Price Formation of Continuous Double Auction Agents using Time as a Strategic Element

Marcel Neumann

June 28, 2011

Abstract

Specialist brokers of the New York Stock Exchange lost their dominance with the introduction of a hybrid system that allows machines to participate in the largest stock exchange of the world. The trading strategies encorporated by the machines became more competitive than their human counterparts. A trading strategy needs intelligent pricing as well as timing. This article enhanced the trading strategies "Zero-Intelligence" and "Zero-Intelligence-Plus" by new timing strategies and placed them in competition with their original complements in order to investigate how time can be incorporated as a strategic element. It is found that agents are outperformed when they do not integrate time as a strategic element at all. The second finding is that despite the already high efficiency of "Zero-Intelligence-Plus" agents, their timing strategy is not optimal. They run into the risk of trading too early and accepting suboptimal deals.

1 Introduction

Brokers of the New York Stock Exchange (NYSE) exchange about 1.6 billion shares worth $45 billion per day. It is the largest stock exchange in the world and a monumental example of real world trading. Specialist brokers base their decisions on experience and available information about the market they work in. Reliance in these human made decisions has decreased ever since the hybrid market of the NYSE took up employment in January 24, 2007. It allows a stock broker to chose in what way his order is to be executed. The options are either via an immediate fully automated electronic exchange or the traditional manual method through a specialist. Already in the first three months of 2007, 82% of the trading volume was automatically executed which shows the user's preference of speed in machine made decision over the judgement of human specialists [1].

These developments raise the question of the competitiveness of computer software agents in electronic markets. Software agents need to demonstrate their ability of making economically intelligent decisions in order to be classified as trustworthy. Gode & Sunder [4] report on experiments with so called "Zero-Intelligence" agents. These agents operate with no memory and limited market information. Still, they extract a high value of allocative efficiency when put solely in a market discipline. However, this type of agent trades far from the so called equilibrium price which would result in great losses when put in real world trading markets. Classical economic theory about markets predicts that the transaction price series approaches an equilibrium price value where the quantity demanded by consumers is equal to the quantity supplied by producers. Experiments conducted by Smith [8] show that even markets with small numbers of human participants possess that property. Hence, Cliff [2] enhanced the idea of the Zero-Intelligence agents which are referred to as "Zero-Intelligence-Plus" agents. These agents are able to adjust their offers in order to trade close to the equilibrium price.

All these trading strategies make the same assumption about the time of activeness in a market. It assumes that it is always better to place offers as soon as their price is being determined. But this does not necessarily have to be the optimal case for profit oriented agents. To give an example, laboratory experiments with Zero-Intelligence agents that are subject to a budget constraint show that the transaction price time series converges to an equilibrium price over time. So, a high bid of a buyer early in the period is quickly accepted by a seller who is pleased with high profit. It would have been more profitable for the buyer to wait and accept a low offer. Of course the same argument holds with a seller shouting a low price early in the period. It can be assumed that early trading is not necessarily the most profitable time of trading for either buyers or sellers.

A similar case can be stated for the trading strategies of Zero-Intelligence-Plus agents. It allows them to observe events in a market and change their subsequent offers accordingly. But since they place their offers constantly, they give away the possiblity of trading at even greater profit at a later date.

The purpose of this paper is to enhance the trading strategies of Zero-Intelligence as well as Zero-Intelligence-Plus agents with alternative timing strategies and compare their performance through a series of laboratory experiments with their original counterparts. Zero-Intelligence agents base their pricing on a random basis. This concept is transferred and similarly applied to time in order to investigate the performance of agents that do not use time strategically. Zero-Intelligence-Plus agents are profit oriented. Therefore, they are enhanced with a sophisticated timing strategy to investigate the performance of agents that try to use time as a strategic element.

The following research questions are investigated:

1. How does the performance of Zero-Intelligence agents change when their deployment of timing is relaxed?
2. How can Zero-Intelligence-Plus agents be more competitve by using time as a strategic element?

The outline of this paper is as follows. Section 2 introduces the reader to auctions, their mechanisms and related work on the topic of economically motivated software agents. Section 3 describes the trading strategies to be investigated together with their enhancements. Section 4 is about the experiments conducted. It gives a description of the simulation environment, performance measures and the specific experimental setups together with their results. The last section gives the conclusions.

2 Background

An auction is a mean to trade goods and services. Participants take the roles of traders which is a general term for someone or something that is either a buyer or a seller. Their intention is to buy or sell a good or service respectively. Hence each member of either party wants to find a match from the other party with whom it is willing to trade. These matches are found during the process of an auction and depend on its mechanism.

The mechanism of an auction is a set of rules. They govern the set of available actions for the traders and the set of information they will receive about the market at each point in time. To give an example, auctions can be set up to have only one seller and run for a certain amount of time but only allow buyers to make bids. So, buyers would be allowed to make a bid at any point in time for as long as the auction lasts. This makes up their set of actions. The seller would usually only be allowed to set a starting price. The information set for all traders could consist of the current highest bid. Another important aspect of the auction mechanism is the process of market clearing. Market clearing may occur one or many times during an auction. When it occurs, it determines the matches of traders that carry out a transaction. The number of matches can be zero or more. The mechanism also states the procedure for unmatched offers. The duration of an auction can be either a fixed time after which no more activity is allowed or subject to a deadline of inactive interval. With the second type of duration the auction is closed after a pre-specified interval of time has passed without any activity. The variety of auction mechanisms is high and only limited by the ingenuity of their executors. The interested reader finds a taxonomy of auctions in Parsons et al. [5].

A common use for auctions is the facilitation of exchange of homogeneous commodities. These kind of goods are equal in characteristic trait which makes their source irrelevant. Hence traders only consider its price as being important. Established examples are stocks, electromagnetic spectrum and oil. The kind of auction mechanism often used to trade these goods is the continuous double auction.

The continuous double auction mechanism consist of one or more periods. During a period traders may make offers at any point in time. An offer is a general term for something that is either a bid or an ask. Buyers bid prices at which they are willing to buy a commodity while sellers ask prices at which they are willing to sell a commodity. Traders may also accept offers. Market clearing depends on whether a trade is determined by the institution or by the traders. Institutional-trade-determination is enforced when a bid and an ask cross and can therefore be matched. The price is then set by the institution which is usually between the prices of the respective offers. Trades can also be determined when a trader accepted an offer. The price is then set to the value of the accepted offer.

Supply and demand are two functions that map quantity to price for a set of sellers and buyers. In other words, for each price there is a quantity of commodity that buyers are willing to buy and sellers are willing to sell. If the quantity demanded is greater than the quantity supplied, the price for one unit of commodity rises. This in turn reduces the quantity demanded because buyers are not willing to pay that price or increase the quantity supplied because sellers are willing to sell more units at that price. Similarly, if the quantity supplied is greater than the quantity demanded, the price for one unit of commodity falls. This in turn increases the quantity demanded because buyers are willing to buy more units or reduce the quantity supplied because sellers are not willing to sell for such cheap prices. Taken this into account, classical economic theory predicts that the price for one unit of commodity will settle at an equilibrium price with a respective equilibrium quantity.

Smith [8] conducted a series of experiments with human traders operating in a market under the continuous double auction mechanism. Each participant was assigned a role of either buyer or seller beforehand. The roles empowered the traders to either buy or sell one unit of a fictious commodity. They also received a number which was only known to them. Smith called it the reservation price. It will now be referred to as the private value. It indicates the value of the commodity for a buyer and the cost of the commodity for a seller. Traders that received a value beyond the equilibrium price will be referred to as "extra-marginal" traders. All other traders will be referred to as "intra-marginal" traders. Buyers were not allowed to buy a unit of commodity at a price higher than their private value. Similarly sellers were not allowed to sell a unit of commodity at a price lower than their private value. However, they should rather try to trade a unit of commodity at a price equal to their private value than not trade at all. The performance of the traders was measured by the absolute value of the difference of the private value and the transaction price if a transaction occurred. As a result of his experiments Smith reported that even markets with a small number of traders have a tendency to trade at prices that converge towards the equilibrium price.

Gode and Sunder [4] extended Smith's work and introduced two types of "Zero-Intelligence" computer software agents. The set-ups of their experiments were similar to Smith's set-up. In these experiments both types of agents submit randomly generated offer prices with the first type being unconstrained within a given price range and the second type being constrained to not make a loss according to its private value. They also conducted experiments with human traders and compared the allocative efficiency for all three types of traders. The allocative efficiency is a measure of market performance which compares the actual profit made by all traders with the theoretical possible profit. Human traders came off best with an efficiency of around 100%. Unconstrained "Zero-Intelligence" agents traded from 48.8% to 90% efficiency. Notably was the impact of the budget constraint on "Zero-Intelligence" traders. These agents were able to achieve an efficiency value in the range of 96% and 99.9% and therefore coming close to the performance of humans. They were even able to trade close to the equilibrium price. The conclusion by Gode & Sunder was that trading strategies do not need to be complex to accomplish efficient trading as long as traders care to not trade at a loss.

Cliff [2] showed that the results by Gode & Sunder were a consequence of their experimental set-up. Their supply and demand functions possess few extra-marginal traders and are therefore designed to result in high allocative efficiency. He also showed that the trading around the equilibrium price by constrained "Zero-Intelligence" traders was a consequence of their private values. They placed a higher probability to offer prices around the equilibrium price. In addition, Cliff introduced a new type of trading strategy going by the name of "Zero-Intelligence-Plus". This trader is able to observe offers and transactions by other traders and accordingly adjust the price of its next offer. His experiments show that it is able to trade close to the equilibrium price after only a few number of periods.

Eventually, Das et al. [3] conducted a series of experiments with agents employing a modified version of the "Zero-Intelligence-Plus" strategy and human counterparts interacting in the same continuous double auction. They report that the computer software agents clearly outperformed the human traders. The agents achieved an average efficiency of more than 100% by exploiting humans while they performed in the range of 92% and 96%.

3 Trading strategies

It follows a description of the basic functionality of the trading strategies that will be used in the experiments documented in section 4. There will be a variety of parameters introduced of which the exact values used in the experiments will be found in sections 4.3 and 4.4 which provide the detailed experimental set-ups.

The purpose of a trading strategy is to determine offers that a trader then submits on a market. Apart from the identity of the trader that calculated it, an offer needs to consitute a price and a time at which the offer is to be submitted. Thus, a trading strategy exists of a pricing and a timing strategy.

The strategies described are "Zero-Intelligence" by Gode & Sunder [4] and "Zero-Intelligence-Plus" by Cliff [2]. Two follow-up strategies are proposed that adopt the pricing strategies of their original complements but introduce new timing strategies. They will be referred to as "Zero-Intelligence Time" and "Zero-Intelligence-Plus Time".

It is assumed that each trader i has a list of limit prices for each good to be traded. A limit price $\lambda_{i,j}$ determines the valuation of a buyer and the cost for a seller of a unit j.

3.1 Zero-Intelligence

Gode & Sunder [4] introduced two trading strategies that they incorporated into so called Zero-Intelligence (ZI) traders. Traders of the first version submit random offers independently, identically, and uniformly distributed over the range of feasible trading prices from 1 to the limit trading price enforced by the market mechanism. This type is referred to as "ZI unconstrained" traders (ZIU). Traders of the second version were subject to a budget constraint. It prevents the traders from engaging in loss making transactions. Thus, buyers submit the uniform random bids between 1 and the valuation of the unit in question. Sellers submit the uniform random asks between the cost of the unit in question and the upper limit trading price of the market. This type is referred to as "ZI with constraint" (ZIC).

3.2 Zero-Intelligence Time

Zero-Intelligence traders are considered to have no rationality. They do not remember past activity in the market and can therefore not learn from experience. Indeed, they do not even try to maximize any kind of profit. However, their timing strategy tells them to place their orders constantly and as soon as possible.

In order to study the behavior of truly "Zero-Intelligence" traders, this practice of timing needs to be relaxed. Hence, "Zero-Intelligence-Time" (ZIT) traders incorporate the pricing strategy of ZI traders but their timing will be randomly distributed over the least remaining time of the auction. The least remaining time of an auction depends on its ending conditions. Laboratory experiments are usually set up to give traders sufficient time to trade. They can have a fixed time span for which the least remaining time is easily determined. The other case is an auction with a deadline of inactive interval. So the least remaining time is at least the size of this interval starting from the last activity in the market. Hence, traders who are placing offers constantly as well as traders who are placing offers over the least remaining time of the auction may do that until they choose not to place offers anymore.

3.3 Zero-Intelligence Plus

Cliff [2] proposed a trading strategy that is carried out by "Zero-Intelligence-Plus" (ZIP) traders. A ZIP trader i maintains a profit margin $\mu_{i,j}$ for each unit j on its list of limit prices. At any time t the offer price $p_{i,j}$ for trader i is then given by $p_{i,j}(t) = \lambda_{i,j}(1 + \mu_{i,j}(t))$ with $\mu_{i,j}(t) \in [0, \infty), \forall t$ for sellers and $\mu_{i,j}(t) \in [-1, 0], \forall t$ for buyers. Hence, the profit margin value for a trader i determines the estimated surplus at time t for unit j. Sellers increase their profit margin by increasing $\mu_{i,j}$ while buyers increase their profit margin by decreasing $\mu_{i,j}$.

ZIP traders are able to observe submitted offers by other traders and the transactions that might result from them. This enables them to adjust their profit margins according to the events in the market. In general, when a trader observes a transaction that reveals the possibility to submit offers with higher profit margins and still secure a deal then the profit margin is increased. To give an example, a seller who indicates to trade for a price of 50 price units may observe a transaction by two other traders at a price of 100 price units. This reveals that the seller could ask more in his next offer than the initial 50 price units and still transact. On the other hand, when a trader observes a transaction that shows that other traders are more competitive than itself then the profit margin is decreased. At last, when a bid stays unmatched buyers adjust their profit margins in the direction of the unmatched bid in order to stay competitive. Similarly, when an ask stays unmatched sellers adjust their profit margins in the direction of the unmatched ask in order to stay competitive. Traders adjust their profit margins for all their limit prices even for units that have already been traded. This information can then be used in subsequent trading periods. However, profit margins are not decreased for units that have already been traded. The underlying assumption is that units already traded may be traded at the same or an even better price in upcoming periods.

The magnitude of an adjustment is based on the Widrow-Hoff "delta rule". In general, the price of the next offer $p_{i,j}(t+1)$ is determined by the previous offer and a change $\Delta_{i,j}(t)$ which gives $p_{i,j}(t+1) = p_{i,j}(t) + \Delta_{i,j}(t)$. The change $\Delta_{i,j}$ depends on a target price $\tau_{i,j}$ and gives asymptotic convergence in the direction of the target at a speed determined by a learning parameter β_i which gives $\Delta_{i,j} = \beta_i(\tau_{i,j}(t) - p_{i,j}(t))$. The target price depends on the last observed event in the market that lead to an adjustment of the profit margin. This is either a transaction or an unmatched offer at price $q(t)$. The target price is not being set equal to $q(t)$ right away in order to leave the possibility to try the market by targeting at an even more profitable price. Hence, two randomly generated coefficients $R_{i,j}$ and $A_{i,j}$ are introduced. They are generated every time a trader wishes to adjust its profit margin. The target price $\tau_{i,j}$ is then determined by $\tau_{i,j} = R_{i,j}(t)q(t) + A_{i,j}(t)$ with $R_{i,j} > 1.0, A_{i,j} > 0$ for price increases and $0 < R_{i,j} < 1, A_{i,j} < 0$ for price decreases. At last, every change in profit margin is subject to a momentum mechanism in order to account for past changes and not run into overhasty decisions. The momentum coefficient γ_i determines how much past changes are considered by the current change in profit margin. The momentum equation is given by $\Gamma_{i,j}(t+1) = \gamma_i \Gamma_{i,j}(t) + (1 - \gamma_i)\Delta_{i,j}(t)$ with $\Gamma_{i,j}(0) = 0$. Eventually, the new profit margin is given by $\mu_{i,j}(t+1) = (p_{i,j}(t) + \Gamma_{i,j}(t))/\lambda_{i,j} - 1$.

3.4 Zero-Intelligence-Plus Time

ZIP traders use events in the market as clues to adjust their profit margins. However, they submit their offers constantly without considering possible trends that could indicate an even larger profit margin in case of withholding an offer and submitting it at a later date. Therefore, "Zero-Intelligence-Plus Time" (ZIPT) traders incorporate the pricing strategy of ZIP traders while introducing a timing strategy based on clues given by the so called bid-ask-spread of a market.

Assuming that buyers and sellers start off with submitting ridiculous low and high prices respectively in order to try the market for high profitable trades, there will not be a transaction taking place right away. This will lead to decreasing ask prices as well as increasing bid prices. Eventually, the highest open bid gets ahead of the lowest open ask and a transaction is carried out. At any time t the difference between the price of the highest open bid and the lowest open ask is called the bid-ask-spread.

The size of the spread is an indicator of how much buyers and sellers need to approach each other in order to carry out a transaction. A buyer for instance may assume that sellers will reduce their asks in order to be more competitive to other sellers. The same holds for sellers. They may assume that buyers will raise their bids in order to be more competitive to other buyers. On the other hand, a small spread does not necessarily mean that trading has reached the equilibrium price. Traders may continue testing the market for higher profit.

Taking this into account, ZIPT traders employ a timing strategy based on the size of the spread. If the size of the spread is large, traders withhold their offers based on the assumption that transaction prices will change in the near future. Traders start submitting offers when the size of the spread falls below a pre-defined threshold. At last, if the size of the spread stays constant for an interval of time, ZIPT traders start testing the market themselves by submitting one offer before checking for the size of the spread again.

4 Experiments

The experiments are conducted using a discrete-event simulation. Simulation time is measured in time units. One time unit is assumed to be the amount of time an agent needs to determine his next offer. The market mechanism is subject to the continuous double auction market rules. This section introduces the simulation environment used for the experiments in section 4.1. The agents are evaluated by a set of performance measures which are described in section 4.2. The explicit setups for the experiments are given in sections 4.3 and 4.4 together with their results.

4.1 Simulation environment

Participants of the experiments are software agents, each employing its own pricing algorithm. Each agent takes a role of either a buyer or a seller. A simulation process runs for a preset number of periods. The deadline for one period is subject to an inactivity interval of 20 time units. Thus, the simulation process closes a trading period when there has not been any activity of any trader for 20 time units. This is meant to give each trader enough time to trade as much units as they are willing and able to. A period is also closed when it is detected that no pair of traders is able to close a deal anymore due to the constraints put on them by their pricing strategies.

Each agent receives a list of limit prices at the beginning of each period that is only known to it. Each item on a list describes the valuation for a buyer or the cost for a seller of one unit of an undefined homogeneous commodity that they are willing to buy or sell respectively during the auction.

The simulation continuously polls offers of each trader. An offer includes a price at which the trader is willing to trade together with a time at which the trader wishes the offer to be activated.

Offers are subject to several rules. All offers are for a single unit and a trader my only have one open offer at a time. Hence, when placing offers, traders consider the limit price for the ith unit before considering the $(i+1)$th unit. Traders that already have open offers may replace them by a new one. Offers are also subject to the NYSE spread improvement rule which states that a new offer has to be superior to the best existing offer of the same type. In other words, the price of a new bid has to be higher than the price of the highest open bid while the price of a new ask has to be lower than the price of the lowest open ask.

The simulation clock is advanced to the time of occurrence of the most imminent set of offers or by a fixed-increment of one time unit if no such set exists. All future offers are now discarded. This enables the traders to re-determine their future offers based on possible events triggered by current offers. The set of current offers are now subject to the clearing process of the auction mechanism. They are processed in a random sequence in order to emulate asychronous activity of the traders as well as stochastic behavior in the connection from the trader's terminal to the auction system. A new incoming offer is attempted to being matched with the best open offer. Thus, in the case of an incoming bid, a transaction occurs when the price of the new bid is higher than the price of the lowest open ask. In the case of an incoming ask, a transaction occurs if the price of a the ask is lower than the price of the highest exisiting bid. The transaction price is set to the price of the earlier offer. New offers that can not be matched are appended to the list of open offers.

4.2 Performance measures

This section covers how the performance of either auction or traders is measured in the experiments. Two measures quantify the performance of individual traders. These are the actual profit of a trader and the coefficient of convergence introduced by Smith [8]. Both can als be used to quantify the performance of a group of traders. The allocative efficiency on the other hand measures the performance of an entire market.

The actual profit of a trader depends on whether his role is that of a buyer or a seller. It considers all his successful transactions and measures them against the private value $\lambda_{i,j}$ of trader i and unit j in question. The actual profit of a buyer and a seller is given in equation 1 and 2 respectively where p_j denotes the transaction price of unit j.

$$pr_i = \sum_{j \in bought} (\lambda_{i,j} - p_j), i \in Buyers \quad (1)$$

$$pr_i = \sum_{j \in sold} (p_j - \lambda_{i,j}), i \in Sellers \quad (2)$$

The actual profit for a group of traders is defined by the sum of the actual profits of each group member.

The coefficient of convergence α introduced by Smith [8] measures how close an individual trader or a group of traders trades to the theoretical equilibrium. It is given in equation 3 and defined as the standard deviation of the actual trade prices p_i from the equilibrium price p_0 as a percentage of it.

$$\alpha = \frac{100}{p_0} \sqrt{\frac{1}{n} \sum_{i=1}^{n} (p_i - p_0)^2} \quad (3)$$

The allocative efficiency e is a measure of performance of an entire market. It is given by equation 4 and defined as the ratio of total actual profit and theoretical profit. Total actual profit is the sum of profits made by each trader while the theoretical profit is the sum of theoretical buyer's profit tbp and theoretical seller's profit tsp (see Figure 1).

$$e = \frac{\sum\limits_{i \in traders} pr_i}{tbp + tsp} \quad (4)$$

4.3 ZI vs. ZIT

It follows the set-up for the experiments involving ZI and ZIT agents. The series of experiments is sub-divided into two sets where the first one involves homogeneous populations of traders while the second one is dedicated to heterogeneous populations of traders.

The traders are given a list of limit prices before the beginning of each experiment. Each list consists of 10 identical values. Any experiment is run with three different distributions of limit prices. The supply and demand functions that account for the distributions of

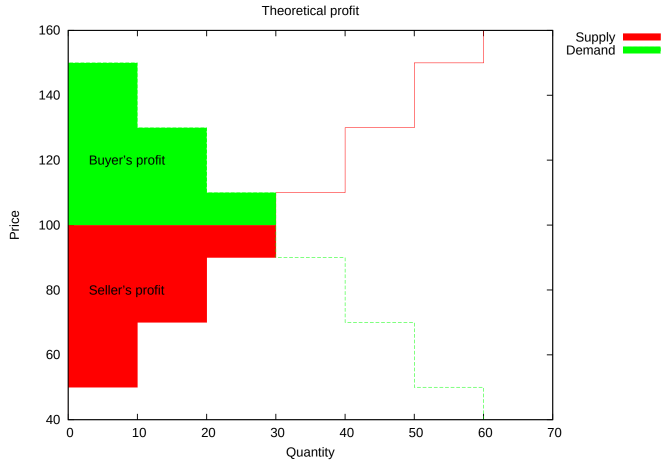

Figure 1: Theoretical profit of buyers and sellers

limit prices are shown in figures 2, 3 and 4. Market 1 and 2 are adopted from the experiments undertaken by Gode & Sunder. They represent markets where the theoretical profit of buyers and sellers is unequal. In addition, a symmetric distribution of limit prices is introduced which will be referred to as market B. The price range within the traders may submit offers is the same for all three markets and set to [0 : 200].

Gode & Sunder [4] made a choice to simplify their implementation which is the deletion of the list of open offers after a transaction has been taken place. This policy is adopted for all experiments involving ZI and ZIT agents in order to compare the results.

The first set of experiments involves homogeneous populations of either ZIU, ZIC, ZIUT and ZICT traders. A population exists of six buyers and six sellers. Any market is run for six periods and the experiments are repeated 1000 times.

Gode & Sunder [4] report that ZIU traders are able to achieve an allocative efficiency of 90% for market 1 and 2. ZIC traders are even able to raise this value to 99.9% for market 1 and 99.2% for market 2. It is expected that the two types of traders will do equally good in this set-up of experiments. The assumption is that ZIT traders will achieve lower values of allocative efficiency because of a higher risk of extra-marginal units to be traded.

The second set of experiments is concerned with heterogeneous populations of ZI and ZIT traders. ZIU traders are confronted with ZIUT traders while ZIC traders are confronted with ZICT traders. A population exists of twelve buyers and twelve sellers, equally divided into the respective type of traders. Any market is run for six periods and the experiments are repeated 1000 times.

The performance of the traders will be measured by the average actual profits earned during any one period.

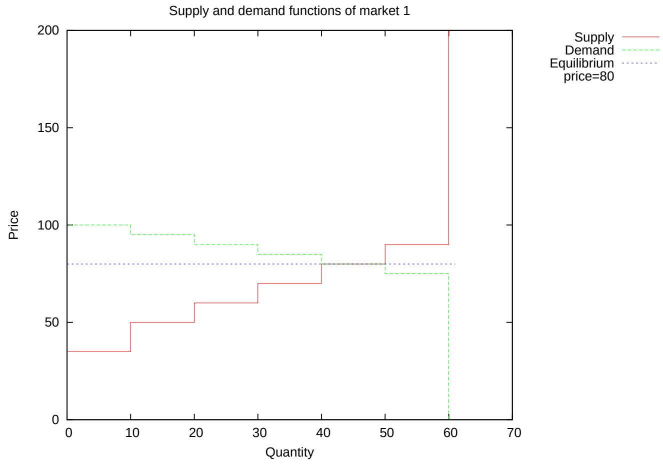

Figure 2: Sellers' theoretical profit = 1050, Buyer's theoretical profit = 500

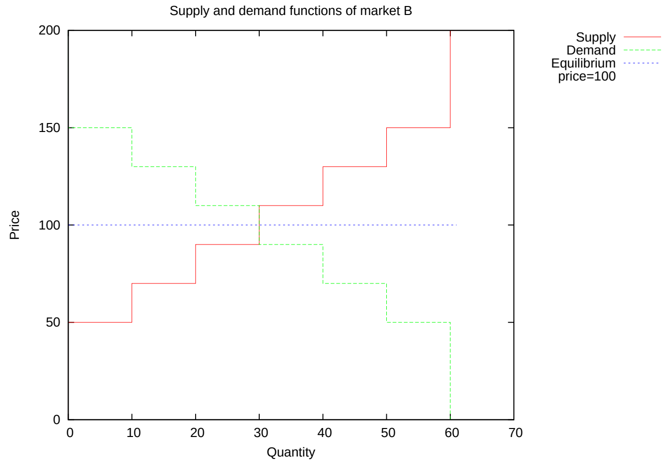

Figure 4: Seller's theoretical profit = 900, Buyer's theoretical profit = 900

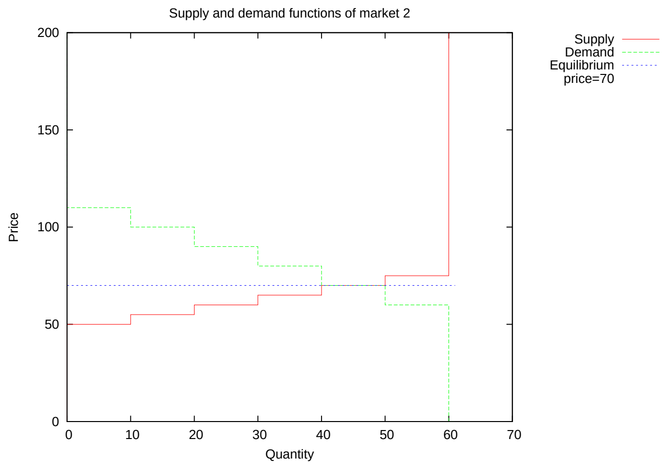

Figure 3: Sellers' theoretical profit = 500, Buyers' theoretical profit = 1000

Figure 5: Transaction price time series for market 1 with 6 buyers and 6 sellers of ZIU traders

Results

The first set of experiments was concerned with homogeneous populations of ZI and ZIT traders.

Figure 5 shows the transaction price time series in market 1 with a population of ZIU traders. The shape of the graph is identical to the one created by Gode & Sunder [4]. The pattern can be stipulated as being random. This is also the case for market 2 and B.

Figure 6 shows the transaction price time series in market 1 with a population of ZIC traders. Again, the shape of the graph is identical to the one created by Gode & Sunder [4]. The price series is less volatile than the price series of the ZIU traders. However, there is no kind of learning observable from period to period. Also, the price series converges towards the equilibrium price within each period. This was explained by Gode & Sunder [4] as follows. The opportunity sets for the traders narrows down as the auction progresses. Offers for units that represent the left end of a supply and demand function are more likely to be accepted because the expected value for their prices are more competitive. After these units have been traded the left end of the supply and demand function shifts to the right which yields to the resulting shape. The observed pattern is similar for market 2 and B.

Figure 7 shows the transaction price time series in market 1 with a population of ZIUT traders. The shape of the graph is similar to Figure 5. A random pattern can be observed with no sign of learning and periods take longer to finish. This is because ZIUT traders do not submit offers as soon as possible anymore but randomly distributed over the least remaining time of the period.

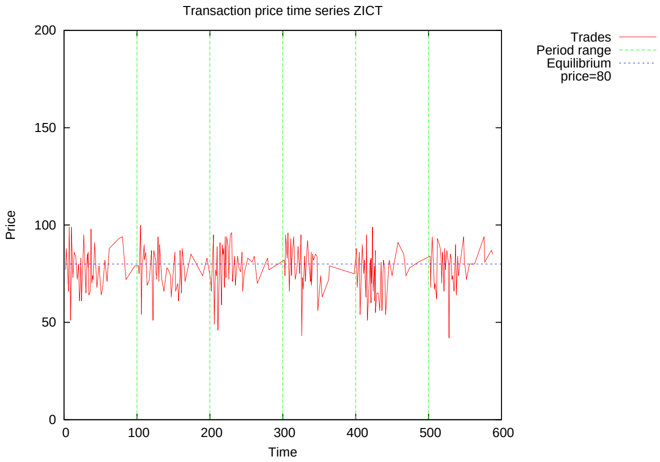

Figure 6: Transaction price time series for market 1 with 6 buyers and 6 sellers of ZIC traders

The pattern is similar for market 2 and B.

Figure 7: Transaction price time series for market 1 with 6 buyers and 6 sellers of ZIUT traders

Figure 8 shows the transaction price time series in market 1 with a population of ZICT traders. The price series looks similar than the one in Figure 6. It still converges to the equilibrium price but not as gradually as it was the case for ZIC traders. Due to the deadline of inactive interval, traders had sufficient time to trade. Two traders with limit prices equal to the equilibirum price needed more time to agree on a price than any other pair of traders. This can be explained by the drop of the list of offers after every transaction. The traders in question had to generate the only price that they could agree on over and over again for all their units to be traded which lead to rather long period lengths.

Parsons et al. [6] noted that the allocative efficiency depends only on the private value of the traders and

Figure 8: Transaction price time series for market 1 with 6 buyers and 6 sellers of ZICT traders

not on the prices at which they trade. It is explained as follows. A transaction at price p between a buyer B with private value λ_B and a seller S with private value λ_S contributes $(\lambda_B - p) + (p - \lambda_S)$ to the numerator of the equation for allocative efficiency. This reduces to $\lambda_B - \lambda_S$. Thus, high allocative efficiency is achieved when buyers that have high private values are matched with sellers that have low private values. Also, every trader on the extra-marginal side of the supply and demand functions is a threat to high allocative efficiency.

Figure 9 shows the sample mean and the sample standard deviation of the allocative efficiency for the mentioned experiments with ZI traders over 1000 runs for each market. Market 1 and 2 were taken from the experiments by Gode & Sunder [4]. Their designs imply efficiency of at least 90 %. This occurs when all available units are being traded which is the case for both types of unconstrained traders. The intra-marginal and the extra-marginal side of the symmetric market B are of equal size. Hence the allocative efficiency of the unconstrained traders in it is 0.

Experiments with ZIC traders did not allow many extra-marginal units to be traded. This is a result of traders constantly submitting offers. Every trader has an open offer at any point in time but as mentioned above offers for units that represent the left end of a supply and demand curve are more competitive. After these units have been traded, the competitive left end shifts to the right. Units representing the right end of a supply and demand curve have the lowest chance of being traded. This was different for the experiments with ZICT traders. The chance of an extra-marginal unit to be traded is higher because it exists the possibility that more competitive traders did not submit an offer yet. This explains the lower allocative efficiency for

ZICT traders in contrast to ZIC traders.

The appropriate α values of the coefficient of convergence shown in figure 10 conform to the findings above. ZIU and ZIUT traders operate far from the equilibrium price. However, the α value for ZIUT traders was up to 5% lower than the α value for ZIU traders. ZIC and ZICT traders on the other hand trade closer to the equilibirum price but again the α value for ZICT traders is lower than the α value for ZIC traders. This again is a result of transactions involving extra-marginal units close to the equilibrium price of which the possibility of occurrence was higher for homogeneous populations of ZICT traders than populations of ZIC traders.

Traders	Market 1 $\bar{x}$ (s)	Market 2 $\bar{x}$ (s)	Market B $\bar{x}$ (s)
ZIU	90.32 ($\pm$ 0.00)	90.00 ($\pm$ 0.00)	0.00 ($\pm$ 0.00)
ZIC	99.17 ($\pm$ 0.23)	98.96 ($\pm$ 0.27)	97.21 ($\pm$ 0.61)
ZIU-T	90.32 ($\pm$ 0.00)	90.00 ($\pm$ 0.00)	0.00 ($\pm$ 0.00)
ZIC-T	99.12 ($\pm$ 0.24)	98.91 ($\pm$ 0.28)	96.95 ($\pm$ 0.67)

Figure 9: Sample mean and sample standard deviation of the efficiency of ZI traders, sample size = 1000

Traders	Market 1 $\bar{x}$ (s)	Market 2 $\bar{x}$ (s)	Market B $\bar{x}$ (s)
ZIU	80.47 ($\pm$ 5.29)	97.34 ($\pm$ 6.41)	60.91 ($\pm$ 3.98)
ZIC	13.37 ($\pm$ 1.77)	17.21 ($\pm$ 1.81)	18.28 ($\pm$ 2.22)
ZIUT	75.44 ($\pm$ 4.80)	92.13 ($\pm$ 6.15)	57.08 ($\pm$ 3.32)
ZICT	12.83 ($\pm$ 1.75)	17.00 ($\pm$ 1.82)	17.91 ($\pm$ 2.18)

Figure 10: Sample mean and sample standard deviation of the coefficient of convergence of ZI traders, sample size = 1000

The second set of experiments was concerned with the comparison of ZI and ZIT traders. Most interesting are the results for the comparison of ZIC and ZICT traders. Figure 11 shows the mean actual profit achieved by the respective traders per period in market B. The mean was calculated over 1000 runs. ZIC traders are able to achieve a mean actual profit of above 920 price units while ZICT traders are only able to achieve a mean actual profit of around 800 price units. Since the supply and demand functions of market B are symmetric the following accounts for buyers and sellers without loss of generality. Given the equilibrium quantity of 60, ZIC buyers and ZICT buyers for instance should be able to trade 30 units each for a theoretical actual profit of 900 price units. The reason why ZIC traders performed above that value while ZICT traders below of it is that ZIC traders were quicker. ZIC traders were able to trade all their intra-marginal units while even some extra-marginal units could be traded with ZICT traders which results in a profit value of above 900. ZICT traders on the other hand were not able to make as many profitable transactions because of less appropriate traders at the time that they were active on the market. The same behavior was observed in market 1 and 2.

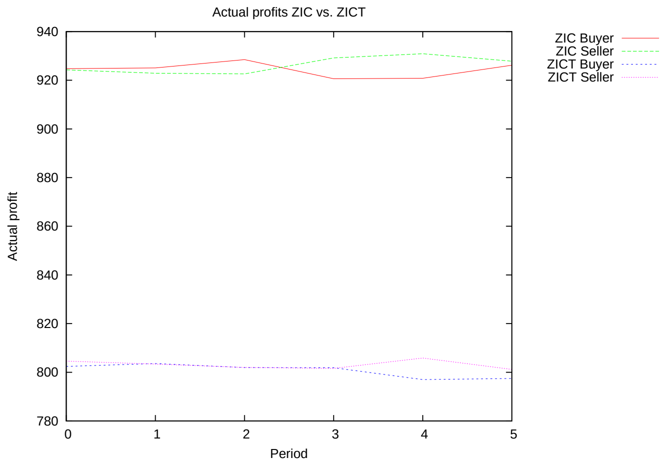

Figure 11: Actual profits earned in market B by a heterogeneous population of ZIC and ZICT traders

4.4 ZIP vs ZIPT

It follows the set-up for the experiments involving ZIP and ZIPT agents. The outline of the experiments is similar to the outline of the experiments in section 4.3. The series of experiments is sub-divided into two sets where the first one involves homogeneous populations of traders while the second one is dedicated to heterogeneous populations of traders.

The traders are given a list of limit prices before the beginning of each experiment. Each list consists of 10 identical values. Any experiment is run with three different distributions of limit prices. The supply and demand functions that account for the distributions of limit prices are shown in figures 2, 3 and 12. Market 1 and 2 are adopted from the experiments undertaken by Gode & Sunder [4]. They represent markets where the theoretical profit of buyers and sellers is unequal. Market A is adopted from the experiments undertaken by Cliff [2] and represents a symmetric distribution of limit prices. The price range for offers is $[0 : 400]$ for this market.

Section 3 introduced a set of parameters used in the pricing strategy for ZIP and ZIPT traders. During all experiments R is uniformly distributed over $[1.0, 1.05]$ for price increases and over $[0.95, 1.0]$ for price decreases. These values are suggested by Cliff [2]. The values for A are changed due to a larger scale for the price units. A is uniformly distributed over $[0.0, 5.0]$ for price increases and over $[-5.0, 0.0]$ for price decreases. The learning rate β is set to 0.3 and the momentum coefficient γ is set to

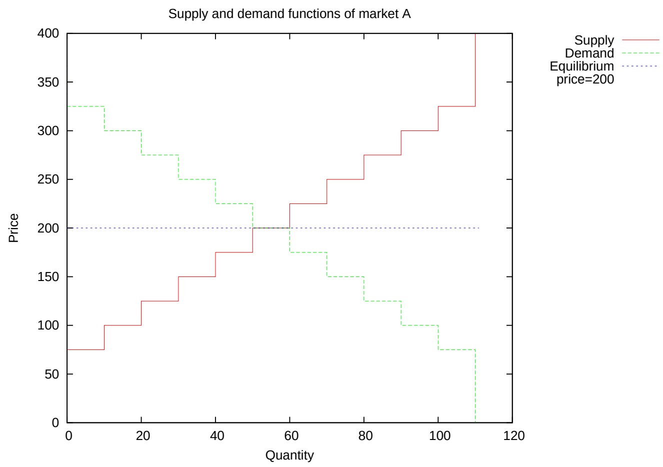

Figure 12: Sellers' theoretical profit = 3750, Buyers' theoretical profit = 3750

0.05 for all traders as suggested by Preist et al. [7]. At last, buyers start with profit margins that give a price of 0 as their first offer. Sellers on the other hand start with profit margins that give the limit price of the market as their first offer.

ZIPT traders start submitting offers when the size of the bid-ask-spreak has fallen below 5.0 price units. If the bid-ask-spread has not changed for 5 time units ZIPT traders start testing the market themselves before considering the size of the bid-ask-spread again. Also, ZIP traders adjust their profit margins in the direction of the best best open respective offer when there has not been a trade for 5 time units. This policy was introduced by Das et al. [3] and is necessary for ZIP traders to be applied in continuous double auctions.

Unlike the set-up in section 4.3 the list of offers persists after the occurrence of a transaction. This is meant to provide a more realistic market place.

The first set of experiments is concerned with homogeneous populations of ZIP and ZIPT traders. Cliff [2] reports that transaction prices of ZIP traders converge to the equilibrium price during early periods and stays stable during later periods. The performance of ZIP traders in this set-up should be equally good. Transaction prices of ZIPT traders should also be able to converge to the equilibrium price since they employ the same pricing strategy as their original complement. However, clues for the profit margin adjustments of the traders are given less frequently which might lead to longer period lengths.

For the second set of experiments ZIP traders are confronted with ZIPT traders in all three markets. The performance of either trader is measured by their actual profit earned during any one period.

Results

The first set of experiment was concerned with homogeneous populations of ZIP and ZIPT traders.

Figure 13 shows the transaction price time series in market A with a population of ZIP traders. As can be seen the transaction prices scallop around the equilibrium price of 200 price units during the first four periods. However, the price series converges towards the equilibrium price where it stays stable for the last two periods. Noteworthy are the prices at which the first transactions are settled. They are below the equilibrium price even tough sellers and buyers started off by submitting offers equally far away from the equilibrium price. This can be explained by the way the target price τ is determined by a trader in the computation of a subsequent offer. Both sellers and buyers do not observe any transaction at the beginning of period one. Therefore, a seller for instance tries to be more competitive than the lowest open ask. Among other things the target price is taken relatively to it by a coefficient R. The same is assumed for buyers. They try to be more competitive than the highest open bid. But since the first bids are rather low while the first asks are rather high the asks decrease faster than the bids increase resulting in first transactions occurring below the equilibirum price.

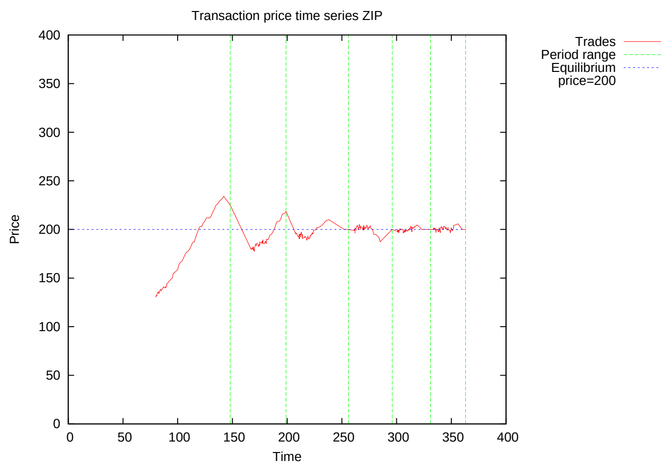

Figure 13: Transaction price time series for market A with 6 buyers and 6 sellers of ZIP traders

Figure 14 shows the transaction price time series in market A with a homogeneous population of ZIPT traders. The shape of the graph looks similar to figure 13. The transaction price time series scallops around the equilibrium price for the first four periods and stabilizes in period five. Noteworthy is that it takes longer for the first transaction to be carried out in contrast to figure 13. This is explained by the fact that ZIPT traders submit less offers and therefore the bid-ask-spread shrinks at a lower pace.

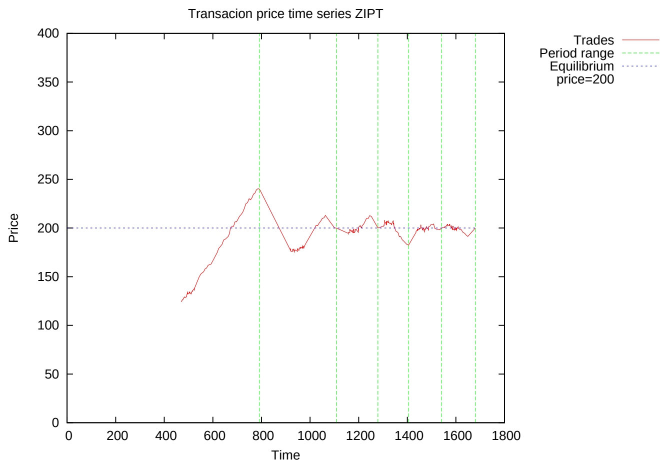

Figure 14: Transaction price time series for market A with 6 buyers and 6 sellers of ZIPT traders

The mean allocative efficiency of markets with homogeneous populations of either ZIP or ZIPT traders are around 99% and summarized in figure 15. The efficiency value for ZIP traders conforms with the findings by Cliff [2].

Traders	Market 1 $\bar{x}$ (s)	Market 2 $\bar{x}$ (s)	Market A $\bar{x}$ (s)
ZIP	99.79 (± 0.24)	99.86 (± 0.19)	99.86 (± 0.30)
ZIPT	99.81 (± 0.22)	99.83 (± 0.22)	98.84 (± 0.31)

Figure 15: Sample mean and sample standard deviation of the efficiency of ZIP and ZIPT traders, sample size = 1000

On average both ZIP and ZIPT traders converge to the equilibrium equally fast for all three investigated markets. Figure 16 shows the the coefficients of convergence for each period for both types of traders operating in market A. The results are the sample mean of α values over 1000 runs.

The second set of experiments was concerned with heterogeneous populations of ZIP and ZIPT traders on all three markets. The findings are explained with market A being the paradigm and then taken over to the other two markets.

Figure 17 shows the transaction price time series for market A. As explained above asks decrease faster than bids increase which results in the first transactions to be taken place below the equilibrium price. At that time ZIP buyers and sellers already submitted offers that now result in transactions. ZIPT buyers and sellers have also started submitting offers due to to a sufficiently small size of the bid-ask-spread. They are also involved in transactions now. However, all sellers are successfully

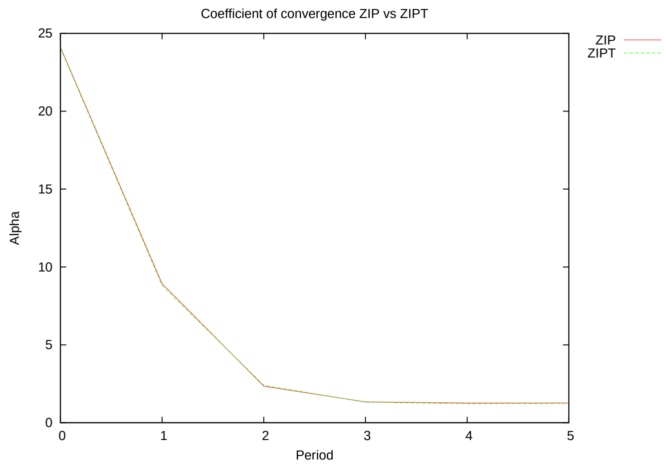

Figure 16: Course of the sample mean of the coefficient of convergence for homogeneous populations of ZIP and ZIPT traders, sample size = 1000

testing the market for higher profit. This results in an increase of the transaction prices. ZIPT traders stop submitting offers because the size of the bid-ask-spread is not sufficiently small. ZIP traders do not detect this trend and keep trading. During this period of time of course ZIP buyers make more profit than they would have made at equilibrium prices. ZIP sellers on the other hand lose potential profit. ZIPT traders also submit offers during this period of time but not as much as ZIP traders do. This is due to the continuously change in size of the bid-ask-spread during the time of increasing transaction prices.

As can be seen in the graph of figure 17 the transaction price time series overshoots the equilibrium price before it slopes down towards the end of period one.

To summarize, it turned out to be more profitable for sellers to wait for the transaction price series to increase. On the other hand, buyers made high values of profit at the beginning of the period. This finding is also illustrated by figure 18 which shows the mean profits per period. During the course of the auction, both type of traders are able to converge to the theoretical total profit of 3750 price units.

Market 1 gives the same findings as described above. The reason for this is that trading starts again below the equilibrium price. Buyers make more profit at the beginning of the period while it is more profitable for sellers to wait for the transaction prices to raise.

In contrast, trading in market 2 starts above the equilibrium price. The transaction price time series is shown in figure 19 and the mean profits per period in figure 20. In this case sellers are better off to sell at the beginning of the period while buyers would look better if they waited for the transaction price series to fall.

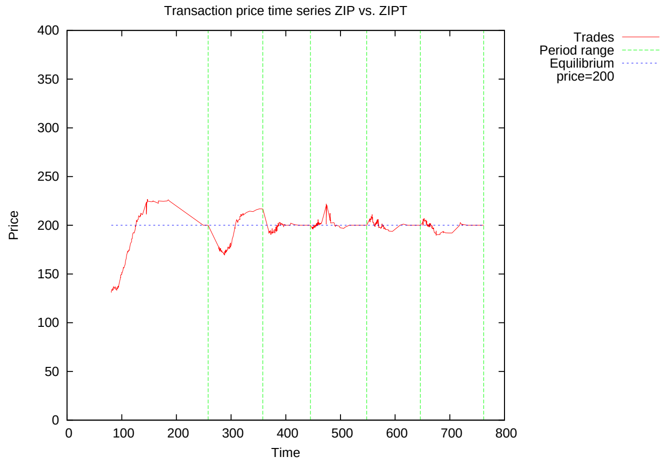

Figure 17: Transaction price time series for market A with a heterogeneous population of ZIP and ZIPT traders

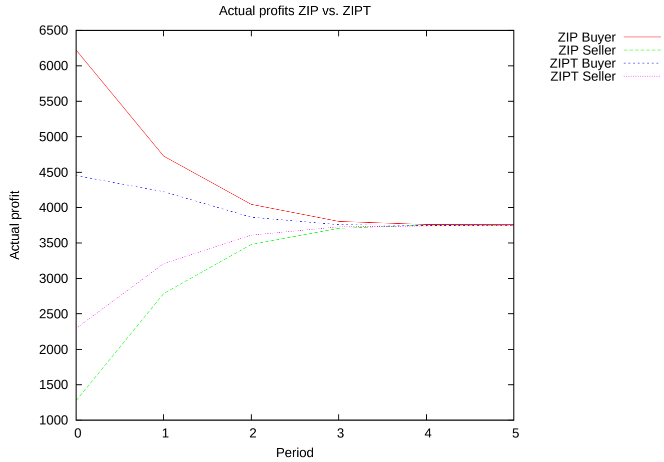

Figure 18: Course of the sample mean of actual profits for a heterogeneous population of ZIP and ZIPT traders in market A, sample size = 1000

5 Conclusions

Two trading strategies were enhanced by new timing heuristics that use time either strategically or not.

"Zero-Intelligence" agents place their offers as soon as their prices are being determined. The enhanced version "Zero-Intelligence Time" does not use time as a strategic element at all and places their offers randomly distributed over the remaining time of the auction. The following considers only the constrained version of ZI traders. Both, ZIC and ZICT traders achieve high values of allocative efficiency when placed with homogeneous populations into a market. However, when confronted with each other, it is shown that ZICT agents are getting outperformed. They are not able to achieve the theoretical actual profit predicted by economic market theory. It is even the case that ZIC traders are able to exploit ZICT traders and achieve actual profits that are even higher than their predicted theoretical profit. The bad timing of ZICT agents does not allow them to be efficient. However, the timing of ZIC agents can not be considered as strategically intelligent either. But the fast timing of ZIC agents is superior to the random timing of ZICT agents.

"Zero-Intelligence-Plus" agents are able to observe the market and adjust the pricing of their next offers accordingly. However, they are not able to detect trends that indicate higher profits at a later date. "Zero-Intelligence-Plus Time" agents use the size of the bid-ask-spread to indicate that prices are going to change in the near future. It is shown that the transaction price time series of ZIP and ZIPT agents converges towards the equilibrium price within a few number of periods when placed with homogeneous populations into a market. During the time of convergence however, either sellers or buyers accept sub-optimal deals depending on whether the price series converges the equilibrium from above or below. When confronted with each other, ZIPT traders are succesfully avoiding these sub-optimal deals by ending the submission of offers as soon as a change of transaction prices is indicated by the size of the bid-ask-spread. Therefore, in the case study of market A for instance, ZIPT buyers were still able to achieve more profit in early periods than their predicted theoretical profit. Also, ZIPT sellers outperformed ZIP sellers due to the avoidance of sub-optimal deals. ZIP sellers were not guarded against them. It is also to be noted that the most profitable group of traders were the ZIP buyers which exploited mainly ZIP sellers for high profitable deals. After transaction prices have been able to converge towards the equilibrium price, profits are equal among all group of traders. But it is unrealistic to assume that an equilibrium price stays stable in real world trading. Adjustments towards the equilibirum price will be made continously.

The outcome of this paper indicates that profit oriented trading agents must not ignore timing as an essential part of their strategies in order to be competitive in real world trading markets. Intelligent timing strategies are indispensable and it can not be made standard practice to only consider timing strategies that place offers as soon as their prices are determined. In order for the ZIPT timing strategy to become more successful, it is not only necessary to avoid sub-optimal transactions but also to detect when relatively high profits can be made.

Sophisticated timing needs to be added to other trading strategies in order to discover how they can be improved further.

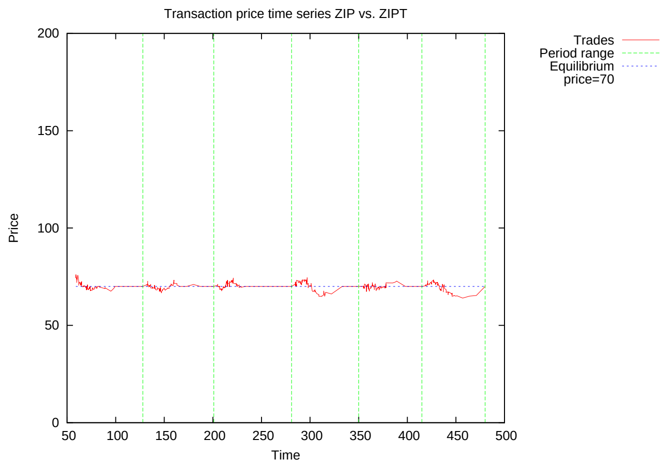

Figure 19: Transaction price time series for market 2 with a heterogeneous population of ZIP and ZIPT traders

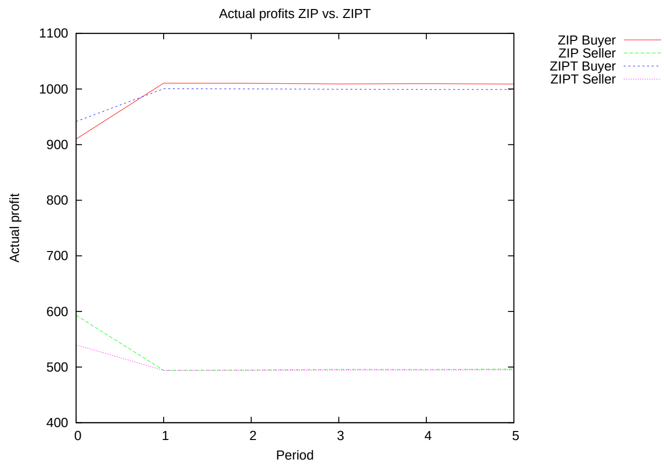

Figure 20: Course of the sample mean of actual profits for a heterogeneous population of ZIP and ZIPT traders in market A, sample size = 1000

References

[1] Adam Shell (2007). Technology squeezes out real, live traders. `http://www.usatoday.com/money/markets/`.

[2] Cliff, Dave (1997). Minimal-intelligence agents for bargaining behaviors in market-based environments. Technical report, School of Cognitive and Computing Sciences, University of Sussex.

[3] Das, Rajarshi, Hanson, James E., Kephart, Jeffrey O., and Tesauro, Gerald (2001). Agent-human interactions in the continuous double auction. *Proceedings of the 17th international joint conference on Artificial intelligence - Volume 2*, pp. 1169–1176, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.

[4] Gode, Dhananjay K and Sunder, Shyam (1993). Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, Vol. 101, No. 1, pp. 119–37.

[5] Parsons, Simon, Rodriguez-Aguilar, Juan A., and Klein, Mark (2011). Auctions and bidding: A guide for computer scientists. *ACM Comput. Surv.*, Vol. 43, pp. 10:1–10:59.

[6] Parsons, Simon, Marcinkiewicz, Marek, Niu, Jinzhong, and Phelps, Steve (in preparation). Everything you wanted to know about double auctions but were afraid to (bid or) ask. Technical report.

[7] Preist, Chris and Tol, Maarten van (1998). Adaptive agents in a persistent shout double auction. *Proceedings of the first international conference on Information and computation economies*, ICE '98, pp. 11–18, ACM, New York, NY, USA.

[8] Smith, Vernon (1962). An experimental study of competitive market behavior. *Journal of Political Economy*, Vol. 70.