Title
Essays on the Continuous Double Auction in General Equilibrium

Author
Williams, Brett

Publication Date
2022

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA
SANTA CRUZ

ESSAYS ON THE CONTINUOUS DOUBLE AUCTION IN GENERAL EQUILIBRIUM

A dissertation submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

ECONOMICS

by

Brett Williams

June 2022

The Dissertation of Brett Williams is approved:

Professor Daniel Friedman, Chair

Professor Natalia Lazzati

Professor Kristian López Vargas

Peter Biehl
Vice Provost and Dean of Graduate Studies

Table of Contents

List of Figures

List of Tables

Abstract

Essays on the Continuous Double Auction in General Equilibrium

by

Brett Williams

The continuous double auction (CDA) has found itself in a state of ubiquity in today's market landscape. Such presence has motivated a vast literature on the double auction, spanning theoretical, experimental and empirical works alike. Despite the century or so of research into the underpinnings of the double auction and its many variations, much is left to discover in terms of convergence and its determinants. This dissertation investigates the driving forces of the continuous double auction institution, as well as mechanisms for equilibrium-achieving trader behavior, all set in simple general equilibrium settings.

Chapter 1 presents a model of zero intelligence trading in the continuous double auction, though set in a general equilibrium economy and unconstrained to the point of reaching a lower level of 'zero' than prior literature. Much like Gode and Sunder (1993), the model's intent is not to give a prescription for what traders do in the world, but to better understand the driving power and mechanisms underlying the CDA in a more complex environment. To fully understand the implications of the model, the institution and the economic setting, I simulate several variations of the market hundreds of times.

Chapter 2 pushes further into the wilderness of bounded rationality, questioning how price accessibility in a CDA impacts trader behavior. Do they respond to orders being posted and traders being made, and if so, how? To investigate, I run a

set of laboratory experiments which vary the accessibility of prices in the orderbook and transaction history. I pair the experiment with another general equilibrium adaptation of a classic CDA agent-based model (Gjerstad and Dickhaut (1998)), one which assumes traders perfectly choose orders based on their beliefs on how acceptable prices are. Chapter 3 postulates a trader behavior model which is able to encapsulate the models of Chapter 1 and 2, and also provide one mapping of the intermediate levels of minimal intelligence between them. In addition to holding beliefs on prices, traders also abide by reservations that allow for disequilibrium trade, and select orders imperfectly via logit choice.

To Reta, Betty, Ed and Ellie,

the motivation you've provided is endless.

Acknowledgments

I am indebted to many brilliant people for their advice, guidance and company in my time as a Ph.D student and candidate, for this dissertation would not have been possible without them.

First, I want to thank my advisor, Daniel Friedman, for his commitment to making me a better economist and for guiding me through the endless 'ocean' that is markets. Dan has been supportive and inspiring from the beginning, as an advisor, a coauthor and a mentor. I also thank Natalia Lazzati and Kristian López Vargas for their guidance in both teaching and research as part of my reading committee. I'm a better economist, experimentalist, theorist, behavioralist, presenter and teacher because of them. With as much as they all helped me directly (and I wish I had reached out to them even more), they all served as professional role models for me (and still do) whether they know it or not.

Next, I want to thank my cohort, coauthors and intellectual colleagues (all of whom I consider friends). To my cohort, I'm grateful for our time together and our camaraderie. From getting through classwork to discussing what would become our dissertations to our cohort dinners, I enjoyed it all. I thank my coauthors for helping me grow as a researcher and a team member. I also thank my colleagues from my IFREE workshop cohort; its been truly helpful being able to discuss ideas with so many experimentalists from all around the world in the same stage of their career as me. To the experimental workshop (now BETE), your comments, patience and willingness to listen were vital in my growth as a researcher and presenter.

Lastly, I thank my family and friends outside of research. They helped ground me throughout this process; anytime I needed a break from study or research or someone to talk to, someone was there. I couldn't have completed this degree without the support of my family, whether mental, financial or emotional. I can't thank them enough.

Chapter 1

Zero Intelligence in an Edgeworth Box

1.1 Introduction

In a discussion on economic modeling and agent behavior, J. Doyne Farmer recently suggested “the main stream (of models) for the last 50 years has entered the wilderness of bounded rationality through the rationality gate... there have been a few who have entered through the other (zero intelligence) gate... though we need ‘tunnelers’ from both directions and we don’t have enough coming from the minimal intelligence direction.” (2020)[1] Agent-based models have been a prominent process by which researchers have entered the zero/minimal intelligence gate. The study of trader behavior and price formation in markets, including the continuous double auction (CDA) institution, can especially benefit from such simple, parsimonious modelling given how easily a market’s underpinnings can become over-complicated.

Gode and Sunder (1993) postulated a model for double auction behavior which

[1]This quote is from The First Conference on ZI/MI Intelligence Agents in 2020, in a discussion on agent-based modelling in economics.

enters through the gate of no intelligence in a quite literal sense, introducing "zero intelligence" traders to the literature.[2] Zero intelligence traders provide a counter to the traditional rational traders who hold perfect utility maximizing capabilities when placing orders. Instead, prices are randomly chosen within the set of feasible prices $[0, M]$, or some no-loss constrained subset of this. Both the partial equilibrium and general equilibrium (Gode, Spear and Sunder(2004)) versions of the model provide glimpses into the natural equilibrating tendencies of the CDA. However, the models rely on a few restrictive assumptions, either institutional or behavioral in nature.

This paper postulates a more generalized lower-bound model of zero intelligence in a continuous double auction set in the general equilibrium environment of an Edgeworth box. In continuity with Gode et al. (2004) (GSS henceforth), agents act as two-way traders in a simple two good economy, participating in a series of market periods where their endowments are reset at the beginning of each period. The main departure of this paper's model lies in the intelligence given to traders when "randomly" placing orders. GSS provides traders with three major choice tendencies: (1) a constraint restricting traders from selecting orders which would allow them to lose utility, (2) an order choice method which has prices chosen via a uniform draw over a set of order vector slopes (in radians[3]), and (3) a set step-size on the length of the order vector. The model presented in this paper relaxes all of the constraints, thus imparting traders with a much lower, truer level of 'zero' intelligence. Traders instead select bundles over

[2]Hurwicz et al. (1975) was likely the first proposal of minimal or zero intelligence in a pure exchange setting. Many thanks to John Ledyard for the reference.

[3]A uniform draw over this set places substantially more weight on lower prices. For example if $(\pi/2, \pi)$ is the feasible set of angle choices, half of the support yields prices between 0 and 1, while the other half covers $(1, \infty)$.

the set of feasible 'next' allocations on the side of entry, via uniform draw from a fine lattice.

A stream of other agent-based models emerged in the late 20th century alongside Gode and Sunder (1993). Wilson's (1987) game theoretic venture, while informative, showed the limitations complexity places on a strategic approach to modelling CDA trader behavior. Friedman (1991) and Easley and Ledyard (1993) both posited non-strategic models with variations on a reservation price mechanism, yet differed in the dynamics of interest. Friedman studied within-period pricing dynamics, while Easley and Ledyard focused on across-period dynamics. Gjerstad and Dickhaut (1998) proposed a belief-based model shortly after, though moving away from zero intelligence towards higher complexities of trader intelligence. Each model in this group is oriented in a partial equilibrium setting; very few models (or even non-theoretical works) have focused on general equilibrium dynamics in these simple CDA markets (Gode and Sunder (2004) and Crockett et al. (2008a) are a couple such papers). This paper adds to the general equilibrium sub-literature, with the hope of growing the discussion moving forward.

Much research has addressed the assumptions made in Gode and Sunder (1993). Cliff and Bruten (1997a) provide insight into the funneling power of the markets containing strictly decreasing demand schedules and increasing supply schedules, while Gjerstad and Shachat (2021) contends the convergence of prices in ZI simulations and calls into question how 'harmless' the no-loss constraint really is. The second main contribution of this paper lies in an expansive investigation into how impactful the model assumptions, market rules and setting attributes are on market efficiency when supposedly unimpeded by trader behavior. At the model level, I test variation in assumptions over the set of feasible orders and the choice process placed over them. Each of the assumptions made, both alone and interacted with one another, crucially impact the behavior of the traders and the performance of the market. Evidence of convergence to equilibrium predictions is increased across the board when enforcing any of the assumptions.

Section 2 defines the institution and environment in focus, while Section 3 recounts the zero intelligence models of Gode and Sunder (1993) and Gode, Spear and Sunder (2004). An alternate general equilibrium model of zero intelligence is postulated in Section 4. Section 5 maps out and analyzes a vast simulation investigation into the underlying determinants of the model and environment. Section 6 concludes.

1.2 Double Auction

The double auction is one of the most ubiquitous institutions used in markets around the world. The most common variety is the continuous double auction. In this section, I'll give a brief formal overview of the institution.

A double auction implies at least one agent attempts to buy some amount of a good, while another agent wishes to sell said good. Each agent provides the price he is willing to pay or receive, respectively. If the payment (or bid) made by the buying agent exceeds the price requested (ask) by the seller, the agents transact the good. There are many conventions for determining the price to be paid; the one selected in this paper, and many other market focused papers, is the crossed price convention. In this case, whoever posted the price first between the two transacting agents has his price granted.

The continuous part of this institution is the ability for any agent in the market to place an order at any given time, regardless of whether an order exists on the other side of the market or not. Orders can also be cancelled or replaced at any moment in time. Two main restrictions are generally imposed, however. First, the agent must be able to completely fill their part of a transaction at the time of crossing, meaning he cannot post an ask for more units than he currently owns or a bid for a price at which he is not liquid. The second restriction is dependent on the timing convention of orders. The CDA studied in this paper does not allow for expiration times on orders, meaning an order only leaves the market if cancelled, replaced or filled. As such, the market restricts the trader to at most one order on each side of the market at any given time.[4]

1.3 Zero Intelligence

Partial Equilibrium

Gode and Sunder (1993) took the pattern of reducing intelligence in non-strategic trader behavior models for the double auction to the limit, creating the "zero intelligence" model. While most of the related models assume some kind of history- or time-dependent driving force behind pricing decisions, ZI agents pay no mind to the state or history of the market. The model can be briefly summarized as follows.

Traders are given roles as either buyers or sellers. Buyers all hold redemption value schedules $\{r_1, .., r_{n_b}\}$ where n_b denotes the number of units the buyers hold at the

[4]In markets with limit orders that contain expiration times, order-shredding (or submitting multiple limit orders with staggered expirations) is feasible.

beginning of the market. Sellers place orders to recover unit costs for their n_s units of the single good, following the cost schedule $\{c_1, .., c_{n_s}\}$. In the base form of the model, all traders choose their order price from $U[0, M]$ via i.i.d. draws. A preferred version, ZI-Constrained (ZI-C), suggests that agents participate in a market which enforces order placement conducive with unit resale at no-loss to the agents. This means buyers draw from $U[0, r_i]$ and sellers draw from $U[c_j, M]$.

Orders are restricted to single unit quantities and must reduce the best bid-ask spread to be placed. Units are transacted in order (i.e. lowest redemption value and highest cost first). After a bid crosses an ask, or vice versa, the orderbook is reset. Traders are held from re-entry until all traders have traded their k^{th} unit. Simulated markets with these traders are run until all intramarginal units have been cleared, eliminating one of the two main drivers of inefficiency.[5]

General Equilibrium

Gode, Spear and Sunder (2004) returned to the ZI paradigm about a decade later to bring zero intelligence to general equilibrium settings (I refer to this model as GSS henceforth). Their environment of choice was the simple two-good Edgeworth box. Adjustments to the original partial equilibrium model are described below.

Traders may participate on either side of the market, and in fact, participate on both sides simultaneously at each entry. Instead of redemption value and cost schedules, traders are induced with Cobb-Douglas preferences. Utility function parameters and initial endowments determine whether traders are more inclined to buy or sell. As in

[5]The other main driver being the trade of extramarginal units in place of intramarginal ones.

ZI-C, traders obey a no-loss constraint in which they will only place orders above their current indifference curve.

The price selection process is predicated on choosing some angle (i.e. relative price) that satisfies the no-loss constraint. For example, with no restriction on quantity, the angle may lie anywhere in $[MRS_c, 1/2\pi]$ in radians, where MRS_c is the marginal rate of substitution at the trader's current allocation. The step-size of the order in the Edgeworth box, however, is restricted. The length of the order vector is set via $r = \sqrt{x^2 + y^2}$, where r is constant for the duration of the market across all traders.[6] This means the price choice space is restricted further, such the lowest price is that of the order vector which lies secant to the trader's current indifference curve on the side of entry.

1.4 Model

I present an alternative version of zero intelligence in the Edgeworth box. While the flavor of the 1993 and 2004 models remains, the majority of their assumptions are relaxed to consider more generalized environments. The model rests in a two-good (X and Y) Edgeworth box with two types of traders, with rule and trader behavior adjustments as follows.

Traders are induced with constant elasticity of substitution (CES) preferences

$$u(x_i, y_i) = c_{\Delta}((a_{\Delta} x_i)^{r_{\Delta}} + (b_{\Delta} y_i)^{r_{\Delta}})^{\frac{1}{r_{\Delta}}} \tag{1.1}$$

which, depending on r, can represent other popular preferences such as Cobb Douglas

[6] x and y are the quantity of x and y to be traded in the event the order is crossed.

($r \to 0$), Leontief ($r \to \infty$) or perfect substitutes ($r = 1$). Each trader can be either a natural buyer or natural seller. As in GSS, such a distinction is determined via the utility parameters (here c, a, b and r) and initial endowments. Note that Δ denotes the type of trader, with $\Delta = b$ for natural buyers and $\Delta = b$ for natural sellers. Alternatively, a natural buyer (seller) can be described as a trader whose marginal rate of substitution is greater (less) than the competitive equilibrium (CE) price when evaluated at his initial endowment.

Traders randomly enter the market one at a time. When a trader enters, he determines the side he will place an order on by flipping a weighted coin. The weights are the relative areas[7] on either side of the market available for order placement.[8] The trader then uniformly randomly chooses an (x, y) bundle from the feasible set. As no no-loss constraint is imposed on the trader, this feasible set can be defined as follows for bids and asks:

$$Buy: \quad [X_{current,i}, X_{endow,b} + X_{endow,s}] \times [0, Y_{current,i}] \tag{1.2}$$

$$Sell: \quad [0, X_{current,i}] \times [Y_{current,i}, Y_{endow,b} + Y_{endow,s}] \tag{1.3}$$

$X_{current,i}$ is the X holding of trader i in his current allocation, and $X_{endow,b} + X_{endow,s}$ denotes the total X holding of a buyer-seller pair at the inception of a market.

The exchange does not enforce a spread reduction rule, allowing any order

[7]Area, here, is the size of the rectangle representing all bundles that satisfy a weak improvement in one good and a weak loss in the other (i.e. a proper bid or ask).

[8]This is admittedly a mild improvement in intelligence over the natural choice of an evenly weighted coin flip. First, this choice is made to help avoid runs of entries on the same side of the market (especially in certain treatments of the simulation panel to be described in Section 1.5). Second, I argue that this decision imbues far less control over the capabilities of the traders and the convergence of the market compared to enforcement of an assumption like a no-loss constraint.

choice to be post-able to the orderbook. In the same vein, the orderbook is not reset upon an order crossing, meaning the book lives the length of the market trading period.

1.5 Simulations

In this section, I present a panel of simulations. Each assumption or rule housed in the model, market, and/or setting is incrementally varied, yielding market outcomes for each potential iteration. The reason for such an expansive investigation is two-fold: (1) it provides a proper test of the model presented in this paper, and (2) it creates the most complete test of the zero intelligence paradigm in a single paper. Below, I give the design for the panel, followed by analysis of each environment at the market level.

1.5.1 Design

Two defining assumptions of the model are primed for variation. First, the set of admissible orders is determined by the inclusion/exclusion of a (budget) no-loss constraint. Prior research (Gjerstad and Shachat (2021)) has pointed out the importance of the budget constraint in funneling ZI-C traders to equilibrium. Second, the selection process is one of the main deviations in this model from Gode, Spear and Sunder. The sequential process of angle (price) and quantity selection is tested against the likely less advantageous lattice choice method.

With regards to the market, Gode and Sunder (1993) made three distinct choices about the rules defining their double auction: (1) orders must be for a single unit,[9] (2) the orderbook is refreshed after any clearing takes place, and (3) trades are priced at that specified in the earlier of the crossing orders. The first two present differences in ZI and the model presented in this paper, and thus are prime variation candidates. The third assumption matches that of this paper, and is also less interesting to test[10]; as such, it is not tested here. A major assumption made in ZI (as well as a vast array of other such models) which could be considered either feature of the institution or choice of the model is the spread reduction rule. I test the (lack of) enforcement of such a rule.

Wholly, these factors combine to create a full factorial (simulated) experimental design with a total of 2^5 treatments. The five main effects and and 27 interactions are tested in Section 1.5.2, with the version described in Section 1.4 representing the control/holdout. Each of the 'factors' in the design are named for analysis as follows: spread reduction (SR), single unit (SU), lattice/angle (LA), orderbook reset (OBR), and no loss (NL).

For each of the 32 variations, I run 250 simulations. Each simulation has 3600 market entries across 12 rounds. Eight computerized traders make up each market, half induced to be natural buyers and half as natural sellers.[11] At the beginning of each round, the endowments of each simulated trader is reset, as is the exchange history.

[9] In GSS (2004), this is amended to have the order vector maintain some uniform length, such that the step-size on the market path is fixed. Assuming single unit order size in general equilibrium can be thought of as a halfway point between full dominion over order choice, and being constrained to a set arc of potential orders. Any results found for the single unit restriction can be expected to inflate if a step-size rule was instead enforced.

[10] The natural alternative here is an even split between the crossing prices, however this just adjusts the split of the gains from trade. While it is feasible this could impact trader behavior and potential shading strategies, zero intelligence traders are incapable of such response, making the factor relatively uninteresting to test.

[11] Buyer endowments are $(x = 3, y = 23)$ with utility parameters $(c = 0.113, a = 0.825, b = 0.175, r = 0.5)$. Buyer endowments are $(x = 11, y = 3)$ with utility parameters $(c = 0.099, a = 0.6875, b = 0.3125, r = 0.5)$.

1.5.2 Analysis

Model Performance

	Mean	St.Dev.
Price	2.27	1.89
$\lvert Price - CE \rvert$	1.45	1.86
Alloc. Eff.	0.66	0.17
Dist. Eff.	0.33	0.13
# Trades	20.13	5.18
RMSE	2.05	8.35
SellerMRS	2.03	0.46
BuyerMRS	3.07	0.63

Table 1.1: Outcome means for the main model in its intended form: no spread reduction, multiple divisible units, lattice choice, no orderbook reset, and no no-loss constraint.

Table 1.1 provides an overview of the model, ZI-G, in its intended form (as described in Section 4). Each estimate shows the average outcome across the 250 simulations ran for this state of the model. As shown, the model performs well in price space. The average round-average price falls within 0.17 units of the general CE price of 2.44 (though rather imprecisely). Average per-trade deviation in price from CE is relatively tight, with the majority of prices falling between 1 and 4 units. The root mean-squared error is moderately acceptable and reports a similar sentiment.

In allocations, the performance of this model variety varies depending on the statistic. Allocative efficiency, measured as the sum of utility gained relative to the sum of gains in equilibrium, shows mild convergence at 0.66, while distance efficiency, or the proportion of the distance from initial endowment to equilibrium traveled by the market, is quite poor at 0.33.[12] The marginal rate of substitution of the traders' final allocations can provide another viewpoint on convergence in allocation space. Here, the

[12]One consolation for these measures is their likely dependence on the length of the market; I report simulations with much longer-lived periods at the end of the section to check this.

seller (buyer) MRS reported is the MRS of the final allocation of an aggregated seller (buyer) agent who aggregates over all natural sellers (buyers) in the market, so as to appear in a standard Edgeworth box. Table 1.1 reports the round-end MRS for these aggregated agents. In equilibrium, both measures should equal the general CE price; in this measure the market performs relatively well, with buyers and sellers lying around 0.5 units away on either side. Thus, while the distance efficiency measure suggests the market has a long way to travel, the allocative efficiency and MRS measures suggest the market has moved enough to have realized a majority of the gains from trade.

Panel Investigation

	Dependent variable:								
	Price	$\lvert Price - CE \rvert$	RMSE	Order Size	# Trades	Trade Size	Seller MRS	Buyer MRS	Alloc. Eff.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Spread Reduction (SR)	−0.298	−0.267	−0.249	−5.930***	7.145***	−0.220***	0.029***	−0.035***	0.007**
	(0.732)	(0.732)	(6.918)	(1.738)	(0.108)	(0.013)	(0.008)	(0.012)	(0.003)
Single Unit (SU)	12.885***	11.272***	12.008*	−13.942***	9.230***	−2.235***	−0.264***	0.636***	−0.014***
	(0.732)	(0.732)	(6.918)	(1.738)	(0.108)	(0.013)	(0.008)	(0.012)	(0.003)
Lattice/Angle (LA)	−0.643	0.208	2.092	23.201***	136.953***	−1.161***	0.042***	0.013	0.012***
	(0.732)	(0.732)	(6.918)	(1.738)	(0.108)	(0.013)	(0.008)	(0.012)	(0.003)
OB Reset (OBR)	0.857	0.817	2.098	−0.139	−5.112***	0.314***	−0.048***	0.041***	0.028***
	(0.732)	(0.732)	(6.918)	(1.738)	(0.108)	(0.013)	(0.008)	(0.012)	(0.003)
No Loss (NL)	0.155	−0.816	−1.368	0.352	−18.100***	0.324***	−0.715***	1.292***	−0.303***
	(0.748)	(0.748)	(7.069)	(1.738)	(0.108)	(0.013)	(0.008)	(0.013)	(0.003)
⋮									
Constant	2.268***	1.447***	2.054	14.942***	20.132***	3.155***	2.035***	3.066***	0.656***
	(0.518)	(0.517)	(4.892)	(1.229)	(0.076)	(0.009)	(0.005)	(0.009)	(0.002)
Observations	95,424	95,424	95,424	96,000	96,000	95,424	95,422	95,424	96,000
R^2	0.024	0.020	0.003	0.019	0.989	0.826	0.464	0.596	0.607
Adjusted R^2	0.023	0.020	0.002	0.019	0.989	0.826	0.464	0.596	0.607
Note:									*p<0.1; **p<0.05; ***p<0.01

Table 1.2: Interaction regression results. First order effects are reported here, the rest are in tables to follow.

Tables 1.2-1.4 present the treatment analysis of the full factorial design (along with Appendix A). Each of the five assumptions/rules being tested is given an indicator, $I(rule)$, with a value of 1 representing the presence of the constraint in the simulations.

The estimation process is represented by the following interaction design:

$$Y_i = \alpha + \sum_{i \in Rules} \beta_i I(i=1) + \sum_{i \in Rules} \sum_{j \in Rules_{/\{i\}}} \beta_{ij} I(i=1) I(j=1) + \cdots + \beta_{ijklm} \prod_{i \in Rules} I(i=1) \tag{1.4}$$

The main effects of the model, i.e. first summation from 1.4, is provided in Table 1.2.[13] Prices seem to be more sensitive to interactions of the main assumptions, as opposed to having only one included. Round-average price, price deviation from CE, and RMSE all show insignificant effects for all assumptions aside from a constraint imposing single unit orders. Collapsing the lattice of possible bundles to a subset of bundles along the line $q = 1$ naturally increases likelihood of higher prices, yielding the massive increase in price and price variation when only the SU assumption is imposed. As might be anticipated, imposing a spread reduction (while leaving all other assumptions unimposed) seems to funnel activity in the market, leading to a very tight MRS spread ($\sim$0.10 on either side of the CE price) and a mild, but significant, improvement in allocative efficiency. From the control, swapping to an angle choice process has a similar impact on activity, though with a slightly larger efficiency gain and no improvement in final buyer allocations. Resetting the orderbook in an otherwise unconstrained market sees the largest gain in efficiency, though a larger spread in final MRS, likely pointing to less even gains across traders with a few seeing larger gains. Adding a no-loss constraint on its own seems to be harmful to the success of the market, however this is likely due to its interaction with the lattice choice method. Forcing bundles to be chosen above the indifference curve naturally imposes higher likelihoods for prices less likely to result in trades; as is reflected in the 90% reduction in trade count.

[13] By construction, the constants report the estimates from Table 1.1.

The second order interactions are reported in Table 1.3. Much like Table 1.2, prices see little adjustment (of any significance) aside from a few interesting interactions. Swapping to an angle choice procedure while SU is imposed (and all other assumptions relaxed) essentially reverses the massive inflation in prices seen in SU estimates from Table 1.2. Weak utility improvement imposition in a single unit order market provides the same regression in prices. Resetting the orderbook in a lattice choice market (with no other constraints) doesn't allow for funneling of prices, leading to price inflation. Adding a second constraint systematically reduces round-average trade counts across the board, with the lone exceptions both involving orderbook resetting.

Allocation adjustments seem to be the main beneficiary of imposing a second assumption. All significant estimates except for seller MRS being positive, paired with all-but-one significant estimates being negative implies convergence in allocation space. A few act as recoveries, with the damage of orderbook resetting, the no-loss constraint and single unit orders being reclaimed by inclusions of a second constraint. Angle choice and spread reduction restrictions are especially effective in progressing NL markets to a more successful final allocation. Similarly, the vast majority of interactions reflect in an increase in efficiency. Larger improvements are reflective of reversals for NL markets mostly, while smaller improvements are most often continuations of efficiency gain in SR and LA markets.

Third and fourth order interactions are reported in Appendix A.1. Tertiary interactions (Table A.1) show mostly decays in market success. Most, if not all, of the improvements seen in Table 1.3 are reversed when adding a third assumption to the market (assuming the remaining two assumptions are relaxed). As most of the measures

	Dependent variable:								
	Price	$\|Price - CE\|$	RMSE	Order Size	# Trades	Trade Size	Seller MRS	Buyer MRS	Alloc. Eff.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
⋮									
SR:SU	0.347	0.336	0.267	5.930^{**}	-4.172^{***}	0.226^{***}	-0.009	-0.017	0.006
	(1.035)	(1.035)	(9.784)	(2.458)	(0.153)	(0.018)	(0.011)	(0.017)	(0.005)
SR:LA	1.082	1.117	7.618	-20.169^{***}	-19.542^{***}	0.159^{***}	0.011	0.020	0.003
	(1.035)	(1.035)	(9.784)	(2.458)	(0.153)	(0.018)	(0.011)	(0.017)	(0.005)
SR:OBR	0.911	0.808	2.823	2.910	-4.815^{***}	0.256^{***}	0.015	-0.028	0.017^{***}
	(1.035)	(1.035)	(9.784)	(2.458)	(0.153)	(0.018)	(0.011)	(0.017)	(0.005)
SR:NL	0.156	0.091	0.091	0.237	-4.473^{***}	-0.083^{***}	0.441^{***}	-0.938^{***}	0.269^{***}
	(1.046)	(1.046)	(9.891)	(2.458)	(0.153)	(0.018)	(0.011)	(0.018)	(0.005)
SU:LA	-13.252^{***}	-11.392^{***}	-14.385	-23.201^{***}	-71.869^{***}	1.134^{***}	0.141^{***}	-0.674^{***}	0.212^{***}
	(1.035)	(1.035)	(9.784)	(2.458)	(0.153)	(0.018)	(0.011)	(0.017)	(0.005)
SU:OBR	-1.097	-0.987	-1.680	0.139	1.625^{***}	-0.315^{***}	0.054^{***}	-0.071^{***}	-0.025^{***}
	(1.035)	(1.035)	(9.784)	(2.458)	(0.153)	(0.018)	(0.011)	(0.017)	(0.005)
SU:NL	-11.234^{***}	-10.236^{***}	-10.908	-0.352	-7.365^{***}	-0.244^{***}	0.174^{***}	-0.219^{***}	0.086^{***}
	(1.047)	(1.047)	(9.899)	(2.458)	(0.153)	(0.018)	(0.011)	(0.018)	(0.005)
LA:OBR	7.055^{***}	7.068^{***}	61.778^{***}	-4.222^{*}	-64.725^{***}	-0.388^{***}	0.004	-0.065^{***}	0.038^{***}
	(1.035)	(1.035)	(9.784)	(2.458)	(0.153)	(0.018)	(0.011)	(0.017)	(0.005)
LA:NL	0.213	-0.186	-2.006	-35.880^{***}	-108.977^{***}	-1.817^{***}	0.355^{***}	-0.976^{***}	0.464^{***}
	(1.047)	(1.046)	(9.893)	(2.458)	(0.153)	(0.018)	(0.011)	(0.018)	(0.005)
OBR:NL	-0.855	-0.812	-2.095	0.178	4.852^{***}	-0.131^{***}	0.014	0.038^{**}	-0.045^{***}
	(1.059)	(1.059)	(10.009)	(2.458)	(0.153)	(0.019)	(0.011)	(0.018)	(0.005)
⋮									
Constant	2.268^{***}	1.447^{***}	2.054	14.942^{***}	20.132^{***}	3.155^{***}	2.035^{***}	3.066^{***}	0.656^{***}
	(0.518)	(0.517)	(4.892)	(1.229)	(0.076)	(0.009)	(0.005)	(0.009)	(0.002)
Observations	95,424	95,424	95,424	96,000	96,000	95,424	95,422	95,424	96,000
R^2	0.024	0.020	0.003	0.019	0.989	0.826	0.464	0.596	0.607
Adjusted R^2	0.023	0.020	0.002	0.019	0.989	0.826	0.464	0.596	0.607
Note:									*p<0.1; **p<0.05; ***p<0.01

Table 1.3: Interaction regression results for second order interactions. This is a continuation of the regression estimates in Table 1.2.

have one or two seemingly negatively-associated assumptions, and each assumption is enforced in six of the ten tertiary interactions, a systematic mild decay is not overly surprising. Quaternary interactions are reported in Table A.2.

Table 1.4, which reports the quinary interaction, provides an interesting connection to the literature. As the GSS model enforces all five assumptions, flipping the signs in Table 1.4 allows the coefficients to represent the average impact of relaxing a single assumption in their model. A slight tightening ($\sim$0.17 reduction) of the MRS

	Price	$\lvert Price - CE \rvert$	RMSE	Order Size	# Trades	Trade Size	Seller MRS	Buyer MRS	Alloc. Eff.
	Dependent variable:								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
⋮									
SR:SU:LA:OBR:NL	-3.235	-3.403	-33.996	10.169	4.865***	0.170***	-0.072**	0.099**	-0.004
	(2.937)	(2.937)	(27.764)	(6.952)	(0.431)	(0.052)	(0.031)	(0.049)	(0.014)
Constant	2.268***	1.447***	2.054	14.942***	20.132***	3.155***	2.035***	3.066***	0.656***
	(0.518)	(0.517)	(4.892)	(1.229)	(0.076)	(0.009)	(0.005)	(0.009)	(0.002)
Observations	95,424	95,424	95,424	96,000	96,000	95,424	95,422	95,424	96,000
R^2	0.024	0.020	0.003	0.019	0.989	0.826	0.464	0.596	0.607
Adjusted R^2	0.023	0.020	0.002	0.019	0.989	0.826	0.464	0.596	0.607
Note:									*p<0.1; **p<0.05; ***p<0.01

Table 1.4: Interaction regression results for fifth order interaction. This is a continuation of the regression estimates in Table 1.2.

spread, along with minimal change in allocative efficiency on average results from a relaxation of one of the five assumptions. As Table 1.5 will show, however, individual comparisons reveal relaxing the angle choice provides most of this variation.

Treatment level averages are reported in Table 1.5; each entry being the summation of the relevant coefficients from Tables 1.2-1.4, A.1 and A.2. For each of the measures representative of convergence or market success, the best performing treatment's mean is bolded in black, the worst performing is highlighted in red. SR:SU:OBR markets performed the worst in price measures, with the mix of lattice-choice-single-unit price inflation and lack of price funneling (as the orderbook can't properly age) results in prices well above CE. A no-loss constraint as the lone assumption produced round-average prices just 0.02 units away from CE. Average price deviation and root-mean-squared error were both minimized by an NL market type, though paired with a spread reduction rule now. NL markets not paired with an angle choice method seemed to perform poorly in allocation measures, likely due to higher variation in prices and thus more frequent deviation from the equilibrium path. Markets that are fully restricted,

					Outcome:									
SR	SU	LA	OBR	NL	Price	$\|Price - CE\|$	RMSE	Order Size	# Trades	Trade Size	Seller MRS	Buyer MRS	Alloc. Eff.	Dist. Eff.
0	0	0	0	0	2.27 (1.89)	1.45 (1.86)	2.05 (8.35)	14.94 (1.04)	20.13 (5.18)	3.15 (0.68)	2.03 (0.46)	3.07 (0.63)	0.66 (0.17)	0.33 (0.13)
0	0	0	0	1	**2.42 (0.61)**	0.63 (0.36)	0.69 (0.38)	15.29 (0.87)	2.03 (1.19)	3.48 (1.39)	1.32 (0.28)	4.36 (0.68)	0.35 (0.19)	0.26 (0.15)
0	0	0	1	0	3.12 (3.09)	2.26 (3.05)	4.15 (11.78)	14.8 (1.02)	15.02 (3.38)	3.47 (0.84)	1.99 (0.41)	3.11 (0.54)	0.68 (0.15)	0.34 (0.13)
0	0	0	1	1	2.42 (0.64)	0.64 (0.38)	0.69 (0.40)	15.33 (0.89)	1.77 (0.98)	3.66 (1.49)	1.29 (0.27)	4.44 (0.66)	**0.34 (0.19)**	0.25 (0.15)
0	0	1	0	0	1.62 (1.78)	1.66 (1.77)	4.15 (21.39)	38.14 (289.29)	157.09 (9.88)	1.99 (0.20)	2.08 (0.57)	3.08 (0.60)	0.67 (0.17)	0.15 (0.16)
0	0	1	0	1	1.99 (0.16)	0.65 (0.11)	0.77 (0.11)	2.61 (0.21)	30.01 (4.69)	0.50 (0.08)	1.72 (0.11)	3.40 (0.18)	0.83 (0.06)	0.60 (0.06)
0	0	1	1	0	9.54 (155.42)	9.54 (155.42)	**68.02 (1476.74)**	33.78 (210.17)	87.25 (5.77)	1.92 (0.26)	2.03 (0.48)	3.06 (0.53)	0.73 (0.14)	0.16 (0.15)
0	0	1	1	1	1.96 (0.19)	0.80 (0.12)	0.93 (0.13)	2.67 (0.26)	24.39 (3.96)	0.53 (0.11)	1.57 (0.12)	3.65 (0.24)	0.76 (0.08)	0.52 (0.07)
0	1	0	0	0	15.15 (1.53)	12.72 (1.53)	14.06 (1.67)	1.00 (0.00)	29.36 (4.58)	0.92 (0.05)	1.77 (0.24)	3.70 (0.69)	0.64 (0.17)	0.24 (0.08)
0	1	0	0	1	4.07 (0.46)	1.67 (0.43)	1.79 (0.39)	1.00 (0.00)	3.90 (1.55)	1.00 (0.00)	1.23 (0.15)	4.78 (0.40)	0.42 (0.14)	0.24 (0.09)
0	1	0	1	0	14.91 (1.75)	12.55 (1.73)	14.48 (1.99)	1.00 (0.00)	25.87 (3.96)	0.92 (0.05)	1.78 (0.25)	3.67 (0.71)	0.64 (0.18)	0.24 (0.09)
0	1	0	1	1	4.07 (0.47)	1.68 (0.43)	1.80 (0.39)	1.00 (0.00)	3.75 (1.51)	1.00 (0.00)	**1.21 (0.15)**	**4.81 (0.39)**	0.41 (0.14)	0.23 (0.09)
0	1	1	0	0	1.26 (0.22)	1.53 (0.11)	1.77 (0.25)	1.00 (0.00)	94.45 (5.23)	0.89 (0.03)	1.95 (0.22)	3.04 (0.22)	0.87 (0.06)	0.53 (0.08)
0	1	1	0	1	1.94 (0.20)	0.72 (0.11)	0.84 (0.12)	1.00 (0.00)	18.41 (1.80)	1.00 (0.00)	1.96 (0.13)	3.05 (0.17)	0.91 (0.04)	0.71 (0.06)
0	1	1	1	0	1.67 (0.28)	1.85 (0.20)	2.43 (0.51)	1.00 (0.00)	71.05 (4.48)	0.90 (0.04)	1.95 (0.22)	3.05 (0.29)	0.84 (0.08)	0.50 (0.10)
0	1	1	1	1	1.99 (0.20)	0.77 (0.11)	0.90 (0.13)	1.00 (0.00)	17.54 (1.73)	1.00 (0.00)	1.91 (0.13)	3.11 (0.17)	0.90 (0.04)	0.69 (0.06)
1	0	0	0	0	1.97 (1.61)	1.18 (1.58)	1.80 (8.77)	9.01 (1.43)	27.28 (5.19)	2.93 (0.52)	2.06 (0.46)	3.03 (0.56)	0.66 (0.17)	0.36 (0.13)
1	0	0	0	1	2.28 (0.32)	**0.45 (0.20)**	**0.53 (0.22)**	9.60 (2.29)	4.70 (1.51)	3.18 (0.82)	1.79 (0.35)	3.39 (0.59)	0.63 (0.15)	0.48 (0.13)
1	0	0	1	0	3.74 (13.55)	2.81 (13.54)	6.72 (46.79)	11.78 (1.21)	17.35 (3.31)	3.50 (0.76)	2.03 (0.41)	3.04 (0.50)	0.71 (0.15)	0.35 (0.13)
1	0	0	1	1	2.33 (0.36)	0.52 (0.22)	0.61 (0.26)	11.28 (1.88)	3.34 (0.97)	3.77 (1.08)	1.64 (0.31)	3.64 (0.58)	0.61 (0.16)	0.46 (0.13)
1	0	1	0	0	2.41 (8.12)	2.51 (8.12)	11.51 (100.31)	12.04 (123.6)	144.69 (8.52)	1.93 (0.21)	**2.12 (0.58)**	3.06 (0.69)	0.68 (0.17)	**0.13 (0.16)**
1	0	1	0	1	2.04 (0.15)	0.61 (0.11)	0.74 (0.12)	1.37 (0.26)	35.39 (4.50)	0.46 (0.06)	1.82 (0.10)	3.23 (0.14)	0.87 (0.04)	0.65 (0.05)
1	0	1	1	0	7.35 (33.47)	7.34 (33.46)	43.49 (299.96)	19.97 (42.80)	82.36 (4.97)	1.89 (0.27)	2.07 (0.44)	3.02 (0.51)	0.75 (0.13)	0.16 (0.15)
1	0	1	1	1	1.93 (0.17)	0.84 (0.11)	0.96 (0.12)	2.28 (0.26)	26.67 (3.33)	0.50 (0.09)	1.59 (0.10)	3.60 (0.18)	0.78 (0.06)	0.53 (0.06)
1	1	0	0	0	15.20 (1.52)	12.79 (1.52)	14.08 (1.59)	1.00 (0.00)	32.33 (4.27)	0.93 (0.05)	1.79 (0.23)	3.65 (0.68)	0.66 (0.17)	0.24 (0.08)
1	1	0	0	1	4.63 (0.37)	2.21 (0.35)	2.27 (0.30)	1.00 (0.00)	4.49 (1.57)	1.00 (0.00)	1.34 (0.15)	4.54 (0.40)	0.52 (0.14)	0.29 (0.09)
1	1	0	1	0	**15.27 (1.73)**	**12.92 (1.71)**	14.97 (1.98)	1.00 (0.00)	28.19 (3.84)	0.92 (0.05)	1.79 (0.25)	3.64 (0.72)	0.65 (0.18)	0.24 (0.09)
1	1	0	1	1	4.53 (0.40)	2.13 (0.36)	2.21 (0.30)	1.00 (0.00)	4.41 (1.55)	1.00 (0.00)	1.32 (0.15)	4.58 (0.40)	0.50 (0.14)	0.29 (0.09)
1	1	1	0	0	1.42 (0.25)	1.66 (0.15)	2.04 (0.37)	1.00 (0.00)	87.40 (4.70)	0.90 (0.03)	1.97 (0.22)	3.03 (0.24)	0.86 (0.06)	0.52 (0.09)
1	1	1	0	1	1.91 (0.20)	0.74 (0.12)	0.86 (0.12)	1.00 (0.00)	19.90 (1.76)	1.00 (0.00)	2.05 (0.13)	**2.93 (0.15)**	**0.94 (0.03)**	**0.75 (0.06)**
1	1	1	1	0	1.73 (0.29)	1.90 (0.21)	2.53 (0.52)	1.00 (0.00)	72.48 (4.29)	0.90 (0.04)	1.96 (0.22)	3.03 (0.29)	0.85 (0.07)	0.50 (0.09)
1	1	1	1	1	1.99 (0.20)	0.78 (0.11)	0.91 (0.12)	1.00 (0.00)	18.78 (1.67)	1.00 (0.00)	2.00 (0.13)	3.00 (0.16)	0.92 (0.03)	0.73 (0.06)

Table 1.5: Outcome averages by treatment. Observations are at the round-average or round-end level. The left panel shows the assumptions enforced. Black bolded estimates are the ‘best’ in the column, while red are the ‘worst’.

aside from orderbook resetting, lead the pack in both measures of efficiency, as well as buyer MRS. GSS markets outperform ZI-G markets in both measures of efficiency and both measures of price variation, with MRS splits being within 0.1 of each other and the ZI-G average price mildly outperforming GSS prices (though insignificantly).

Panel Main Effects

While a full factorial analysis is informative of the incremental response to variation in the market's design, the main effect of relaxing or enforcing an assumption may be hard to discern. Tables 1.6 and 1.7 show such an effect for efficiencies and prices, respectively.

	Allocative Efficiency			
	Mean (I=0)	Mean (I=1)	Difference	P-Value
Spread Red.	0.67	0.72	0.06	0.00
Single Unit	0.67	0.72	0.05	0.00
Angle Choice	0.57	0.82	0.25	0.00
OB Reset	0.70	0.69	-0.01	0.00
No Loss	0.72	0.67	-0.05	0.00
	Distance Efficiency			
Spread Red.	0.37	0.72	0.35	0.00
Single Unit	0.36	0.72	0.36	0.00
Angle Choice	0.30	0.82	0.52	0.00
OB Reset	0.41	0.69	0.29	0.00
No Loss	0.31	0.67	0.36	0.00

Table 1.6: Efficiency. T-tests to show main effect of each treatment factor on allocative and distance efficiency.

In first differences, allocative and distance efficiencies show quite different trends, though with a unifying mechanism. Spread reduction and single unit assumptions show small, yet significant improvements in allocative efficiency while orderbook resetting and no-loss constraints lead to mild reductions. Angle choice leads to a rather large improvement, likely driven by the more CE-localized trade prices and increased trade count. Distance efficiency on the other hand sees substantial improvements across the board. Such a mild response in one measure and improvement in the other points towards a more equitable redistribution of gains from trade in utility-terms so as to bring the path traveled by the market closer to the equilibrium path.

	Average Trade Price			
	Mean (I=0)	Mean (I=1)	Difference	P-Value
Spread Red.	4.42	4.42	-0.00	0.99
Single Unit	3.10	5.74	2.64	0.00
Angle Choice	6.19	2.67	-3.52	0.00
OB Reset	3.92	4.92	1.00	0.00
No Loss	6.16	2.66	-3.51	0.00

Table 1.7: Price Main Effects.

Prices are less consistently impacted by the enforcement or relaxation of the tested assumptions. Spread reduction inclusion sees zero adjustment in average price; Figure A.1 confirms the distributions are nearly identical. The other four assumptions are split in their direction of change. Single unit order and orderbook reset constraints both inflate prices on average, though with much different responses in their distributions. A massive redistribution, or flattening, occurs with single-unit-order constrained markets, while orderbook resetting markets see a small shift right in price (as shown in Figure A.1). Markets with traders who either use an angle choice process for their orders or follow a no-loss constraint report nearly identical adjustments in average round-average price, over halving the estimates from above 6 units of y per unit of x, to relative prices just over 0.2 units above the CE prediction. Distributional changes, however, are quite different between these two sets of markets. Figure A.1 shows the mass of the distribution funnelling quite close to the mean for angle-choice markets, however the support for the distribution remains unchanged. No-loss constrained markets also show tighter mass near the CE-price price. The massive reduction in the size of the support is perhaps the more remarkable result when giving traders the intelligence to always obey their own preferences.

1.6 Conclusion

Understanding both the implications of the rules implied by a market institution and the underlying behavior defining trader behavior in said market are crucial tasks in economic research. The main issue plaguing such an endeavor is the clear entanglement between the two. One way to isolate the first study is to place traders with no strategic behavior in the market, so as to let market outcomes be only guided by the rules of the institution. Gode and Sunder (1993) proposed such a model and test in a partial equilibrium setting, and then brought the model to a general equilibrium via the Edgeworth box (Gode, Spear and Sunder (2004)). The proposed zero intelligence traders however either abided by potentially influential behavioral assumptions (such as a no-loss constraint or an order choice process giving more weight to less aggressive prices), or participated in markets with rules that may guide the allocation path. This paper provides (1) a more generalized, lower-'zero' version of zero intelligence in an edgeworth box, and (2) a test of the assumptions made in this model and those mentioned.

Agents participate as two-way traders in a CDA, where there are two goods and two types of traders (thus constructing an Edgeworth box). Despite being induced with CES preferences, and predisposed to prefer one order type to another based on the curvature of said preferences, traders enter the market (and particular side) randomly. Orders are uniformly drawn from a fine lattice placed over the set of all bundles satisfying the properties of a transaction occurring (a gain in one good and loss in the other), with no regard for gains or losses in utility from trade.

Empirically, I test the major assumptions made in ZI models, as well as those made in the models' respective CDA markets, via a novel, expansive simulation procedure. All combinations of the five assumptions either relaxed or enforced are simulated with traders from this paper's model. Each variation was simulated 250 times with each run containing 3600 entries, yielding a data set of 28.8 million market entries and order placements across 96,000 trading periods in 8000 simulated markets. First differences show an improvement across the board when imposing one of the five restrictions, as well as a reorganization of gains from trade resulting in systematic improvements in distance efficiency. An interactions model reports the incremental impact of these assumptions in the full factorial design.

Compared to the traders and setting of Gode, Spear and Sunder (2004), this model is shown to provide a much lower level of zero intelligence in a much less constrained version of general equilibrium (despite also residing in an Edgeworth box). Allocative and distance efficiencies are 0.28 and 0.42 units lower than those found in Gode, Spear and Sunder, while average price is slightly (though insignificantly) closer to CE and price volatility is twice as large in the lower-zero model. When paired with another price-funneling-prone assumption, both spread reduction and no-loss rules provide large improvements in market performance. Interestingly, markets with either pairs of enforced rules or at most one relaxed assumption exhibit the most equilibrating tendencies. Rather unsurprisingly, markets in which traders' order choice behavior is dictated by an angle-choice process or a no-loss tendency benefit immensely in market performance from the added intelligence.

Hopefully, this project displays the benefits of and need for modelling ventures into the wilderness of bounded rationality via the minimal intelligence gate. Many subfields of economics outside of markets could benefit from such practice. Additionally, investigations in more complex, general equilibrium settings are needed and should be pursued moving forward. The world doesn't usually operate in stylized environments with well-known supply and demand schedules or traders content with the exchange on one indivisible unit of a single good.

Chapter 2

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

2.1 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, rolled out in three waves to three separate lake regions, allowed for easier access to on-shore prices across multiple small fishing markets. In each case, the roll-out 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 unit type(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 trade 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 prices 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 those 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), pays 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's (2012), 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 (1997b)), 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. (2008b), 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.

2.2 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.

2.2.1 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 $\hat{m} \equiv m/p_y$.

$$\max_{(x_i,y_i)} u_i(x_i, y_i) \quad s.t. \quad px_i + y_i = \hat{m} \tag{2.1}$$

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

$$Z_i^X(p|(x_{i,o}, y_{i,o})) = \frac{a^{\gamma_i}(y_{i,o} + px_{i,o})}{p(a^{\gamma_i} + p^{\gamma_i} b_i^{\gamma_i})} - x_{i,o} \tag{2.2}$$

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

$$Z^X(p|(x_o, y_o)) = \sum_{i=1}^{N} Z_i^X(p|(x_{i,o}, y_{i,o})) = 0 \tag{2.3}$$

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.

2.2.2 Continuous Double Auction

The market type of choice in this paper is likely the most prolific, 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 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 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.

2.2.3 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.

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 bounded 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 if buying.

The set of orders that have been posted to the orderbook or transacted over the duration of the market (up until the current state of the market) 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 next entry.

2.2.4 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 (p, q) choice is multiplied by the probability that 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

[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.

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

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.

2.2.5 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 traders; 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 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

$$TA(\rho) \equiv \sum_{k=1}^{|\tau_A(\rho)|} \sqrt{q_k}\frac{q_{k,traded}}{q_k} \tag{2.4}$$

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:

$$p_a(a) = \frac{\sum_{\rho \geq a} TA(\rho) + \sum_{\rho \geq a} TB(\rho)}{\sum_{\rho \geq a} TA(\rho) + \sum_{\rho \geq a} TB(\rho) + \sum_{\rho \leq a} RA(\rho)} \tag{2.5}$$

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

$$p_b(b) = \frac{\sum_{\rho \leq b} TB(\rho) + \sum_{\rho \leq b} TA(\rho)}{\sum_{\rho \leq b} TB(\rho) + \sum_{\rho \leq b} TA(\rho) + \sum_{\rho \geq b} RB(\rho)} \tag{2.6}$$

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

[3]See Gjerstad and Dickhaut (1998) for discussion on characteristics of p(a), including monotonicity.

beliefs are over the domain [0,M], with $p_a(0) = 1$ and $p_a(M) = 0$.

2.2.6 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

$$S_i^k = \max\{\max_{\rho \in P}(u_i(x_{i,k}, y_{i,k}) - u_i(x_{i,k-1}, y_{i,k-1})) \cdot p(\rho), 0\} \tag{2.7}$$

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 from the lower bound to 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

[4]The entrant's marginal rate of substitution in this sense can be considered a reservation price.

from having the order fully 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.

2.2.7 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]

2.3 Design

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

2.3.1 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

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

most recent trade in the market, updated and replaced with each new trade).

2.3.2 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).

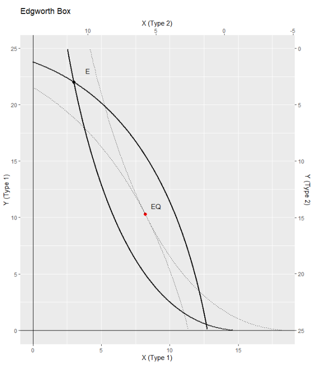

Figure 2.1: Edgeworth box displaying natural buyer and seller preferences. Type 1 refers to buyers, 2 to sellers.

	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 2.1: CES parameters, starting endowments and equilibrium allocations.

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 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 2.1 and Table 2.1.

2.3.3 Laboratory Realization

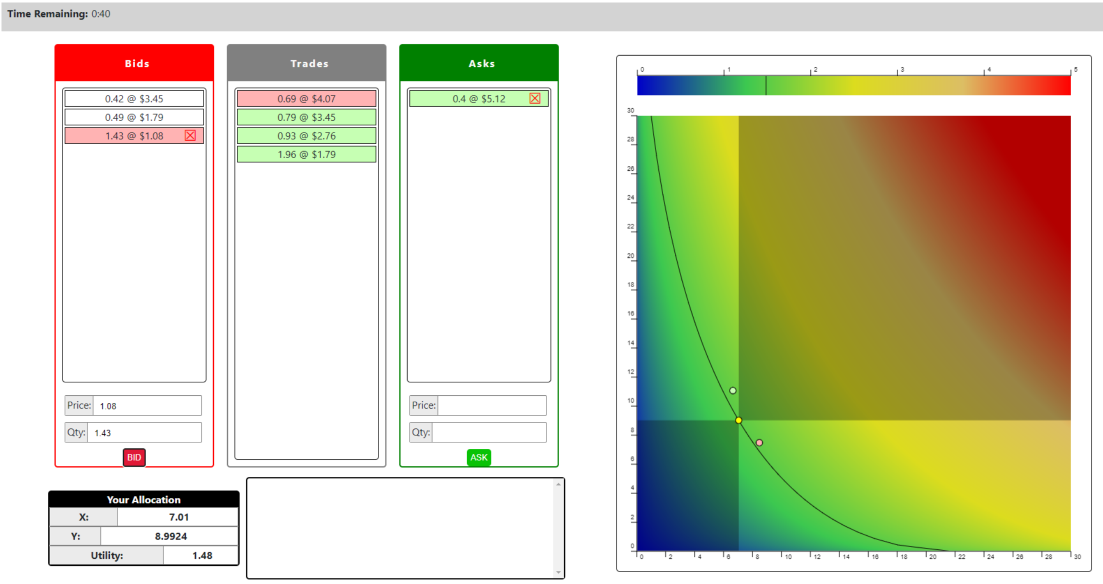

Figure 2.2: User interface for laboratory market sessions. Contains an orderbook, preferences heatmap, allocation box and error log box.

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 2.2. 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.

2.3.4 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]

2.4 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.

[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.

2.4.1 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.

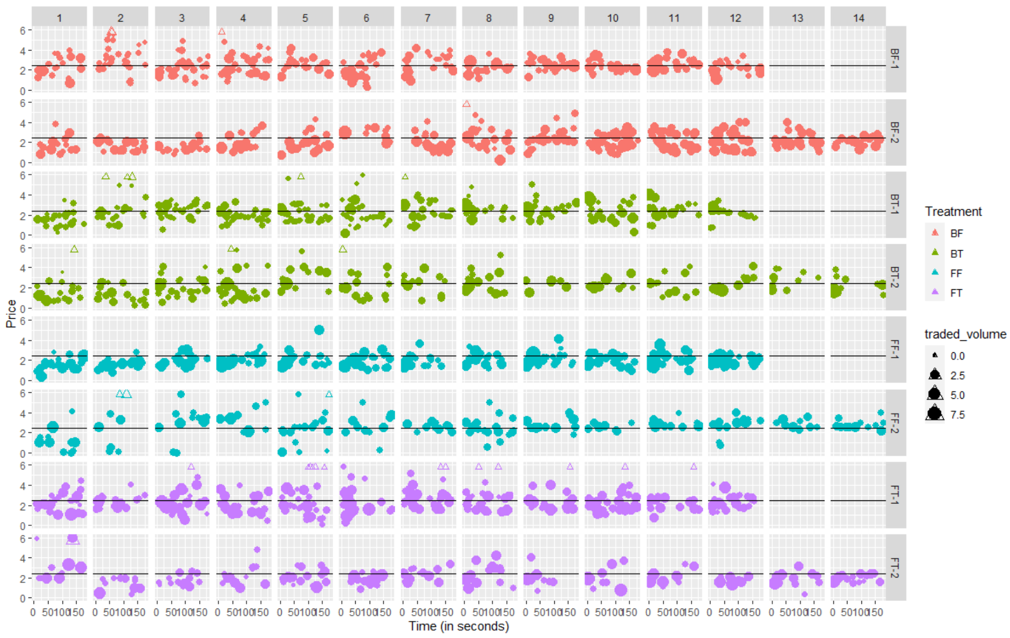

Figure 2.3: Individual transaction prices. Trades with prices above 5 units of y per unit of x are plotted as outliers (triangles along the $p = 5$ line). The black line plots the round-start competitive equilibrium price.

Figure 2.3 plots all trades[9] in each session against the time the exchange marked the 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

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

bargaining advantage, with a majority of prices appearing below the equilibrium price line.

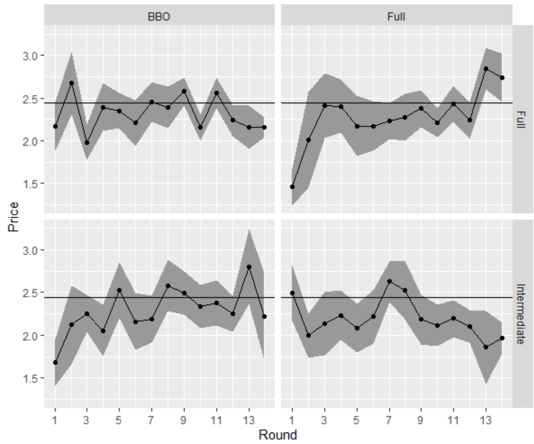

Figure 2.4: Round-average transaction prices. Shaded region shows 95% confidence interval. Columns show levels for orderbook factor. Rows show levels for transaction history level.

A round-average depiction of the prices seen in Figure 2.3 can be found in Figure 2.4. 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
Full	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)
	# 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)
TT	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)
	# 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 2.2: 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 2.2 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 2.4).

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 2.2 (thus from BBO-TT to Full-TT to Full-Full) as a potential improvement path, markets show a monotonic progression

[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.

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.

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 2.4.4, ideally the MRS of these aggregated agents (and each individual trader) will be equal to the competitive equilibrium price. Table 2.2 suggests 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.2.

2.4.2 Allocations

The reallocation of goods throughout a market's 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 2.5 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.

Three criteria for performance in Figure 2.5 are: tightness of cluster, deviation

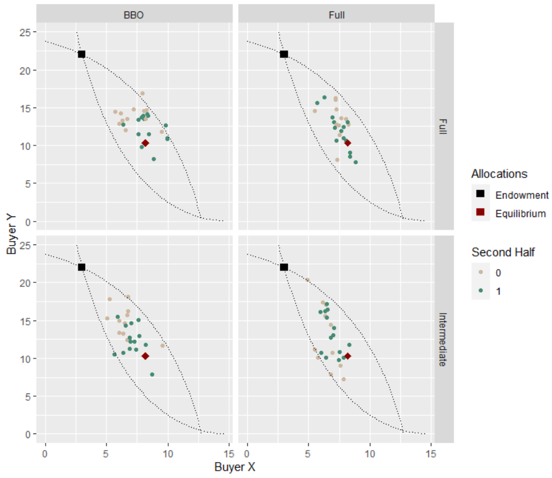

Figure 2.5: Round-end allocations.

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 2.4.1. 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 2.5 is compressed and re-imagined in terms of distance from the equilibrium allocation in Figure 2.6 and Table 2.3. 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 2.3 and via a best fit line in Figure 2.6.

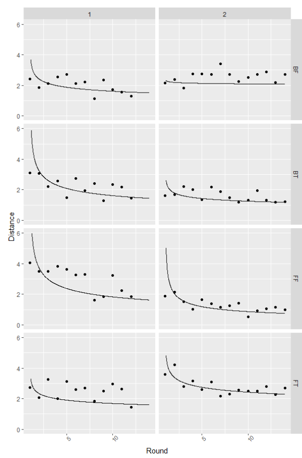

Figure 2.6: Round-end distance of final allocations from equilibrium allocations.

The plots show linear decay in the majority of sessions, with session FF-2

	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 2.3: 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.

being the closest to being exponential. FF sessions also displayed the steepest decay. Table 2.3'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 2.3 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.

2.4.3 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 Williams (2021).

2.4.3.1 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 (2021b) 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 (ZI-G from here) 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 in ZI-G:

- 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,

[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).

the above rules were chosen to match those enforced on the laboratory traders.

2.4.3.2 Estimates

Table 2.4 presents the average allocative efficiency 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) 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 2.4: 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-G markets.[13] 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 agents. 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 agents. Traders in these markets are substituting efficiency in an effort to accommodate more aggressive price discovery behaviors.

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 2.5: 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.

[13]Traders in the version of ZI-G run here place orders subject to a no-loss constraint. Also, no market level spread reduction rule is enforced (to match the rule choice in the experiments). However, ZI-G traders did abide by an internal 'spread reduction' process in that they would only place a new draw if the draw improved their current order in the market.

Much like with final allocation distances (Table 2.3), end-of-round estimates may not be telling the whole story. Table 2.5 displays exponential decay coefficients for the gap in allocative efficiency at the round level, as well as decay for time of least inefficiency across 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.

Cumulative density functions for the round-end trader-level difference in utility gained and expected utility gain in equilibrium are reported in Figure B.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.*

2.4.4 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 tabl". 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 2.7 displays the MRS for the aggregated natural buyer and natural seller across rounds averaged between sessions for each treatment. As discussed in Section 2.4.1 and summarized in Table 2.2, markets with higher transaction history accessibil-

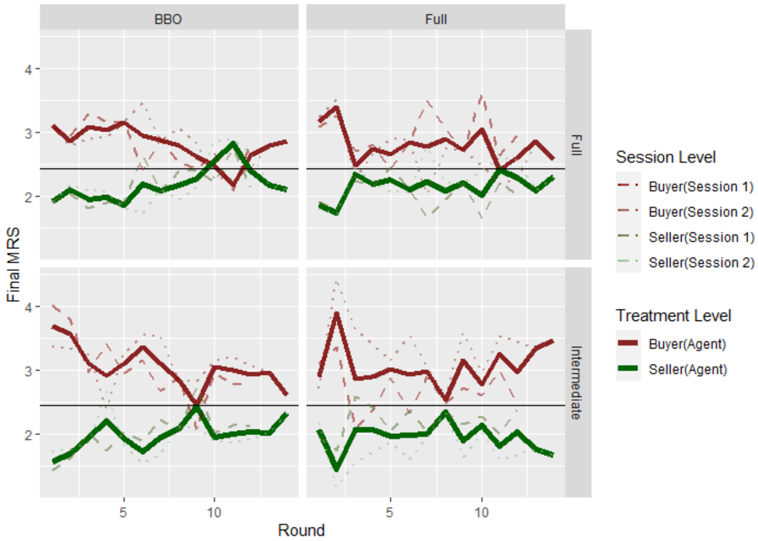

Figure 2.7: Marginal rate of substitution at the end of each round, averaged across session within treatment. Estimates are for aggregated buyer and seller traders who, after each transaction occurs, adjust their allocations by the average individual adjustment made by the natural buyers and sellers, respectively. Dashed lines show the session level round-end estimates.

ity 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 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.*

2.4.5 Treatment Effects

Table 2.6 provides regression analysis for the main outcomes discussed in sections 2.4.1-2.4.4. Main effects in the price regressions corroborate the story told in Section 2.4.1, 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 2.2, with the large positive increase of 0.160 matching the jump from Full-TT to Full-Full.

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

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.

2.5 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.

Chapter 3

Expanding Minimal Intelligence in the Double Auction

3.1 Introduction

For markets to converge, must traders be perfectly rational agents choosing orders subject to some expected utility paradigm? Or, can traders who intentionally deviate for the sake of future gains from trade still guide a market to equilibrium predictions?

Classical agent-based models have provided several mechanisms for how traders achieve equilibrating play and market convergence, though with nearly all assuming some form of utility-improvement restriction. Traders are presumed to perfectly kite along the utility-improving side of their indifference curves (or, analogously, remain on the surplus improving side of their cost or redemption value schedules). Laboratory experiments, however, have repeatedly shown evidence, especially in more complex settings, that traders in a continuous double auction routinely break this assumption. There are few potential reasons for this. The mechanism privy to the model in this paper is that traders occasionally intentionally taking utility losing trades in order to set themselves up for future trades on both sides of the market.

Such a mechanism, when modelled, can find itself in the oft-mentioned 'wilderness of bounded rationality.' A vast expanse of deviations from perfect rationality have been explored, with many more yet to be charted. Two gates to this wilderness, or maze, are typically recognized: that which assumes perfect rationality and that which imparts no (or very minimal) intelligence upon the economic agents.

Models and implications at both entrances are numerous, though admittedly with a larger mass at the rational end. A mapping between gates, however, is less so explored. The main goal of this paper is to present a model that provides one such mapping. I propose a new agent-based model of CDA trader behavior in an Edgeworth box economy. The flavor of two influential assumptions on trader activity are incorporated: (1) beliefs on the acceptability of prices (Gjerstad and Dickhaut, 1998), and (2) reservation prices which adjust within-period (Friedman, 1991). Traders place orders by applying logit choice probabilities to each admissible[1] order; the logit choice parameter allows a trader's choice to capture precision anywhere between the random choice of zero-intelligence traders and the traders of Gjerstad and Dickhaut.

Section 2 presents a summary of past CDA trader behavior models. A wave of influential models at the end of the 20th century provide the design motivation for much of this paper's model (as well as a vast experimental and applied literature over

[1]Nomenclature taken from Friedman (1991). Here admissible will mean satisfying a utility analog of reservation price.

the last twenty years). Section 3 proposes the setting of this paper. The double auction institution and general equilibrium setting are invoked to provide a test bed for the theorized agents.

A tractable model of trader behavior is presented in Section 4. Agents make use of reservation prices and logit choice probability to select and/or accept orders as two-way traders conscious of their positioning for both sides of the market. Section 5 proposes a simple design for a test of the model via simulations. Measures of efficiency are impressively high, which, when paired with convergence in allocations, hints at relatively equitable reallocations near a point on the contract curve. Section 6 concludes the paper.

3.2 Prior Theory

A few literatures are relevant for providing a background for this project, as well as framing its contribution overall. Theoretical work surrounding the continuous double auction price dynamics has arrived in two distinct waves over last fifty years. First, a batch of partial equilibrium models were proposed in the late 80's and throughout the 90's. Then, a newer wave of more generally applicable model were given in the '00's and '10's. Here I'll present a synopsis of both waves, as well as a selection of related research in the intermediate period.

Wilson (1987) began the first wave with a game theoretic model, positing a strategic multilateral setting where each trader's actions directly impact the pricing strategies of the other traders. Though highly plausible in smaller markets, strategic

Easley and Ledyard (1993) took a less complex route to defining double auction play, entirely removing the strategic interaction. Traders participated in the market under the assumption their own decisions have no impact on the order placement/acceptance of others, all the while guiding their own orders via an across-period deterministic reservation price. Friedman (1991) similarly took on a game against nature stance, however with traders administering a more sophisticated within-period reservation price bidding/selling strategy. Gode and Sunder (1993) simplified trader behavior even further (hence its running name of "zero intelligence") by having order price be randomly chosen, supposedly leaving the only driving factor of price formation being the underlying rules of the double auction itself. Much closer to the perfectly rational gate, Gjerstad and Dickhaut (1998) models traders who develop beliefs on the acceptability of prices, and then select the price which yields the maximum expected surplus.

The literature from this point split over the last couple decades. A batch of parsimonious, tractable, often heuristic-driven models entered the learning model literature. Though not directly designed for markets, the following models can naturally be bent to account for the more complex setting. Roth and Erev (1995) provide a reinforcement learning model designed for dynamic games, with an emphasis on testing performance and convergence in the intermediate term. Agents develop choice propensities for each strategy, with successful outcomes increasing a strategy's propensity to be chosen in future decisions.[2] Fudenberg and Levine (1995) postulate a theory of 'cautious fictitious play', which places beliefs over the probability of opponent's playing given strategies. Agents use these beliefs to make their own strategy, each strategy

[2]The model was tested across three game types, with one being a simple market.

being chosen with some logit choice probability.[3] Camerer and Hua Ho (1999) house reinforcement-based learning and belief-based learning as special cases of a more complex experience-weighted attraction learning model; some flexible convex combination of the two is shown to generally be a better fit to game data than either of the two as stand-alone models.

A strain of models imposing higher levels of complexity in behavior or more complex market settings, or both, have also been proposed recently. One such model that is highly malleable in terms of its application and setting is the individual evolutionary learning model (IEL) of Arifovic and Ledyard (2011). Economic agents maintain an evolving pool of potential choices which they draw from subject to a probability distribution that is constantly updating via experimentation and replication stages. A few years later, Anufriev et al. (2013) applied IEL to the continuous double auction setting, in a partial equilibrium environment.[4] General equilibrium adaptions of the ZI model were promoted by Gode et al. (2004) and Crockett et al. (2008a) a decade or so after the original model was published. The former features an price-angle order choice process, while the latter proposes a learning process by which the allowable subset of the contract curve is restricted round after round. Williams (2021a,b) bring two models from the first wave to a general equilibrium setting. A competing model to that of Gode, Spear and Sunder (2004), Williams (2021b) postulates a GE-based zero intelligence (from here, ZI-G) model with a lower sense of 'zero' intelligence. Williams (2021a)

[3]Feltovich (2000) tested these two models in the laboratory where subjects played a two-stage game with asymmetric information. The reinforcement model better predicted choice probability of a subject's next action, while the belief-based model proved better more often for aggregate trends in play.

[4]The timing in the paper lends itself more to an analysis of multiple iterations of a call market as opposed to a double auction in continuous time. A timing adjustment suggested in van de Leur and Anufriev (2018) better aligns the model with continuous time.

brings the belief-based process of Gjerstad and Dickhaut (1998) to general equilibrium to better understand impacts price information may have on market convergence.

3.3 Environment

3.3.1 Message Space

Here, I lay out the space encompassing all exchange related information the traders are given: the message space. This space is a crossing of several one dimensional sets (yielding information in the form of order n-tuples) to be described below.

Prices: $P \subset R_+$ s.t. elements $p \in P$ denote per unit prices

Quantities: $Q \subset R_{/\{0\}}$ s.t. elements $q \in Q$ denote desired unit adjustment

Time: $\mathcal{T} \equiv [0, T] \subset R_+$ s.t. elements $t \in \mathcal{T}$ denote time a market period

The above create the 3-dimensional space which defines the standard elements of orders: the per unit price, the desired adjustment in units of the commodity, and the time in the markets life at which the order was placed. Here I will assume orders are infinitely-lived, only expiring if the market's duration ends or if a trader replaces their order with a new one.

This typical 3-tuple order will be augmented to store information about the trader's involved. The set of traders, $\mathcal{N}$, is split into two subsets: natural buyers, $\mathcal{B} = \{1, ..., N_B\}$, and natural sellers, $\mathcal{S} = \{N_B + 1, ..., N_B + N_S\}$. The set N satisfies

$N = B \cup S$ and $B \cap S = \{0\}$.[5]. Natural buyers are traders whose marginal rate of substitution at the inception of a market is greater than the competitive equilibrium price for the initial allocation of goods across the market. Natural sellers are defined similarly, with their MRS residing below the initial CE price.

With trader identities defined, we can define the standard notation for orders in this framework:

$$o_{\Delta,a} = (p_a, q_a, t_a, b_a, s_a)$$

where a denotes the action number (the a^{th} action taken in this market), with $a \in \mathcal{A} \subset N$. Actions are defined as an order placement or order acceptance. Also, Δ denotes the side of the market the order is being placed, with

$$\Delta = \begin{cases} b, & q > 0 \\ s, & q < 0 \end{cases}$$

For any ordering being placed in the orderbook, b_a or s_a must be 0.

The set of bids and asks can thus be defined as follows:

$$\textit{Buys}: \quad \Omega_B := \mathcal{P} \times \mathcal{Q} \times \mathcal{T} \times \mathcal{N} \times \{0\} \equiv \{o_{b,a} : p_a \in \mathcal{P}, q_a \in \mathcal{Q}, t_a \in \mathcal{T}, b_a \in \mathcal{N}, s_a \in \{0\}\}$$

$$\textit{Sells}: \quad \Omega_S := \mathcal{P} \times \mathcal{Q} \times \mathcal{T} \times \{0\} \times \mathcal{N} \equiv \{o_{s,a} : p_a \in \mathcal{P}, q_a \in \mathcal{Q}, t_a \in \mathcal{T}, b_a \in \{0\}, s_a \in \mathcal{N}\}$$

Note that b_a and s_a can lie in all of $\mathcal{N}$, as opposed to $\mathcal{B}$ or $\mathcal{S}$, as all traders have the capacity to place orders and trade on either side of the market (we refer to them as two-way traders). The set of orders is defined as $\Omega \equiv \Omega_B \cup \Omega_S$.

[5] $\|\mathcal{N}\| = \|\mathcal{B}\| + \|\mathcal{S}\| = N_B + N_S = \bar{N}$

3.3.2 Exchange Definitions

Every message which enters the market, or exchange, lives in the message space laid out above. Below, I define the characterizing elements/processes that make up the exchange.

Orders: An order is a single message sent to the exchange by a trader, of the form $o_{\Delta,a}$. Orders may be submitted at any time $t \in \mathcal{T}$.

Asks/Bids: Asks satisfy $q_a < 0$, $b_a = 0$, $s_a \neq 0$ s.t. $o_{s,a} \in \Omega_S$. Bids similarly satisfy $q_a > 0$, $b_a \neq 0$, $s_a = 0$ s.t. $o_{b,a} \in \Omega_B$.

Orderbook (Ω_O): Any order $o_{\Delta,a}$ which satisfies $b_a = 0$ or $s_a = 0$ exists in the exchange's orderbook.

To accommodate the idea of trading units or filling orders, an amendment to the notation of orders is needed, as well as a definition of how orders are filled. An order, $o^{\kappa}_{\Delta,a}$, has $\kappa = 0$ if the order is newly posted to the orderbook, $\kappa \in (0, 1]$ if the order was (partially) filled. The value of κ in the latter case is equal to the proportion of units that were filled out of q_a. If $\kappa = 1$, order $o^{\kappa}_{\Delta,a}$ has been fully filled, and thus is removed from the orderbook.

Cross/Accept: An order $o_{\Delta,a}$ crosses order $o_{\Delta',a'}$ (for $a' < a$) if:

1. $\Delta \neq \Delta'$ and

2. $p_a \leq p_{a'}$ if $\Delta = s$ or $p_a \geq p_{a'}$ if $\Delta = b$

Also, define $a_T : N \rightarrow \mathcal{A}$ as the mapping from the ordering of crossings/traders to the action at which the crossing/trade occurred.

3.3.3 Histories

While the orderbook provides a snapshot of the present state of the exchange, a system for (1) referencing older orders no longer in the book, and (2) providing context for the expanse of trader's memories within the market must be defined to track the adjustment of the market.

History: The set of all orders (past and present) and trades in the lifetime of the market. The full history can be split into three types of orders:

- Trades (Ω_T) $\longrightarrow$ The set of orders $o^{\kappa}_{\Delta,a}$ which satisfy $b_a, s_a \in \mathcal{N}$ and $\kappa = 1$.
- Cancelled Orders (Ω_C) $\longrightarrow$ The set of orders $o^{\kappa}_{\Delta,a}$ which satisfy $b_a = 0$ or $s_a = 0$, and $\kappa = 1$
- Orderbook (Ω_O; as defined above)

A useful union of these subsets for the purpose of understanding trader behavior is $\Omega_T \cup \Omega_C$, which houses all past, or closed, orders in the history of the market. Also, note that this union is the definition of *history* (Ω_H) in Gjerstad and Dickhaut (1998).

The history, as just defined, shows the entirety of the exchange's life, from the perspective of the exchange (in the sense that the history is full/complete). Each trader, however, may or may not have the capacity (or desire) to maintain as complete of a history. In the vein of Gjerstad and Dickhaut (1998), traders can remember all past orders within the last L successful trades[6]:

Memory: The set of traded orders and cancelled orders that have occurred in the last L orders (where the most recent trade was the j^{th} trade in the market) , $\Omega_{M(L)} \equiv \{o^{\kappa}_{\Delta,a} \in \Omega_H : a < a_T(j-L)\}$.

3.3.4 Trader Preferences

Much like the ZI-G and GD-G general equilibrium models, traders are motivated via utility functions. This is opposed to the cost and redemption-value schedules driving traders in more classical partial equilibrium settings. Generally, the standard assumptions on the utility function of trader i, u_i, are assumed: u_i is twice differentiable, decreasing and quasi-concave. For the remainder of this paper, I'll focus on the constant elasticity of substitution functional form:

$$u_i(x,y) = c_i((a_i x)^r + (b_i y)^r)^{\frac{1}{r}} \tag{3.1}$$

For simplicity, I normalize relative preference parameters a and b such that they sum to one and are both non-negative. The curvature parameter r is also assumed to lie in $(-\infty, 1]$ to satisfy the quasi-concavity requirement.

6

I add the flavor of reservation prices to trader's, though through an avenue more appropriate for general equilibrium. Traders maintain reservations around the utility gained at each price (or quantity change in x). As agents are two-way traders, they develop these reservations as both buyers and sellers. These reservation utilities are captured via parameter η, where $\eta(t)$ is a function of within-period time t.

The reservation adjustment η takes the form:

$$\eta(t) = \left(\frac{T-t}{T}\right)\frac{min\{a,b\}}{max\{a,b\}}min\{a,b\} \tag{3.2}$$

and enters into the trader's utility function as follows

$$u_{i,b}(x,y|\eta) = c_i((a_i+\eta)^r x^r + (b_i-\eta)^r y^r)^{\frac{1}{r}} \tag{3.3}$$

$$u_{i,s}(x,y|\eta) = c_i((a_i-\eta)^r x^r + (b_i+\eta)^r y^r)^{\frac{1}{r}} \tag{3.4}$$

Here $u_{i,b}$ is the buyer reservation utility for trader i, and $u_{i,s}$ is the seller reservation utility. Some natural (and necessary) comparative statics arise from for $\eta(t)$. First, η is decreasing in t. This implies that both $u_{i,s}$ and $u_{i,b}$ collapse to the trader's true indifference curve u_i as a period progresses. Second, the adjustment is decreasing in relative "side", or good, preference, as $\min\{a,b\}/\max\{a,b\}$ is decreasing in $|a-b|$.[7] Thus, traders greatly preferring one good to the other will have weaker reservation adjustments.[8] Finally, the last part of equation 3.2 implies that preferences for both goods remain positive (as η can't surpass a or b in magnitude).

[7]Holding $|a+b|$ constant.

[8]This means $u_{i,\Delta}$ will be more accepting on the Δ side, and $u_{i,-\Delta}$ will be less influential. The trader is less inclined to submit increasingly utility-reducing (relative to u_i) offers in this case.

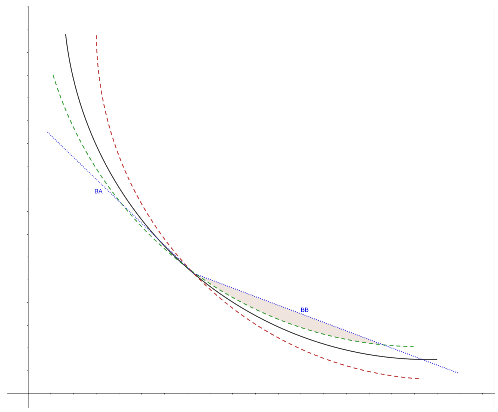

Figure 3.1: Reservation Utility. The black curve shows an indifference curve (IC) of u_i. The green and red dotted curves show IC's for $u_{i,b}$ and $u_{i,s}$, respectively. BA and BB (blue dotted lines) show best ask and bid prices in the market. The shaded region shows the space in which this trader would automatically accept a posted ask, were he to enter as a buyer.

A few special cases should be mentioned as well. First, with respect to functional form, perfect substitutes ($r = 1$), perfect complements ($r \to -\infty$) and Cobb-Douglas ($r \to 0_{-,+}$) are naturally folded into CES preferences. Both perfect substitutes and perfect complements provide interesting responses/interpretations when including η as in equations 3.3 and 3.4. The former, graphically, mimics reference prices from the partial equilibrium literature of the 90's, as the slope of the IC gives a natural reservation price. The latter is analogous to an adjustment in the desired complement ratio.

3.4 Agent-Based Model

This section lays out the details of the model, now that the environment has been established. Much like the GD-G model from Williams (2022), four main processes determine the flow of the market and trader behavior in this model. These are entry, belief updating, market interaction, and re-entry determination.

Entry refers to the actions taken and snapshot of the market received by the trader who enters the market in time t. In all times aside from the inception of the market, entry is actually the second step of a two-part market entry/exit flow process along with the re-entry determination phase. The belief-updating phase takes the snapshot of the market in the entry phase and allows the entrant to readjust his interpretation on which prices my potentially be successful moving forward. Market interaction defines the order selection and submission process, as well as potential clearing. The re-entry determination phase sees all traders briefly evaluate their holdings, beliefs and the state of the market to evaluate their desire for re-entry. Below, each of these will be fleshed out in much greater detail.

3.4.1 Entry

Entry (and re-entry) into this environment's markets can take a couple of different forms depending on the age of the market and the potential entrant's previous participation in the market. The inception of the market (i.e. the first entry in the first iteration, or period, of the market) is unique in that no prior history exists. As such, this is the only instance in which entry is entirely random. Similarly, the first entrant of any period after the first is uniformly drawn.

The second (and far more common) entry situation is any entry after the first in any market period. To foreshadow the re-entry process discussed in section 3.5.1, the trader who wins the re-entry draw (with re-entry probabilities being dependent on average utility gain above a given trader's reservation utility) enters the market next.

In this case, the trader drawn to (re)enter checks the market's best bid and ask against their own current reservation utilities and begins the belief updating process before making a decision on how they wish to use their entry.

3.4.2 Belief Updating

First, lets recall the belief formation and updating process of Gjerstad and Dickhaut (1998). Here, traders establish beliefs over the acceptability of certain prices on either side of the market. Traders recall a portion of the history, Ω_H, and tally the success and failure rate of each price, ρ, seen for each side of the market, $TA(\rho)$ for asks and $TB(\rho)$ for bids.

In Gjerstad and Dickhaut's original setting, these tallies were defined as counts with a count of 1 given to each order that satisfied the criteria (traded or cancelled) of interest. This was appropriate as each order in their partial equilibrium setting was required to be for a single indivisible unit. However, uniform counts are not attuned to settings with multiple and/or divisible units. Williams (2021) provides a general-equilibrium-adjusted version of Gjerstad and Dickhaut's model (from here referred to as GD-G), in which each order is given weight equal to the proportion of the original quantity successfully traded, $\sqrt{q_k}\frac{q_{k,traded}}{q_k}$. A similar weighted count is defined for the rejected (cancelled) portions of orders, $RA(\rho)$ and $RB(\rho)$.

Traders aggregate over the success of orders at less desirable (to the rest of the market) prices than one they may be considering. This is assessed relative to the success of these worse prices along with the failure of prices placed on the desired side at more desirable prices. For some bid ρ,

$$p_b(\rho) = \frac{\sum_{\rho \leq b} TB(\rho) + \sum_{\rho \leq b} TA(\rho)}{\sum_{\rho \leq b} TB(\rho) + \sum_{\rho \leq b} TA(\rho) + \sum_{\rho \geq b} RB(\rho)} \tag{3.5}$$

represents the probability of acceptance. Each trader holds such a belief for each price represented in Ω_H. Note that beliefs are over the domain [0,M], with $p_b(0) = 0$ and $p_b(M) = 1$ for bids and the reverse for asks.

3.4.3 Market Interaction

Contrary to ZI-G and GD-G, this model considers the use of two types of order placement strategies, accepting orders directly and placing orders in the book. While both of the prior models can achieve both strategies via only the latter (as orders placed in the book which cross, essentially accept another order directly), a couple of distinctions should be made. First, in a setting where orders can have multiple and/or partial unit quantities, crossing orders won't always interact as cleanly in the orderbook as an accept. Second, it seems natural to consider the two actions as responding to separate lines of intent for the trader, with accepts being very short-term, heuristic driven choices and orderbook additions being more long-term plays. Establishing such distinctions between the two also provides a nice analog to the ideas of market orders and limit orders in the financial literature.

The market interaction, in concert with the above, is a two part process: checking for and interacting with orders that may be desirable immediately, and submitting an order to the exchange to add to the existing book. Note that the second step is only reached if the trader does not satisfy the "interacting" portion of the first step. Below are the processes of accepting and placing orders explained in detail.

Accepting Orders

Upon entry, even before the belief updating process has occurred, the trader has an idea of their reservation utility on their selected side of entry. Consistent with previous reservation price models, traders have an incentive and desire to accept with certainty an order on the contra-side of the market which is greater than their reservation. This means the trader would be checking first if the current book leaves any room between the best order on the entered side and his current reservation utility, or:

$$|BP_{\Delta} - MRS_{u_i}| - |MRS_{u_{i,\Delta}} - MRS_{u_i}| > 0 \tag{3.6}$$

A check with evidence of a contra-side order in this region induces the entrant to accept the order outright. If multiple orders exist in this region, the order with the highest resultant utility is chosen. A null result from the check leads the trader to stage two of their market interaction, described below.

Placing an Order

While on side Δ, the trader has three indifference curves to consider: the curve for $u_{i,\Delta}$, the curve for u_i and the curve for $u_{i,-\Delta}$. Functionally, only two of these will be considered. The more restrictive reservation utility, $u_{i,\Delta}$, has already been shown to be used as a bound for immediately-acceptable orders. The weaker reservation utility, $u_{i,-\Delta}$, provides a lower bound for the bundles necessary to be at least as happy in future entries (especially if entering on side $-\Delta$ in their next interaction). The curve associated with u_i is left to serve as a target for activity very late in the market's life, with $u_{i,\Delta}$ and $u_{i,-\Delta}$ providing "goal posts" moving over time as a reflection of a trader's continuation value.

Using this lower goal post as a criterion for utility-improving orders, the trader considers any bundle on the Δ side that is weakly better than their current reservation utility $u_{i,-\Delta}$. This set of orders lies in

$$\mathcal{P}_{u_{i,-\Delta}} \times \mathcal{Q}_{u_{i,-\Delta}} \equiv (min\{|MRS_{u_{i,-\Delta}}|, Boundary_{\Delta}\},\ max\{|MRS_{u_{i,-\Delta}}|, Boundary_{\Delta}\}] \times (0, \bar{x}(\Delta)] \tag{3.7}$$

$\mathcal{P}_{u_{i,-\Delta}}$ takes an open lower bound at the marginal rate of substitution at the trader's current endowment on the contra-side reservation utility $u_{i,-\Delta}$ and a closed upper bound at the boundary price on that side (0 if $\Delta = b$ and M if $\Delta = s$). $\mathcal{Q}_{u_{i,-\Delta}}$ is more tedious to define, as the upper bound must take both current allocation and the non-zero[9] intersection point between the indifference curve and line associated with the best price on that side, BP_{Δ}, into account. The upper bound on $\mathcal{Q}_{u_{i,-\Delta}}$ is dependent on both the intersection between $u_{i,-\Delta}$ and the price vector extending from the trader's current allocation (call this $\hat{x}$) and the trader's current holdings of x. When $\Delta = b$, $\bar{x}$ is generally equal to $\hat{x}$; however, if $\hat{x}$ is non-existent or sufficiently large, then $\bar{x}$ is bounded above by the total x remaining in the market. For $\Delta = s$, $\bar{x}$ is the minimum of the total y remaining in the market divided by the price of the order and $\hat{x}$.

For each potential bundle, $o_z \in O_z := \mathcal{P}_{u_{i,-\Delta}} \times \mathcal{Q}_{u_{i,-\Delta}}$, the trader considers their belief on the acceptability of the given price. Each bundle thus has an expected level of utility improvement. The trader considers the possible bundles with logit choice

[9] The "zero" intersection here would be at the trader's current allocation.

probability:

$$Pr(o_z|x_k, y_k, \Omega_M) = \frac{exp[\lambda p_\Delta(o_z)(u_{i,-\Delta}(o_z) - u_{i,-\Delta}(x_k, y_k))]}{\sum_{o_z' \epsilon O_z} exp[\lambda p_\Delta(o_z')(u_{i,-\Delta}(o_z') - u_{i,-\Delta}(x_k, y_k))]} \tag{3.8}$$

The parameter λ implies some preciseness over the trader's ability to choose the expected utility-gain maximizing order. Reservation adjustment aside, $\lambda = 0$ would yield a uniform distribution over the orders, much like ZI-G. Similarly, $\lambda \to \infty$ would imply perfect choice as in GD-G.

3.4.4 Re-entry Determination

Now that the current entrant has entered, updated and (attempted to) place their order, and the exchange has updated the book and/or processed a transaction, the rest of the market (and the entrant herself) can individually reflect and gather their potential gains on either side of the market were they to enter next. To do so, a trader checks equation (3.6) for each side of the market , though with $u_{i,\Delta}$ replacing $u_{i,-\Delta}$. If (3.6) fails, the trader gives a value of 0 for that side of the market. Otherwise, the trader performs the same process as was taken by the entrant when placing an order. A couple minor adjustments to the process are needed however.

For side Δ, each trader considers $u_{i,\Delta}$ when determining the set of admissible orders (those which are immediately acceptable) $\mathcal{P}_{u_{i,\Delta}} \times \mathcal{Q}_{u_{i,\Delta}}$. Each admissible order is given an expected utility gain using the trader's developed beliefs for price acceptability. The trader averages over the expected gains of all admissible orders, giving them an idea of the expected gain for entering on that side. Each trader-side is treated as a separate draw for the next entry into the market, with each draw's probability being the draw's expected gain divided by the sum of all trader-side expected gains.

3.5 Simulations

3.5.1 Implementation

The performance of the model presented here is demonstrated via a set of simulated markets. A group of eight computerized traders are placed in a simulated CDA, playing in multiple periods of a single market. This multi-period-life market is simulated many times, completely refreshed at the inception of each simulation.

The main assumptions of the model, institution and equilibrium are applied to the traders; a series of 40 markets are simulated under these conditions (and with the parameters described below). Each market lives twelve periods of identical length. A market period is comprised of 200 market entries, with the entrant being determined via the draw described in Section .[10]

Each computerized trader has CES preferences over two goods, with parameter sets:

	c	a	b	r	(x_{Endow}, y_{Endow})
Buyers	0.113	0.825	0.175	0.5	(3,23)
Sellers	0.099	0.6875	0.3125	0.5	(11,3)

Table 3.1: Simulated Agent Parameters.

All traders maintain a memory of $L = 5$, implying they can perfectly recall

[10] As in the model, no market level spread reduction rule is enforced. However, traders have 'internal' spread rules, only replacing their own order if its better than one currently in the market. As this still allows for order placement at prices worse than the best bid and ask, I don't feel such a restriction is overly influential in market success. These internal rules are the only impediment on orders not being placed in the book. In the batch of simulations discussed here, 79.4 out of 200 orders were placed per period.

all transactions and order cancellations in the market within the last five transactions. This memory may span across market periods within the same run, however may not carry over between runs.[11] Additionally, when conducting his logit choice procedure over the set of feasible orders, each trader will have a logit choice parameter, γ, of 0.75. This places traders' choice precision between uniformly random and perfect, leaning more on the side of random.

3.5.2 Performance

	Mean	St. Dev.	Range
I. Prices			
Price	2.32	0.22	(1.68, 3.48)
$\lvert Price - CE \rvert$	0.56	0.14	(0.23, 1.23)
RMSE	0.69	0.19	(0.26, 2.15)
Final 5 Prices	2.35	0.23	(1.68, 3.48)
II. Allocations			
Final Distance	0.76	0.48	(0.01, 5.78)
Seller MRS	2.41	0.18	(1.09, 3.16)
BuyerMRS	2.53	0.21	(2.04, 4.77)
III. Efficiencies			
Allocative	0.96	0.04	(0.39, 1.00)
Distance	0.79	0.06	(0.26, 0.92)
Observations	480	480	480

Table 3.2: Simulation Outcomes. Observations at the round-average level. Panel I shows price related estimates. RMSE is the root-mean-squared error. Panel II reports outcomes in allocation space. MRS here is the marginal rate of substitution at the final allocation of aggregated representative agents. Panel III lists estimates for two measures of efficiency.

Table 3.2 records the main performance measures for the simulated markets. A quick glance shows evidence of surprisingly successful markets. Estimates show promis-

[11]This assumption is tested in a batch of simulations with trader memories that refresh at the beginning of every period.

ing levels of convergence in both allocation and price space, with markets tracking remarkably well around the equilibrium path.

All estimates are means of round-average (in the case of all price measures) or round-end (in the case of allocation and efficiency measures) level observations. Average price lies just 0.12 below the CE prediction from market inception, which, when accompanied with low volatility, implies transaction prices lying in an impressively tight band around the CE price. Figure 3.2 confirms not only the round-averages, but the individual transaction prices across the markets are closely bound. Convergence within period, however, requires tighter bounds on the time in focus. The final batch of transactions in a period provide an idea of traders' desire to trader and urgency to lack thereof to reap more gains from trade. I find an estimate even tighter to the CE prediction, suggesting prices not only lie close to the equilibrium, but tighten and converge in some smaller bound as the period ends.

Even so, prices can only provide so much of a picture of the full success of the market. Panel II gives two distinct pictures of how these simulated traders reallocate the two goods among themselves. The first is how far away the market is as a whole from the equilibrium set of allocations. To examine this, I collapse[12] the two types of traders into representative agents. These agents can aptly be fully represented in the Edgeworth box. On average, the final distance[13] the pair lies away from equilibrium

[12]For example, the four natural buyers can be aggregated into a single agent by averaging over each transaction made by one (or two) of the traders. If a buyer transacts with a natural seller, the adjustment in the representative buyer's allocation will be a quarter of that realized by the individual trader. If two natural buyers transact, the representative sees no adjustment in his allocation.

[13]The Euclidian distance that the representative buyer (and equivalently, seller) is away from the equilibrium allocation in the Edgeworth box. The y contribution to the distance is deweighted by the equlibrium price. The distance function is thus $dist(\cdot) = \sqrt{(x_i - x_{CE})^2 + (\frac{1}{p_{CE}}(y_i - y_{CE}))^2}$.

allocation bundle pair is within a unit radius of the final. Allocations approach the contract curve, on average lying in nearly Pareto optimal final resting places.

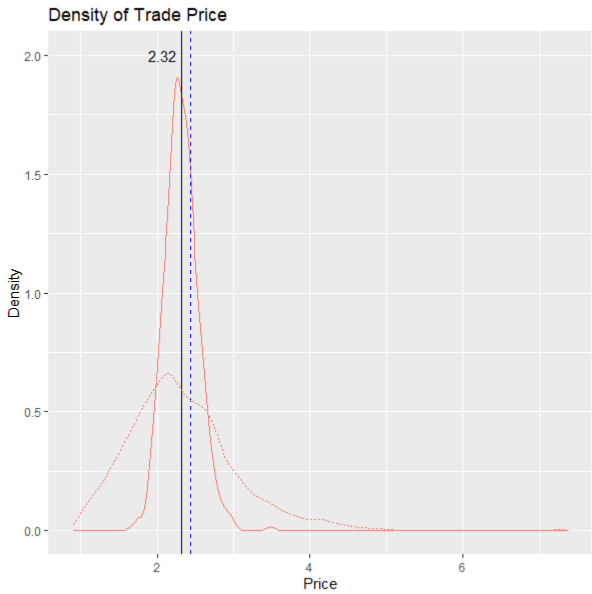

Figure 3.2: Kernel density for prices. Red line shows round-averages, while red-dotted shows individual transaction prices. The black vertical line is the mean of the round-averages, and the blue dotted line is the CE price of 2.44.

The marginal rate of substitution of market participants presents gives a proxy of convergence in allocative efficiency, as a trader's MRS should equal the CE price in equilibrium. Natural buyers are characterized by their initial MRS being above the equilibrium prediction; natural sellers lie on the other side of the price. As such, the traders, and their representative agents, should reallocate resources throughout the market period to collapse their MRS to the CE-price. The average final allocations of the representatives approach encouragingly close to 2.44, with sellers 0.03 below and buyers 0.08 above. Despite large ranges of round-end estimates for this spread, tight standard deviations suggest poorer MRS spreads are rather uncommon.

Two measures of efficiency are estimated. Allocative efficiency in this GE

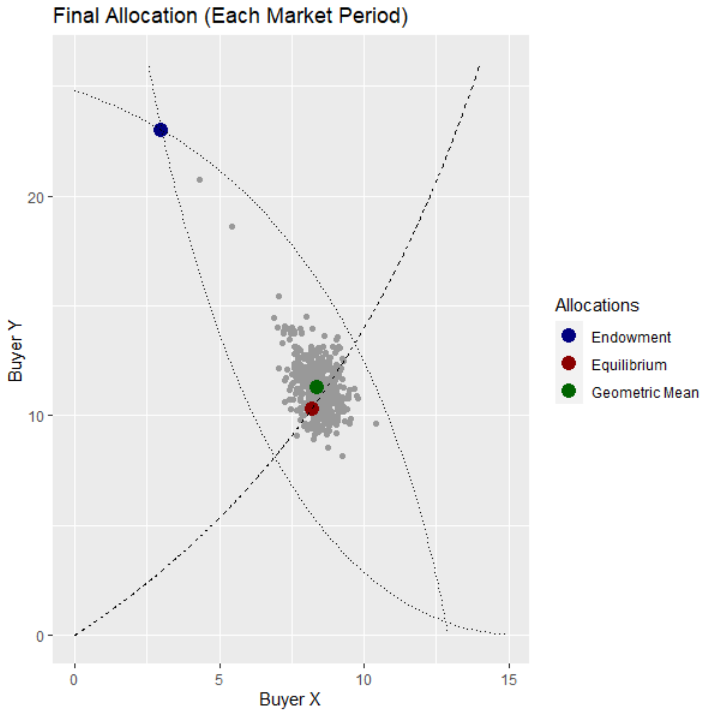

Figure 3.3: Final Allocations. Each grey dot represents the final allocation of the representative agents in the Edgeworth box. The red dot shows the equilibrium bundles, while the green dot represents the geometric mean of the scattered grey dots. The CE-price de-weighted distance between the red and green dots is 0.47 units. The dotted lines show the indifference curves of the representative agents evaluated at the endowment allocation. The dashed line shows the set of Pareto optimal allocations.

setting is adapted from the sum of profits over expected sum in equilibrium to the sum of utility gains relative to expected gains in equilibrium. The measure captures whether the market as a whole could reap the full gains from trader available, but is subject to ignoring disparities in gains between traders. This is less so true in GE, however, thanks to properties of the trader's preferences. Clearly, these traders are capable of capturing most of what the market has to offer, as average allocative efficiency is 0.96. For reference, simulated traders from the general-equilibrium zero intelligence models of Gode, Spear and Sunder (2004) and Williams (2021), as simulated in Williams (2021), yield allocative efficiencies of 0.92 and 0.66 in markets with even longer periods (300 entries).

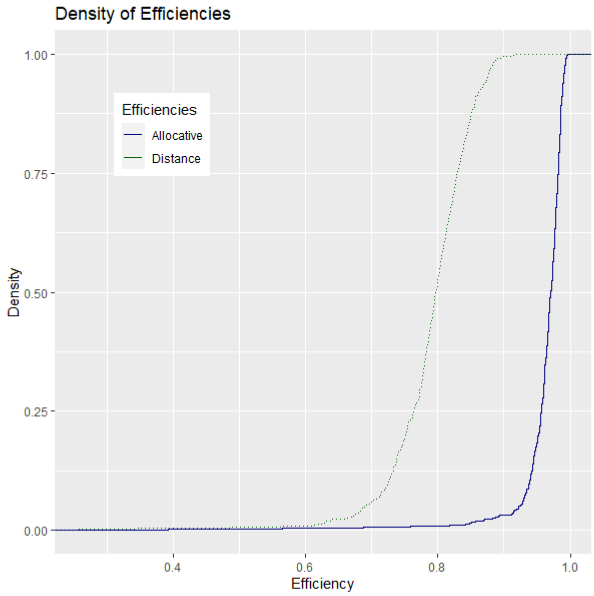

Figure 3.4: Cumulative Density Functions for Efficiency.

Distance efficiency is measured by the difference in the distance between the equilibrium and endowment allocations and distance between the final and equilibrium allocations, divided by the distance between the equilibrium and endowment allocations.[14] Geometrically, this means any point lying on the surface of the ball $\mathcal{B}_{Dist.Eff.}$ has the same value, thus penalizing deviation from the equilibrium path in much the same way as under or over trading along the path would be. The simulated markets perform quite well, with an average distance efficiency of 0.79, and max efficiency of 0.92. Estimates of both efficiency measures are tightly grouped, as shown in Figure 3.4, with allocative efficiency first order stochastically dominating distance efficiency. The Spearman rank correlation between the estimates is 0.80, which, given such high

[14]Distance here is the average of the 2-space GE-price-de-weighted Euclidean distance for each of the eight traders.

measures in both, suggests most markets en relatively close to a Pareto optimal allocation, but deviate to lie closer to other points on the contract curve than the equilibrium allocation. Figure 3.3 gives some weight to this claim.

3.6 Concluding Remarks

This paper models market dynamics in an Edgeworth box where trader's have 'imperfect' choice procedures when placing orders in a CDA. Traders have the capacity to remember a portion of the history of the market, developing beliefs over the acceptability of order prices. Beliefs account for the relative success of each past price based on order size. Agents recognize that they may participate on both sides of the market, and develop reservations depending on which side that enter. As traders maintain some utility preferences over their holdings, these reservations are held in terms of utility (as opposed to reservations on price as in Friedman (1991)). A curvature parameter η (which is a function of the time remaining in the market) determines what orders are immediately acceptable on the entered side and what orders satisfy the trader's reservation were they to enter on the contra-side in their next entry. Such a process allows traders to maker order selections that appear to be utility-reducing relative to their true preferences, though allow the trader to position themselves as to better perform as a two-way trader.

A set of simulations test the performance of markets with computerized traders imbued with the behavior described in the model. Prices near the equilibrium prediction consistently. Round-averages remain slightly below equilibrium, creating tight bounds but not quite converging. Allocations, both in 16-space and 2-space, regularly lie in impressively close to Pareto optimal allocations along the contract curve at round's end, often very close to the equilibrium allocation bundle. Given rounds are not run intentionally until allocations are Pareto optimal, achieving nearly this so consistently is promising. Seller and buyer marginal rates of substitution provide supporting evidence for convergence in allocations as well. Efficiencies, both allocative and distance, are repeatedly high, suggesting gains from trade are often equitably spread and mostly drawn from the market.

The major implication of the findings of this project is the feasibility of individually irrational (though deliberate) order placement decisions in markets that show convergent tendencies. Strategic repositioning in the orderbook, and in anticipated holdings, is a legitimate consideration traders may be making in double auctions. This paper confirms such a consideration is not as harmful as some perfectly rational purists may suspect; in fact, estimates here perform near or level with some more complex models. Furthermore, the model provides a mapping from the zero intelligence gate (beginning with ZI) through the wilderness to a model fit much closer to the perfectly rational gate (this being Gjerstad and Dickhaut's belief-driven model).

A few natural adjustments to this model exist. First, individualized η functions, dependent on arguments such as current holdings, within-round and market-life earnings, and overal time in the market (aggregated across periods), is an interesting adaptation. Estimation of functional form for variations on η via laboratory experimentation could be illuminating for the external validity of this model's mechanism. Given the results of Williams (2021a), an inclusion of prices in the orderbook in the belief updating process would likely improve fit. Finally, and more in the vein of a substitute instead of an adjustment, a heuristic-based version of this model that maintains the driving mechanism without the tenuous book-keeping by the traders could provide a more tractable and lab-friendly testbed.

Appendix A

Chapter 2 Appendix

A.1 Regression Table Continued

	Dependent variable:								
	Price	$\lvert Price - CE \rvert$	RMSE	Order Size	# Trades	Trade Size	Seller MRS	Buyer MRS	Alloc. Eff.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
⋮									
spreadRed:singleUnit:lattAng	-0.972	-1.059	-7.362	20.169^{***}	9.522^{***}	-0.163^{***}	-0.017	0.019	-0.018^{***}
	(1.464)	(1.464)	(13.837)	(3.476)	(0.216)	(0.026)	(0.015)	(0.025)	(0.007)
spreadRed:singleUnit:obReset	-0.599	-0.503	-2.349	-2.910	4.158^{***}	-0.258^{***}	-0.026^{*}	0.048^{*}	-0.028^{***}
	(1.464)	(1.464)	(13.837)	(3.476)	(0.216)	(0.026)	(0.015)	(0.025)	(0.007)
spreadRed:singleUnit:noLoss	0.352	0.381	0.376	-0.237	2.091^{***}	0.077^{***}	-0.350^{***}	0.757^{***}	-0.192^{***}
	(1.473)	(1.473)	(13.920)	(3.476)	(0.216)	(0.026)	(0.015)	(0.025)	(0.007)
spreadRed:lattAng:obReset	-3.887^{***}	-3.856^{***}	-34.719^{**}	9.380^{***}	12.323^{***}	-0.230^{***}	-0.022	0.003	-0.010
	(1.464)	(1.464)	(13.837)	(3.476)	(0.216)	(0.026)	(0.015)	(0.025)	(0.007)
spreadRed:lattAng:noLoss	-0.888	-0.983	-7.494	24.618^{***}	22.250^{***}	0.101^{***}	-0.374^{***}	0.789^{***}	-0.236^{***}
	(1.472)	(1.472)	(13.914)	(3.476)	(0.216)	(0.026)	(0.015)	(0.025)	(0.007)
spreadRed:obReset:noLoss	-0.861	-0.745	-2.747	-1.269	3.711^{***}	0.153^{***}	-0.130^{***}	0.206^{***}	-0.018^{***}
	(1.481)	(1.481)	(13.998)	(3.476)	(0.216)	(0.026)	(0.016)	(0.025)	(0.007)
singleUnit:lattAng:obReset	-6.404^{***}	-6.580^{***}	-61.530^{***}	4.222	44.816^{***}	0.394^{***}	-0.008	0.100^{***}	-0.065^{***}
	(1.464)	(1.464)	(13.837)	(3.476)	(0.216)	(0.026)	(0.015)	(0.025)	(0.007)
singleUnit:lattAng:noLoss	11.544^{***}	10.423^{***}	13.354	35.880^{***}	58.409^{***}	1.844^{***}	0.190^{***}	-0.090^{***}	-0.203^{***}
	(1.472)	(1.472)	(13.919)	(3.476)	(0.216)	(0.026)	(0.015)	(0.025)	(0.007)
singleUnit:obReset:noLoss	1.092	0.992	1.689	-0.178	-1.509^{***}	0.132^{***}	-0.035^{**}	0.031	0.028^{***}
	(1.482)	(1.482)	(14.009)	(3.476)	(0.216)	(0.026)	(0.016)	(0.025)	(0.007)
lattAng:obReset:noLoss	-7.090^{***}	-6.924^{***}	-61.623^{***}	4.236	59.364^{***}	0.238^{***}	-0.120^{***}	0.244^{***}	-0.092^{***}
	(1.481)	(1.481)	(13.998)	(3.476)	(0.216)	(0.026)	(0.016)	(0.025)	(0.007)
⋮									
Constant	2.268^{***}	1.447^{***}	2.054	14.942^{***}	20.132^{***}	3.155^{***}	2.035^{***}	3.066^{***}	0.656^{***}
	(0.518)	(0.517)	(4.892)	(1.229)	(0.076)	(0.009)	(0.005)	(0.009)	(0.002)
Observations	95,424	95,424	95,424	96,000	96,000	95,424	95,422	95,424	96,000
R^2	0.024	0.020	0.003	0.019	0.989	0.826	0.464	0.596	0.607
Adjusted R^2	0.023	0.020	0.002	0.019	0.989	0.826	0.464	0.596	0.607
Note:									*p<0.1; **p<0.05; ***p<0.01

Table A.1: Interaction regression results for third order interaction. This is a continuation of the regression estimates in Table 1.2.

	Dependent variable:								
	Price	$\lvert Price - CE \rvert$	RMSE	Order Size	# Trades	Trade Size	Seller MRS	Buyer MRS	Alloc. Eff.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
⋮									
spreadRed:singleUnit:lattAng:obReset	3.474*	3.474*	34.067*	−9.380*	−3.183***	0.228***	0.029	−0.019	0.028***
	(2.070)	(2.070)	(19.568)	(4.916)	(0.305)	(0.036)	(0.022)	(0.035)	(0.010)
spreadRed:singleUnit:lattAng:noLoss	0.197	0.399	6.771	−24.618***	−11.336***	−0.098***	0.363***	−0.714***	0.185***
	(2.076)	(2.076)	(19.628)	(4.916)	(0.305)	(0.037)	(0.022)	(0.035)	(0.010)
spreadRed:singleUnit:obReset:noLoss	0.454	0.348	2.205	1.269	−2.985***	−0.150***	0.139***	−0.228***	0.029***
	(2.083)	(2.083)	(19.694)	(4.916)	(0.305)	(0.037)	(0.022)	(0.035)	(0.010)
spreadRed:lattAng:obReset:noLoss	3.757*	3.869*	34.709*	−10.169**	−14.317***	−0.169***	0.057***	−0.073**	−0.014
	(2.082)	(2.082)	(19.683)	(4.916)	(0.305)	(0.037)	(0.022)	(0.035)	(0.010)
singleUnit:lattAng:obReset:noLoss	6.501***	6.474***	61.427***	−4.236	−40.185***	−0.244***	0.097***	−0.258***	0.121***
	(2.083)	(2.083)	(19.691)	(4.916)	(0.305)	(0.037)	(0.022)	(0.035)	(0.010)
spreadRed:singleUnit:lattAng:obReset:noLoss	−3.235	−3.403	−33.996	10.169	4.865***	0.170***	−0.072**	0.099**	−0.004
	(2.937)	(2.937)	(27.764)	(6.952)	(0.431)	(0.052)	(0.031)	(0.049)	(0.014)
Constant	2.268***	1.447***	2.054	14.942***	20.132***	3.155***	2.035***	3.066***	0.656***
	(0.518)	(0.517)	(4.892)	(1.229)	(0.076)	(0.009)	(0.005)	(0.009)	(0.002)
Observations	95,424	95,424	95,424	96,000	96,000	95,424	95,422	95,424	96,000
R^2	0.024	0.020	0.003	0.019	0.989	0.826	0.464	0.596	0.607
Adjusted R^2	0.023	0.020	0.002	0.019	0.989	0.826	0.464	0.596	0.607
Note:									*p<0.1; **p<0.05; ***p<0.01

Table A.2: Interaction regression results for fourth order interaction. This is a continuation of the regression estimates in Table 1.2.

A.2 Round-Average Price Densities

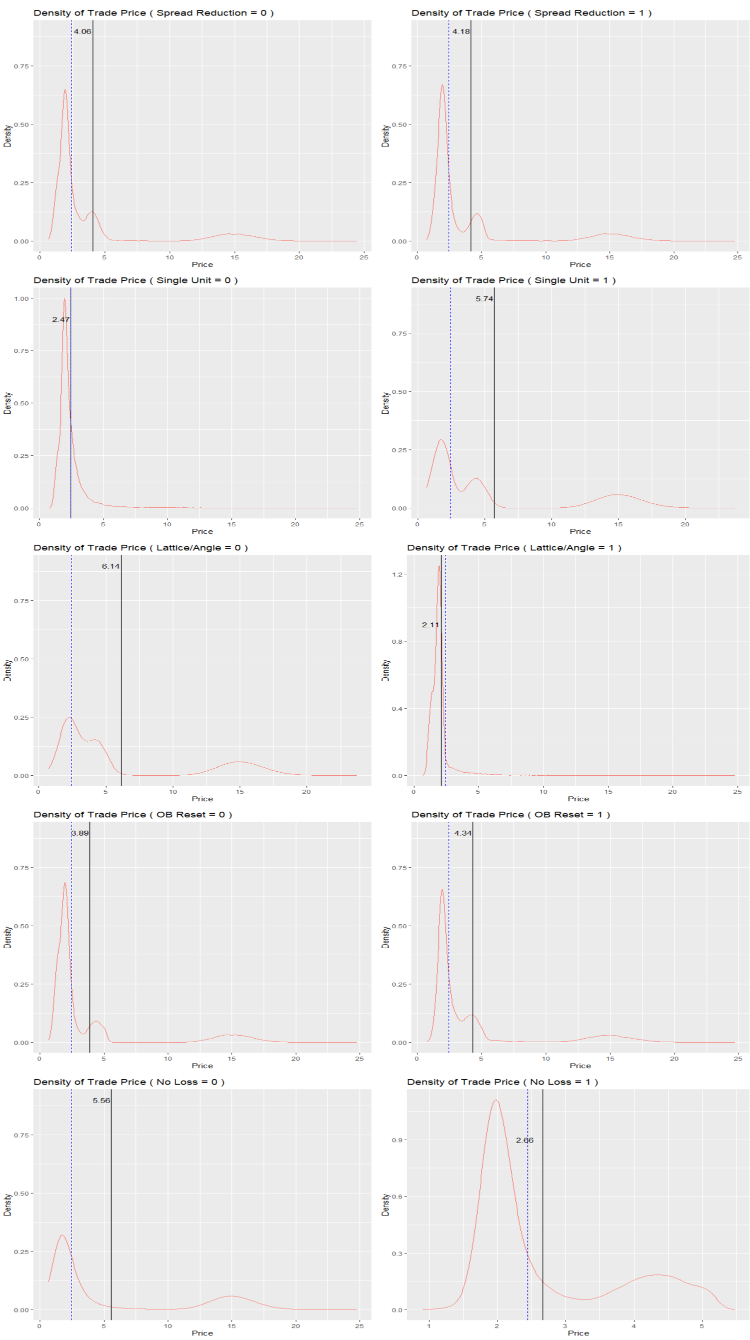

Figure A.1: Round-average price densities. Blue line is CE price and black line is subset average.

Appendix B

Chapter 3 Appendix

B.1 Distributional Tests

B.1.1 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**(+)
$\lvert Price - CE \rvert$	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 B.1: Wilcoxon test p-values for outcomes in Table 2.2. () denote the direction of change when moving from Treatment A → Treatment B in the column.

B.1.2 Utility Gain CDFs

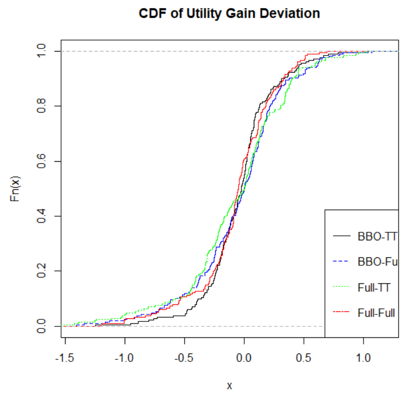

Figure B.1: Cumulative density functions for the round-end trader-level difference in utility gained and expected utility gain in equilibrium.

B.2 Across and Within Period Price Dynamics

B.2.1 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.2: Estimates for one period lagged price deviation (from time 0 competitive equilibrium price, p^*).

B.2.2 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.3: Within-period price adjustment.

B.3 Within-Period Allocation Adjustment

Figure B.3 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. 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 B.3 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.

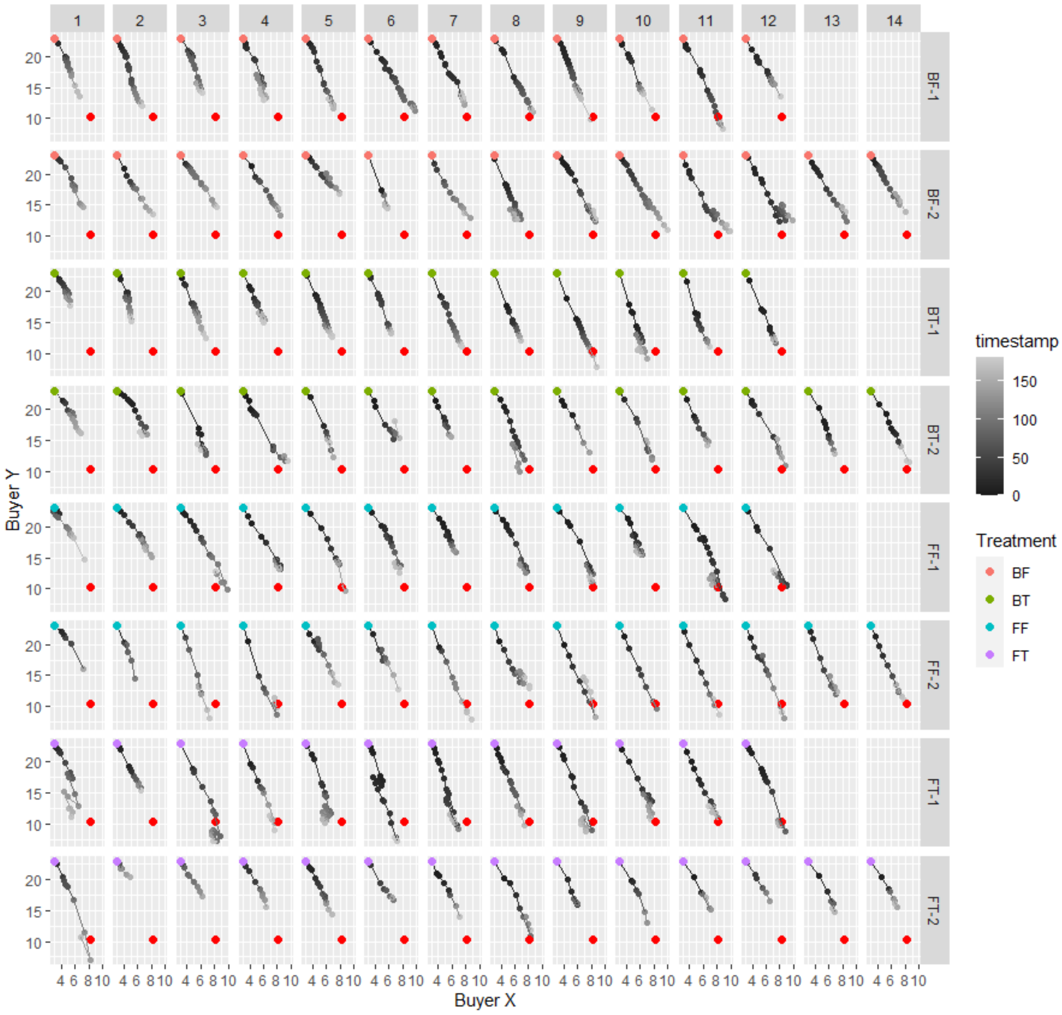

Figure B.2: Within-period allocation adjustment plotted in Edgeworth boxes. The shade of the allocation dot fades later in the period.

B.4 Utility-Losing Behavior

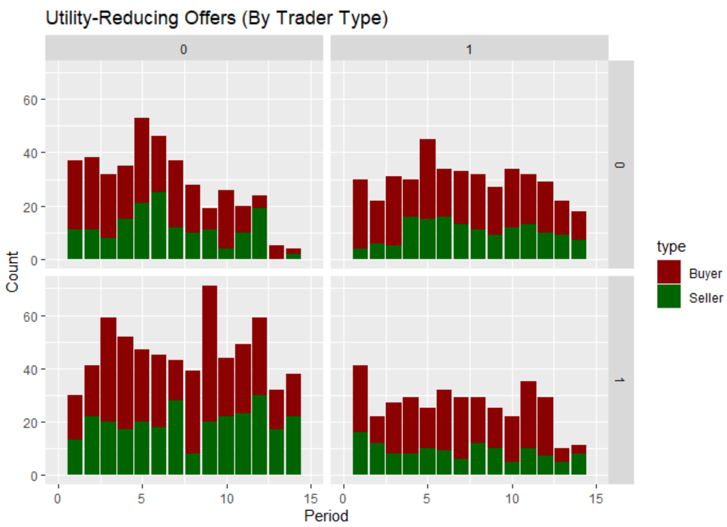

Figure B.3: Counts of utility-losing orders placed, partitioned by trader type.

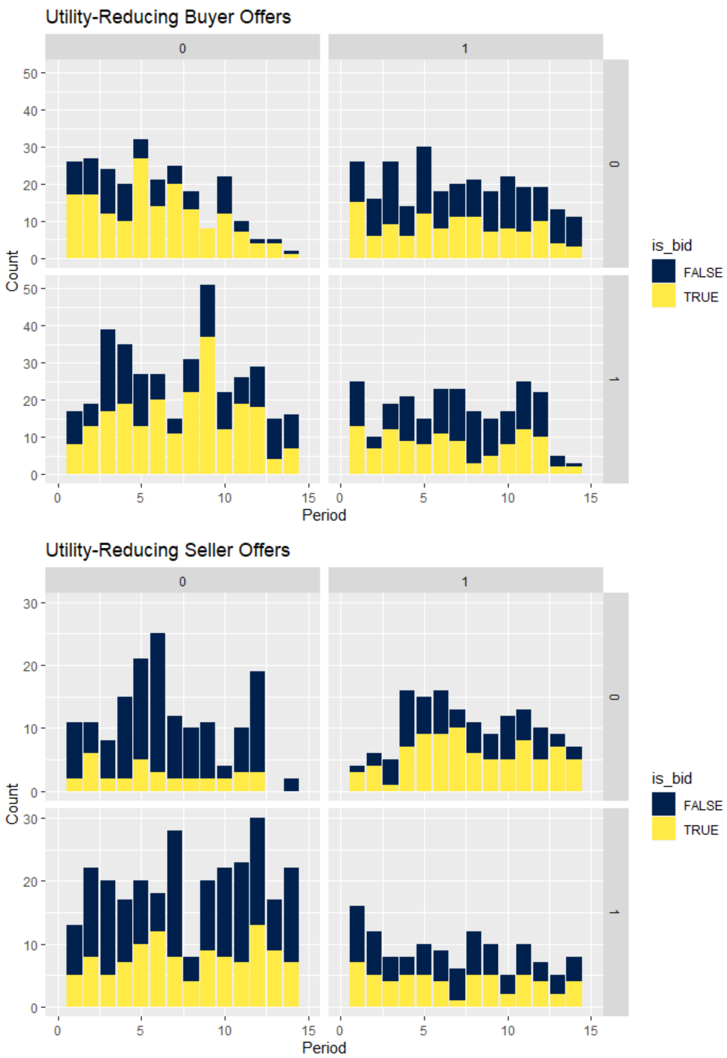

Figure B.4: Counts of utility-losing orders placed, partitioned by trader type and market side.

	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 B.4:

Appendix C

Chapter 4 Appendix

C.1 Robustness Checks

Table C.1 shows simulation results for markets with various adjustments for the purpose of robustness. As seen in the left panel, allowing memories to straddle periods is not driving the impressive results in the paper. In fact, resetting the history (and thus memories) each period yields slight improvements in most outcomes relative to the markets examined in the main text. Means for round-average prices and final prices fall just a few tenths short of the main simulations, though with tighter ranges. Measures of final distance and both efficiencies are just slightly improved in the markets with resetting memories; buyer and seller MRS actually shows a much tighter spread.

Where history-resetting markets see mild improvements in market success, even milder regressions (or very often, pushes) are seen in markets with no internal spread reduction rule. While this is not surprising as lower quality draws are forced into the orderbook, meaning a few will find their way into trades when the market is very thin. On average however, there is essentially no impact on market success, just a forced increase in the number of order placements and cancellations.

	History Reset			No Internal SR		
	Mean	St. Dev.	Range	Mean	St. Dev.	Range
I. Prices						
Price	2.28	0.17	(1.84, 2.92)	2.34	0.26	(1.64, 3.47)
$\|Price - CE\|$	0.52	0.12	(0.20, 0.99)	0.62	0.17	(0.28, 1.14)
RMSE	0.65	0.17	(0.24, 1.74)	0.77	0.24	(0.37, 2.07)
Final 5 Prices	2.32	0.20	(1.77, 2.94)	2.34	0.28	(1.64, 3.69)
II. Allocations						
Final Distance	0.73	0.36	(0.07, 1.84)	0.89	0.53	(0.05, 3.97)
Seller MRS	2.46	0.15	(2.01, 2.89)	2.38	0.22	(1.49, 3.06)
BuyerMRS	2.48	0.14	(2.15, 2.97)	2.56	0.24	(2.00, 3.87)
III. Efficiencies						
Allocative	0.97	0.02	(0.91, 1.00)	0.95	0.04	(0.64, 0.99)
Distance	0.80	0.05	(0.66, 0.92)	0.75	0.06	(0.43, 0.87)
Observations	240	240	240	240	240	240

Table C.1: Simulation Outcomes for Markets with Robustness Checks. Left Panel: markets with history reset at period start. Right Panel: markets where traders have no internal spread reduction rule

Bibliography

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.

Colin Camerer and Teck Hua Ho. Experience-weighted attraction learning in normal form games. *Econometrica*, 67(4):827–874, 1999.

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

Dave Cliff and Janet Bruten. More than zero intelligence needed for continuous double-auction trading. 1997a.

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

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, 2008a.

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

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.

Nick Feltovich. Reinforcement-based vs. belief-based learning models in experimental asymmetric-information games. *Econometrica*, 68(3):605–641, 2000.

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.

Drew Fudenberg and David K Levine. Consistency and cautious fictitious play. *Journal of Economic Dynamics and Control*, 19(5-7):1065–1089, 1995.

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.

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

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

Dhananjay K Gode and Shyam Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of political economy*, 101(1):119–137, 1993.

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.

Alvin E Roth and Ido Erev. Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. *Games and economic behavior*, 8 (1):164–212, 1995.

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. Opening the book: Price information's impact on market efficiency in the lab. 2021a.

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

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