Learning Paretian Equilibrium

Jacob K. Goeree and Brett Williams[1]

April 11, 2025

Abstract

We introduce a parameter-free model of trader behavior in multiple interlinked markets. Transaction prices inform traders about the likelihood certain prices are acceptable. Traders submit limit orders accordingly or accept others' limit orders. We demonstrate this simple learning model reproduces data patterns from prior general equilibrium experiments, whether observed prices converged, kept cycling, or led to odd outcomes. In each simulated replication, LEAPE predicts the more-general Paretian equilibrium outcomes recently proposed by Goeree and Williams (2024).

Keywords: *General equilibrium, Paretian equilibrium, learning, zero intelligence*

[1]AGORA Center for Market Design, UNSW. We gratefully acknowledge funding from the Australian Research Council (DP190103888 and DP220102893).

1. Introduction

Nearly a century and a half ago, Leon Walras 1874 put forth his theory on competitive equilibrium. A century of theoretical advancement followed (e.g. Pareto (1906)), championing general equilibrium as a pillar of mathematical economics. Theories of equilibrium and optimum converged to form the Fundamental Welfare Theorems, while others diverged from Walrasian equilibrium as disequilibrium trade was ignored in Walras' central theory. This theoretical pursuit continued into the mid 1900's, where a shift towards empirical and experimental works took place. Over the last 50 years, general equilibrium has declined in favor, to the point where Crockett (2013) summarizes the current state of the theory as "General equilibrium theory has fallen out of fashion".

However, given a growing market landscape paired with rising complexity in goods, market rules and preferences, general equilibrium is positioned to (perhaps slowly) move it's way back to center stage. Partial equilibrium exploration assumes independence across markets with respect to prices and excess demand, yet disequilibrium trade and real-world price dispersion refute this. As such, research in markets with more complex preferences, including those which are nonconvex, has remained in general equilibrium (Starr, 1969; Milgrom and Watt, 2022). Advancements in laboratory market engines and user interfaces have allowed for recent experimental works to explore more complex preferences in general equilibrium as well (Gjerstad, 2013; Crockett et al., 2020; Williams, 2023a; Friedman et al., 2024).

Parallel to the decline of GE was the rise of agent-based modelling, a simple computational philosophy concerned with observing the relationship between the behavioral rules or characteristics of agents and phenomena that occur in social or group settings.[1] Perhaps a less known or renowned medium for economic research, agent-based models have flourished in computational fields, and represent a way to help guide theory and experiments in these more complex general equilibrium spaces. In fact, the 90's already displayed the impact ABM can have on economic theory, as several ABMs appeared for partial equilibrium accounts of trade in the double auction (Wilson, 1987; Friedman, 1991; Easley and Ledyard, 1993; Gode and Sunder, 1993; Gjerstad and Dickhaut, 1998).

[1] See Axtell and Farmer (2022) for a new survey on the ABM literature in economics and finance.

To provide a simple medium for understanding complex markets, we present LEAPE, a minimal intelligence agent-based model of trader behavior. LEAPE equips traders with simple price adjustment rules based the market's trade history, placing more emphasis and intelligence on the quantity selection portion of order placement. In it's simplest form, LEAPE is nearly parameter-less, with only the distributional shape of the price distribution and the bounds on those prices parametrizing the model. Otherwise, trader's enter the market one at a time, view the market, decide if anything is acceptable (and act if so), draw a price (if not), and find the quantity at that price that yields the highest utility reallocation. Importantly, the decision (and ability) to accept orders is crucial to (1) incorporate the idea of utility-chasing individuals trading while gains from trade can be made, and (2) deal with economies where traders especially need trades filled entirely at once. The sole behavioral assumption in this base version of the model is that trader's are IR, though this follows naturally given they are perfect-utility-maximizing. Several adjustments to the model are made to consider interesting behavioral phenomenon, such as imperfect choice or across-period learning, or to accommodate more complex market settings, such as multiple markets (i.e. more that two goods).

We draw theoretical inspiration in this paper from our equilibrium contribution: *Paretian equilibrium* (Goeree and Williams, 2024). Paretian equilibrium recognizes that disequilibrium trade (relative to a Walrasian paradigm) exists in modern markets, as does price dispersion. As such, a Paretian equilibrium is defined as a Pareto optimal final allocation paired with a sequence of prices that are consistent with individual rationality (IR). In the base version of LEAPE, traders act as perfect utility maximizers, which paired with a minimal intelligence price-learning process produces Paretian equilibrium outcomes. Such is true even under more strenuous conditions, including multiple markets, imperfect quantity choice, or even nonconvex preferences.

A handful of prominent experimental and theoretical results provide a stress test for the model. We replicate the price and allocation dynamics of CES economies with income effects (Gjerstad, 2013), non-Walrasian economies with extreme price divergence (Gale et al., 1963; Crockett et al., 2011), and even multiple markets with price cycles (Scarf, 1960; Anderson et al., 2004; Goeree and Lindsay, 2016). In each case, LEAPE agents are able to provide mirror the final price and allocation outcomes as well as the dynamics of human subjects. Even the final stress test, two pure exchange economies with either one or two trader types with $max(\cdot,\cdot)$ preferences, does not hamper the LEAPE traders' ability to reap gains from trade and reach Pareto optimal outcomes, despite trading in a space no typically covered by standard general equilibrium theory.

The paper unfolds across six sections. Section 2 recounts prior literature on general equilibrium theory and experiments, agent based models, and their intersection. Section 3 follows with the model and its various adjustments. We stress test the model in Section 4, followed by a discussion in Section 5 on intricacies and design choices both of the model and moving forward in practice. We conclude the paper in Section 6.

2. Prior Experiments and Agent-Based Models

General equilibrium theory has grown for over a century since the likes of Walras, Marshall and Pareto studied the intricacies of existence, stability and (less so) dynamics of equilibria in markets. Experimental work followed decades later with the work of Chamberlin (1948) and Smith (1962), finding convergent tendencies in laboratory markets. However, these markets were of a much simpler, highly stylized framework, i.e. partial equilibrium, than those of the founding theory. Simplifying away from utility theory and price dependence across markets allowed experiments to focus on pricing relative to given demand and supply schedules, often over one unit of some indivisible good. Hundreds of laboratory markets have been run in such a way, with the consensus generally being mild to strong convergence to the economy's Walrasian equilibrium. The general consensus in these simple markets was an agreement that tatonnement, the dynamic process led by some Walrasian auctioneer, could support the majority of the analysis in partial equilibrium.

Following major improvements in computing, laboratory markets slowly gained the support needed to carry out more complex markets. Naturally, this offered a prime opportunity to move away from partial equilibrium and towards general equilibrium; however, the transition has been slow. Interestingly, the majority of GE market experiments over the last 20 years have been tests of non-Walrasian example-economies from 1950's and '60's; namely, the price divergence example of Gale et al.

(1963) and the price cycling example of Scarf (1960). Anderson et al. (2004) conducted the first such experiment, producing the across-period price cycles of Scarf's three-good-three-trader economy. Goeree and Lindsay (2016) replicated Anderson et al.'s findings, while also proposing a schedule-based order process, which eliminated the cycling and found convergence to the unstable equilibrium price pair. The most recent test of Scarf's economy, Gillen et al. (2021), again replicate clear price cycling in the multiple-unit double auction (MUDA), and claim partial equilibrium analysis is likely sufficient in characterizing GE price dynamics. Book-ended by the Scarf tests was the laboratory recreation of Gale's (1963) economy by Crockett et al. (2011). The study again reproduced the price divergent tendencies of Gale's setting, though only after enforcing a price floor (to gain the high corner solution) or price ceiling (low corner solution).

A few other experimental studies on general equilibrium laboratory markets emerged in the same time period with focuses outside of non-Walrasian dynamics. Williams et al. (2000) was the first test of GE CDA markets in the lab, testing behavior of CES-induced buyers across two markets. Crockett (2008) studied reallocations in a market consisting of real subjects and simulated agents to study adjustments in final allocations along the contract curve across trading periods. While reaching the contract curve was quite common, and learning was apparent across periods in terms of relative equality in final utility gains, not much convergence was found in across-period final allocation adjustment. Gjerstad (2013), like the prior two works, studied traders with CES preferences, however a more extreme curvature was considered, inducing stronger income effects.[2] The results focused much more on prices, finding a lack of convergence to a Walrasian competitve equilibrium price as the income effects accommodated price stickiness away through disequilibrium trade. More recently, Crockett et al. (2020) studied CDA laboratory markets with subjects induced by their (elicited) naturally occurring preferences.

A closely related literature to experiments on markets is that on agent-based modelling. Since the late 20th century, agent-based models have been used to simulate markets under specific behavioral rules followed by the traders. Naturally, such a practice pairs well with laboratory experiments as the ABM simulations can be used

[2] An inner exponent of -1 was used in both trader types' utility functions, which marks the existence of stronger income effects and the point at which multiple equilibria can begin to exist.

to replicate, substitute, or even predict human behavior in the lab. Specific to the CDA institution, the sub-literature began with a trio of papers, namely, Wilson (1987), Friedman (1991) and Easley and Ledyard (1993). The first studied a game-theoretic approach to double auction trade, the second proposed a model of within-period price adjustment and reallocation through the use of reservation prices, and the third took an across-period approach to price adjustment. Perhaps the most well-read of the sub-literature, Gode and Sunder (1993) introduced 'zero intelligence' (ZI) traders to the landscape. These traders bear no behavioral intricacies, other than the abidance of individual rationality (IR).[3] Several adaptions of the ZI model have been put forth, with Cliff and Bruten (1997) suggesting traders slightly increase their intelligence to prioritize profit margins, and Gode et al. (2004) and Crockett et al. (2008) bringing ZI agents to a general equilibrium exchange economy setting.[4]

More complex minimal-intelligence models also appear in the literature, with the first being Gjerstad and Dickhaut (1998). In their partial equilibrium model, agents maintain beliefs over the acceptability of prices based on the prior trade and order-book history, and have preferences over their next order-placement time dependent on the maximum expected surplus to be potentially received upon order placement and acceptance. Arifovic and Ledyard (2011) take a more evolutionary approach, imbuing agents with experimentation, replication and selection stages of their strategy choice and learning process. Anufriev et al. (2013) explicity applies this more general learning model and testbed to the continuous double auction. More recently, Williams (2023b) considers a logit-choice process where agents hold reservation utilities (akin to reservation prices) over potential bid and ask orders, and may consider utility-losing orders as a means to manipulate their future allocation bundles and marginal rate of substitution for better future trade.

[3]The original framing in Gode and Sunder (1993) was the traders follow a 'budget constraint; however, both Cliff and Bruten (1997) and Gjerstad and Shachat (2021) refute the lack of intelligence and the latter explicitly shows the constraints equivalence to IR.

[4]Williams (2023c) tests, via an expansive simulation exercise, the impacts each behavioral and market rule from Gode et al. (2004) (and to some extent Gode and Sunder (1993) and Crockett et al. (2008)) has on ZI agents and market performance in this general equilibrium context.

3. The Model

We present a simple model of learning and order selection in the continuous double auction. A set $\mathcal{N}$ of traders trade goods X and Y in an Edgeworth box. Traders have preferences, $u_i(\cdot)$, over these goods which may or may not follow classical assumptions in general equilibrium, such as convexity. Additionally, traders (1) hold beliefs over prices via distribution-updating learning process, and (2) choose order sizes according to their offer curves. Traders enter one at a time, uniformly randomly drawn from the set of all traders.

Prices. Order selection follows a two step process. First, after a trader has decided she wishes to enter a new order in the orderbook, she must decide on a price. We formulate this decision as being a random draw from a truncated normal distribution, with mean μ and standard deviation θ. The distribution is truncated below at 0 and above at some maximum allowable price (as set by the market designer) M. Here μ and θ follow a simple updating process, informed by the per-unit average price of the trades that have already happened in the market, $\tilde{p}$.[5] Equations 1 and 2 detail said process:

$$\mu_t = \frac{\frac{M}{2} + t \cdot \tilde{p}}{1+t} \tag{1}$$

$$\theta_t = \frac{\frac{M^2}{12} + t \cdot \theta(\tilde{p})}{1+t} \tag{2}$$

$\theta(\tilde{p})$ represents the standard deviation of the prices of all previous trades. The draw of trader i in entry time r, $p_{i,r}$, determines the price of the trader's order. Graphically, this creates a line through the trader's current endowment (x_i, y_i) with slope $p_{i,r}$, along which anticipated utilities can be compared. The number of orders, t, provides the weight for the trade history's statistic in the weighted averages defining μ_t and θ_t.

Quantities. The market designer in this setting allows for multiple unit orders, where units are divisible. Given this feature, as well as the trader's desire and ability to both buy and sell each good, the second step of the order process is selecting a

[5] Per-unit average price is defined as $\tilde{p} = \frac{\langle p,q \rangle}{\sum q}$

quantity that may either be negative (sell) or positive (buy). The lower and upper bounds on this quantity support, $[\underline{s}, \overline{b}]$, can be defined as:

$$\underline{s} \equiv (-1) \cdot min\{\frac{\overline{Y} - y_i}{p_{i,r}}, x_i\} \tag{3}$$

$$\overline{b} \equiv min\{\frac{y_i}{p_{i,r}}, \overline{X} - x_i\} \tag{4}$$

$\overline{X}$ and $\overline{Y}$ represent the total amount of X and Y in the market. The trader then selects the point along the order's price-line, bounded by $[\underline{s}, \overline{b}]$, which maximizes their expected utility. This choice both determines the side of the market the trader adds an order to, as well as the quantity field in her order, $q_{i,r}$.

3.1. Imperfect Order Choice

A natural adjustment to the non-parametric model described above is to allow for error in the order selection process, or more specifically the quantity selection process. As such, we replace deterministic utility maximization with a logit-choice process, thus adding one parameter to the model.

Given her price draw $p_{i,r}$, the trader considers a set of points along her order's price-line, associated with a lattice over quantities of X. For each reallocation a in the set of reallocations, A, along the drawn price vector, the trader can determine her expected utility of moving along the price-line to the potential allocation. She then associates each allocation/quantity, $q^a_{i,r}$, with a probability of being chosen:

$$Pr(q^{a^*}_{i,r}) \equiv \frac{exp[\lambda u_i(x_i + q^{a^*}_{i,r},\ y_i - p_{i,r} \cdot q^{a^*}_{i,r})]}{\sum_{a \in A} exp[\lambda u_i(x_i + q^{a}_{i,r},\ y_i - p_{i,r} \cdot q^{a}_{i,r})]} \tag{5}$$

Varying the added parameter λ allows the trader choice precision to range from something akin to a variation of zero-intelligence (Gode and Sunder, 1993; Gode et al., 2004; Williams, 2023c) to being perfect utility maximizers. The base version of the model assumes $\lambda = \infty$.

3.2. Acceptance through Market Orders

As posited in Goeree et al. (2024), the inclusion of market orders by the designer can positively enforce trade when gains are available but an equilibrium in the traditional sense may not exist. To such an end, we allow for the placement of market orders.

Upon entry, if a trader notices an order in the orderbook which would improve her utility upon acceptance, she will ignore her regular order placement process. Instead, she will place an order on the contra-side of the market matching the price of the desired order (to guarantee a cross occurs). The quantity of the market order placed is determined in the same way as placing a limit order, with the quantity yielding the highest utility-gain being selected. In the case of imperfect order selection, the quantity is selected as described in Section 3.1. In both cases, a cap is placed on the possible quantity choice, placed at the order size of the order being crossed.

3.3. Across-Period Learning

The base version of our model takes into consideration within-period price learning, however ignores the possibility of this learning continuing across trading periods. A sleek way to update the model, free from adding another parameter, is to replace the starting μ for the traders' price distribution in all periods after the first. More specifically, the starting mean of $M/2$ in μ_t is replaced by the final trade price of the previous period.

3.4. K-dimensional Edgeworth Box

Perhaps the most ambitious adjustment to the model, we maintain that the order selection and learning process can be tractably extended to an n-dimensional Edgeworth box. Two assumptions over the structure of the market, common among the previous waves of general equilibrium theory, are made: (1) the k^{th} good is a numeraire to be used as a cash of sorts (cite the Shapley 61 paper on this), and (2) all trades must include the numeraire, essentially separating the marketplace into $k-1$ markets each with their own orderbook. To maintain simplicity, a trader will only place one order at a time (i.e. one order per entry).

The price and quantity processes are adjusted slightly. For each of the $k-1$ goods, a separate truncated normal distribution over prices is known to all traders. Thus, there are $k-1$ processes $\mu_{\kappa,t}$, and associated $\theta_{\kappa,t}$ processes. For each of the $k-1$ goods, the trader draws a potential price $p_{i,r,\kappa}$. Each associated price-line is given their own support $[\underline{s}_{\kappa}, \overline{b}_{\kappa}]$. The definition of these lower and upper bounds becomes slightly more complex as the constraints arise over the amount of available numeraire. Given a trader's orders currently posted in other orderbooks across the $k-1$ goods, the amount of numeraire she can offer in a buy order is capped weakly below by her current holdings.

$$\underline{s}_{k'} \equiv (-1) \cdot min \left\{ \frac{\overline{Y} - y_i - \sum\limits_{j \in K \setminus \{k'\}} p^s_{i,j} q^s_{i,j}}{p_{i,r,k'}},\ x_{i,k'} \right\} \tag{6}$$

$$\overline{b}_{k'} \equiv min \left\{ \frac{y_i - \sum\limits_{j \in K \setminus \{k'\}} p^b_{i,j} q^b_{i,j}}{p_{i,r,k'}},\ \overline{X} - x_{i,k'} \right\} \tag{7}$$

Notice that the expressions for her current holdings of x and y have adjusted compared to (3) and (4). The summations encapsulate the holdings tied up in orders placed in the other $k-2$ markets, where the superscripts on p and q denote whether they are the trader's buy (e.g. p^b) or sell (p^s) order and K is the set of good indexes.[6] The trader considers all allocations along the set of $k-1$ pricelines. If she is a perfect utility maximizer, she will precisely choose the optimal quantity (and thus the orderbook and side of the market she will enter). However, if she has the capacity to err in her decision-making, she will use a logit-choice process as in Section 3.1.

4. Replications

Given the simplicity of our model and its prioritization of reaping gains from trade (as opposed to abiding by a Walrasian pricing principle), we provide a series of replication stress tests. Each replication exercise analyzed below reprises the main findings of seminal experimental and theoretical general equilibrium works on continuous double auction trade. In each case, we show our model can (1) match the experimental

[6]Note that the order in the k' market is not included in the summation as the order about to be potentially placed by the trader will cancel and replaced her existing order in that market.

findings contingent on adjustments in λ and other model adjustments specified in section 3, and (2) support the theoretical equilibrium predictions when the simulated agents are equipped to perfectly follow their offer curves.

4.1. Test 1: Standard CES Preferences w Income Effects

Our first testbed is that of Gjerstad (2013). Agents are induced with constant elasticity of substitution (CES) preferences

$$u_i(x_i, y_i) = c_i((a_i x_i)^{r_i} + (b_i y_i)^{r_i})^{\frac{1}{r_i}} \tag{8}$$

with agents partitioned into natural buyers or sellers. All buyers have the same set of CES parameters and the same starting endowments; all sellers are the same as well. The r_i used for both agent types is set to -1, which generally implies strong income effects. In general, this parameter can range across $(-\infty, 0) \cup (0, 1]$ in order to satisfy the usual convexity assumption made in general equilibrium. The rest of the parameters are as follows: $(c_i, a_i, b_i) = (0.695,\ 109.89,\ 0.362)$ for buyers and $(0.256,\ 109.89,\ 2.982)$.[7] The starting endowments are (0, 1800) and (18, 0) for buyers and sellers, respectively. With respect to these starting conditions, the competitive equilibrium price for any market with equal number of each agent type is 91.[8]

Figure 1 reports time-average price trends for agents in Gjerstad's setting, where all traders are perfect maximizers (in the sense that given a price, an agent will chose the corresponding quantity on his offer curve). Price trends converge on average nicely to the CE prediction. Despite being the first test in the exercise, disequilibrium trade in the current setting is difficult to rebound from; thus, such strong convergence is promising. Final allocations, as shown in Figure 2, lie on or close to the contract curve. While dispersed somewhat widely along the contract curve, the geometric mean of the sample lies very close to the predicted equilibrium allocation. Two important notes can be made in this regard. First, this spread near or along the contract curve is reminiscent of the final allocations in Gjerstad (2013). Second, the model clearly follows our focus on Paretian as opposed to Walrasian dynamics, as

[7]For the sake of consistency across all simulations, we assume Y to be the numeraire. As such, in the Gjerstad setting, we swap the parameters impacted x and y in both the endowments and utility functions. This does not change the CE price.

[8]The equilibrium allocation for buyers is $(7, 1163)$. For sellers, the allocation is $(11, 637)$.

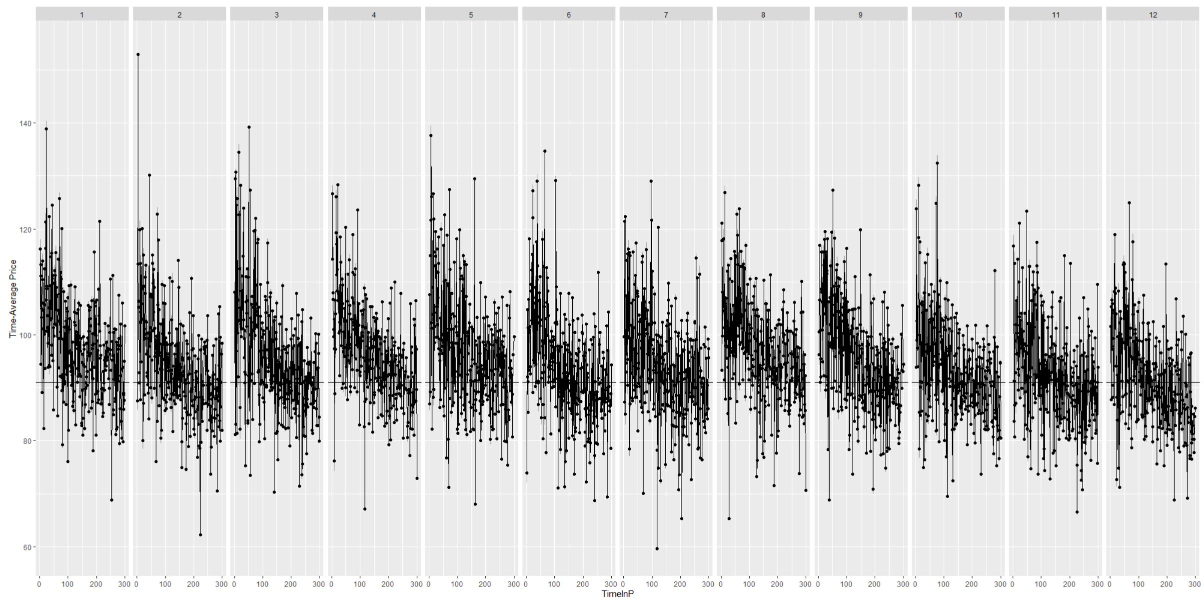

Figure 1: Time-average prices for CES markets as parametrized in Gjerstad (2013). The CE price is 91. No across-period price adjustment is assumed in this set of simulations, i.e. each period is a full reset of endowments as well as memory.

nearly all markets reach a Pareto optimal allocation in the aggregate.

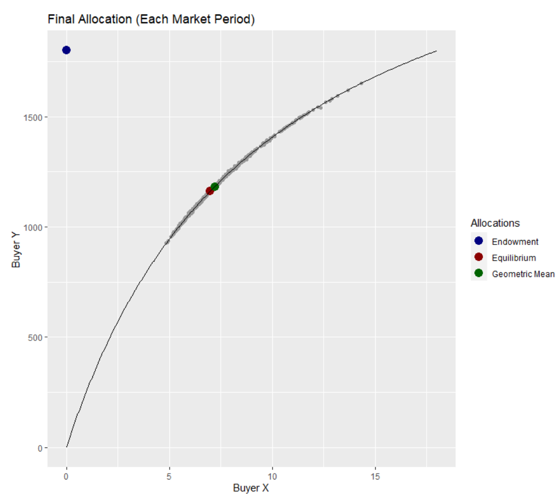

Figure 2: Final allocations of aggregated agents for each simulated market period (600 in total). The allocations are plotted against the contract curve.

4.2. Test 2: Gale's Example Economy

The second level of our replication exercise is the first real stress test the model encounters, with Gale's (1963) economy creates strong divergent price tendencies. Namely, Walrasian dynamics driven by excess demand lead disequilibrium prices to diverge rapidly from the unstable Walrasian equilibrium. The stylized economy induces agents with altered Leontief preferences of the form $min\{a_i x, y_i - b\}$. We follow the parametrization specific to the experimental work of Crockett et al. (2011), with twelve agents participating in a market, split evenly to naturally prefer either selling or buying X. All agents of each type are given the same (a_i, b_i) as well as the same start-of-market endowment. For natural buyers, $(a_i, b_i) = (658,\ 3947)$ and the starting endowment is (5, 5600); for sellers, $(a_i, b_i) = (35.5,\ -1349)$ and the starting endowment is (15, 400). Such parametrization yields a competitive equilibrium price (contingent on the starting endowment) of 158.

As Gale's original example predicts and the experimental findings of Crockett et al. (2011) support, this equilibrium is highly sensitive to disequilibrium trades, to the point of strong divergence. Crockett et al. (2011) use price floors and ceilings to induce such disequilibrium trade; we mimic this strategy, though through only the use of ceilings, by adjusting M to be either below the CE price or well above it. Specifically, we use M values of 120 and 1000 to induce divergence to prices well below or well above the equilibrium prediction.

Figure 3 provides the time-average price trends for each round within each M choice. Clearly, the divergent patterns confirmed in Crockett et al. (2011) are replicated by the LYQ agents. Given a price ceiling of 120, prices quickly plummet to a band oscillating around 20. A ceiling of 10000, however, matches the price climb elicited by a price floor of just above 200 in Crockett et al. (2011). Interestingly, a price ceiling in the range of 300-400 (here 350) can reign in the prices enough to keep them relatively stable near 158, as seen in the middle plot of Figure 3.

Figure 4 shows round-end allocations across all simulated rounds for aggregated representative agents and individual agents.[9] The bow-tie shape created in the Edge-

[9]These representative agents are the result of tracking type-average allocation adjustments. For example, the representative natural buyer moves $1/n$ of the adjustment made the natural buyer who made a trade. However, if this trade was with another natural buyer, the average adjustment for all natural buyers is null, and the representative buyer's allocation does not move.

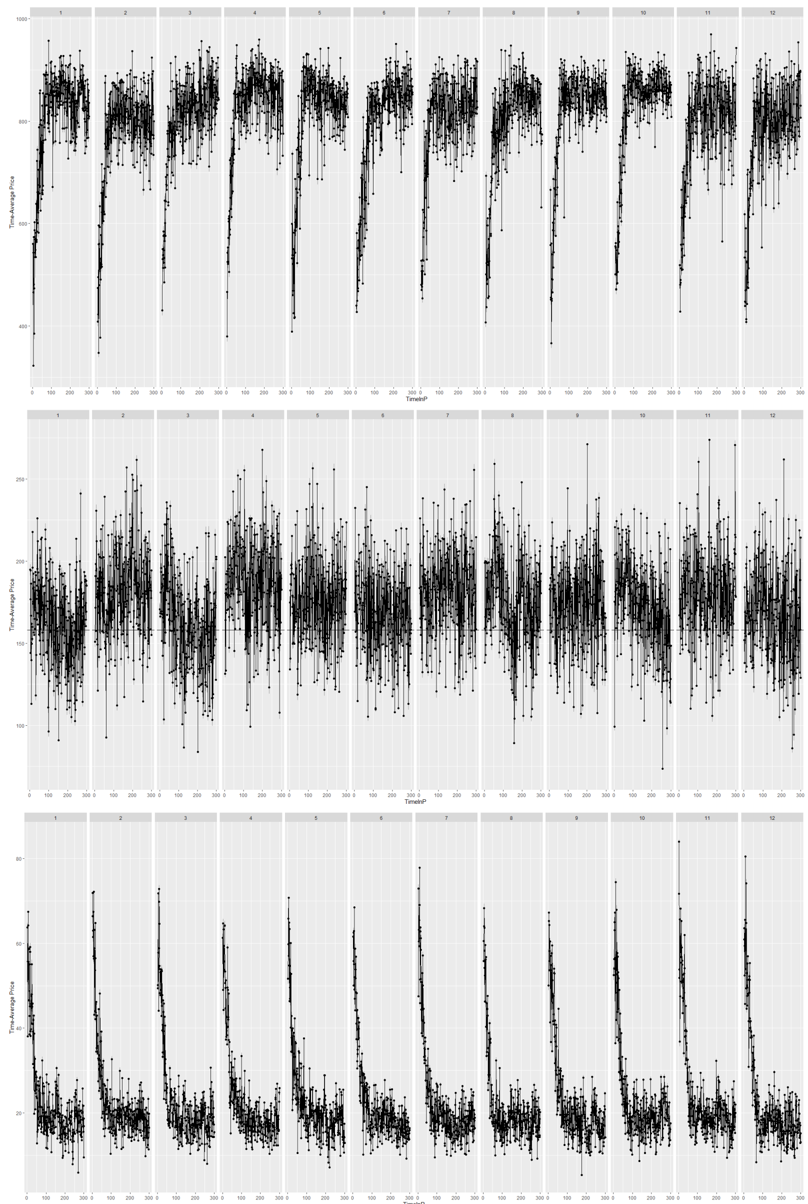

Figure 3: Time-average prices. The upper row is for M=1000 markets, the middle row is for M=350 markets, and the lower row is for M=120 markets.

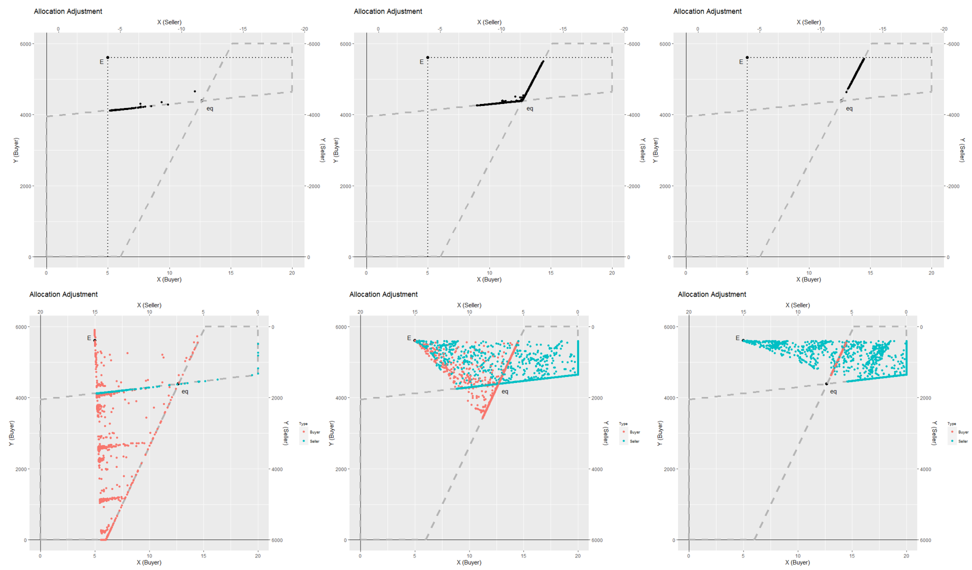

Figure 4: Aggregated Agent Final Allocations. M=10000 markets on the left, M=350 in the middle, and M=120 on the left.

worth box shows the expansion paths of the two trader types, which houses the Pareto optimal allocations for each trader type. Aggregated agent final allocations fall impressively along these paths, with the M=10000 and M=120 markets approaching the corner solutions and M=350 markets centering around the unstable Walrasian equilibrium final allocation. Individual traders create fans from the initial allocation to the traders' respective expansion paths, with the fans spanning between the $p = 0$ and $p = M$ vectors from the initial allocation. Even in relatively shortrun markets, a large portion of the traders successfully reach Pareto optimal outcomes.

4.3. Test 3: Scarf's Example Economy

Perhaps the most difficult stress test we address, Scarf's (1960) three good economy (two commodities and one numeraire) adds a new element to capture: cyclical price dynamics between the two commodity markets. The mechanism behind such price movement can be briefly summarized as excess demand pressure from one market to the other. In other words, when prices rise in one market, demand falls for that good and relative demand for the other good rises causing a delayed increase in that good's price, creating a cycle of across-market excess demand pressure. A few experimental studies have tested Scarf's model (Anderson et al. (2004); Goeree and Lindsay

(2016); Gillen et al. (2021)), all replicating the cyclical phenomenon as a portion of their analyses.

As with the first two stress tests, we use the same parameters as the experimental studies to induce our simulated agents. Each of the three trader types has Leontief preferences over two of the three goods, with initial holdings in one of their two preferred goods. Different permutations of the endowments below result in changes of direction of the expected price cycle.

Utility	Endowment (ω_A,ω_B,ω_C)
$u_1(x,y) = 40min\{q_A/10, q_B/20\}$	(10,0,0)
$u_2(x,y) = 40min\{q_B/20, q_C/400\}$	(0,20,0)
$u_3(x,y) = 40min\{q_C/400, q_A/10\}$	(0,0,400)

Table 1: Simulated Agent Parameters in Scarf Economy.

As suggested in Section 3.3, we adjust the above simulation model to allow for across period price trends. This is done simply by replacing the starting price distribution means with the final trade prices in the X and Z markets from the prior period. To reduce the early-round volatility in adjustments to θ_t, we impose a tighter support to the price draw distribution that shifts with the moving average price within the full support. Figure 5 displays cyclical price trends reminiscent of those found in the experimental data of Anderson et al. (2004) and Goeree and Lindsay (2016). With 1000 entries per round, not only are cycles clearly formed and persistent, but the center of the cycles stays relatively close to the equilibrium price pair of (40,20).

4.4. Test 4: Nonconvexities

Though ranging considerably in their equilibrium predictions and prescribed utility functions, all of the previous tests have an important feature in common: convexity. Namely, all of the above scenarios satisfy the (crucial) assumption in general equilibrium theory that the utility functions are convcave (i.e. the preferences are convex). We thus present a test of simple environments discarding this assumption.

We consider two economies, both with two goods and two trader types. The first contains two agents with $max(\cdot,\cdot)$ preferences, and the second replaces one of these agent types with one whose preferences are Leontief ($min(\cdot,\cdot)$). The former takes

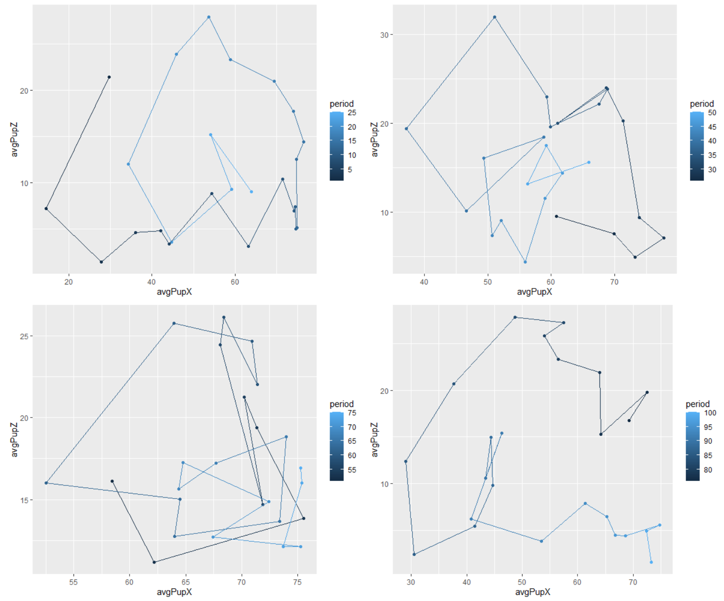

Figure 5: Round-average prices; 1000 entries per round. Starting price mean pair was set to (20,30).

our disregard of the convexity assumption at face value, with two simple yet very nonconvex agents participating in trade. The latter aims to place pressure on the exposure problem that arises from nonconvex preferences, i.e. that moves made by the nonconvex trader are “all or nothing”, and partial fills of their orders (driven by the competing interests of the Leontief trader) can lead to severe losses.

First, the two nonconvex agents scenario. The agents are induced with $max(2x, y)$ and $max(3x, y)$ preferences depending on type. We consider several initial endowments across a range of replica counts from one to four of each type. Figure 6 reports final allocations at the aggregate agent and individual trader levels for three different starting allocations. Depending on the starting allocation, the trader types clearly face varying pressure to split amongst themselves at their preferred prices. For endowments with $x_b > 2/3$, final allocations in the aggregate flock to the $y_b = 0$ and $x_b = 1$ corner, while the rest approach the $y_b = 1$ and $x_b = 0$ boundaries. Though similar in aggregate outcomes, the $(1/4, 2/3)$ and $(1/6, 1/6)$ economies yield large differences in sellers’ willingness to finish without holdings of X. In all three cases, the aggregate outcomes finish in weakly Pareto optimal allocations.[10]

[10]These also happen to lie on the predicted equilibrium allocations of Yquilibrium, as in Goeree et al.

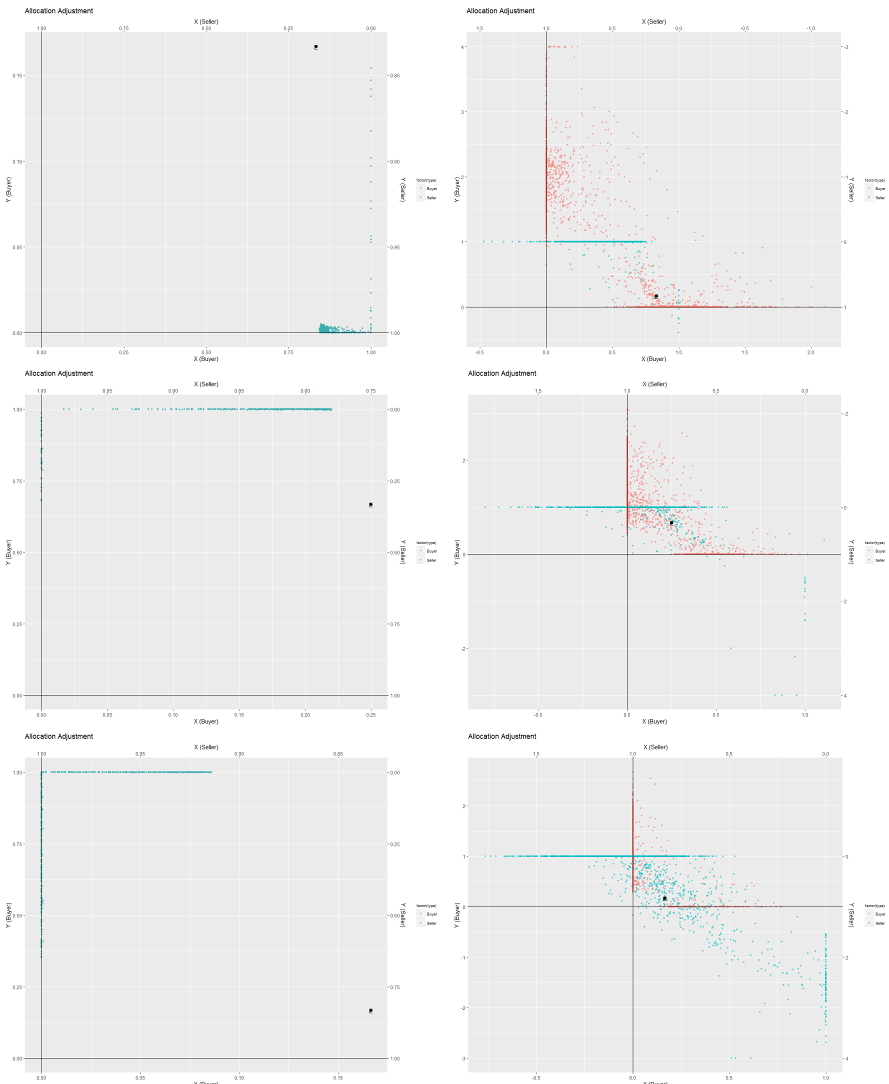

Figure 6: Allocations for max versus max economy (n=4). Left column: aggregate agent final allocations. Right column: individual agent final allocations. First Row: buyer endowment is (5/6,1/6). Second Row: buyer endowment is (1/4,2/3). Third Row: buyer endowment is (1/6,1/6).

Moving to the mixed economy with Leontief traders trading with max preference traders, more variety appears in the final allocation (Figure 7). Leontief traders clearly prioritize moving to line which equalizes their min(,) arguments, in this case $y = 3x$. When the starting endowment lies close to this line, Leontief traders are proficient at achieving this goal, while max traders once again split themselves between holding only x or y. However, with Leontief traders arriving at their destination so quickly, max traders have difficulty fully exhausting trade and moving to holdings of only a single good. When the endowment is relatively far from $y = 3x$, Leontief traders, instead, are less efficient in their reallocation, as max traders decide to all trade to their originally desired good rather than splitting.

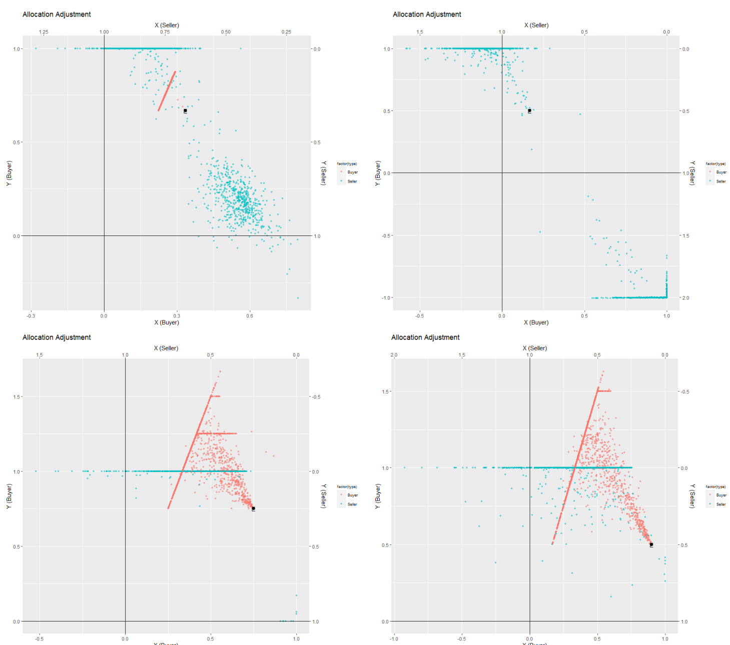

Figure 7: Individual final allocations for min versus max economy (n=4). Top Left: buyer endowment is (1/3,2/3). Top Right: buyer endowment is (1/6,1/2). Bottom Left: buyer endowment is (3/4,3/4). Bottom Right: buyer endowment is (9/10,1/2).

(2024).

5. Discussion

5.1. Design Choices

Looking forward to future work on markets where the existence of partial orders may be harmful to some, or all, of the traders, i.e. those with traders following non-convex preferences, a design choice regarding market orders is interesting to consider. In the main body of this paper, we impart a 'best-response' style of accept rule, where traders look at the current state of the book, check if any reallocations along the best-priced order's price vectors are utility improving and place such an order if so. This BR-style rule matches naturally with the utility-maximizing sentiment of Paretian equilibrium, and the quantity dynamic discussed in Goeree and Williams (2024). However, in economies where traders place importance on fully filling orders (or, avoiding partially-filled orders), a BR-style accept rule may prove harmful to some traders. An alternative accept rule then for LEAPE traders with such preferences (i.e. nonconvex), is an all-or-as-much-as rule.

In a convex world, such a rule change does not change the distribution over final allocations, only the time it takes to get to them. However, in the nonconvex world, such a rule safeguards those with orders already placed in the book, while still ensuring the trader accepting the order is also better off (as they would only accept the order if it were utility improving). In a follow-up project, Goeree et al. (2024) pair this adjusted accept rule with the concept of Yquilibrium (Goeree et al., 2024) to address design questions in combinatorial auctions with nonconvexities.

5.2. Nonconvexity

A key advantage of LEAPE is the non-reliance on a convexity assumption. Agents simply consider their utilities at different allocations and choose (sometimes noisily) their highest option. Given this freedom, LEAPE agents are able to participate in markets when induced by non-convex preferences, yielding near-equilibrium price and allocation predictions in a setting often ignored or considered unexplorable. In future work, we intend to explore such freedom, using LEAPE agents paired with the notion of Paretian equilibrium to provide theoretical and experimental insights even past those explored in this paper.

6. Conclusions

We provide a novel agent-based model for traders in general equilibrium, by which laboratory market data is consistently replicated and outcomes in typically difficult markets can be predicted. Final allocations determined by the base version of the model are characterized as Paretian equilibrium, as agents reach the contract curve guided by an IR sequence of trade prices. Non-Walrasian price trends proposed by prominent theoretical counters and found in laboratory markets are replicated with traders in the LEAPE paradigm. Given the applicability to more general market settings, and it's ability to replicate (and likely predict) market outcomes, we intend to follow this paper with an experimental exploration into nonconvex markets.

References

Anderson, C. M., C. R. Plott, K.-I. Shimomura, and S. Granat (2004). Global instability in experimental general equilibrium: the scarf example. *Journal of Economic Theory 2*(115), 209–249.

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

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

Axtell, R. and J. Farmer (2022). Agent based modeling in economics and finance: past, present, and future. *Journal of Economic Literature*.

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

Cliff, D. and J. Bruten (1997). Zero is not enough: On the lower limit of agent intelligence for continuous double auction markets.

Crockett, S. (2008). Learning competitive equilibrium in laboratory exchange economies. *Economic Theory 34*(1), 157–180.

Crockett, S. (2013). Price dynamics in general equilibrium experiments. *Journal of Economic Surveys 27*(3), 421–438.

Crockett, S., D. Friedman, and R. Oprea (2020). Naturally occurring preferences and general equilibrium: A laboratory study. *International Economic Review*.

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

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

Easley, D. and J. O. Ledyard (1993). 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*.

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

Friedman, D., V. Zheng, and B. Williams (2024). Visual Markets. *UCSC working paper In preparation*.

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

Gillen, B. J., M. Hirota, M. Hsu, C. R. Plott, and B. W. Rogers (2021). Divergence and convergence in scarf cycle environments: experiments and predictability in the dynamics of general equilibrium systems. *Economic Theory 71*, 1033–1084.

Gjerstad, S. (2013). Price dynamics in an exchange economy. *Economic Theory 52*, 461–500.

Gjerstad, S. and J. Dickhaut (1998). Price formation in double auctions. *Games and economic behavior 22*(1), 1–29.

Gjerstad, S. and J. Shachat (2021). Individual rationality and market efficiency. *Nonlinear Dynamics, Psychology & Life Sciences 25*(4).

Gode, D. D. K., S. Sunder, and S. Spear (2004). Convergence of double auctions to pareto optimal allocations in the edgeworth box.

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

Goeree, J. K. and L. Lindsay (2016). Market design and the stability of general equilibrium. *Journal of Economic Theory 165*, 37–68.

Goeree, J. K., L. Lindsay, and B. Williams (2024). YCE: A Yquilibrium Combinatorial Exchange. *AGORA working paper, in preparation*.

Goeree, J. K., J. Tayawa, and B. Williams (2024). Duality and Equilibrium. *AGORA working paper*.

Goeree, J. K. and B. Williams (2024). Paretian Equilibrium. *AGORA working paper*.

Milgrom, P. and M. Watt (2022). Walrasian Mechanisms for Non-convex Economies and the Bound-Form First Welfare Theorem. *working paper, Stanford University*.

Pareto, V. (1906). *Manuate di Economia Politica*. Società Editrice Libraria, Milano.

Scarf, H. (1960). Some examples of global instability of the competitive equilibrium. *International economic review 1*(3), 157–172.

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

Starr, R. (1969). Quasi-Equilibria in Markets with Non-Convex Preferences. *Econometrica 37*, 25–38.

Walras, L. (1874). *Elements d'Économie Pure, ou Théorie de la Richesse Sociale*. Paris:

Guillaumin et cie.

Williams, A. W., V. L. Smith, J. O. Ledyard, and S. Gjerstad (2000). Concurrent trading in two experimental markets with demand interdependence. *Economic Theory 16*(3), 511–528.

Williams, B. (2023a). Minimal Intelligence in the Double Auction: Logit-Choice and Reservation Utility. *AGORA working paper*.

Williams, B. (2023b). Minimal Intelligence in the Double Auction: Logit-Choice and Reservation Utility. *AGORA working paper*.

Williams, B. (2023c). Zero Intelligence in an Edgeworth Box. *AGORA working paper*.

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

For online publication only

Appendix A Scarf Cycles

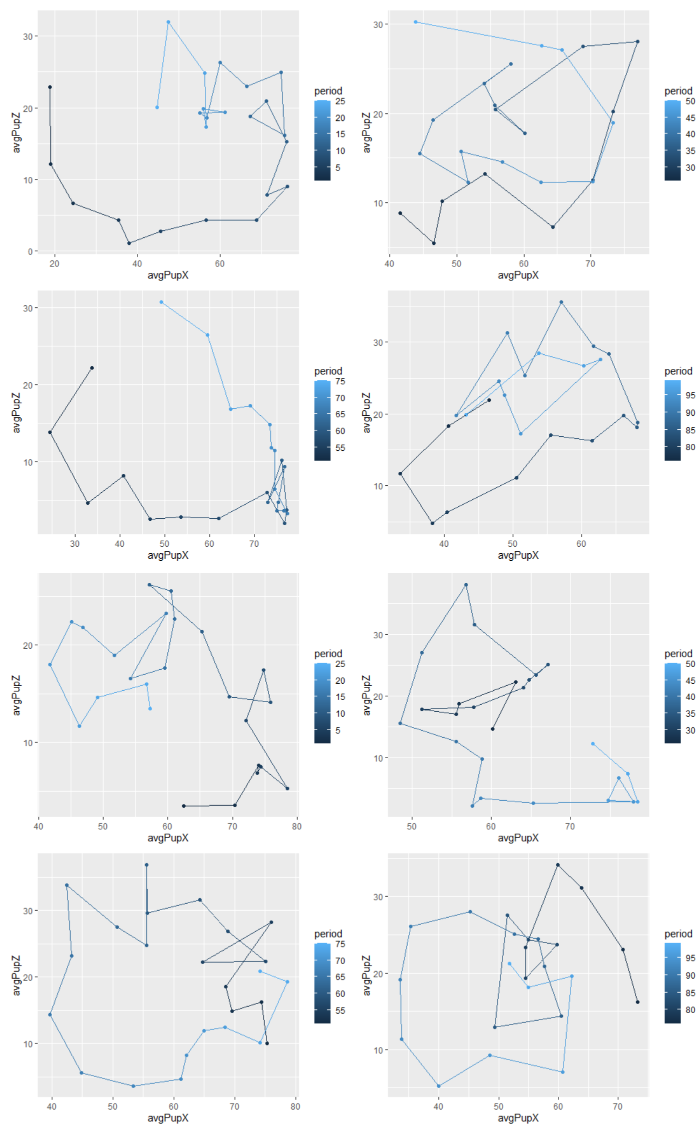

Figure A.1: Price trends for Scarf markets in the counterclockwise configuration. The upper four graphs depict trends with a starting price pair of $(p_x = 20, p_z = 30)$, while the lower four start with the pair $(p_x = 60, p_z = 10)$

.

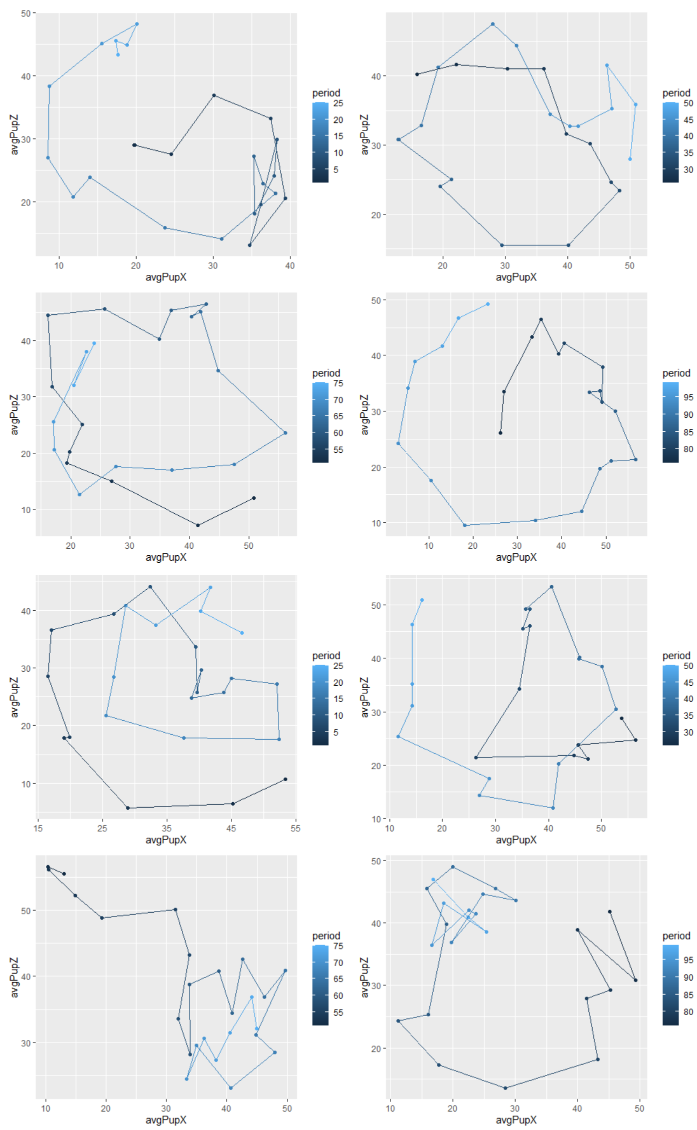

Figure A.2: Price trends for Scarf markets in the clockwise configuration. The upper four graphs depict trends with a starting price pair of ($p_x = 20, p_z = 30$), while the lower four start with the pair ($p_x = 60, p_z = 10$).

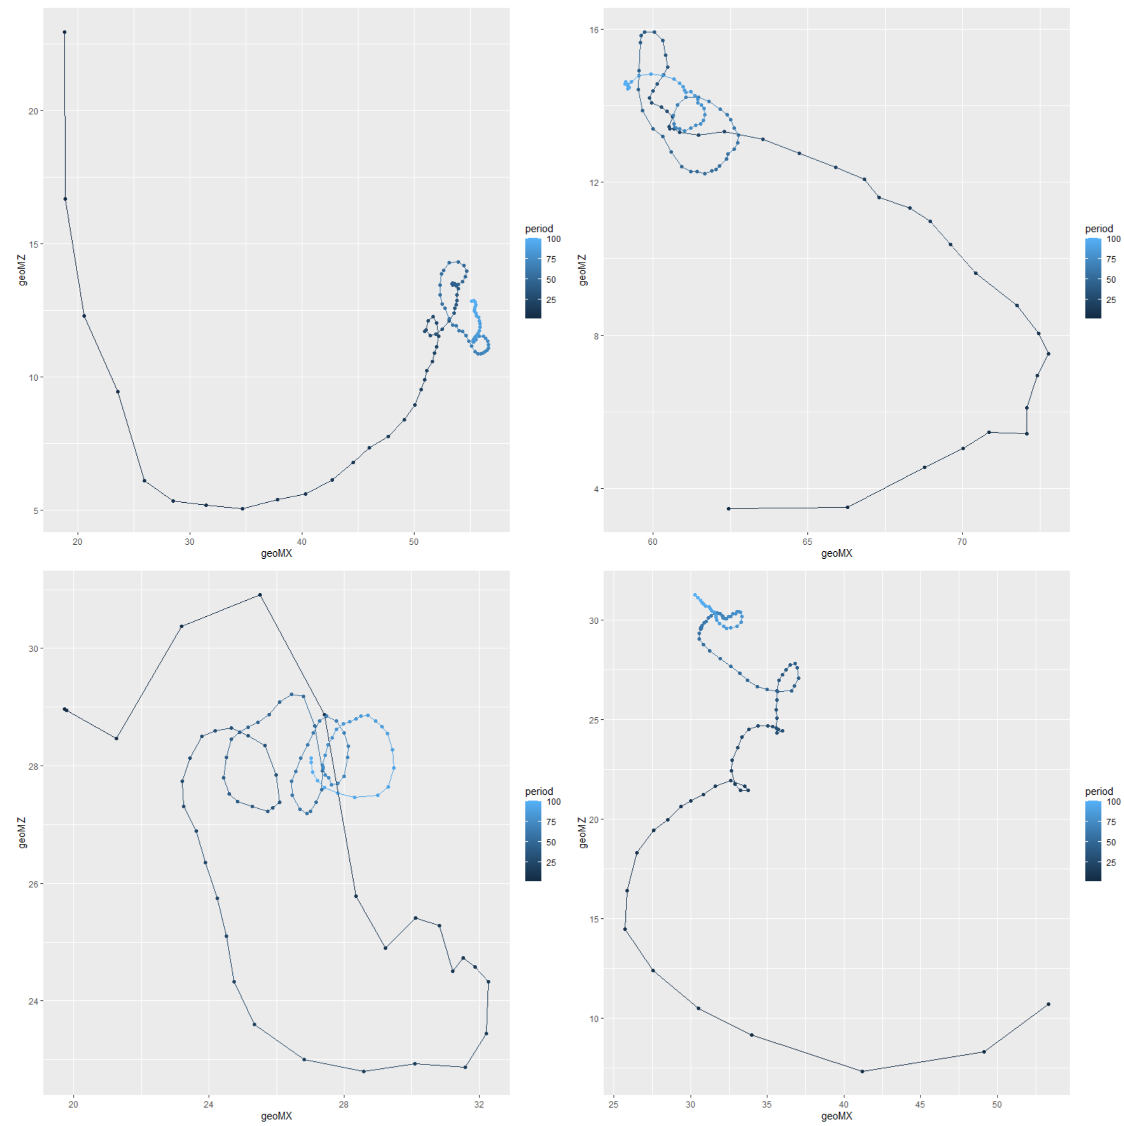

Figure A.3

Appendix B MinMax Accept Impact

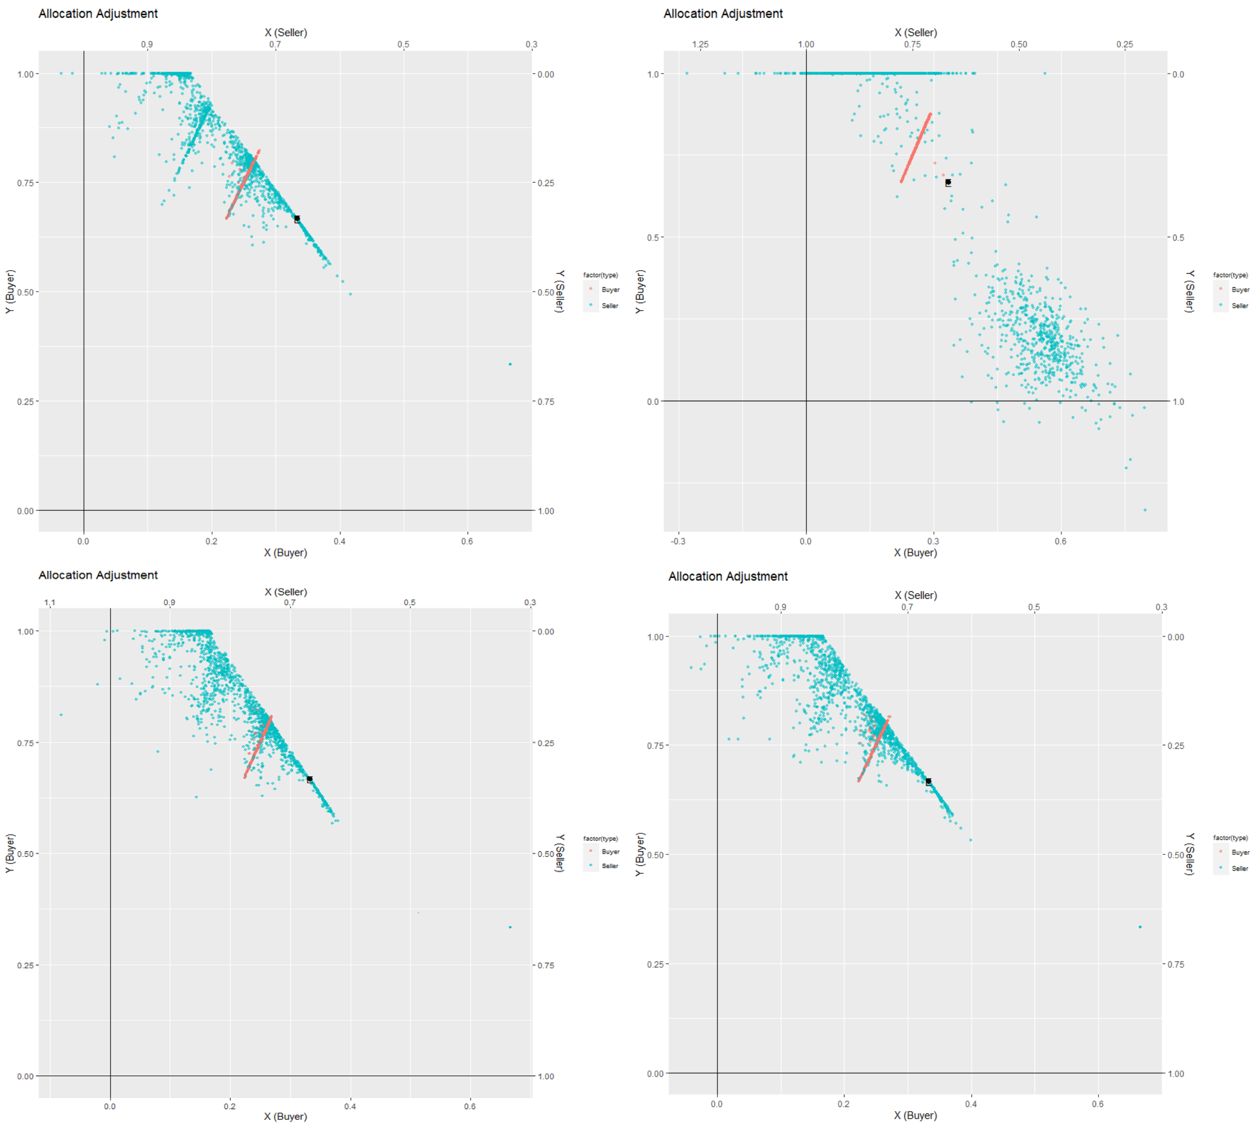

Figure B.1: Individual final allocations for markets with min(3x,y) and max(2x,y) traders, with starting endowment for the min trader being (1/3, 2/3). Rows denote accept rule use: no accept (bottom), or accept (top). Columns denote concavification use: 2x+y (left) or max(2x,y) (right).

Appendix C Accept Rule: All-or-as-much-as

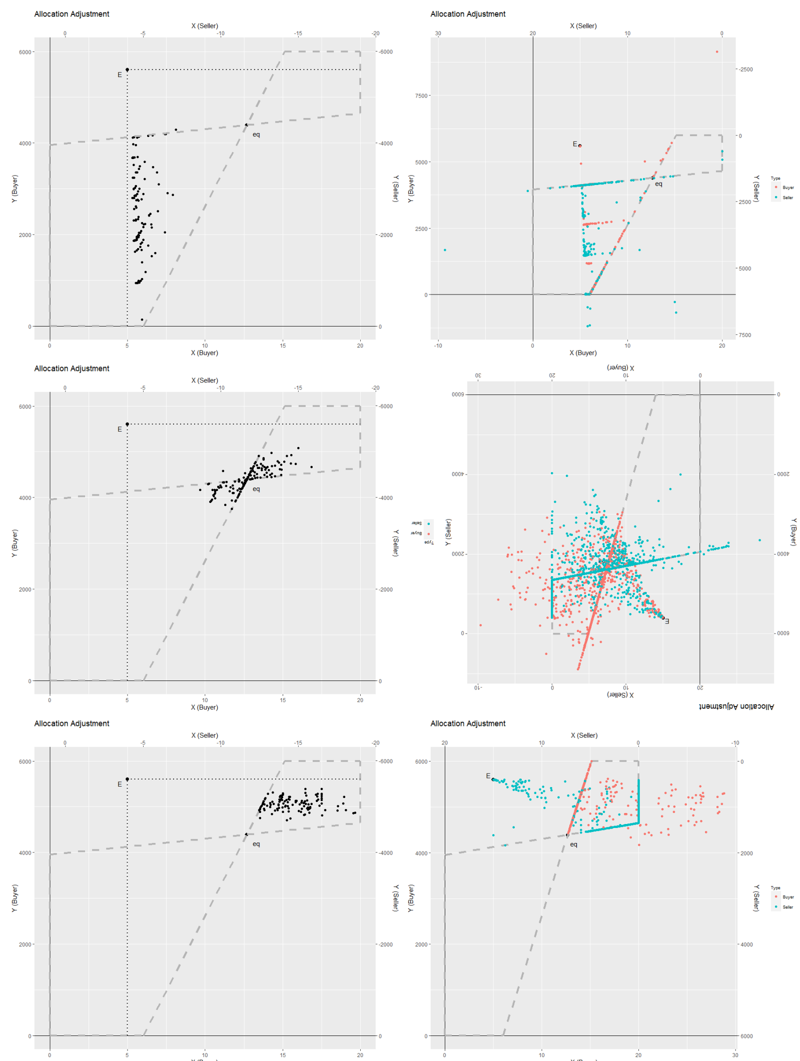

Figure C.1: Aggregated Agent Final Allocations. M=10000 markets on the left, M=350 in the middle, and M=120 on the left.