

Minimal Intelligence in the Double Auction: Logit-Choice and Reservation Utility

Brett Williams*

Abstract

I propose a new boundedly-rational trader-behavior model for the continuous double auction, fit in the general equilibrium (GE) framework. Traders' order placement and acceptance strategies are driven by GE-adaptions of two classic mechanisms: within-period reservation prices (here utilities), and price-acceptability beliefs. Reservation utility associated with the side of entry is used to gauge immediate acceptability of contra-side orders while contra-side reservation utility informs the trader of admissible orders for placement. Traders place logit choice probabilities over their set of admissible orders, based on beliefs formed about the acceptability of each order. The logit choice parameter allows the precision of trader order placement to range from uniformly random to perfectly maximizing given beliefs. Simulations report market performance for this model, finding promising measures of efficiency and evidence of convergence to nearly Pareto optimal allocations. The model is then applied to experimental data from Williams (2025a) and Friedman et al. (2025), capturing an encouraging amount of the utility-reducing behavior and order placement in general.

Keywords: Continuous Double Auction, General Equilibrium, Exchange Economy, Logit Choice, Reservation Utility

JEL Classifications: C63, D44, D51, D83

*AGORA Centre for Market Design, UNSW Sydney; brett.williams2@unsw.edu.au
UNSW Business School (E12) Level 4 West Lobby, Cnr Union Rd & College Rd, UNSW Sydney, Kensington
NSW 2052, Australia

I thank Daniel Friedman, Jacob Goeree, Natalia Lazzati, Kristian López Vargas and Jason Tayawa for their comments and discussions.

Declaration of interest: none.

1 Introduction

Markets imbue agents with a centralized setting to reallocate commodities amongst themselves in an effort to arrive at a utility-improving state. Generally, the thought is that through markets this reallocation process is efficient. How agents in these markets traverse the landscape of even simple general equilibrium spaces such as an Edgeworth box is still to be determined to some extent. Many questions have been, are currently being and will continue to be formed about this process. This paper presents and investigates the following. Must every step on a markets path to equilibrium be utility-improving? Likely not, but what motivation guides traders to deviate and how can they do so in a productive manner?

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 move 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 (CDA) routinely break this assumption. There are a few potential reasons for this. The mechanism privy to this paper's model is that traders occasionally intentionally take utility losing trades in order to set themselves up for future trades on both sides of the market.

Such a mechanism places the model in the 'wilderness of bounded rationality.' (Farmer, 2003) A vast expanse of deviations from perfect rationality have been explored, with many more yet to be charted. Two gates to this wilderness 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 often attempted. 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 admissible¹ order; the logit choice parameter allows a trader's choice to capture precision anywhere between the random choice of zero-intelligence traders (Gode and Sunder (1993); Williams (2025b)) and the surplus-maximizing 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 the last twenty years). Advancements in behavioral learning models and a second wave of agent-based models of the double auction are also recounted. 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. The set of agents follow an order decision process and entry/re-entry rule that (1) loosen traditional restrictions on the placement of strictly-improving orders, (2) impart a more holistic view on market participation, and (3) incorporate the idea of adjusting an allocation for the sake of maneuverability (i.e. moving to a bundle with a more accommodating marginal rate of substitution for both directions of trade). Agents make use of time-dependent reservation utilities and logit choice 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 and prices, hints at relatively equitable reallocations near a point on the contract curve. Three new or updated measures of efficiency are presented, as well as two

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

penalized variants. The combination affords the ability to capture performance in utilities, prices and allocations and thus a fuller view on convergence.

I provide extensions and applications of the model in Section 6. Variations on the use of reservation utilities characterize alternative trader types, some of which are more 'minimal' in their intelligence than others. These types are then brought to the data of two recent laboratory market experiments, Williams (2025a) and Friedman et al. (2025). I find encouraging levels of categorization, with up to 90 percent of non-individually-rational (non-IR) orders being captured by the model. In a test of full-period (or market-duration) categorization, I find up to 78 percent of the trader-period pairs satisfy the criteria for at least one of the defined types (and up to 93 percent if the traders are allowed to make one error as that type). Section 7 concludes the paper.

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 the last fifty years. First, a batch of partial equilibrium (PE) models were proposed in the late 80's and throughout the 90's. Then, a newer wave of more generally applicable models 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.

A trio of models set the scene for the first wave. Wilson (1987) proposed a game theoretic model, positing a strategic multilateral setting where each trader's actions directly impact the pricing strategies of the other traders. 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. Two more influential models followed, providing bounded-rationality bookends for the wave. 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 of couple decades. A batch of parsimonious, tractable, often heuristic-driven models entered the learning 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 being chosen with some logit choice probability.³ Camerer and 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

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

³Feltovich (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.

(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 adaptations of the ZI model were promoted by Gode et al. (2004) and Crockett et al. (2008) a decade or so after the original model was published. The former features a 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. An alternative model to that of Gode, Spear and Sunder (2004), Williams (2025b) postulates a GE-based zero intelligence (from here, ZI-G) model with a lower sense of ‘zero’ intelligence. Williams (2025a) (henceforth ‘GD-G’) 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 Environment

I consider a market over two goods, X and Y . Traders in this market begin each trading day, or period, with some endowment of each of these goods, with the market totals of these endowments helping to define the Edgeworth box the traders are placed in. One of these goods is considered a standard commodity (let this be X), and the other considered as a numeraire (Y), as is the case with most if not all such settings since Shapley and Shubik (1967). Such an assumption allows the traders to trade quantities of X at prices represented by units of Y per unit of X in traditional auction settings.

⁴The timing in the paper lends itself to both an analysis of multiple iterations of a call market as well as a double auction in near continuous time. van de Leur and Anufriev (2018) extends the model via a more complex timing problem.

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.

In typical fashion, orders are comprised of a price, a quantity and a time of placement. Quantities are real numbers excluding 0, while prices are non-negative real numbers and times of entry are natural numbers. In practice, all three elements of the order have finite support. Prices are bounded above by some maximal price M , times are bounded above by the length of a trading period T , and quantities are bounded by either the allocation bundle of the trader or the total amount of each good in the economy.

In addition to the standard elements, each order holds information on associated traders (via placement or trade). The set of traders, $\mathcal{N}$, is partitioned into natural buyers ($\mathcal{B}$) and natural sellers ($\mathcal{S}$). Traders in $\mathcal{B}$ have marginal rates of substitution (MRS) above the competitive equilibrium price (calculated at the starting endowment bundle), while those in $\mathcal{S}$ have a starting MRS below the CE price.⁵ While these monikers hint at the preferred order type, traders are not restricted in their trade direction, i.e. they operate as two-way traders.

3.2 Exchange Definitions

A few subsets of the message space above define the key objects in the market. First, an **order** o is an element of $\mathcal{P} \times \mathcal{Q} \times \mathcal{T}$ placed by some trader in $\mathcal{N}$. The sign of the quantity element in an order determines whether it is deemed an ask (-) or bid (+). After an order is placed to the exchange it exists in what is called the **orderbook**, which is a ledger organizing orders into asks and bids and then by price and time. The best bid in the orderbook is the highest priced; if multiple bids exist at this price then the best is the earliest placed. A

⁵Note that MRS changes throughout the life of the market, potentially enough to transition a trader from one side of the competitive equilibrium price to the other.

similar process holds for asks, except price priority is instead from lowest to highest.

Trade in this exchange can occur in two ways. First, if an order is placed that **crosses** another order(s) already in the orderbook, i.e. a new bid (ask) that is priced higher (lower) than the best ask (bid), then a transactions will occur at the price of the already placed order. The amount of the good transacted is equal to the lesser of the two orders' quantities. The larger order will continue to live in the orderbook (with its updated desired quantity), while the smaller order will be removed as it has been entirely filled. Additionally, if the newer order crosses multiple (k) orders and has a large enough quantity to fill the first k-1 of them fully and the kth partially, then each of these individual transactions will occur. Each order retains the placing trader as well as the set of transacted traders. The second trade process is via directly **accepting** an existing order. In this case, the trader placing an order may accept an existing order by placing an order of the same price and the quantity's additive inverse.

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. The **history** of the market is the set of all orders and trades in the lifetime of the market. This history can be partitioned into three main subsets: **trades**, **cancelled orders** and the **orderbook**. Trades are the set of orders which have been placed in the orderbook, interacted with by at least one other trader than the trader who placed it, and filled some quantity either via crossing or acceptance. Cancelled orders are the set of orders which had a non-zero quantity left to be filled that were removed from the orderbook. Thus, all orders that were placed by traders and no longer exist in the orderbook are either in the set of trades or the set of cancelled orders.

Agents' recollection of the history of the market may or may not be complete. In this

sense, each agent has some **memory** of the history. Consistent with Gjerstad and Dickhaut (1998), this memory is defined as the set of trades and cancelled orders within the last L trades.

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 and quasi-concave. For the remainder of this paper, I'll focus on the constant elasticity of substitution functional form:

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

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.

I add the flavor of reservation prices to trader's preferences, 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. For example, an agent who enters as a buyer may set a reservation utility by considering an altered version of her utility function, one which now requires more X to satisfy a utility-improving order. Similarly an agent electing to sell may consider an adjusted utility function which readjusts her relative preferences to favor Y less than her true preferences. For CES preferences, this would mean adjusting the relative sizes of a_i and/or b_i .

One natural consideration would be to have a piece-wise, kinked indifference curve satisfying the above adjustments. Such a model would be an attempt at a GE-version of the reservation price model of Friedman (1991). While this is an interesting option, I take the




Figure 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. The two blue dotted line segments show best ask (right) and bid (left) 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.

position that this may be a waste of the ‘other halves’ of these adjusted indifference curves. Agents, instead, use each of these preference sets for two different order choice rules.

Figure 1 provides an example of the adjustment in curvature, and thus relative preferences over X and Y . The two rules, to be discussed at length in Section 4, dictate accepting orders and placing orders. Depending on the side of entry, the reservation utility IC above the true IC is used for the acceptance rule and the reservation IC below the true IC (the portion thrown out usually) guides the order placement rule.

These reservation utilities are captured via parameter η , where $\eta(t)$ is a function of within-period time t , and enters into the trader's utility function as follows

$$u_{i,b}(x, y|\eta) = c_i((a_i - \eta)^r x^r + (b_i + \eta)^r y^r)^{\frac{1}{r}} \quad (2)$$

$$u_{i,s}(x, y|\eta) = c_i((a_i + \eta)^r x^r + (b_i - \eta)^r y^r)^{\frac{1}{r}} \quad (3)$$

Here $u_{i,b}$ is the buyer reservation utility for trader i , and $u_{i,s}$ is the seller reservation utility.

A few desirable statics arise when determining an appropriate functional for η . First, as gains from trade decline over the life of the market and a smaller range of prices become competitive or desirable, $d\eta/dt < 0$ should be satisfied. Second, the size of the adjustment should be related to the trader's relative preference between the two goods. Namely, a trader who strongly prefers one good to the other may be less inclined to consider large deviations in their reservations away from their true preferences. Third, a trader's reservation preferences shouldn't change her outlook on a product from a 'good' to a 'bad', i.e. $\eta \leq \min\{a_i, b_i\}$. Thus, the form of $\eta(t, a, b)$ I consider in this paper is

$$\eta(t, a, b) = \left(\frac{T - t}{T} \right) \frac{\min\{a, b\}}{\max\{a, b\}} \min\{a, b\} \quad (4)$$

An final important note, however, is that the role of a and b 's relative sizes reverses when crossing r changes sign. Given r 's support lies on both sides of 0, an adjustment to η is needed to maintain the relationship between the three indifference curves as portrayed in Figure 1 (or, maintain the ordering of marginal rates of substitution across them). The ordering can be maintained by simply considering

$$\eta(t, a, b, r) \equiv \eta(t, a, b) * \text{sgn}(r) \quad (5)$$

A few special cases should be mentioned as well. First, with respect to functional form, perfect substitutes ($r = 1$), perfect complements ($r \rightarrow -\infty$) and Cobb-Douglas ($r \rightarrow 0_{-,+}$) are naturally folded into CES preferences. Both perfect substitutes and perfect complements provide interesting responses/interpretations when including η as in equations 2 and 3. The former, graphically, mimics reference prices from the partial equilibrium literature of the 1990's, as the slope of the IC gives a natural reservation price. The latter is analogous to an adjustment in the desired complement ratio.

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 (2025a), 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 of which prices may 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.

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 4.4, 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.

4.2 Belief Updating

First, 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 (2025a) provides a general-equilibrium-adjusted version of Gjerstad and Dickhaut's model, in which each order is given weight equal to the proportion of the original quantity successfully traded, $\sqrt{q_k} \frac{q_{k, \text{traded}}}{q_k}$. 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. In the notation of Gjerstad and Dickhaut (1998), for some bid ρ ,

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

represents the probability of acceptance. The analogous definition for sell price acceptability is

$$p_s(s) = \frac{\sum_{\rho \geq s} TA(\rho) + \sum_{\rho \geq s} TB(\rho)}{\sum_{\rho \geq s} TA(\rho) + \sum_{\rho \geq s} TB(\rho) + \sum_{\rho \leq s} RA(\rho)} \quad (7)$$

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.

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 crossing orders 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.

4.3.1 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 whose price is better 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, $BP_{-\Delta}$, and his current reservation utility, or:

$$|BP_{-\Delta} - MRS_{u_i}| - |MRS_{u_{i,\Delta}} - MRS_{u_i}| > 0 \quad (8)$$

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. Figure 1 displays such a check, where an entry on the buy side could yield an auto-accept as the best ask price vector lies above the stricter reservation utility. If the quantity associated with the ask yields a utility improvement, then the trader will accept the order.

4.3.2 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

$$\mathcal{P}_{u_{i,-\Delta}} \times \mathcal{Q}_{u_{i,-\Delta}} \equiv (\min\{|MRS_{u_{i,-\Delta}}|, Boundary_{\Delta}\}, \max\{|MRS_{u_{i,-\Delta}}|, Boundary_{\Delta}\}) \times (0, \bar{x}(\Delta)] \quad (9)$$

$\mathcal{P}_{u_{i,-\Delta}}$ 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 ($Boundary_\Delta$; 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-zero⁶ 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 probability:

$$Pr(o_z | x_k, y_k, \Omega_M) = \frac{\exp[\lambda p_\Delta(o_z)(u_{i,-\Delta}(o_z))]}{\sum_{o'_z \in O_z} \exp[\lambda p_\Delta(o'_z)(u_{i,-\Delta}(o'_z))]} \quad (10)$$

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, $\lambda \rightarrow \infty$ would imply perfect choice as in GD-G.

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. 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

⁶The "zero" intersection here would be at the trader's current allocation.

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.

5 Simulations

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 4.4.⁷

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 1: Simulated Agent Parameters.

⁷As 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, 146 out of 200 orders were placed per period on average. See Appendix A.4 for simulations without such a reduction rule.

Given these parameters, the natural buyer's and seller's η adjustments begin at 0.035 and 0.142, decaying to 0 by the end of each period. The competitive equilibrium price associated with the traders' true preferences, u_i , at the endowment point is 2.44. If traders trade solely based on their stricter preferences, $u_{i,\Delta}$, the CE price calculated at time 0 and for the starting endowments is 2.82.⁸ Following only lower reservation utilities, $u_{i,-\Delta}$, the starting CE price would be 2.24.

All traders maintain a memory of $L = 5$, implying they can perfectly recall all transactions and order cancellations in the market within the last five transactions.⁹ This memory may span across market periods within the same run, however may not carry over between runs.¹⁰ Additionally, when conducting the logit choice procedure over the set of feasible orders, each trader will have a logit choice parameter, λ , of 5. This places traders' choice precision between uniformly random and perfect, leaning more on the side of random. See Appendices A.1 for robustness runs testing history choices.

5.2 Measuring Efficiency

Traditionally, the measure coined as 'allocative efficiency' has been used as the main efficiency measure in this type of simple market. The measure comes from the partial equilibrium literature, defined as the sum of profits made by each of the traders on each of their units over the expected gain across all units in equilibrium. More formally, this is written as

$$\frac{\sum_{b=1}^B \sum_{i=1}^{P_b} (p_{i,b} - v_{i,b}) + \sum_{s=1}^S \sum_{j=1}^{P_s} (c_{j,s} - p_{j,s})}{\sum_{b=1}^B \sum_{i=1}^{P_b} (p_{CE} - v_{i,b}) + \sum_{s=1}^S \sum_{j=1}^{P_s} (c_{j,s} - p_{CE})} \quad (11)$$

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. The sets $\{v_i\}$ and $\{c_j\}$ are

⁸This falls to 2.59 at the halfway point of a period.

⁹This is the memory length used in Gjerstad and Dickhaut (1998). I test two other memory lengths, 0 and 10, book-ending this choice for robustness. See Appendix A.2.

¹⁰This assumption is tested in a batch of simulations with trader memories that refresh at the beginning of every period.

the buyers' and sellers' redemption values and unit costs, while $\{p_{i,b}\}$ and $\{p_{j,s}\}$ are the buy and sell trade prices and p_{CE} is the competitive equilibrium price.

One natural analogue to this measure in GE is to replace the main PE outcome, cash profit, with the main GE outcome, utility gain. As such, I term the sum of realized utility gains over expected utility gains, seen in equation (12), as allocative efficiency in a two-good Edgeworth box economy.

$$E^{Alloc} \equiv \frac{\sum_{i=1}^N (u_i(x_{i,T}, y_{i,T}) - u_i(x_{i,0}, y_{i,0}))}{\sum_{i=1}^N (u_i(x_i^{CE}, y_i^{CE}) - u_i(x_{i,0}, y_{i,0}))} \quad (12)$$

While this provides a nice outlook on the gains from trade reaped by the market in utility terms, the measure is indiscriminating in terms of relative gains across individual traders. For example, for two market realizations both not reaching competitive equilibrium, it is entirely possible the two vectors of utility gains are noticeably different. In fact, the two vectors need not even have equivalent lengths. I provide a penalized version of allocative efficiency as well, where the value in (12) is multiplied by a scaling penalty defined in (13). For interpretation of the following notation, $\vec{v}_{\omega \rightarrow T}$ is the vector from the tuple of utilities at a period's starting endowment to the tuple of final utilities at time T in the period and $\vec{v}_{\omega \rightarrow CE}$ is the vector from the tuple of utilities at a period's starting endowment to the tuple of expected utilities in competitive equilibrium.

$$\frac{\|proj_{\vec{v}_{\omega \rightarrow CE}} \vec{v}_{\omega \rightarrow T}\|}{\|\vec{v}_{\omega \rightarrow T}\|} \quad \text{where} \quad proj_{\vec{v}_{\omega \rightarrow CE}} \vec{v}_{\omega \rightarrow T} \equiv \frac{\vec{v}_{\omega \rightarrow T} \cdot \vec{v}_{\omega \rightarrow CE}}{\|\vec{v}_{\omega \rightarrow CE}\|^2} \vec{v}_{\omega \rightarrow CE} \quad (13)$$

Thus, if the tuple of final utilities deviates greatly from the path determined by $\vec{v}_{\omega \rightarrow CE}$, the estimate is scaled down proportional to its deviation.

I also provide a second interpretation of partial equilibrium's allocative efficiency in general equilibrium, named 'profit efficiency'. For each trade made by a pair of traders, each trader i has a price, call it $p^{Ind}(i)$, at which she could have moved the same number of units of X on the side she traded while remaining on the same indifference curve. The difference

in this price and the actual trade price, multiplied by the traded quantity,¹¹ is the GE analog of a cash gain, i.e. the gain (or loss) in the numeraire. Such a gain summed across all trades across all traders divided by the gain when replacing each trade price with the CE price defines this profit efficiency, as in (14). For notation, let A_τ be the set of actions a resulting in a trade, i.e. the a represented by the set of orders o comprising Ω_τ .

$$E^{Profit} \equiv \frac{\sum_{a \in A_\tau} [\kappa_a q_a (p^{Ind}(b_a) - p_a) + \kappa_a q_a (p_a - p^{Ind}(s_a))]}{\sum_{a \in A_\tau} [\kappa_a q_a (p^{Ind}(b_a) - p_{CE}) + \kappa_a q_a (p_{CE} - p^{Ind}(s_a))]} \quad (14)$$

Though these two efficiency measures provide a meaningful account of the market's ability to capture gains from trade in utility and price terms, neither seems to capture the market's path and proximity to the equilibrium allocation. The third measure of efficiency I examine, 'distance efficiency', aims to capture the market's performance in allocations. Equation (15) defines the statistic as a deviation from one hundred percent. The deviation is measured as the distance to the equilibrium allocation bundle, $\{x^{CE}, y^{CE}\}$, at the end of a trading period relative the total distance traveled in equilibrium.

$$E^{Dist} \equiv 1 - \frac{d_{CE}(\{x^T, y^T\})}{d_{CE}(\{x^\omega, y^\omega\})} \quad (15)$$

with distance measure

$$d_{CE}(\{x, y\}) \equiv \frac{1}{N} \sum_{i=1}^N \sqrt{(x_i - x_{CE})^2 + \left(\frac{1}{p_{CE}}(y_i - y_{CE})\right)^2} \quad (16)$$

The endowment bundle and final bundle are denoted as $\{x^\omega, y^\omega\}$ and $\{x^T, y^T\}$.

Much like with the allocative efficiency measure, this measure may not fully account for large variation in allocative gains across traders. Thus, I provide a penalized version of distance efficiency as well. In line with the notation of the penalized allocative efficiency, I denote the vector from the endowment bundle to the equilibrium bundle as $\vec{\alpha}_{\omega \rightarrow CE}$ and the

¹¹For a trade occurring in action a , this is $\kappa_a q_a$, or the proportion of the order filled multiplied by the quantity desired.

vector from the endowment to the final allocation as $\vec{\alpha}_{\omega \rightarrow T}$. As the set of bundles that satisfy the same level of distance efficiency all lie on the surface of the same ball around $\{x^{CE}, y^{CE}\}$, the penalty should increase in intensity as the final bundle deviates from $\vec{\alpha}_{\omega \rightarrow CE}$ and as the market begins over-trading. The penalty is thus defined as in equation (17), and is applied as a scaling term multiplying the distance efficiency measure in practice.¹²

$$\frac{\|\vec{\alpha}_{\omega \rightarrow CE}\| - \|\vec{\alpha}_{T \rightarrow CE}\|}{\|proj_{\vec{\alpha}_{\omega \rightarrow CE}} \vec{\alpha}_{\omega \rightarrow T}\|} \quad (17)$$

Notice the numerator refers to the length of the vector from the endowment bundle to the point lying on $\vec{\alpha}_{\omega \rightarrow CE}$ which is equidistant from $\{x^{CE}, y^{CE}\}$ as $\{x^T, y^T\}$.

5.3 Performance

Table 2 records the main performance measures for the simulated markets. A quick glance shows evidence of surprisingly successful markets. Estimates show promising 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 0.16 above the CE prediction from market inception, which, when accompanied with a relatively low average deviation, implies transaction prices lie in a tight band around the CE price. Figure 2 confirms not only the round-averages, but the individual transaction prices across the markets are closely bound. The per-unit average price (total units of y traded divided by total units of x traded across the period) is less sensitive to high outlier prices as they are accompanied with small trade quantities; de-weighting as such yields an estimate just 0.01 unit away from CE. Convergence within period, however, requires tighter bounds on the time in focus. The final batch of transactions in a period provide an idea of the traders' desire to trade and urgency to reap more gains from trade. I find an estimate

¹² $\|\cdot\|$ here is the standard Euclidean norm.

	Mean	St. Dev.	Range
<i>I. Prices</i>			
Price	2.60	0.22	(2.07, 3.51)
Per-Unit Avg.	2.43	0.26	(1.83, 3.55)
$ Price - CE $	0.55	0.17	(0.24, 1.34)
RMSE	0.95	0.34	(0.31, 2.10)
Final 5 Prices	2.47	0.14	(2.13, 3.30)
<i>II. Allocations</i>			
Final Distance	0.86	0.58	(0.03, 3.79)
Seller MRS	2.41	0.17	(1.56, 2.79)
BuyerMRS	2.54	0.20	(2.18, 3.80)
<i>III. Efficiencies</i>			
Allocative	0.99	0.01	(0.91, 1.00)
- Penalized	0.77	0.10	(0.43, 0.97)
Distance	0.82	0.05	(0.65, 0.95)
- Penalized	0.68	0.08	(0.41, 0.88)
Profit	0.91	0.08	(0.67, 1.13)
Observations	480	480	480

Table 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 three measures of efficiency.

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 reveal a portion 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 collapse¹³ the two types of traders into representative agents. These agents can aptly be represented in the Edgeworth box. On average, the final distance¹⁴ the pair lies away from equilibrium allocation bundle pair is within a unit radius of

¹³For 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.

¹⁴I.e. the Euclidian distance that the representative buyer (and equivalently, seller) is away from the




Figure 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 final. Allocations approach the contract curve, on average lying in nearly Pareto optimal final resting places. As displayed in Figure 3, the geometric mean of the final allocations is very close to the equilibrium bundle, and even closer to the contract curve in general. In fact, the vast majority of the final allocations lie on the contract curve or quite close.

The marginal rate of substitution of market participants gives a proxy for 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. Figure 4 reinforces such a claim, as around 90% of the periods yield an MRS

equilibrium allocation in the Edgeworth box. The y contribution to the distance is de-weighted by the equilibrium price. The distance function is thus $dist(\cdot) = \sqrt{(x_i - x_{CE})^2 + (\frac{1}{p_{CE}}(y_i - y_{CE}))^2}$.




Figure 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 0.478 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.

spread within 0.5 units.

Three measures of efficiency are estimated: allocative, distance and profit. Each is meant to capture the performance of the market in a different convergence indicator: utility, allocations, and prices. The simulations report promising results in each.

First, the base measure of allocative efficiency reports an average of 0.99 with a minimum estimate of 0.91. The market clearly is capturing essentially all of the gains from trade achievable in utility-terms. However, is this being driven by one or two traders dominating the market or are gains seen by all traders? The penalized average of 0.77 reflects some mild deviation from the path given by $\vec{v}_{\omega \rightarrow CE}$, with the minimum estimate drops from 0.91 to 0.43. A Spearman rank correlation of 0.35 between the two measures suggests high efficiency periods are not entirely dependent on markets with over-equitable or inequitable utility gain distributions.

While gains from trade in utility terms are mostly being realized, is the market ar-




Figure 4: The left figure shows the densities for round-end representative agent Buyer (red) and Seller (green) MRS. The black line represents the CE price. The right figure shows the CDF for the round-end gap between the Buyer and Seller MRS.

iving at the correct final allocation? Figure 3 provides a hint in two dimensions via the representative agents. The base measure for distance efficiency corroborates these findings in space represented all traders individually, achieving an average value of 0.82. While an 18% average deviation may seem large, the localization along the contract curve in Figure 3 suggests these deviations are likely lateral. The penalized estimate of 0.66 supports this, as the penalty is $0.68/0.82$, or 0.83. This implies the average distance traveled is 20% greater than the numerator ($\|\vec{\alpha}_{\omega \rightarrow CE}\| - \|\vec{\alpha}_{T \rightarrow CE}\|$) which is equivalent to the base measure estimate of 0.82. Therefore, the average distance traveled essentially matches that of the equilibrium path, matching the idea that the majority of the markets are finishing near a point on the contract curve.

The final efficiency measure, profit efficiency, examines the market's ability to maximize gains from trade via price selection. As over-trading is possible in this model, the value can exceed 1. The average profit efficiency for the simulations is 0.91. Despite the ability to intentionally place utility-losing offers, the markets capture nearly all of the gains in the numeraire compared to equilibrium predictions. Sellers appear to be the group of traders compromising on price more as they are capturing only 60% of their expected profit in




Figure 5: Cumulative Density Functions for Efficiency.

Y , while buyers are overperforming with 124%. While pronounced, this difference is not particularly surprising as sellers' η is much larger throughout the period, leading them to post much more aggressive orders regardless of the side they enter. Additionally, the lower η of the buyers leads to a higher likelihood of auto-accepting those aggressive orders.

5.4 Adjustments in λ

The left panel of Table 3 shows estimates for markets with a logit parameter of 0, meaning traders are placing uniform probability over their order choices. Unsurprisingly, nearly all estimates are worse than the $\lambda = 5$ markets reported in Section 5.3. What is surprising is how small some of the differences are. Buyer and Seller MRS still sit about 0.1 unit on either side of the CE price, though the supports suggest overtrading occurs more often when taking away the use for price beliefs. The deviations in efficiency estimates match are similar in magnitude to those of the $L=0$ markets in Appendix A.2; this suggests having no order history is as harmful to market performance as random selection. Price estimates show the largest deviation, with average price over 0.3 away from CE and some periods having an average nearly double the CE price.

	$\lambda = 0$			$\lambda = 20$		
	Mean	St. Dev.	Range	Mean	St. Dev.	Range
<i>I. Prices</i>						
Price	2.78	0.42	(1.65, 4.65)	2.54	0.21	(0.96, 3.33)
Per-Unit Avg.	2.52	0.34	(1.55, 3.53)	2.46	0.24	(1.03, 3.18)
$ Price - CE $	0.94	0.35	(0.34, 2.86)	0.37	0.17	(0.18, 1.48)
RMSE	1.40	0.60	(0.42, 3.67)	0.61	0.26	(0.25, 1.92)
Final 5 Prices	2.56	0.31	(1.87, 3.66)	2.44	0.17	(0.96, 3.07)
<i>II. Allocations</i>						
Final Distance	0.81	0.50	(0.01, 2.72)	0.80	0.79	(0.03, 6.63)
Seller MRS	2.36	0.19	(1.75, 3.02)	2.35	0.23	(0.90, 2.61)
Buyer MRS	2.55	0.21	(2.07, 3.34)	2.59	0.31	(2.29, 5.31)
<i>III. Efficiencies</i>						
Allocative	0.94	0.06	(0.44, 1.00)	0.99	0.07	(0.16, 1.00)
-Penalized	0.76	0.12	(0.15, 0.97)	0.87	0.11	(0.06, 0.98)
Distance	0.75	0.08	(0.39, 0.90)	0.86	0.08	(0.09, 0.95)
-Penalized	0.56	0.12	(0.11, 0.80)	0.76	0.10	(0.08, 0.91)
Profit	0.86	0.09	(0.15, 1.07)	0.92	0.08	(0.21, 1.08)
Observations	240	240	240	239	239	239

Table 3: Simulation outcomes for markets with logit parameter, λ , values of 0 and 20. One round in the $\lambda = 20$ simulations had zero trades, hence the lower observation count.

The right panel shows estimates for markets with much more intelligent traders. With a logit parameter of $\lambda = 20$, small adjustments in expected utility across order options are more discernible to the trader, allowing for more precise order choice. While estimates in these markets do outperform the $\lambda = 0$ estimates, they do not outperform the main $\lambda = 5$ results. In some cases, estimates are marginally worse, in fact. Penalized distance efficiency is noticeably better (base Distance efficiency only marginally) in these more intelligent markets, while all Panel II estimates are quite close to the Table 1 estimates. Markets are thus arriving at final allocations a similar distance away from the CE bundle as the $\lambda = 5$ markets, however with these points lying closer to $\vec{v}_{\omega \rightarrow CE}$.

6 Extensions and Applications

6.1 Types

A natural extension of this framework is to consider traders of different types, or with alternative motives for holding reference utilities. I introduce three interesting types, in addition to the trader introduced in Section 3, that make use of these side-specific reservation utilities in ways that intuitively fit different behavior profiles. First, consider traders' patience, or impatience, over participating in trades and reallocating. Two profiles or types can be quickly defined by considering the inner and outer "kinked" preferences created by a trader only considering one reservation utility for each side of the market for both order types. For instance, consider a trader who is patient, only willing to reallocate at very advantageous prices to them both immediately and in the future. Then, that trader may consider only $u_{i,\Delta}$ for both accepts and limit orders; this is equivalent to considering a piecewise kinked utility function where $u_{i,s}$ is considered for sells and $u_{i,b}$ for buys.¹⁵ This is also the most natural adaption of the reservation price behavior considered in partial equilibrium (Friedman (1991)).

Alternatively, a trader may be exceedingly impatient. If the trader is operating under the intent to accept orders, or reallocate, as often as possible (and maintain a sense of bettering their reservation), they may consider only their more permissive reservation utilities (i.e. $u_{i,-\Delta}$). Using same-side reservation utility yields weakly more opportunities to accept an order upon entry, as the subset of the Edgeworth box (given entry side) that is same-side reservation-utility-improving contains the subset that is improving when considering the contra-side reservation as a proper subset for any $t < T$. Additionally, since the beliefs over price acceptability are weakly monotonic in the direction of orderbook ordering, the limit orders they place are more likely to be accepted in less time. I'll refer to these two types as

¹⁵Crockett and Oprea (2012) also consider a kinked utility function in a general equilibrium setting, though under different behavioral considerations.

Patient (type-P) and *Impatient (type-I)* traders, respectively.

The last alternative type I'll define here is essentially the reverse of the trader defined in Section 3. This type of trader similarly wants to consider different reservation utilities based on side and order type in an effort to reposition. However, they are naive in their implementation, in the end switching which reservations are used for which order type. Namely, the trader considers their same-side reservation for limit orders, while using their contra-side reservation for market orders. I'll refer to these types as *Naively Repositioning (type-NR)*. Similarly, I'll refer to the traders defined in Section 3 as *Strategically Repositioning (type-SR)*.

	Naive Repositioning			Impatient			Patient		
	Mean	St. Dev.	Range	Mean	St. Dev.	Range	Mean	St. Dev.	Range
<i>I. Prices</i>									
Price	2.46	0.29	(1.83, 3.84)	2.30	0.14	(1.92, 2.74)	2.70	0.38	(1.05, 3.71)
Per-Unit Avg.	2.37	0.27	(1.80, 3.47)	2.25	0.15	(1.83, 2.79)	2.68	0.38	(1.06, 3.72)
$ Price - CE $	0.74	0.15	(0.38, 1.40)	0.52	0.12	(0.20, 0.90)	0.73	0.26	(0.26, 1.62)
RMSE	0.91	0.18	(0.50, 1.92)	0.68	0.16	(0.29, 1.44)	1.18	0.51	(0.38, 2.82)
Final 5 Prices	2.36	0.35	(1.48, 3.84)	2.24	0.17	(1.68, 2.64)	2.47	0.25	(1.05, 3.37)
<i>II. Allocations</i>									
Final Distance	0.92	0.65	(0.06, 4.86)	0.66	0.40	(0.03, 2.33)	0.74	0.80	(0.03, 7.30)
Seller MRS	2.34	0.24	(1.34, 2.91)	2.46	0.16	(1.93, 3.17)	2.29	0.20	(0.78, 2.53)
Buyer MRS	2.59	0.30	(2.02, 4.32)	2.45	0.15	(2.05, 3.10)	2.61	0.36	(2.32, 5.95)
<i>III. Efficiencies</i>									
Allocative	0.93	0.07	(0.53, 0.99)	0.97	0.06	(0.60, 1.00)	0.97	0.10	(0.02, 1.00)
-Penalized	0.41	0.11	(0.04, 0.74)	0.39	0.11	(-0.06, 0.72)	0.88	0.11	(0.01, 0.97)
Distance	0.77	0.10	(0.28, 0.96)	0.79	0.11	(0.34, 0.95)	0.84	0.10	(0.01, 0.93)
-Penalized	0.33	0.11	(0.02, 0.71)	0.32	0.11	(-0.04, 0.63)	0.73	0.11	(0.00, 0.88)
Profit	0.81	0.09	(0.26, 0.97)	0.85	0.11	(0.25, 1.02)	0.93	0.10	(0.02, 1.10)
Observations	240	240	240	240	240	240	238	238	238

Table 4: Simulation outcomes for alternative trader types, Naive Repositioning (left), Impatient (middle) and Patient (right).

In Table 4, I report performance estimates for simulated markets of each of the above variants. Under the same moderate λ tested in type-SR markets, I find each alternative arrives at roughly similar base allocative efficiency levels, though via noticeably different processes. In price, averages arrive from the anticipated directions. Impatient traders trade at below base-CE prices, driven by larger concessions in price being made on the sell side than the buy side. Similarly, Patient traders approach CE price from above, with the average final

five prices lying closer to CE than the per-unit average price. These are consistent with the starting CE prices when taking the reservation utilities being used as given, with a CE price of 2.82 for a market of all type-P traders and 2.24 for an all-type-I market. As shown later in Figure 6, these ranges of CE prices over t correspond to and support experimental findings of along considerable subsections of the contract curve (e.g. (Crockett, 2008); (Gjerstad, 2013); Friedman et al. (2025)), and thus variation in trade price across market period and session. Much like with type-SR markets, type-NR markets report price averages centered near the base CE price.

While the allocative efficiencies at base are quite high (and similar to type-SR levels), penalizing the measure reveals significant dispersion in utility gain at the individual level. Type-NR and Type-I markets plummet to less than half of their base estimates, with average penalized allocative efficiencies of 0.41 and 0.39 respectively. Penalized estimates for type-P markets come in at over twice those of the other alternate type markets, falling by only 0.1 roughly. A similar story appears in distance efficiency measures, with base estimates all roughly near 0.8, yet penalized values showing a sizable gap growing between type-NR/type-I performance and type-P. Estimate comparison between Patient trader markets and type-SR markets from Section 5.3 reveals type-SR markets are able to capture utility losing behavior, while maintaining similar efficiency values (penalized or not) and even yielding prices nearer to the base CE price.




Figure 6: Final allocations for individual traders (dots) and aggregate agents (level sets). Green dots denote seller final allocations while red dots denote buyers. The large blue dot reports the geometric mean of the aggregate agent final allocations.

As hinted by the penalized efficiency estimates, final allocations at the individual and aggregate levels, shown in Figure 6, vary greatly across alternative type. Perhaps unsurprisingly, the individual and aggregate final allocations are worst for the type-NR markets. Impatient trader allocations are able to gather closer to the contract curve, though spread wider along said curve. Both market types yield several utility-reducing final allocations at the individual level. Patient trader markets, in the aggregate cluster close to the base CE allocation, though with individuals falling short of the curve. This is an artifact of lower trader count relative to the other two alternatives; given all orders are placed relative to the less permissible reservation utility, prices are less desirable for larger portions of the market’s duration.

6.2 Experimental Data

As natural application of the model, I test data from two recent general equilibrium market experiments. Williams (2025a) uses the same parametrization as in Section 5.3, testing variation in orderbook and trade history transparency in a continuous double auction. I test and categorize play in the full transparency sessions from the experiment. The second batch of laboratory markets tested are from Friedman et al. (2025), where four user interfaces with increasing market visualization are implemented. Session layout and parametrizations for each of the papers are summarized below in Table 5.

Williams (2025)					
a. Parameters	a	r	Endowment	CE p	CE Δx
Buyer/Seller	0.825 / 0.6875	0.5 / 0.5	(3,23) / (11,3)	2.44	5.2
b. Sessions	# sessions	Traders per	Periods per	Treatments Tested	
	4	8	6	Full Transparency	
Friedman et al. (2025)					
a. Parameters	a	r	Endowment	CE p	CE Δx
Buyer/Seller (LoNeg)	0.65 / 0.779	-0.5 / -0.5	(20, 250) / (140, 20)	1.4	74.6
Buyer/Seller (HiPos)	0.35 / 0.843	0.25 / 0.25	(20, 250) / (140, 20)	3.3	74.6
b. Sessions	# sessions	Traders per	Periods per	Treatments Tested	
Inexperienced	8 (4)	10	6	No Frontier, Combo	
Experienced	4 (2)	10	9	No Frontier, Combo	
Block Treatments	LoNeg then HiPos (order reversed in even-numbered sessions)				

Table 5: Summary of Experimental Sessions tested from Williams (2025) and Friedman et al. (2025). In parametrizing CES preferences, both papers set $b = 1 - a$.

The three parametrizations tested across Williams (2025a) and Friedman et al. (2025) span the range of CES preferences which guarantee a unique competitive equilibrium, i.e. those which have $r \in [-1, 1]$. Varying this parameter both influences the strength of income effects in the economy, as well as how permissive reservation utilities are for traders. In addition to the relevance of the tested parametrizations, the user interfaces used in these laboratory markets (see Appendix B) are best fit to test this paper’s model. Laboratory traders interact via heatmaps over the space of feasible allocations, with colors denoting the utility associated with each bundle of X and Y (up to the pixel). Overlaid on the heatmap is also the trader’s current indifference curve, and in some cases even a frontier visualizing the current state of the orderbook. Thus, these interfaces make utility-reducing orders and accepts highly visible to the trader, meaning these actions are likely more intentional in this setting than other market experiments using simple grids or calculators to induce preferences.

As a first check of the data against the model, I check individual orders and accepts

against the reservation utilities predicted of each trader. Four different η 's are checked; two functional forms for the portion associated with time and two for the portion dependent on relative good preference (i.e. a and b). The functional forms are listed in Table 6, all of which satisfy the desired statics described in Section 3.4. Each has η decaying with time, and being bounded by the relative preference of the trader's less desired good (guaranteeing both goods remain 'goods').

Tested η 's	
Base (B)	$\left(\frac{T-t}{T}\right) \frac{\min\{a,b\}}{\max\{a,b\}} \min\{a,b\}$
Slow Base (SB)	$\sqrt{\left(\frac{T-t}{T}\right) \frac{\min\{a,b\}}{\max\{a,b\}}} \min\{a,b\}$
Wide (W)	$\left(\frac{T-t}{T}\right) \min\{a,b\}$
Slow Wide (SW)	$\sqrt{\left(\frac{T-t}{T}\right)} \min\{a,b\}$

Table 6: Functional forms of η tested against experimental data.

Table 7 summarizes this first check, reporting percentages of utility-reducing orders and accepts that are captured by the reservation utilities driven by η 's listed above. So long as the order or accept is captured for at least one of the trader types from Section 6.1, the count captures it. The set of non-IR actions in general is not insignificant across these experiments, with roughly 15 to 40 percent of all actions by inexperienced traders being utility-reducing (roughly 10 percent for experienced), depending on parametrization. Among these, Table 7 reports promising model-capture percentages. Using the most restrictive η , Base, between 26 and 69 percent of all non-IR actions are rationalizable by the model.¹⁶ Non-IR orders directly placed to the orderbook are effectively represented in the model, as up to 90 percent are captured when considering the most permissive reservation utilities. Non-IR accepts appear more difficult to capture, with levels ranging from 0.10 to 0.62 percent across η 's. However, experience appears to improve this, as every count in the Friedman et al. (2025) data improved, holding η constant. These leftout accepts are generally either late in the

¹⁶The lower counts reported for Williams (2025) can likely be attributed in part by the sessions being run online during Covid shutdowns in 2021.

round, after the reservation utilities have converged close to the trader’s true preferences, or are at more extreme prices and small quantities (or, nearly directly below or directly left of the trader’s current allocation in the Edgeworth box).

	Williams (2025)				Friedman et al. (2025)				Friedman et al. (2025)			
	r=0.5				r=0.25				r=-0.5			
	B	SB	W	SW	B	SB	W	SW	B	SB	W	SW
Inexperienced												
Orders	0.36	0.56	0.43	0.66	0.86	0.89	0.87	0.90	0.80	0.84	0.81	0.86
Accepts	0.14	0.38	0.20	0.50	0.10	0.27	0.14	0.38	0.12	0.32	0.19	0.44
All	0.26	0.48	0.30	0.59	0.69	0.75	0.70	0.79	0.59	0.68	0.60	0.74
Experienced												
Orders	-	-	-	-	0.78	0.83	0.79	0.87	0.81	0.87	0.83	0.89
Accepts	-	-	-	-	0.16	0.51	0.23	0.62	0.27	0.45	0.33	0.56
All	-	-	-	-	0.51	0.64	0.52	0.70	0.51	0.64	0.52	0.70

Table 7: Percent of Non-IR Orders and Non-IR Accepts that are captured by the model using various η ’s

For a stricter test of the model and how data is categorized by the proposed types, I check labelling at the trader-period level.¹⁷ For each trader in each period, all actions taken by that trader forms a data set where each row is an action; if every action matches a criteria of a specific type, that trader can be categorized as that type. In other words, every order placed and accept made is checked against the predicted reservation utilities for that trader at that time in the period. Figure C.1 reports the proportion of trader-period data points that can be categorized by the model (benchmarked against IR), and which types they are labelled as within the set that are categorizable. In every parametrization and experience level of the data tested, the models can predict between 7 and 19 percent points more of the data than the IR benchmark under the most permissive η tested, with 48 (53) to 71 (78) percent of the trader-period pairs being categorizable by the model overall under the least (most) permissive η . If one error is allowed in the categorization of trader-period pairs (i.e.

¹⁷For example, in the Williams (2025) markets there are eight traders per period and 12 periods per session, meaning there are 96 trader-periods to categorize per session.

$k-1$ of the k orders of a trader represent a type), then these numbers improve to 69 (75) and 89 (93) respectively (see Appendix C).

Within categorizable trader-period pairs, the vast majority fall within the SR type described in Section 4, and the type nested within it, namely Patient (or type-P). Roughly 90 percent of the data across parametrizations is covered by type-SR and type-P when considering the Base form for η . Relaxing the functional form with respect to relative good preference, i.e. moving to the Wide form, results in a shift away from type-P categorization to type-I or type-SR. While, moving from Base to Wide creates more room for non-IR orders and accepts, it similarly creates less room for orders and accepts that satisfy the upper reservation utility. As a result, type-I and type-SR catch the traders who aren't as restrictive in their play. Reducing the pace at which traders collapse their reservation utilities to their true preferences (or considering the Slow versions of each form) shows a similar displacement when holding Base or Wide constant.

7 Concluding Remarks

This paper models market dynamics in an Edgeworth box where traders have ‘imperfect’ choice procedures when placing orders in a continuous double auction. 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 and fill. Agents recognize that they may participate on both sides of the market, and develop reservations depending on which side they 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 make order selections that appear to be utility-reducing relative to their true preferences, though allow




Figure 7: Type Categorizations at the Trader-Period level. Left plots reports percentage of data points covered by IR versus all types for each η . Right plots show type breakdown, within data that is covered by the model.

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. Unlike many previous studies, our simulations stop at a preset time and do not automatically continue until gains from trade are exhausted. Therefore it is especially promising that final allocations in our simulation model are so consistently so close to Pareto optimal. 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.

When brought to the data of two recent laboratory experiments, Williams (2025a) and Friedman et al. (2025), the model categorizes an impressive proportion of the non-IR behavior, and order placement in general. The model provides an improvement in action and trader categorization than just assuming individual rationality; up to 79 percent of the actions taking (within one of the parametrizations) can be captured by the model. Four behavioral types are then used to label traders at the period level. Roughly half to three quarters of the trader-period pairs are rationalizable as one of the four types. Patient and Strategically Repositioning types dominate the landscape, capturing up to 90 percent of the traders that the model can reconcile.

The major implication of the findings of this project is the feasibility of boundedly rational order placement decisions in markets that show convergent tendencies. Specifically, 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 preferring perfectly rationality may suspect; in fact, estimates here perform near or level with some more complex models. Furthermore, the model pro-

vides a mapping from the zero intelligence gate (beginning with ZI) through the wilderness to a model fit much closer to the rational gate (this being Gjerstad and Dickhaut's (1998) 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 overall time in the market (aggregated across periods), provide 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 (2025a), an inclusion of prices in the orderbook in the belief updating process would likely improve fit.

References

- 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. A behavioral model for mechanism design: Individual evolutionary learning. *Journal of Economic Behavior & Organization*, 78(3):374–395, 2011.
- Colin Camerer and Teck Hua Ho. Experience-weighted attraction learning in normal form games. *Econometrica*, 67(4):827–874, 1999.
- Sean Crockett. Learning competitive equilibrium in laboratory exchange economies. *Economic Theory*, 34(1):157–180, 2008.
- Sean Crockett and Ryan Oprea. In the long run we all trade: Reference dependence in dynamic economies. *University of California Santa Barbara Working Paper*, 2012.
- Sean Crockett, Stephen Spear, and Shyam Sunder. Learning competitive equilibrium. *Journal of Mathematical Economics*, 44(7-8):651–671, 2008.
- 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.
- D Farmer. Avoiding getting lost in the wilderness of bounded rationality: the path from zero intelligence to no arbitrage. *QUANTITATIVE FINANCE*, 3(4):C64–C65, 2003.
- Nick Feltovich. Reinforcement-based vs. belief-based learning models in experimental asymmetric-information games. *Econometrica*, 68(3):605–641, 2000.
- Daniel Friedman. A simple testable model of double auction markets. *Journal of Economic Behavior & Organization*, 15(1):47–70, 1991.

- Daniel Friedman, Brett Williams, and Vivian Juehui Zheng. Market visualizations. *Available at SSRN 5480969*, 2025.
- Drew Fudenberg and David K Levine. Consistency and cautious fictitious play. *Journal of Economic Dynamics and Control*, 19(5-7):1065–1089, 1995.
- Steven Gjerstad. Price dynamics in an exchange economy. *Economic Theory*, 52(2):461–500, 2013.
- Steven Gjerstad and John Dickhaut. Price formation in double auctions. *Games and economic behavior*, 22(1):1–29, 1998.
- Dhananjay Dan K Gode, Shyam Sunder, and Stephen Spear. Convergence of double auctions to pareto optimal allocations in the edgeworth box. 2004.
- 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.
- 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.
- 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.
- Brett Williams. Opening the book: Price information’s impact on market efficiency in the lab. 2025a.
- Brett Williams. Zero intelligence in an edgeworth box. *Journal of Behavioral and Experimental Finance*, page 101098, 2025b.
- Robert B Wilson. On equilibria of bid-ask markets. In *Arrow and the ascent of modern economic theory*, pages 375–414. Springer, 1987.

Appendix A Robustness Checks

A.1 History Reset

	History Reset		
	Mean	St. Dev.	Range
<i>I. Prices</i>			
Price	2.48	0.16	(2.05, 3.04)
Per-Unit Avg.	2.37	0.16	(1.97, 2.78)
$ Price - CE $	0.50	0.15	(0.23, 1.12)
RMSE	0.76	0.32	(0.31, 1.93)
Final 5 Prices	2.46	0.16	(2.07, 3.64)
<i>II. Allocations</i>			
Final Distance	0.63	0.33	(0.07, 1.80)
Seller MRS	2.44	0.13	(1.98, 3.00)
BuyerMRS	2.47	0.13	(2.13, 3.02)
<i>III. Efficiencies</i>			
Allocative	0.98	0.03	(0.74, 1.00)
- Penalized	0.82	0.09	(0.37, 0.97)
Distance	0.81	0.04	(0.61, 0.93)
- Penalized	0.67	0.08	(0.27, 0.84)
Profit	0.90	0.07	(0.54, 1.09)
Observations	240	240	240

Table A.1: Simulation outcomes for markets with the trader history reset at the beginning of every period.

Table A.1 shows simulation results for markets with trader’s history entirely reset at the start of each period. 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 above of the main simulations, though with tighter ranges. Measures of final distance

and penalized allocative efficiency are just slightly improved in the markets with resetting memories; buyer and seller MRS actually shows a much tighter spread. Distance and profit efficiencies are essentially unchanged.

A.2 Other Memory Lengths

	L=0			L=10		
	Mean	St. Dev.	Range	Mean	St. Dev.	Range
<i>I. Prices</i>						
Prices	2.39	0.15	(2.01, 2.86)	2.57	0.17	(2.16, 3.12)
Per-Unit Avg.	2.31	0.16	(1.89, 2.76)	2.40	0.26	(1.76, 3.18)
$ Price - CE $	0.45	0.12	(0.22, 0.96)	0.53	0.14	(0.24, 1.14)
RMSE	0.64	0.24	(0.27, 1.73)	0.93	0.30	(0.35, 1.94)
Final 5 Prices	2.45	0.15	(2.11, 3.11)	2.45	0.17	(1.93, 3.68)
<i>II. Allocations</i>						
Final Distance	0.66	0.35	(0.03, 2.10)	0.80	0.51	(0.00, 2.39)
Seller MRS	2.44	0.14	(2.04, 3.12)	2.43	0.14	(1.83, 2.72)
Buyer MRS	2.48	0.13	(2.09, 2.93)	2.50	0.16	(2.19, 3.19)
<i>III. Efficiencies</i>						
Allocative	0.97	0.03	(0.78, 1.00)	0.99	0.01	(0.96, 1.00)
-Penalized	0.82	0.09	(0.43, 0.97)	0.77	0.10	(0.51, 0.97)
Distance	0.81	0.05	(0.64, 0.92)	0.82	0.05	(0.69, 0.93)
-Penalized	0.67	0.09	(0.36, 0.85)	0.68	0.08	(0.43, 0.87)
Profit	0.90	0.07	(0.57, 1.06)	0.92	0.08	(0.63, 1.14)
Observations	240	240	240	240	240	240

Table A.2: Simulation outcomes for markets with memory lengths of 0 and 10 [20 runs each].

Table A.2 reports simulation results for markets with traders holding memories of L=0 and L=10. Given the tightness of the reservation utilities driven by the η in the main paper, results appear relatively robust for lower values of L (0, 5 or 10 here). L=10 reports mildly better estimates in prices (aside from RMSE), while L=0 narrowly nudges ahead in allocation outcomes. Efficiencies display a mild separation in average estimates in favor of L=10 markets and supports are tighter on nearly all measures. Relative to the L=5 estimates reported in Section 5.3, L=10 performs nearly identically, though with a negligible lead in

MRS estimates and slightly tighter supports across most estimates. $L=0$ seems to mildly outperform the $L=5$ allocation estimates, while systematically lower in prices.

A.3 No Accept Rule

	$\lambda = 0$			$\lambda = 5$		
	Mean	St. Dev.	Range	Mean	St. Dev.	Range
<i>I. Prices</i>						
Price	2.42	0.43	(1.44, 3.62)	2.43	0.20	(1.83, 3.09)
Per-Unit Avg.	2.32	0.42	(1.39, 3.47)	2.35	0.27	(1.65, 3.07)
$ Price - CE $	0.70	0.22	(0.22, 1.37)	0.45	0.17	(0.10, 1.23)
RMSE	0.87	0.27	(0.26, 1.65)	0.67	0.23	(0.12, 1.60)
Final 5 Prices	2.26	0.36	(1.45, 3.70)	2.43	0.12	(2.06, 2.87)
<i>II. Allocations</i>						
Final Distance	1.22	0.73	(0.05, 3.71)	0.81	0.52	(0.05, 3.47)
Seller MRS	2.25	0.28	(1.56, 2.95)	2.39	0.17	(1.63, 2.70)
Buyer MRS	2.70	0.33	(1.98, 3.76)	2.54	0.19	(2.27, 3.66)
<i>III. Efficiencies</i>						
Allocative	0.91	0.07	(0.59, 1.00)	0.99	0.01	(0.96, 1.00)
-Penalized	0.74	0.12	(0.33, 0.96)	0.82	0.09	(0.51, 0.99)
Distance	0.71	0.09	(0.36, 0.88)	0.85	0.05	(0.68, 0.94)
-Penalized	0.51	0.13	(0.14, 0.78)	0.72	0.08	(0.38, 0.88)
Profit	0.85	0.09	(0.51, 1.07)	0.95	0.08	(0.76, 1.15)
Observations	240	240	240	240	240	240

Table A.3: Simulation outcomes for markets with the accept rule removed from trader's decision process. Estimates are shown for markets with logit parameter values of 0 and 5.

Table A.3 reports simulation results for markets with traders who strictly place limit orders. This is more in-line with the majority of the trader behavior and market theoretical literature as limit orders are generally the only means of market participation in these simpler models. I test this simplification of the traders' order placement process in markets with logit parameters of 0 and 5.

First, in both the right and left panel, price-related estimates are systematically lower than the markets with the an accept rule. This is supported by (1) the lower expected CE price of 2.24 if traders only trade on their natural side while being guided by their $u_{i,-\Delta}$, and (2) the lack of an accept rule means more aggressive prices (those posted closer to the lower reservation IC) are not immediately taken. A similar explanation reconciles the improvement

in profit efficiency in the $\lambda = 5$ panel.

The $\lambda = 0$ results are considerably lower than those of the $\lambda = 5$ markets. The largest difference appears in the final cluster of prices, with $\lambda = 0$ markets clearly tapering towards the $u_{i,-\Delta}$ CE price prediction, while $\lambda = 5$ markets make an effort to stay around the true CE prediction. Efficiency estimates are considerably lower regardless of measure for the $\lambda = 0$ sessions compared to both the right panel and the main results.

A.4 No Internal Spread Reduction Rule

	No Internal SR		
	Mean	St. Dev.	Range
<i>I. Prices</i>			
Price	2.56	0.28	(2.01, 4.07)
Per-Unit Avg.	2.35	0.29	(1.70, 3.67)
$ Price - CE $	0.69	0.25	(0.28, 1.94)
RMSE	1.13	0.41	(0.43, 2.36)
Final 5 Prices	2.50	0.25	(2.05, 4.37)
<i>II. Allocations</i>			
Final Distance	0.87	0.49	(0.03, 2.41)
Seller MRS	2.48	0.13	(1.89, 2.78)
BuyerMRS	2.47	0.13	(2.14, 3.20)
<i>III. Efficiencies</i>			
Allocative	0.99	0.02	(0.68, 1.00)
- Penalized	0.76	0.11	(0.47, 0.96)
Distance	0.81	0.06	(0.49, 0.94)
- Penalized	0.67	0.09	(0.23, 0.85)
Profit	0.94	0.07	(0.61, 1.11)
Observations	240	240	240

Table A.4: Simulation outcomes for markets with no internal spread reduction rule imposed.

Table A.4 tests whether imposing an internal spread reduction rule impacts the performance of the model. An internal spread reduction rule restricts the trader from posting an order with a worse price than one she has already posted. This is different from the standard

spread reduction rule, in which traders can only post orders that tighten the best bid-ask spread.

While price estimates appear to be negligibly worse than the main results, efficiency estimates are essentially equivalent, and MRS estimates converge (though slightly above CE). As the MRS estimates are measured for the representative agents, the improvement in convergence while maintaining the same efficiency levels points to slightly more trading across trader type (e.g., a trade between two natural sellers leads to no change in the representative sellers allocation).

Appendix B Experiment User Interfaces




Figure B.1: Williams (2025) user interface. Left side shows an open orderbook, reporting bid and ask orders, as well as completed traders. Right side displays a heatmap representing the preferences being induced. Also the No Frontier interface from Friedman et al. (2025).




Figure B.2: Friedman et al. (2025) user interface. Left side shows an open orderbook, reporting bid and ask orders, as well as completed traders. Right side displays a heatmap representing the preferences being induced.

Appendix C Fuzzy Categorizations (One Error Allowed)




Figure C.1: Type Categorizations at the Trader-Period level. Plots show type breakdown, within data that is covered by the model, with one action error allowed per trader-period. UD stands for "Undefined", i.e. data points not capture by the model.