0.00**(-)	**0.00**(-)	0.93(+)
Trade Size	0.63(∼)	**0.00**(+)	**0.03**(+)	**0.00**(+)	**0.01**(+)	**0.08**(-)
Seller MRS	**0.01**(+)	0.45(∼)	0.19(+)	**0.01**(-)	0.43(∼/+)	**0.10**(+)
Buyer MRS	**0.06**(-)	0.54(∼/-)	0.18(-)	**0.01**(+)	0.77(∼)	0.16(-)

Table A.1: Wilcoxon test p-values for outcomes in Table 2. () denote the direction of change when moving from Treatment A → Treatment B in the column.

A.2 Utility Gain CDFs

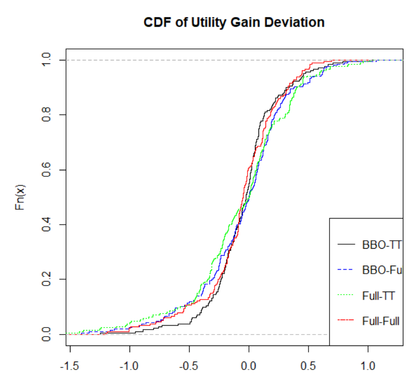

Figure A.1: Cumulative density functions for the round-end trader-level difference in utility gained and expected utility gain in equilibrium.

Appendix B Across and Within Period Price Dynamics

B.1 Across Period Estimates

Session	Δp_1	γ	s.d.(γ)
BF-2	-0.17	0.38	0.18
BF-1	0.10	-0.49	0.15
BT-2	-0.06	0.31	0.27
BT-1	-0.16	-0.27	0.28
FF-2	0.36	-0.17	0.31
FF-1	-0.21	0.43	0.05
FT-2	-0.51	-0.23	0.26
FT-1	-0.07	0.45	0.16

Table B.1: Estimates for one period lagged price deviation (from time 0 competitive equilibrium price, p^*).

B.2 Within Period Estimates

		BBO-TT			BBO-Full			Full-TT			Full-Full		
Period	Session	Δp_1	γ	s.d.(γ)	Δp_1	γ	s.d.(γ)	Δp_1	γ	s.d.(γ)	Δp_1	γ	s.d.(γ)
1	(1)	0.21	-0.35	0.15	-0.82	0.12	0.56	-0.66	0.01	0.98	-0.26	0.43	0.03
	(2)	-0.68	0.11	0.72	-2.17	-0.39	0.32	-0.70	0.17	0.55	0.56	0.26	0.46
2	(1)	0.36	0.49	0.01	-0.67	0.04	0.85	0.36	0.15	0.45	-0.07	-0.04	0.89
	(2)	-0.94	-0.05	0.84	1.28	-0.13	0.83	-0.55	0.18	0.46	-0.99	-0.27	0.54
3	(1)	-0.00	0.48	0.00	-0.40	0.40	0.08	-0.04	0.10	0.60	-0.02	0.18	0.43
	(2)	-0.47	0.14	0.48	0.76	-0.06	0.80	-0.34	-0.07	0.78	-0.73	0.05	0.87
4	(1)	0.08	0.05	0.80	-0.26	0.10	0.69	-0.55	0.03	0.89	-0.40	-0.07	0.79
	(2)	-0.33	0.21	0.47	1.33	-0.67	0.40	-0.30	0.04	0.84	0.10	0.04	0.93
5	(1)	0.26	0.08	0.70	-0.31	0.12	0.60	-0.06	0.03	0.85	-0.31	0.13	0.51
	(2)	-0.58	0.09	0.69	-0.12	-0.59	0.07	0.79	-0.11	0.77	-0.51	-0.32	0.26
6	(1)	-0.15	0.66	0.00	-0.77	-0.38	0.10	-0.33	0.06	0.72	-0.11	0.15	0.39
	(2)	0.29	-0.07	0.85	0.27	0.14	0.74	-0.43	0.08	0.64	-1.09	-0.59	0.10
7	(1)	0.12	0.19	0.27	-0.52	-0.11	0.72	-0.17	0.06	0.75	0.36	-0.01	0.96
	(2)	-0.30	0.27	0.35	0.28	-0.37	0.39	-0.25	-0.25	0.47	-0.03	0.17	0.66
8	(1)	-0.29	0.16	0.41	-0.45	-0.27	0.30	0.27	-0.45	0.15	0.18	-0.19	0.36
	(2)	0.11	0.35	0.12	0.06	-0.24	0.36	0.20	0.40	0.17	0.19	-0.06	0.86
9	(1)	0.17	0.42	0.01	-0.16	0.27	0.23	0.14	0.05	0.83	-0.19	-0.27	0.29
	(2)	0.03	0.07	0.78	0.45	0.60	0.09	-0.18	0.48	0.28	-0.66	0.05	0.91
10	(1)	-0.09	0.67	0.00	-0.32	0.18	0.46	-0.20	-0.20	0.38	-0.41	-0.06	0.73
	(2)	-0.59	-0.04	0.83	0.03	-0.09	0.85	0.06	0.04	0.92	-0.46	-0.41	0.44
11	(1)	0.14	0.13	0.47	-0.12	0.16	0.45	0.07	-0.30	0.17	-0.31	-0.27	0.31
	(2)	-0.04	-0.13	0.51	0.91	-0.65	0.16	-0.42	-1.18	0.01	-0.26	0.55	0.29
12	(1)	-0.12	0.43	0.03	-0.37	0.39	0.08	-0.11	0.08	0.61	-0.22	-0.17	0.41
	(2)	-0.32	-0.36	0.07	0.18	0.15	0.68	0.38	0.75	0.22	-1.14	-0.09	0.87
13	(2)	-0.23	0.31	0.16	0.66	0.13	0.74	0.67	0.02	0.95	-1.62	-1.21	0.01
14	(2)	-0.34	0.07	0.70	0.52	-0.68	0.06	-0.20	-0.21	0.56	-0.45	-0.21	0.53

Table B.2

Appendix C Within-Period Allocation Adjustment

Figure C.1 presents Edgeworth box depictions of each market (trading period) for each treatment. The box shows the average movement of natural buyers in their x and y holdings after each trade, with the lower horizontal axis and left vertical axis marking each respectively.

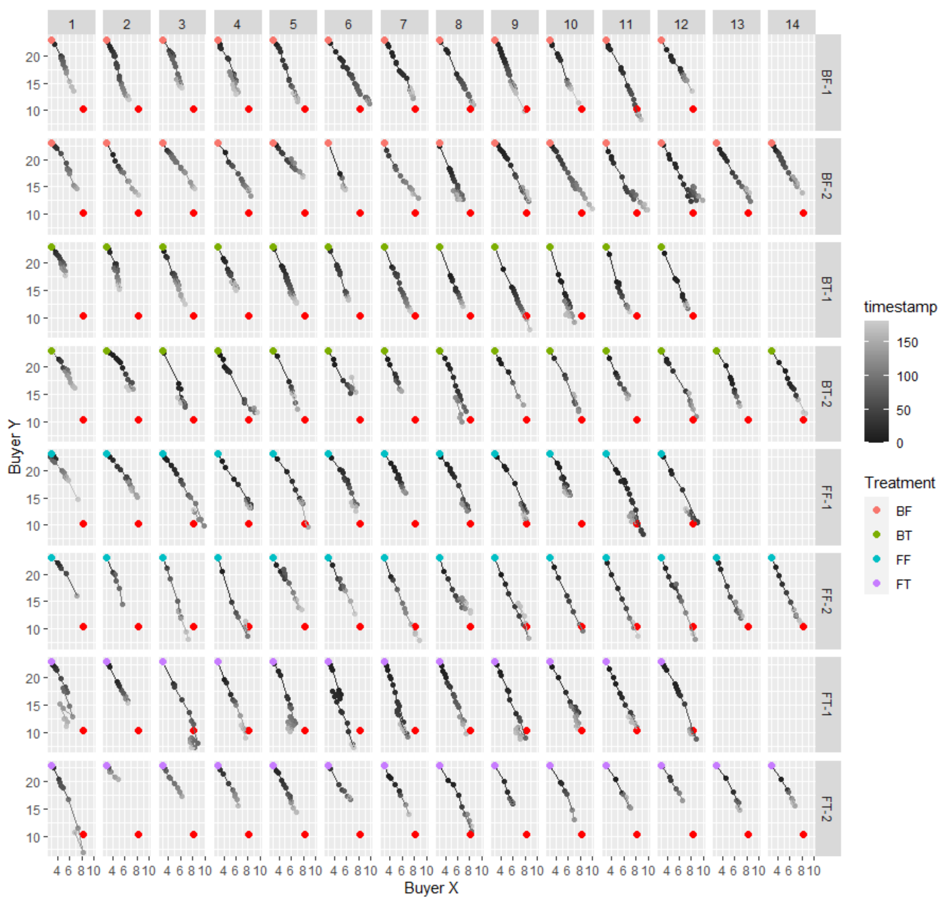

Figure C.1: Within-period allocation adjustment plotted in Edgeworth boxes. The shade of the allocation dot fades later in the period.

The average movement of natural sellers is mapped similarly, as the average seller allocation is the average amount of x and y left in the market.

Figure C.1 plots the allocation adjustment path for each period in all eight sessions. Much like with prices, a trend of slow improvement appears for most sessions. Included within this trend is a tendency to improve quickly and then revert back to poor progression, as if traders are readjusting their price discovery patterns to achieve greater gains. More occurrences of this in sessions with lower accessibility supports this potential mechanism. Outside of session FT-2, however, the markets appear to converge relatively well in at least half of the periods.

Appendix D Utility-Losing Behavior

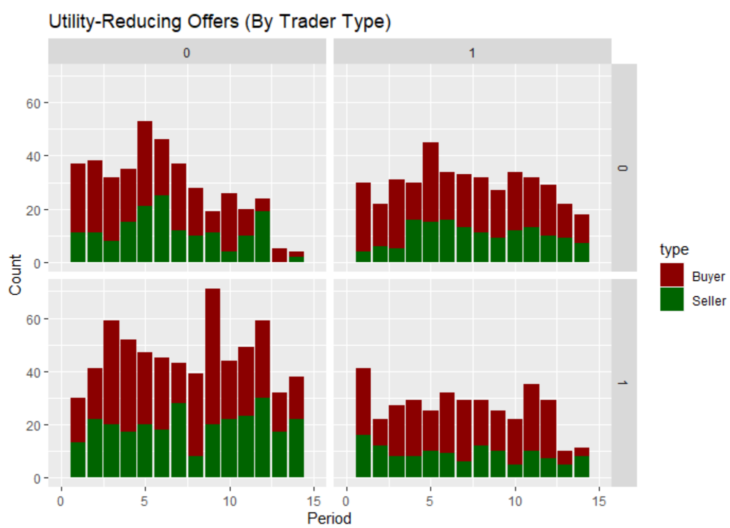

Figure D.1: Counts of utility-losing orders placed, partitioned by trader type.

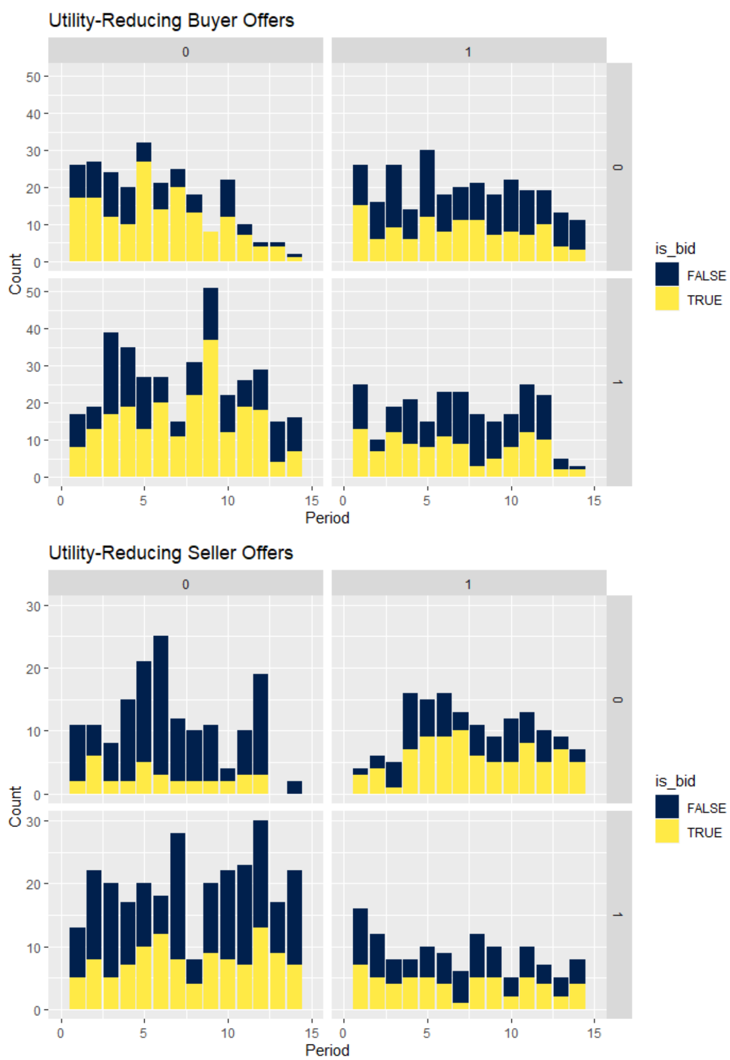

Figure D.2: Counts of utility-losing orders placed, partitioned by trader type.

	Dependent variable:
	utilLoss
distFromCenter	0.019***
	(0.008)
FullT	−0.044
	(0.192)
FullOB	−0.306***
	(0.105)
FullT:FullOB	−0.156
	(0.181)
Constant	−1.305***
	(0.150)
Observations	9,208
Log Likelihood	−4,566.272
Akaike Inf. Crit.	9,142.545
Note:	*p<0.1; **p<0.05; ***p<0.01

Table D.1