How Does Learning Makes Market Efficient? The Role of Market Microstructure

Misbah Hameed [1]

Maaz Javed[2]

Kinza Qureshi [3]

Abstract

The Zero-Intelligence (ZI) agents in the double auction market ensure efficiency by providing the highest surplus and trade volume but have been unable to ensure price convergence. Afterwards, focusing on the importance of rationality, the ZI agents are allocated limited intelligence to learn from the market and improve their behaviour over time. These agents with limited learning power are known as Zero-Intelligence Plus (ZIP) agents, and they ensure allocative and volume efficiency and make price convergence possible in the presence of the Marshallian trading sequence. However, these ZIP agents possess assumed learning rules to guarantee the desired efficiency level, and these rules are not based on how humans learn in real markets. To overcome this downside, this paper replaces the assumed learning rule of ZIP agents with new learning rules extracted from real human behaviour, i.e., introducing ZIP_H agents. In addition, the market efficiency of these ZIP_H agents is examined for different trading sequences and varying demand and supply schedules, i.e., symmetric, asymmetric, and box-shaped demand and supply schedules.

Keywords: Market efficiency, agent-based model, ZI agents, learning mechanism, market microstructure, price convergence, and competitive market.

Introduction

For over a century, competitive equilibrium has been a central concept in economic theory, suggesting that rational economic agents seeking profit naturally lead markets to equilibrium. This idea is encapsulated by Adam Smith's metaphor of the "invisible hand," which illustrates how markets can achieve maximum efficiency and surplus for buyers and sellers without external intervention. The pursuit of profit by individual agents drives the market toward perfection, making it efficient on its own.

The assumption that rational economic agents seek profit maximization is believed to drive market efficiency, aligning with the invisible hand theory in competitive markets. However, this theory faced criticism from non-cooperative game evidence (Becker, 1962)

---

[1] Independent Research (PhD). [Email: misbah.kureshii@gmail.com]
[2] University of Prince Mugrin (UPM), Saudi Arabia. [Email: m.javed@upm.edu.sa]
[3] Independent Research (PhD). [Email: kinzahameedqureshi@student.usm.my]

until Chamberlin (1948), Becker (1962), and Smith (1962 and 1976) provided empirical support demonstrating how the invisible hand fosters market efficiency.

Experiments in Double Auction (DA) markets have surprisingly supported competitive equilibrium theory, showing that buyers and sellers, driven by profit maximization within budget constraints, naturally lead the market to efficiency. These experiments confirm the invisible hand phenomenon, where market outcomes align with competitive market predictions over time as economic agents learn and transaction prices converge.

Experiments in Experimental Economics (EE) have tested the assumptions of competitive market theory, revealing that human traders learn and adapt to market environments, leading to outcomes consistent with competitive market predictions. This learning is influenced by experience, available information, and market mechanisms. The results are robust across different trader numbers, cultural contexts, and even in markets with extreme inequality. Notably, DA market experiments show that market efficiency can be achieved even when some competitive market assumptions, like the need for a large number of buyers and sellers, are relaxed.

Initial experiments in Double Auction (DA) markets show robust efficiency across various conditions, leading to skepticism about the necessity of human learning for achieving market efficiency. This skepticism arises from the consistent performance of DA markets in diverse environments, prompting researchers to question the role of learning and other assumptions, such as full information availability. To address this, learning models like Zero-Intelligence Plus (ZIP) and Snipers have been developed, but they often focus on outcomes rather than actual learning behaviours. Researchers have extracted learning rules from real human experiments and applied them to Zero-Intelligence agents to better simulate human learning in markets, aiming to more accurately replicate human behavior in simulated markets.

To test the validity of competitive market theory using Zero-Intelligence (ZI) agents in Agent-Based Modeling (ABM), we replicated human experiments in a simulated market. We explored differences between ZI agents and human traders and assessed the robustness of our learning model across various market environments, including symmetric and non-symmetric supply and demand schedules. By applying learning rules derived from actual human behaviour, we aimed to closely imitate human learning in these simulations and evaluate the model's consistency with real-world outcomes.

Considering the literature gap for not implementing the human-based learning rules in the double auction market, this paper aims to achieve the following research objectives.

I. To investigate the importance of learning mechanisms and rationality in simulated Double Auction markets.
II. To build a learning model for ZI agents and examine if it replicates the results of experiments with human traders.

Research Methodology

Gode and Sunder's studies (1992, 1993) showed that markets following the Marshallian trading sequence could achieve efficiency without human learning, suggesting that rationality isn't necessary at the micro-level. However, when alternative trading sequences like Reverse Marshallian or Randomized sequences are used, market efficiency may decline, highlighting the need for rationality. Cliff and Bruten (1997) supported this by demonstrating that merely relying on market forces isn't enough for efficiency, leading them to develop the Zero-Intelligence Plus (ZIP) agents with more complex learning mechanisms. This research further tests the theory by using a Randomized trading sequence to evaluate the robustness of ZI market results, emphasizing that "zero intelligence" alone is insufficient for achieving competitive market outcomes.

To evaluate the importance of learning in Double Auction (DA) simulations, this study replicates Cliff and Bruten's (1997) model, which builds on Gode and Sunder's (1993, 1994) work but introduces learning-capable agents. Both models use a Marshallian trading sequence. Cliff (2021) further explored how limited trading time impacts market surplus, noting a trade-off: as time runs out, traders lower their profit expectations to ensure they can trade before the period ends. This behaviour mirrors human traders in DA markets, who initially seek high profits but adjust their bids and asks as they learn about the market and face time constraints.

Cliff (2021) emphasizes that the initial limit prices set for agents—minimum prices for buyers and maximum for sellers—significantly affect their learning behavior. For example, a high seller limit price expands the range for shout prices, slowing down the agent's learning process as they require more time to determine optimal trading prices. This creates an inverse relationship between learning speed and the price range available to agents. Cliff also updates ZIP agents, considering the impact of this inverse relationship and the time allowed for trading in simulations, focusing on the exogenously set limit prices.

While exogenously set limit prices affect agent learning behaviour, the primary focus here is on the trading sequence used in these artificial markets. The first step is to replicate the ZIP agents' learning model, which uses the same Marshallian trading sequence as ZI agents. The emphasis is on the learning rules ZIP agents follow to optimize their trades and maximize gains (see Appendix II).

By incorporating learning rules into ZI agents, we create ZIP agents capable of improving their shout prices based on their environment. Using the demand and supply model from Cliff and Bruten (1997), we replicate a DA market with ZIP agents and introduce variations in the trading sequence. The goal is to test whether ZIP agents, who learn and adapt, can achieve market efficiency across different trading sequences.

The ZIP agents in Cliff and Bruten's (1997) and Cliff's (2021) models, using a Marshallian Path microstructure, achieve higher market efficiency than ZI agents from Gode and Sunder's studies. However, these results are sensitive to factors like the shout price range and trading sequence. To address this, we propose using an alternative randomized trading sequence while allowing ZI agents to learn about the market, instead of the Marshallian sequence with limited randomization.

This is the second phase of testing the importance of market microstructure in the DA market environment. The first attempt is to apply an alternative trading sequence for ZI agents with no intelligence; hence, there is no rationality or learning mechanism. If the results of an alternative randomized trading sequence differ from the trading sequence used by Gode and Sunder (1993 and 1997) we reach to settle on the point that the market itself doesn't lead the market toward the efficient outcome. Once these results are obtained, the introduction of intelligence to the agents is mimicked, as in Cliff and Bruten (1997) and Cliff (2021). Afterwards, we intend to test the DA market populated with agents who can learn about the market environment. It helps to figure out if the learning mechanism of traders proposed in the literature is helpful in always achieving market efficiency.

An Alternative Learning Mechanism of ZI Agents

ZI agents, while initially promising, can be made more efficient with learning mechanisms. Cliff and Bruten (1997) introduced a complex learning algorithm, but Preist and Tol (1998) highlighted that their simpler heuristic is also effective. Both approaches involve algorithms and coefficients aimed at maximizing surplus and improving agent performance.

Gjerstad and Dickhaut (1998) introduced a complex learning mechanism where agents improve their shout prices based on communication with others and past experiences. Their approach includes strategies like wait-watch-grab, where agents exploit others by waiting for optimal prices before acting. Kaplan agents use this sniping strategy but may be driven out once others adapt. Genetic Programming (GP) agents use a pool of strategies, switching strategies to enhance performance, while Bayesian Game Against Nature (BGAN) agents update their beliefs based on the distribution of others' shouts. All these strategies are assumed by the experimenters and not derived from real-world heuristics.

The number of learning mechanisms introduced in the literature is even greater than that discussed above. However, there are a few major concerns regarding these mechanisms. First, these markets with learning agents assume learning rules by themselves instead of borrowing them from empirical results, experiments, or real market data. Second, all the learning mechanisms are so complex that they are very hard to understand, let alone follow in the actual trades. Therefore, learning rules were extracted from the experimental data. Although there are separate theories on how human traders learn about the market, there is no evidence that applies the rule extracted from human traders to ZI agents. The major theories on human learning are Regret Theory and Directional Learning theory.

Empirical research on learning mechanisms in DA markets often falls into two categories: advanced theories based on human learning and complex fictional models for ZI agents. However, no studies have applied human-derived learning rules to ZI models. The key to bridging experimental economics and agent-based modeling is to extract real human behavior rules from experiments and apply them to artificial markets. The goal should be to create markets based on these real-world rules, rather than merely aligning artificial intelligence models with competitive equilibrium predictions.

In this study, we extracted rules from real-world experimental data to develop our model. The necessary data included participant numbers, redemption values, costs, shout prices, transaction prices, and surpluses from trades. This information, provided by DeYoung (1993) from a double auction market experiment, allowed us to define how traders adjust their bid-ask spreads to become eligible for trade. These learning rules are then applied in an artificial market with the same market setting as of DeYoung (1993) to confirm if these rules offer the same results as provided in the human-populated market experiment[4].

Box 1: Learning Rule for ZI Populated Simulated Double Auction Market

The learning rule for the simulated double auction market can be customized in a way that mirrors the results of human traders in the same market with similar market conditions.

Before going into details of how the model works, it is necessary here to explain some assumptions regarding the behavior of agents in the market.

Assumptions:

Agents are risk averse. Implication of this assumption is the learning behavior of agents emerges over the trades as they witness the successful trades at the initial shout prices with lower profit margin. So, over the periods the agents start with high profit margins but when they see that they are unable to trade at first iteration then they decrease their profit margin on second iteration and so on.

Following steps are followed for the DA market populated with artificial agents and model is developed in Python to reconfirm if the results from artificial agents can replicate the results of human traders.

Step 1:

---

[4] Different extracted learning rules are applied many times in an artificial market to come up with the same results as DeYoung (1993), and lastly, the rules are finalized that provide the same outcome.

- Randomly any of the buyer or the seller is selected from the market.
- If one of the buyers and one of the sellers is selected then go to step 2 otherwise keep repeating unless market has at least one of buyers and sellers.

Step 2:

- Record the value (cost) for buyer/buyers (seller/sellers) selected in step 1.
- Following approach is used to allocate the ask price to each selected seller and bid price to each selected buyer.
  For the selected Buyer:
  P_{Bid} = Random Value [(Value * Profit Margin), Value]
  For the selected Seller:
  P_{Ask} = Random Value [Cost, (Cost * Profit Margin)]
  Profit Margin for all Traders = Trading Period * 10%

Step 3:

- Once the ask and bid prices are allocated to selected buyers and sellers then, over the iteration within each period and across the period as well, these trades have to be improved depending on the agent behavior.
- In our model, these two shout improvement rules are as follow.
  Over the periods: As the agents are risk averse so at the start of trading period the lower profit margins is set but as the period passes and they witness trades to be successful they start improving the profit margin over the periods.
  Over the iterations: Within each period, the agents who have not been successful in the making trade in past iteration they decrease their profit margin to be able to make trade.

Step 4:

Once the buyers and sellers are selected randomly at first step then the following steps are to be followed.

- After the first step of selecting the traders to trade from the list of all buyers and sellers, next step is to compare if these traders are eligible to trade.
- If any of the buyer has bid price greater than the ask price of any of the seller, then they both gets eligible to trade.
- Here if more than two buyers (sellers) have bids (asks) greater (lower) than the ask (bid) price of seller (buyer) then greatest (lowest) of the two will be selected.

Step 5:

- If any of the buyers does not has bid price greater than the ask price of any of the seller (as happened in iteration 1) then the active agents will revise their shout by decreasing and replacing it with new profit margin.
  Revise profit margin (Over the Iteration)$_t$ = $\frac{Profit\ Margin_{t-1}}{2}$

- Any of buyer or the seller is selected randomly, as in step 1.
- This process will keep going until any of the buyer has bid price greater than ask price of seller.

Step 6:

- If the bid price of any of the buyer is greater than the ask price of any of the seller, then they get eligible for trade.

Step 7:

- When the trade will happen between buyer of highest bid and seller of lowest cost then both of these agents will go out of market for that period.
- Trade price is equal to the shout price of trader who enters the market first.

Alternative Learning Mechanism: Marshallian VS Randomized Trading Sequence

Once the learning rules are defined for trading in an artificially simulated market, the next step is applying this learning mechanism in the double auction market. It is applied in a number of steps to test the required double auction models.

1. Once the learning rules are established, then these rules are applied to the number of double auction market environments having different demand and supply functions, i.e., symmetric, asymmetric, and box-shaped demand and supply functions.
2. These learning rules are applied in a double auction market following the Marshallian sequence and then applied in the market with a Randomized trading sequence. It helps to find out which of these two trading sequences is helpful to gain maximum surplus, to have maximum trade volume, or to get the prices to converge to the equilibrium level.

After each experiment, we assess performance by measuring allocative efficiency, price variability, convergence coefficient, and total trade volume. We evaluate successful trades to determine agent surplus and compare it with overall market gains. According to competitive price theory, optimal trade combinations should maximize market surplus. We can calculate the market efficiency, E, after extracting the value of maximum theoretical surplus[5], ε. Formula to measure efficiency is given below.

$$\mathrm{E} = \frac{\Sigma\,(\mathrm{Vi-Pi}) + \Sigma\,(\mathrm{Pi-Ci})}{\varepsilon} * 100 \tag{3.1}$$

Here V_i is the value of one unit of good allocated to each buyer and C_i is the cost of producing one unit for each seller. P_i denotes the transaction price of for each transaction and summation sign is applied only on successful trades as unsuccessful trades do not have a price. The efficiency index, E, is sensitive to microstructure of the market and must

---

[5] Maximum surplus shows the total profit of buyers and sellers from all the trades.

be interpreted accordingly, for instance, if extra-marginal units gets traded then it'll have an negative effect on market efficiency.

Another market performance measure taken from Davis and Holt (1993) is the coefficient of convergence, α, which measure the price variability in the market as well as the deviation of transaction prices from competitive equilibrium prices. It's important in measuring the gap between the actual market price and the competitive equilibrium price level. It shows in what direction the market is going. It can be measured as the square root of the variance of prices around the theoretical competitive equilibrium level of price and is calculated as,

$$\alpha = \sqrt{\frac{\Sigma_{n=1}^{N}(P_t - \mathrm{P}_e)^2}{N}} \tag{3.2}$$

Where N is the total number of successful transactions in one trading period, P_t is the price at which a transaction takes place and P_e is the competitive equilibrium price level. If all transactions are done at a competitive equilibrium level of prices then the value of α will be zero showing no difference between transaction price and equilibrium price. The value of α increases with deviation of mean prices from competitive equilibrium or with an increase in price volatility. It is usually observed in various researches that the value of α drops substantially with each trading period showing the convergence of prices towards competitive equilibrium (LiCalzi, Milone and Pellizzari, 2011)[6]. The reason why its value decreases is not analyzed in the literature.

The last measure is the trade volume which is nothing but the sum of all successful trade in one trading period. In some cases, the averaging trading volume is also considered to compare the market performance in the literature. Actual trade volume is compared with the theoretical maximum possible trade that can happen in the market in a trading period. For the market to be efficient, the actual trade volume should be equal to theoretical trade volume.

All of these performance appraisal measures are to be used in evaluating the market performance of all the simulations. Purpose of using the same performance measures as used in studies earlier is to compare the results.

Empirical Results

Before going into the results and their discussion, it would be convenient to summarize what these results are about. The empirical results of all the simulations are divided into three main parts and each of these is subdivided in two sub-strands, as below.

1. Market performance when agents can learn about the market:
    a. Duplication of ZI populated market with learning mechanism introduced by Cliff and Bruten (1997) while following Marshallian trading sequence.

---

[6] This phenomenon of price convergence is evident in a market following Marshallian trading sequence.

    b. Results of alternative Randomized trading sequence in ZI populated market who follows learning rule of Cliff and Bruten.
2. Market performance with new learning rule:
    a. Application of new extracted learning rules (from DeYoung (1993)) in Marshallian trading sequence.
    b. Replication of new learning rules in market mechanism with an alternative trading sequence.

The simulations are run across various demand and supply functions to ensure that the results are robust. This purpose is attained by implementing three main types of demand and supply functions as follow.

1. Symmetric (with equal but opposite) demand and supply functions
2. Asymmetric (with unequal) demand and supply functions
3. Box-shaped (with relatively constant) demand and supply functions

To achieve our objective, we use five types of demand and supply schedules, drawn from literature, to test market efficiency. These schedules include symmetric, asymmetric, and box-shaped types, as detailed in Appendix-I. Symmetric markets have matching demand and supply curves, asymmetric markets have differing values and costs, and box-shaped markets feature values and costs arranged in a box-like shape.

Market outcomes are influenced by specific setup details. Competitive equilibrium predictions are not independent of market microstructure. For example, a market with ZI agents may not achieve efficiency or convergence, depending on its structure. A Marshallian trading sequence is optimal for maximizing surplus, while a Randomized trading sequence is better for maximizing trader participation. Thus, market outcomes are sensitive to the details of the market setup.

This section examines whether market efficiency is universal across different microstructures, using a ZI-populated DA market framework. It evaluates market performance with individual rationality, focusing on agent learning. Initial results show that even without learning or rationality, ZI markets can achieve allocative efficiency similar to markets with human traders.

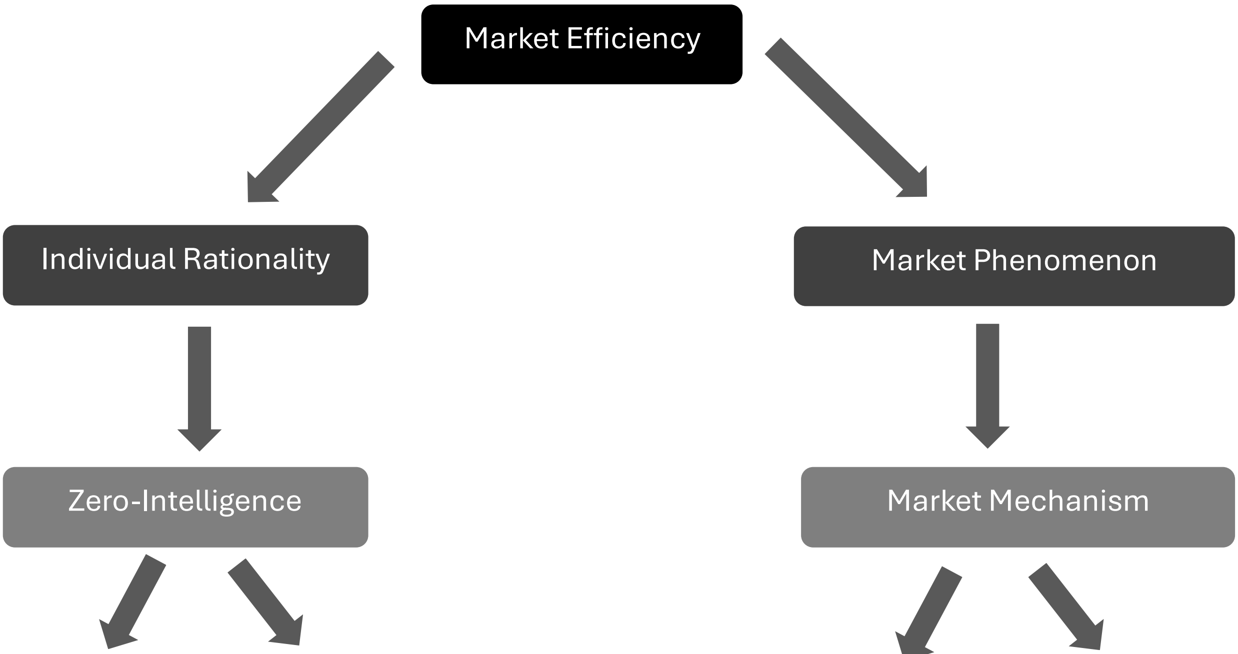

Figure 1: Market mechanism with alternative learning mechanisms

The previous section questioned the efficiency of ZI-populated markets, so this analysis focuses on how learning affects agent performance. Cliff and Bruten (1997) showed that incorporating learning mechanisms into ZI agents significantly improves market efficiency compared to ZI agents without learning. This indicates that learning has a notable impact on market outcomes. The effect of learning is further explored by comparing efficiency from simulated learning rules and rules extracted from human experiments.

Incorporating learning mechanisms into ZI agents, termed ZIP agents by Cliff and Bruten (1997), significantly improves market performance. However, the learning rules used for ZIP agents are assumed rather than derived from human trader experiments. For more accurate and efficient outcomes, learning rules should be based on actual human trading behavior rather than hypothetical models.

This purpose is served in the following steps.

a. Come up with a ZI model based on learning rules of ZIP agents that provide the same outcome as provided by ZIP traders.
b. The next step is to use extract the learning rules from the experiment with human traders.
c. Now these learning rules, that come from the real human experiment, can be substituted in the ZI-populated simulated market in the first step. The benefit of replacing only the learning rules is to control all the other factors used by Cliff and Bruten (1997) except the learning rules. It will help to investigate the impact of change in the learning mechanism on the performance of ZI agents.

The results for all these three forms of simulations are provided in the coming section. The next section is dedicated to looking into the results of ZI traders with learning mechanisms and how these results are achieved.

ZI-Plus Agents in DA Market

To imitate the results of Zero-Intelligence Plus agents, the same demand and supply schedules are required to be followed. The demand and supply schedules are illustrated in Figure 18 which exhibits symmetric demand and supply as used in market 1 of G&S. Here 200 is the equilibrium price level with an equilibrium quantity of 6 units.

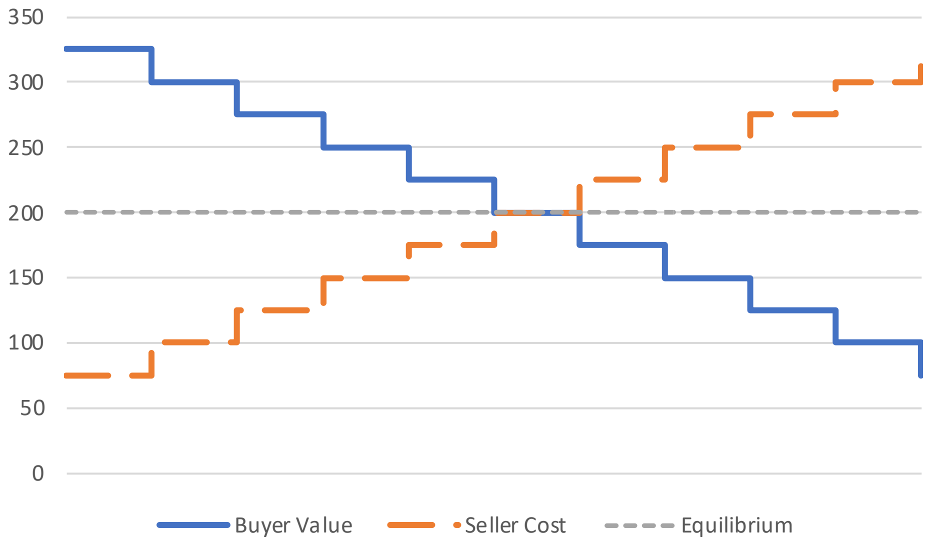

Figure 2: Market Demand and Supply Schedule by Cliff and Bruten (1997)

With the implementation of the ZIP learning rules in the Marshallian sequence, the results from trade across 10 periods are given in Figure above[7]. The main objective of C&B was to show the price convergence process because of the learning process. Their results just provide period-wise an average transaction price for the same market.

Market Efficiency of ZIP_H Traders with Simulated Learning Mechanism

After applying the same learning rules the results of ZIP_H traders in a Marshallian trading sequence demonstrate the price convergence, as witnessed in C&B[8]. The prices are not only converging towards market equilibrium value but also converging to equilibrium across the periods. These results contrast with earlier results with ZI agents where no price convergence was observed. It shows that the learning mechanism has a significant effect in leading the transaction prices to their theoretical equilibrium value[9].

---

[7] The selection of 10 periods is done to imitate the results of Cliff and Bruten (1997).

[8] Here we name the zero-intelligence agents as ZIP_H agents as they possess the learning abilities and these learning rules come from the human-populated experiments.

[9] As Cliff and Bruten (1997) doesn’t provide analysis on market efficiency, it is why here only the results of price convergence are provided to make sure that these results are in line with the price convergence of ZIP agents.

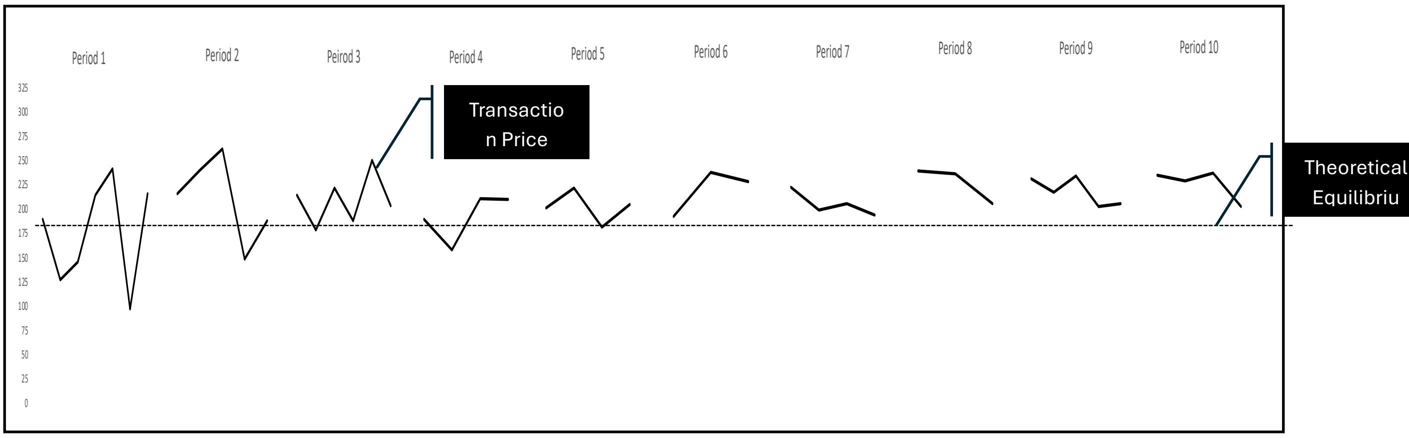

Figure 3: Transaction prices with ZIP agents

The convergence of transaction prices towards equilibrium highlights the importance of the learning mechanism in simulated markets. Without learning, ZI-populated markets fail to achieve price convergence. Convergence ensures that transaction prices align with the competitive equilibrium, thereby reflecting effective market efficiency and a closer match to theoretical outcomes.

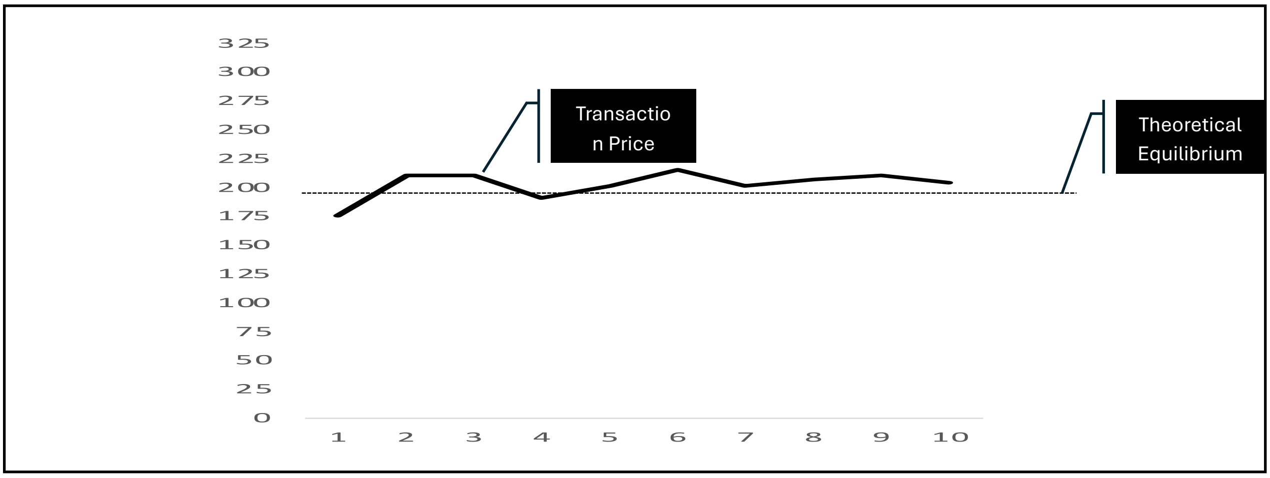

Figure 4: Period-wise average transaction price

Average transaction prices converge to the theoretical equilibrium, indicating market efficiency, driven by the learning mechanism in ZIP agents compared to simple ZI agents. This challenges the idea that market microstructure is irrelevant. However, the learning rules used by Cliff and Bruten (C&B) are complex and not reflective of actual human trading behaviours, raising concerns about the simulation results. This prompts a reevaluation using learning rules based on real human trading experiments to better understand if different learning mechanisms could yield varying results.

ZIP Agents with Learning Mechanism of Human Traders with MS: The Baseline Model

The previous section highlighted how learning mechanisms enhance the convergence of transaction prices toward competitive equilibrium. This section focuses on applying learning rules from real human trader experiments to ZIP agents, rather than using assumed rules. DeYoung (1993) is utilized for its detailed data on bid and ask prices and transaction prices, crucial for deriving accurate learning rules and analyzing simulated market outcomes.

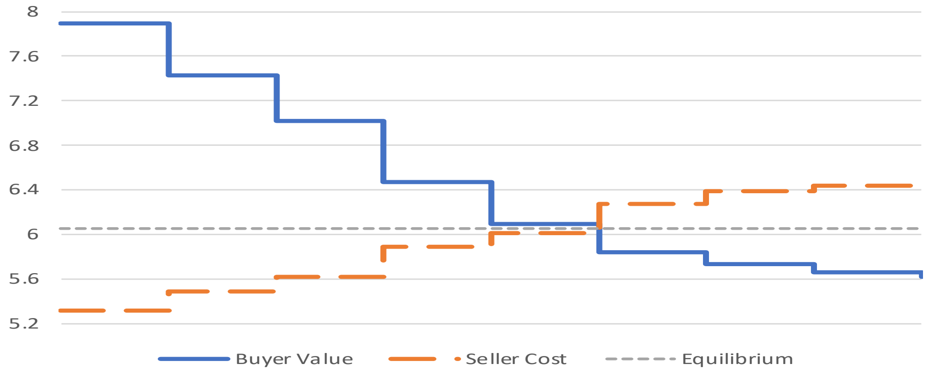

Figure 5: Market Demand and Supply Schedules by DeYoung (1993)

The demand and supply schedules used in DeYoung (1993) are asymmetric, as can be seen in Figure 21. The theoretical competitive equilibrium price is 6.05, and the maximum achievable surplus is 6.57 in this market structure.

Price Convergence of ZIP Agents with Experiment-Based Learning Rules

Now the purpose of using this market structure is to extract the learning rule from the results of DeYoung (1993) and come up with a ZI-populated simulated market that reaches the same outcome as DeYoung (1993). It will help to ensure the effectiveness of learning rules and the reliability of these rules in the artificially simulated market. Later, this simulated market can be used to check if it provides better results as compared to a market populated by ZI agents with no learning mechanism. As a first step, these learning rules will be applied to a simulated market with the Marshallian trading sequence. The focus here is on examining if the transaction prices converge toward the competitive equilibrium price level.

Different parameters are used to assess market efficiency. While earlier research focused on surplus and trade volume, this study examines price convergence in markets with learning mechanisms. ZI agents without learning don’t achieve price convergence, whereas ZIP agents with learning mechanisms do, as they adjust prices toward equilibrium. The study uses learning rules from DeYoung (1993) to analyze ZIP agent behaviour, with transaction prices reflecting their learning progress. Results show that ZIP agents’ transaction prices approach competitive equilibrium over time, aligning with human-populated market results from DeYoung (1993). This indicates that ZIP agents effectively learn and adapt to market conditions.

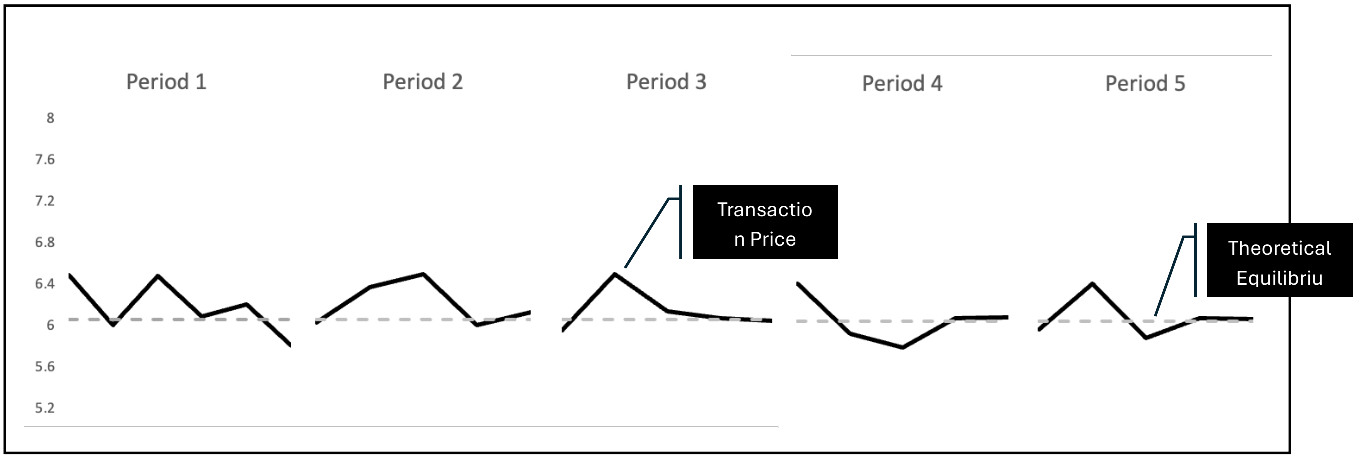

Figure 6: Transaction price over the period with ZIP agents with learning mechanism

Allocative Efficiency of ZIP Agents with Experiment-Based Learning Rules

In terms of allocative efficiency, the learning process shall be supportive to increase it. When the agents learn about the market, they shout the price close to the competitive equilibrium price level, and at this price level, only the intra-marginal agents are more likely to trade[10]. The trade between the intra-marginal buyers and sellers leads to higher allocative efficiency. It is why the agents with the learning mechanism should have experienced increasing allocative efficiency over the periods as they learn about the market mechanism.

Table 1: Period-wise surplus of ZIP agents over the transactions (in percentage)

Period	**Period 1**		**Period 2**		**Period 3**		**Period 4**		**Period 5**	
Transaction	**Buyer**	**Seller**	**Buyer**	**Seller**	**Buyer**	**Seller**	**Buyer**	**Seller**	**Buyer**	**Seller**
T-1	17.9	8.8	23.7	12.8	23.1	14.7	22.9	15.4	22.9	14.8
T-2	19.2	7.8	17.9	15.5	18.6	11.8	18.7	10.7	18.4	10.7
T-3	7.8	6.8	9.3	15.4	12.7	7.8	13.7	8.0	13.5	8.1
T-4	0.2	9.7	7.3	0.0	5.6	0.9	5.6	3.2	5.3	3.1

---

[10] Intra-marginal agents are the market agents who have position at the left side of competitive equilibrium level. Or they are buyer who have redemption values higher than the competitive equilibrium price and seller who have cost less than the competitive equilibrium price level.

T-5	2.3	3.2	0.3	0.1			0.0	1.5	0.3	1.0
T-6	0.7	0.5								

The ZIP-populated agent market results demonstrate the effect of learning on allocative efficiency. Using learning rules from DeYoung (1993), allocative efficiency is assessed by summing individual surpluses of buyers and sellers. As transaction prices converge, individual surpluses decrease. The table shows a declining trend in individual surpluses over time, with allocative efficiency measured as a percentage of the total possible surplus, aligning with competitive market standards.

Table 2: Period-wise allocative efficiency of ZIP traders imitated learning rules (%)

	Period 1	**Period 2**	**Period 3**	**Period 4**	**Period 5**
Allocative Efficiency	86.1	97.6	97.8	100	100

Allocative efficiency improves over time, rising from 86.1% in period 1 to 100% in period 5. This increase reflects how agents learn market dynamics, adjusting their prices closer to the competitive equilibrium. As a result, only those agents near the equilibrium trade, maximizing market surplus. With learning mechanisms similar to human traders, ZIPH agents achieve both price convergence and higher allocative efficiency.

Trade Volume Efficiency of ZIP_H Agents with Human-Based Learning Rules

The third measure of market efficiency is trade volume efficiency which exhibits the number of trades that happen in a period as a percentage of the maximum possible trade that can happen. In a market with demand and supply schedules borrowed from DeYoung (1993), the trade volume efficiency is quite higher than it is for the ZIP_H agents with no learning.

The theoretical competitive number of trades is 5 in each period for this market. Whereas the actual trades that happen in each of these 5 periods are 6, 5, 4, 5, and 5. So, the overall average trade volume efficiency is around 100 percent which is the maximum the market can offer.

ZIP Agent with Learning Mechanism: Randomized Trading Sequence

Till now, the newly introduced learning model has been implemented in a market with a Marshallian trading sequence and demonstrates that when the learning is introduced to ZIPH agents then, their performance becomes better, and they achieve higher market efficiency in terms of higher surplus, a greater number of trades and the convergence of transaction prices. This section now attempts to introduce a learning model in a ZIP_H-populated simulated market that follows the Randomized trading sequence.

The purpose here is to explore the market efficiency when the ZIP$_H$ traders follow human-induced learning rules in a Randomized trading sequence instead of following the Marshallian one. For this purpose, the first step is to introduce three different market demand and supply schedules. These three markets have symmetric, asymmetric, and box-shaped demand and supply schedules are shown below in Market 1, Market 2, and Market 3.

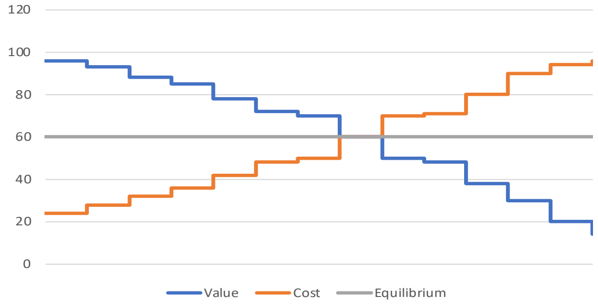

Figure 7: Market 1 with symmetric demand and supply schedule by Isaac and Plott

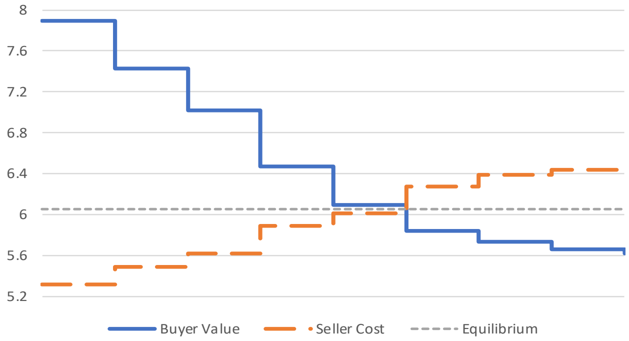

Figure 8: Market 2 with asymmetric demand and supply schedule by Gjerstad

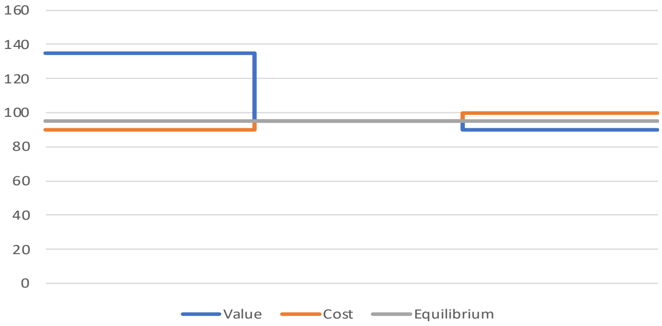

Figure 9: Market 3 with box-shaped demand and supply schedule by Gode and Sunder

The next section presents the results for ZIPH agents using three market schedules with a Randomized trading sequence and human-inspired learning rules. Market efficiency will be evaluated based on transaction price convergence, allocative efficiency, and trade volume efficiency.

Allocative Efficiency of ZIP$_H$ Traders with Randomized Trading Sequence

The allocative efficiency of a DA market with a Randomized trading sequence will test if human-inspired learning rules are more effective than assumed ones. While ZIP agents use assumed learning rules, ZIPH agents use human-inspired rules. Unlike ZIP agents, whose efficiency improves with the Marshallian Path, ZIPH agents combine human-inspired learning with Randomized trading.

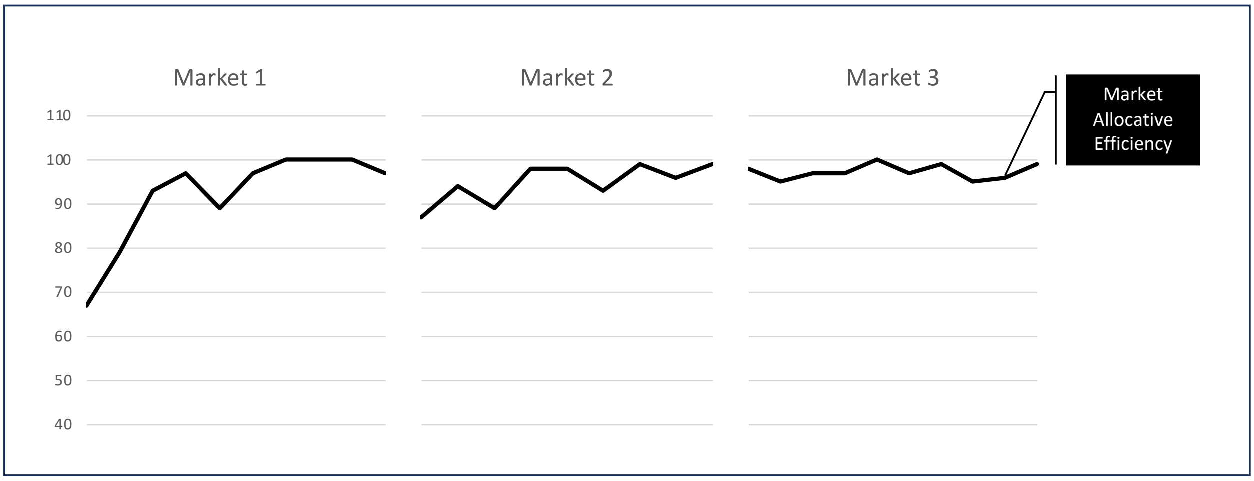

Figure 10: Allocative efficiency of market with ZIP$_H$ agents following RS

Allocative efficiency is highest when agents use human-inspired learning rules, with box-shaped schedules showing the best results (97%), followed by asymmetric (95%) and symmetric schedules (92%). This indicates that learning improves market efficiency compared to ZI agents and ZIP agents with assumed learning. These findings align with human experiments, highlighting that rationality and learning are crucial for achieving high market efficiency and surplus. The next section will address price convergence as another measure of efficiency.

Price Convergence in ZIP-populated Market with Randomized Trading Sequence

To assess the market performance of ZIP$_H$ agents in terms of price convergence in a randomized trading sequence all three types of market (symmetric, asymmetric, and box-shaped) are evaluated again. If the transaction prices converge toward the competitive equilibrium price level, then it means the actual outcome of the simulated market is getting closer to the theoretically predicted outcome. The transaction prices for these markets are illustrated below.

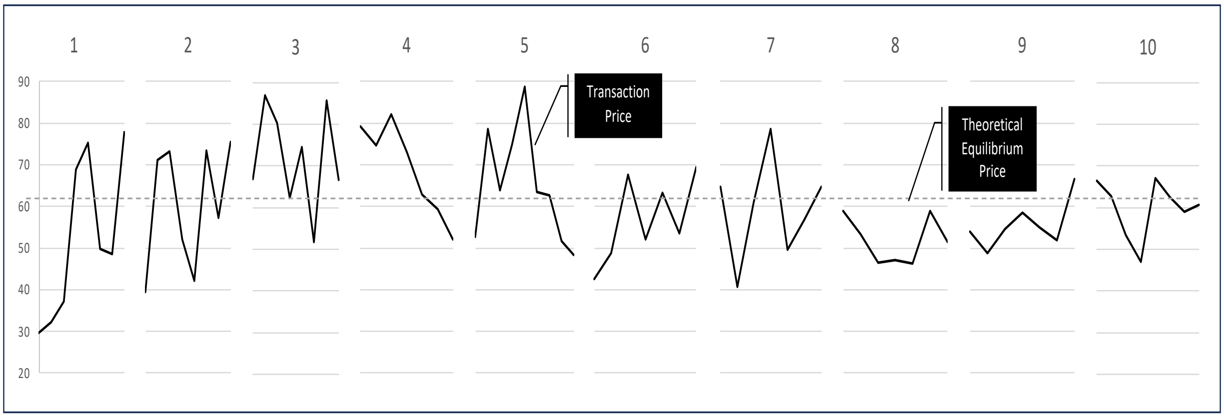

Figure 11: Period-wise TP of the market with symmetric D&S schedule

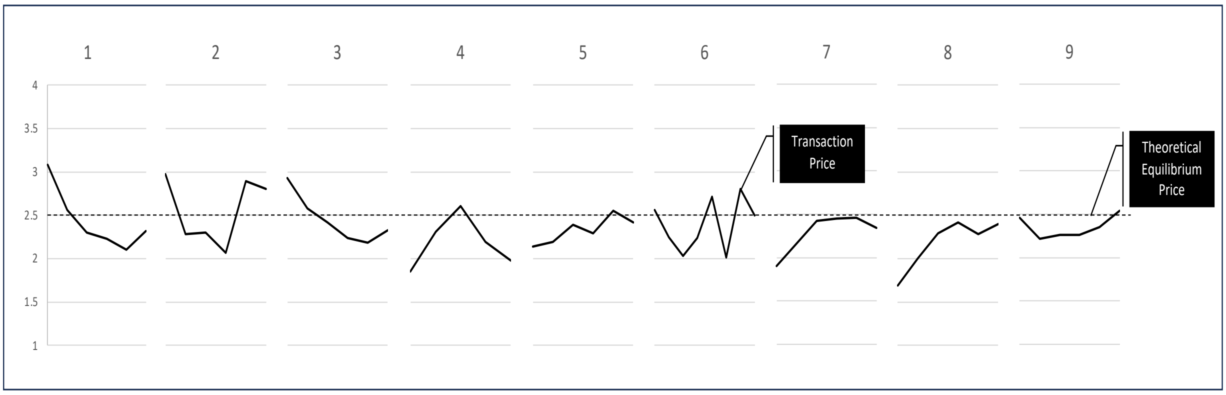

Figure 12: Period-wise TP of the market with asymmetric D&S schedule

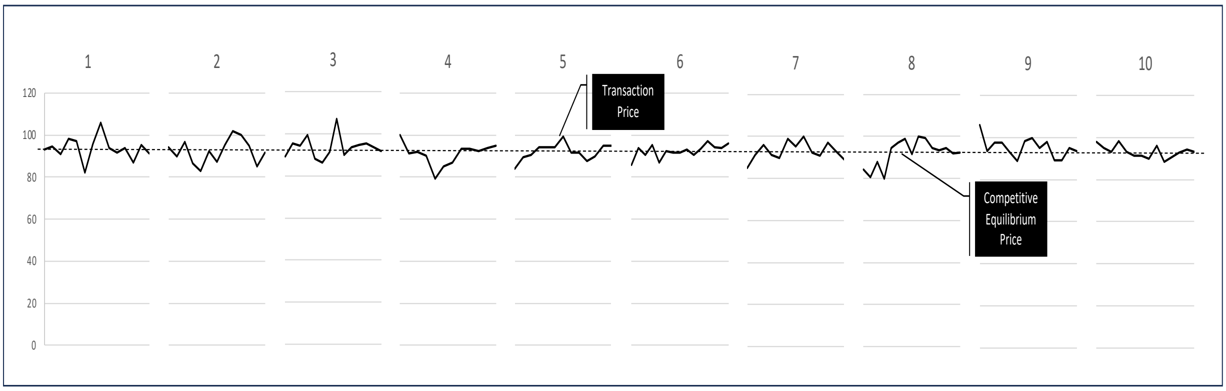

Figure 13: Period-wise TP of the market with box-shaped D&S schedule

Transaction prices show strong convergence to the competitive equilibrium price in each period, particularly in asymmetric and box-shaped markets. However, this convergence is weaker compared to ZIPH markets with Marshallian trading sequences. The Marshallian sequence results in higher market efficiency and better price convergence, but efficiency decreases somewhat with randomized sequences, though not as drastically as with ZI agents lacking learning mechanisms. The next section will address trade volume efficiency.

Trade Volume Efficiency of ZIP$_H$-populated Market with RS

Trade volume efficiency is another measure of market efficiency along with allocative efficiency and price convergence. Trade volume efficiency is also compared for the ZIP$_H$-populated market with varying demand and supply schedules.

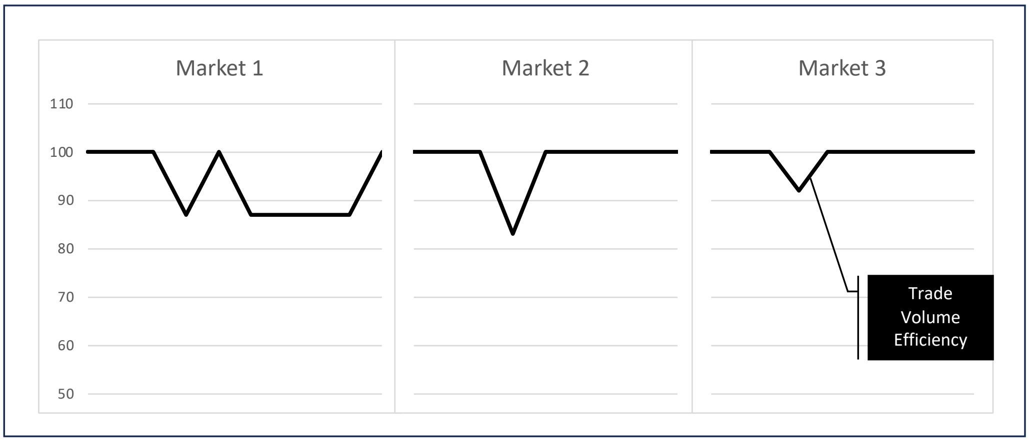

Figure 14: Volume efficiency of market with ZIP$_H$ agents following randomized sequence

Trade volume efficiency shows a high level of efficiency as in both markets (asymmetric and box-shaped), the efficiency is 100 per cent except for one transaction. However, the trade volume efficiency is lower in symmetric demand and supply scheduled markets because the extra-marginal traders trade with intra-marginal traders, and some of the potential intra-marginal agents have not been able to get into the trade. It leads to a decrease in the number of trades as a percentage of the competitive number of trades in each of these markets.

When compared with the results of ZIP$_H$ agents with Marshallian trading sequence, the difference is not huge as in that market trade volume efficiency was also higher as it is here. These results show that the trade volume is a measure of market efficiency that provides high-efficiency values irrespective of the trading sequence.

Whereas the market efficiency for the ZIP$_H$-populated market is visibly higher for the Marshallian trading sequence than it is for the Randomized trading sequence. These results explain the sensitivity of market outcomes and their efficiency with respect to the microstructure. It means changing the details regarding how the simulated market is set up; the market efficiency keeps altering. The efficiency of the market is sensitive to the trading sequence, type of D&S schedules, and the learning behavior of the agents involved in the market.

Conclusion

This study underscores the significant impact of learning mechanisms and market microstructures on market efficiency within double auction (DA) markets. The results reveal that Zero-Intelligence Plus agents with human learning mechanisms achieve higher market efficiency compared to basic Zero-Intelligence Plus and Zero-Intelligence

(ZI) agents. This finding is consistent with earlier research demonstrating that learning mechanisms enhance both allocative efficiency and price convergence.

Specifically, markets populated with ZIP agents exhibit strong convergence of transaction prices towards competitive equilibrium, particularly when using a Marshallian trading sequence. In contrast, while still efficient, markets with randomized trading sequences show slightly weaker convergence and efficiency. The learning mechanisms based on human-inspired rules further improve allocative efficiency, with box-shaped demand and supply schedules yielding the highest efficiency, followed by asymmetric and symmetric schedules. Trade volume efficiency remains high across different trading sequences, reflecting the robustness of market efficiency measures. However, the sensitivity of market outcomes to microstructure details—such as trading sequences and learning behaviours—indicates that market efficiency is not universal but rather contingent on these factors.

In summary, incorporating learning mechanisms and carefully selecting market microstructures are crucial for optimizing market performance. This study highlights the importance of aligning artificial market models with human trading behaviours to enhance market efficiency and provides insights into how different trading sequences and demand-supply schedules affect market outcomes.

References

Becker, G. S. (1962). Investment in human capital: A theoretical analysis. *Journal of political economy*, *70*(5, Part 2), 9-49.

Becker, G. S. (1962). Irrational behavior and economic theory. *Journal of political economy*, *70*(1), 1-13.

Chamberlin, E. H. (1948). An experimental imperfect market. *Journal of political economy*, *56*(2), 95-108.

Cliff, D., & Bruten, J. (1997). Zero is Not Enough: On The Lower Limit of Agent Intelligence For Continuous Double Auction Markets ,Ä.

Cliff, D. (2021). Parameterised-Response Zero-Intelligence (PRZI) Traders. *CoRR*.

Davis, D. D., & Holt, C. A. (1993). Experimental economics: Methods, problems, and promise. *E-studios Economics*, 179-212.

DeYoung, R. (1993). *Determinants of cost efficiencies in bank mergers* (Vol. 93, No. 1). Office of the Comptroller of the Currency.

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

Gode, D. K., & Sunder, S. (1997). What makes markets allocationally efficient?. *The quarterly journal of economics*, *112*(2), 603-630.

Gode, D. K., & Sunder, S. (1992). *Lower Bounds for Efficiency of Surplus Extraction in Double Auctions* (No. 1992-17). Carnegie Mellon University, Tepper School of Business.

Gode, D. K., & Sunder, S. (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.

Gode, D. K., & Sunder, S. (1994). Human and artificially intelligent traders in computer double auctions. *Computational organization theory*, 241-262.

LiCalzi, M., Milone, L., & Pellizzari, P. (2011). Allocative efficiency and traders' protection under zero intelligence behavior. *Computational Methods in Economic Dynamics*, 5-28.

Preist, C., & Tol, M. V. (1998). *Economic agents for automated trading*. Bristol: Hewlett Packard Laboratories.

Smith, V. L. (1962). An experimental study of competitive market behavior. *Journal of political economy*, *70*(2), 111-137.

Smith, V. L. (1976). Experimental economics: Induced value theory. *The American Economic Review*, *66*(2), 274-279.

Appendices

Appendix-I: Market Demand and Supply Schedules by Gode and Sunder (1993)

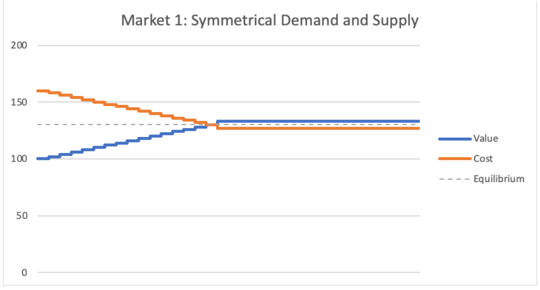

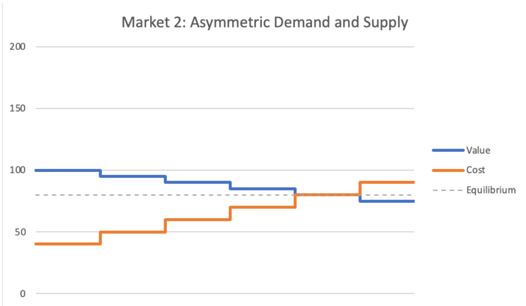

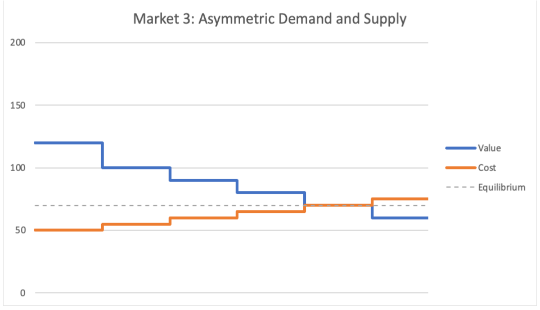

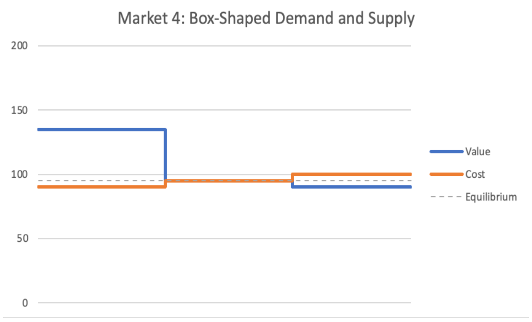

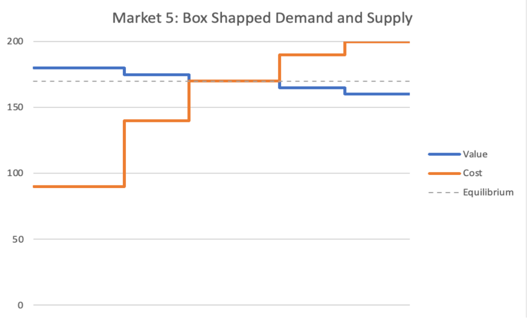

Appendix-II: Description of Zero-Intelligence Plus Agents by Cliff and Bruten (1997)

Similar to ZI agent populated market of Gode and Sunder (1993), the microstructure of ZIP populated agent also works under the Marshallian sequence. So, the whole market mechanism is same except the condition applied on the agents to learn from their surroundings. In ZIP populated market, agents are allowed to learn from the other traders and improve their shout prices accordingly.

1. Allocate limit prices (i.e. redemption values and costs to buyers and sellers) given already by the experimenter.
2. Allocate the profit margin for buyers and sellers that is in range of 5% and 35%. This profit margin is randomly selected from this range to start the experiment in period-1 and remains fixed for 10 periods. With this allocated profit margins the initial shout prices are calculated for all the traders.

3. Randomly any one trader (buyer or seller) is selected from the list of all traders. This process keeps repeating until at least 1 buyer and 1 seller is selected.
4. Transaction condition: If the shout price (bid) of buyer >= shout price (offer) of seller then the trade will happen between the two at shout price of trader who enters the market first.
5. When the trade will happen, both buyer and seller will become inactive as now they can learn about the market and improve their shout price accordingly but cannot involve in the transaction.

For the learning purpose, the

ZIP-agent use the market response and observable activity of other agent in the market to update its profit margin with each passing trade. ZI-agents adapt their price using the Widrow-Hoff rule each time they hear a shout from another trader in the market. This rule includes learning rate which influence how quickly an agent can learn. Cliff and Bruten defines this learning rate to be random number from a uniform distribution [0.1, 0.5]. So, every ZIP trader is learning positively with each passing trade. Then there is a random momentum value drawn from a uniform distribution [0.2, 0.8]. If the agent had a successful (unsuccessful) transaction in the previous round, it then increases (decreases) the profit margin.

Set of rule for ZIP by Cliff and Bruten (1997):

i. ZIP Seller has following set of rules.

   If the last shout is accepted at price q[11]:

   a. If last shout is a ask price, any seller who asked a price lower or equal to last transaction price (q), raises its profit margin in current period.
   b. If last shout is a bid price, any seller who asked a price higher or equal to transaction price lowers its profit margin in current period.

   Else:

   a. If last shout is an ask price, any seller who asked a price higher or equal to q lowers its profit margin.

ii. ZIP Buyer has the following set of rules.

   If the last shout is accept at price q:

   a. Any buyer who bid a price higher or equal to transaction price, raises its profit margin in the current time period.
   b. If the last shout is an offer, any buyer who asked a price lower than or equal to transaction price, lower its profit margin.

   Else:

---

[11] Here ‘q’ is the transaction price.

a. If last shout is a bid, any buyer who asked a price lower than or equal to q, lowers its profit margin in current period.

The profit margin is updated using the Widrow-Hoff with momentum learning rule:

$$\Delta_i(t) = \beta_i * (\tau_i(t) - p_i(t))$$

Here ‘ß’ is learning rate, ‘p’ is shouting price of agent, and ‘τ’ is target price calculated based on recent submitted shout price. After the trader observes any submitted shout it updates its profit margin ‘μ’ according to following equation.

$$\mu_i(t+1) = \frac{p_i(t) + \Gamma_i(t+1)}{l_i - 1}$$

with ‘l_i’ representing the trader’s limit price and ‘Γ_i(t+1)’ being calculated through following equation.

$$\Gamma_i(t+1) = \gamma_i(t) + (1 - \gamma_i(t)) * \Delta_i(t)$$

where ‘γ_i’ represents the so-called momentum coefficient. Finally, when the trader submit a shout at time ‘t’ it will calculate a shout price through equation

$$p_i(t) = l_i * (1 + \mu_i(t))$$

For ‘β’, each trader generates a random value from a uniform distribution U (0.1, 0.5). ‘γ’ is randomly generated for each trader from a uniform distribution U (0, 0.1).

‘μ’ is randomly generated for each trader from a uniform distribution U (5%, 35%).

‘c’ is perturbations that also randomly generates a value for each trader from a uniform distribution U (0, 0.05).

All participating traders have a limit price, in this scenario again randomly generated from a uniform distribution U (0.75, 3.5).

6. All the ZIP agents, whether active or inactive, learn about the market with the learning rules describes above. In the subsequent periods, the profit margin is updated but all the agents start with same shout prices as used in the last trading period. So, ZIP traders keep learning across the trading periods and this learning process does not stop even if they are inactive or not involved in trading as they can still observe what’s happening in market.