Simulated Double Auction-markets with production and storage*

Johan Dahlin†
Department of Economics
Umeå University

January 17, 2011

Abstract

Earlier agent-based models simulate Double Auction (DA)-markets by using trading periods, each lasting until all agents have traded a certain amount of commodity. In real-world markets, agents are entering and exiting at random thus changing the amount of tradable commodity. In this paper we therefore study the effects of removing trading periods from earlier proposed agent-based models. We present a method of removing trading periods from DA-markets by introducing storage and production. The resulting transaction price series is shown to have closer resemblance to price series found in real-world markets, than previous models with trading periods. The effects of storage on markets populated by machine agents are also studied. We present results showing that the Zero-Intelligence Plus agents (ZI-P) performs poorly even with very small storage fees. Finally we show how a simple modification of the ZI-P agents results in a significant better performance on markets with storage costs.

*This paper was written as a Bachelor thesis in Economics for a Degree of Bachelor in Business and Economics at Umeå University. The thesis is a part of the course *Research project*, 15 ECTS-credits given during the fall 2010. The paper was made possible by the guidance and suggestions by thesis supervisor Ulf Holmberg, Department of Economics, Umeå University.

†E-mail: joda0009@student.umu.se.

Contents

1 Introduction

Pricing is a subject of everyday importance; individuals, producers and other economic agents engages daily in bargaining and price setting. Usually this subject is studied using the concept of perfectly competitive markets[1], assuming fully rational economic agents in an equilibrium environment. Recently the requirement of rational agents have been challenged by studies of markets populated by machine agents, equipped with none or little intelligence. Simulations of such artificial markets yields adjustment patterns to the equilibrium price with obvious resemblance experiments of markets populated by human agents. These studies raise the question of how much rationality is needed for the market price to converge to the Walrasian equilibrium price. Perhaps the complex real-world market may be explained by a simpler mechanism, that relaxes the assumption of fully rational economic agents.

The Walrasian equilibrium price is guaranteed existence by a large number of assumptions. Including that the market contains some amount of tradable commodity some period and a large number of active economic agents. In addition, the Walrasian price is derived assuming no additional friction, i.e. no asymmetric information and transaction costs. These assumptions are far from the real world markets that often experience different kinds of information problems and problems with few sellers and many buyers. Studies of pricing without these strict assumptions are therefore important for understanding these markets. Attempts to soften the assumptions of the Walrasian equilibrium model has resulted in impracticable models and therefore no complete theoretical model exists of real-world markets. This have motivated the use of less traditional economic methods, such as agent-based computational economics[2].

Agent-based modelling[3] has long been used in the fields of mathematics, physics and biology to model the emergent properties of a system by using small agents. The main thought behind the method is the use of many agents, which (in general) are quite simple constructs, that interact using

[1]For a thorough presentation of perfectly competitive markets and the Walrasian theory, see Mas-Colell et al. (1995). More accessible introductory presentations are found in Varian (2006) and Gravelle and Rees (2004).

[2]For motivation for the use of agent-based methods in economics and discussions on the difficulties in using the Walrasian model in real-world markets, see Tesfatsion (2006).

[3]An accessible introduction to the use of agent-based modelling in the Social Science is found in Gilbert (2007).

some rules and creates some emergent macroscopic behaviour. This corresponds quite well with the structure of the connection between microeconomics and macroeconomics, where the interactions between the small parts of the economy aggregated creates macroeconomic phenomenon. The main advantage of agent-based modelling is the ability to study dynamic system behaviour and therefore analyse non-equilibrium situations in economics. A drawback is the difficulty in finding the simple rules/behaviours of agents that creates the desired emergent properties.

Previous studies have shown that the Walrasian equilibrium price may be found by using Double Auction(DA)-markets populated by machine agents. In a DA-market, sellers and buyers are able to submit offer and bid prices to a central auctioneer which announces a trade if a bid and offer price intersects. In this paper, we will use machine agents and DA-markets to study more complex markets then previously done. We will generalize the previous studies by removing trading periods[4] using production and storage (with/without a storage fee).

The subject of studying the competitive market with experiments was first introduced by Smith (1962). The author investigates the hypotheses of Walrasian competitive market theory by the use of human agents in market experiments. The conclusion of the experiments was that human agents quickly find the theoretical equilibrium price, for a range of different sets of demand and supply functions. Gode and Sunder (1993) continued the work of Smith in a simulation-based setting[5], replacing humans agents with simple computer-based agents submitting random bids and offers. The hypothesis was that the actions of these irrational *Zero-Intelligent agents* (ZI) would create an aggregated competitive market with high performance[6]. The main conclusion of the study was that the Walrasian equilibrium price was found by simple machine agents (with no reasoning, learning or adaptability) populating a DA-market[7]. The work of Gode and Sunder (1993) was heavily

[4] Earlier studies have allowed agents to trade a number of units of commodity during a trading period. Each period lasts until all agents have traded all units at their disposal. After each trading period all agents are given the same amount of units to trade and the simulation is continued.

[5] For a discussion of the use of machine agents in market experiments, see Duffy (2006).

[6] The market performance measure used by Gode and Sunder (1993) is the part of the total possible profit (the sum of the consumer and producer surpluses) that the agents manages to *extract* in trading. As discussed by Becker (1962) the use of a simple rule (the budget restriction) would generate an effective market.

[7] This work is also discussed and analysed in Brewer et al. (2002).

criticized by Cliff et al. (1997). They demonstrated that the convergence towards the theoretical equilibrium price was the result of market symmetry. This was verified by simulations on asymmetric markets where the ZI-agents preformed poorly (converging to the calculated *incorrect* price and not to the theoretical *correct* price). To resolve this problem, Cliff et al. (1997) introduced some adaptability to the ZI-agent and named these ZI-P agents. The use of ZI-P agents resulting in the transactional prices more similar to observed prices from markets with human agents.

This paper continues the work of Gode and Sunder (1993) and Cliff et al. (1997) by removing trading periods and suggesting a future improvement of the ZI-P agents to handle markets with storage fees. We generalize the previous work in five different steps; (i) agents may trade multiple units, (ii) agents may store units between periods with some fee, (iii) sellers may trade with loss to avoid storage cost, (iv) adding production of units and finally (v) ZI-P agents may add an expected storage cost onto submitted ask prices, these new agents are named ZI-Pe agents. Thereby creating a complex artificial market more similar with real-world markets, which is compared to the previous existing simpler market models. We also study how the ZI-P by Cliff et al. (1997) and ZI-Pe agents react to an increasing storage fee.

We present the reader with three main results; (i) simulations with production and free storage creates transaction price series more similar to market experiments with human agents. (ii) Adding a small storage fee generates poor performance by ZI-U, ZI-C and ZI-P agents. Most sellers goes bankrupt with a resulting drop in the equilibrium price. (iii) An enhancement of the ZI-P agents allows the sellers to add the storage cost onto the offered prices. These enhanced agents have some success (for small storage fees) contributing with more transactions and a more stable equilibrium price series, compared with the original ZI-P agents.

The outline of this paper is as follows, the papers continues with the first theoretical part presenting and discussing the work of Gode and Sunder (1993) and Cliff et al. (1997) in technical terms for reference. The second theoretical part describes the additional factors and rules unique for this paper and the motivation for adding them into the model. Continuing with a presentation, analysis and discussion of the results. Finally we present some conclusions, implications and questions for future study.

2 Market models and agents

This section presents the model environment of the agent-based simulations. We begin by formally introduce the Double Auction (DA) market and proceed by describing the markets (and agents of) Gode and Sunder (1993) and Cliff et al. (1997). Finally comparing the behaviour and performance of markets populated by human agents and each of the three types of machine agents.

2.1 Double Auction Market

The price setting mechanism used in this paper is the Double Auction(DA)-market[8]. The DA-market is a very robust and simple method to create markets, which efficiently and quickly approaches the theoretical Walrasian equilibrium price. DA-markets are therefore commonly used in stock exchanges, commodity markets and when selling government bonds. Gjerstad and Dickhaut (1998) defines a DA-market as a *microeconomic system*, which in turn is defined in *Definition 1* found in Smith (1982).

Definition 1 (Microeconomic system) *A microeconomic system,* $S = (\mathbf{e}, \mathbf{I})$*, is comprised of the environment* $\mathbf{e}$ *and an institution* $\mathbf{I}$*. The environment* $\mathbf{e}$ *is defined as,*

$$\mathbf{e} = \prod_{i \in \mathcal{A}} \mathbf{e}_i,$$

where $\mathcal{A} = \{1, 2, \dots, n\}$ *is the set of agents and* $\mathbf{e}_i$ *is the characteristics of agent* i *(i.e. preferences, technology and endowment). The institution* $\mathbf{I}$ *consists of a message space* M_i *for each agent* i*, an adjustment rule (specifying the sequence of agent messages) and an outcome function,* $h(m_t)$*,*

$$h(m_t) = (h_1(m_t), h_2(m_t), \dots, h_n(m_t)),$$

where $m_t = (m_{1,t}, m_{2,t}, \dots, m_{n,t}) \in \mathcal{M}_t$ *where* $\mathcal{M}_t$ *is the vector of the agent's messages defined as,*

$$\mathcal{M}_t = \prod_{i \in \mathcal{A}} \mathbf{M}_{i,t}.$$

[8] There exits a large number of introductions to DA-markets. A comprehensive treatment is given by Friedman (1993). The formation of prices in DA-markets, using experiments and simulations is investigated in Friedman (1984), Cason and Friedman (1996) and Gjerstad and Dickhaut (1998).

We use two different kinds of agents; *sellers* and *buyers*. Sellers manufacture some commodity and have a *production cost*, $c_i = (c_{i,1}, c_{i,2} \dots c_{i,n})$, for each agent $i \in \mathcal{A}_1$ and n units of commodity. We assume that the seller must sell the units in such order, that the cost is an increasing function of quantity. The sellers' *profit* is the difference between the price and the total cost of the commodity, $\pi_{i,k} = p_k - c_{i,k}$, where p_k is the transactional price for good k. We assume that sellers attempts to maximize its own total profit, $\pi_i = \sum_k \pi_{i,k}$.

The buyers buy some commodity and have a *redemption value*, $r_j = (r_{j,1}, r_{j,2}, \dots, r_{j,n})$, for each agent $j \in \mathcal{A}_2$ and m units of commodity. We assume that the buyers must buy the units in such order, that the redemption value is an decreasing function of quantity. The buyers seek to maximize the *net redemption value*, the difference between the redemption value and the transactional price, $\eta_{j,k} = r_{j,k} - p_k$, motivated by the assumption that the buyers are equipped with some *monotonically increasing* utility function, $U(\eta_j, \cdot)$, where $\eta_j = \sum_k \eta_{j,k}$.

In our experiments there exists two goods, that are traded between the agents; a currency and a commodity. The total set of agents is denoted $\mathcal{A} = \mathcal{A}_1 \cup \mathcal{A}_2$, where $\mathcal{A}_1$ is the set of *sellers* and $\mathcal{A}_2$ the set of *buyers*. The environment $\mathbf{e}$ is defined as the combined vectors of producers' profits π_i and buyers' net redemption values η_j,

$$\mathbf{e} = \{\pi_i\}_{i \in \mathcal{A}_1} \cup \{\eta_j\}_{j \in \mathcal{A}_2}.$$

An important remark is that the ordered vector of individual redemption values $r = (r_1, r_2 \dots r_n)$ and producer costs $c = (c_1, c_2 \dots c_m)$, for n buyers and m sellers, are some points of the demand and supply functions of the market. These ordered vectors are used as the empirical supply and demand functions to find the Walrasian equilibrium price.

Continuing with the institution of the DA-market, we have buyers shouting (making/announcing) bids and sellers offers. Formalizing this procedure, we introduce the notion of a *message space* in *Definition 2* and rules for sellers in *Definition 3* and buyers in *Definition 4* as in Gjerstad and Dickhaut (1998).

Definition 2 (Message space) *Let $\mathcal{N} = \{x : x = \frac{n}{10^m}$ for $n \in \mathbb{N}\}$, i.e. all positive valued number with m decimals. Every seller $i \in \mathcal{A}_1$ and buyer $j \in \mathcal{A}_2$ has a message space $\mathcal{M}^{(s)}_{i,t}$ and $\mathcal{M}^{(b)}_{j,t}$ at any time t as,*

$$\mathcal{M}^{(s)}_{i,t} \subset \{i\} \times \{0\} \times \{N\} \qquad \mathcal{M}^{(b)}_{j,t} \subset \{0\} \times \{j\} \times \{N\}.$$

That is if a seller $i \in \mathcal{A}_1$ would like to submit an ask for $a \in \mathcal{N}$, the message space would be $(i, 0, a)$, if a buyer $j \in \mathcal{A}_2$ would like to shout a bid of $b \in \mathcal{N}$ the message would be $(0, j, b)$.

Definition 3 (Best (outstanding) ask and acceptance) *The best ask (or outstanding ask) is denoted as $a^* = a_k > a_i$ for any $i \in \mathcal{A}_1 \cap \{k\}$. A trade (accept) occurs when buyer $j \in \mathcal{A}_2$ shouts bid $(0, j, b)$ such that $b > a^*$, the transaction price p is the best ask $p = a^*$. When an accept occurs we set the best prices $a^* = b^* = 0$. If ask $(i, 0, a)$ is shouted, by seller $i \in \mathcal{A}_1$, and $a^* = 0$ or $a^* > a$, we set $a^* = a$ as the new best ask.*

Definition 4 (Best (outstanding) bid and acceptance) *The best bid (or outstanding bid) is denoted as $b^* = b_k < b_j$ for any $j \in \mathcal{A}_2 \cap \{k\}$. A trade (accept) occurs when seller $i \in \mathcal{A}_1$ shouts ask $(i, 0, 1)$ such that $a < b^*$, the transaction price p is the best bid $p = b^*$. When an accept occurs we set the best prices $a^* = b^* = 0$. If bid $(0, j, b)$ is shouted, by buyer $j \in \mathcal{A}_2$, and $b^* = 0$ or $b^* < b$, we set $b^* = b$ as the new best bid.*

2.2 ZI-U and ZI-C agents

The simplest possible agents for use in DA-markets was presented by Gode and Sunder (1993). They introduced two types of *Zero-Intelligence* (ZI) agents, which makes purely random bids and offers to the market. Each ZI-agent have some predetermined maximum production cost, Λ_c, or a maximum redemption value, Λ_v. Although that the two different versions of the ZI-agent uses the same DA-market arrangement, summarized in *Algorithm 1* in *Appendix A*, they use two different methods of generating shouts. The first type of ZI-agents, generates uniformly distributed random shouts using,

$$\text{offers: } a_{i,j} \sim \mathcal{U}[0, \Lambda_c] \qquad \text{bids: } b_{i,j} \sim \mathcal{U}[0, \Lambda_v]. \tag{1}$$

Since this type of agent faces no budget constraints, i.e. submit offers/bids that may generate losses, it is named *Zero-Intelligence Unconstrained agents* (ZI-U). The other version of the ZI-agent uses a budget constraint and is named *Zero-Intelligence Constrained agents* (ZI-C). The constraint is implemented by reducing the set of possible bids and offers, such that shouts creates non-negative profits,

$$\text{offers: } a_{i,k} \sim \mathcal{U}[c_{i,k}, \Lambda_v] \qquad \text{bids: } b_{j,k} \sim \mathcal{U}[0, r_{j,k}], \tag{2}$$

where $c_{i,k}$ the *production cost* for agent $i \in \mathcal{A}_1$ and good $k \in \{1, 2, \ldots, n\}$ and $r_{j,k}$ denotes the *redemption value* for agent $j \in \mathcal{A}_2$ of good $k \in \{1, 2, \ldots, m\}$. The value and cost for each agent is calculated as an uniformly distributed random variable as $c_{i,k} \sim \mathcal{U}[0, \Lambda_c]$ and $r_{j,k} \sim \mathcal{U}[0, \Lambda_v]$. As a result, ZI-C agents are heterogeneous each having different cost/redemption values.

2.3 ZI-P agents

Cliff et al. (1997) introduces *Zero-Intelligence Plus agents* (ZI-P), a more advanced version of the ZI-C agents including the ability of adaptive behaviour. These agents observes the shouts and transactional prices on the market and uses this information to adapt the mark-up. In such manner, that the agents own shouts approaches the observed transactional prices and thereby also the theoretical price. The ZI-P agents therefore do not make purely random shouts and also rely on an open order book, i.e. all shouts are public to all agents.

As the ZI-P agents uses a mark-up onto the random cost/value, the shout generating function differs from function used by the ZI-C agents. Given a cost $c_{i,k}$ or redemption value $r_{j,k}$, the shout prices are generated as,

$$\text{offers: } a_i(t) = c_{i,k}(1 + \mu_i(t)) \qquad \text{bids: } b_j(t) = r_{j,k}(1 + \mu_j(t)), \tag{3}$$

where $\mu_i(t)$ is the *mark-up* calculated by the Widrow-Hoff *delta rule*, discussed in detail below. Due to the budget restriction introduced into the ZI-C agents in the previous section, the mark-up must be negative for the buyers, $\mu_j(t) \in [-1, 0]$, and positive for the sellers, $\mu_i(t) \in [0, \infty)$. The mark-ups are updated after each shout using the following expressions,

$$\text{sellers: } \mu_i(t) = \left[\frac{p_i(t-1) + \psi_i(t-1)}{c_{i,k}}\right] - 1, \tag{4}$$

$$\text{buyers: } \mu_j(t) = \left[\frac{p_j(t-1) + \psi_j(t-1)}{r_{j,k}}\right] - 1, \tag{5}$$

where ψ_i is a *momentum-based correction* defined below. The mark-up correction itself is calculated as the difference in mark-up needed to bring the price to some target price level. Cliff et al. (1997) discusses the problem with using the *target price*, $\tau_i(t)$, as the most recent transactional price. This could result in that the simulated mean transactional price largely differs from theoretical equilibrium price. To resolve this problem, Cliff et al. (1997) includes a mechanism imposing sellers to always strive to increase the price (to earn more profit) and buyers to always seek to decrease the transactional price (to increase net utility). This is modelled as a stochastic function of the most recent shout price as,

$$\tau_i(t) = \mathcal{R}_i(t)q(t) + \mathcal{B}_i(t), \tag{6}$$

where $q(t)$ is the most recent shout price with $\mathcal{R}_i(t)$ and $\mathcal{B}_i(t)$ as some uniform random variables (with signs depending on if the agent would like to increase or decrease the mark-up). An agent wishing to increase the price would therefore use positive values of the random variables. This would create a target price that is somewhat higher then the most recent shouted price. This target price is used by the Widrow-Hoff delta value which calculates the adaptive change in the shout price by,

$$\Delta_i(t) = \beta_i \left(\tau_i(t) - p_i(t)\right), \tag{7}$$

where β_i is the *learning rate coefficient*, with $\beta_i \in [0,1]$ for each agent $i \in \mathcal{A}$. To remove any high-frequency oscillations around the price level, a momentum-based update procedure is also added. This dampens any big changes in the mark-up between two updates. The value of the *momentum-based mark-up update*, $\psi_i(t)$, is calculated using,

$$\psi_i(t) = \gamma_i \psi_i(t-1) + (1-\gamma_i)\Delta_i(t-1) \text{ with, } \psi_i(0) = 0, \tag{8}$$

where γ_i is called the *momentum coefficient*, with $\gamma_i \in [0,1]$ for each agent $i \in \mathcal{A}$. This mark-up update is implemented in *Algorithm 4*, together with some rules for when to increase and decrease the mark-up as shown in *Algorithm 3* (in *Appendix A*).

2.4 Results and conclusions from previous work

We proceed by presenting some results from the ZI-agents, ZI-P agents and human experiments from previous work. A comparison[9] between human traders and the two ZI-agent types introduced by Gode and Sunder (1993) is presented in *Figure 1*.

[9]The simulation is based on the use of 12 traders (6 buyers and 6 sellers) with the maximum redemption value $\Lambda_v = 200$ and maximum production cost $\Lambda_c = 200$. 6 trading periods are used with the duration of 30 seconds per period, and 2 minutes per period in the experiments with human agents.

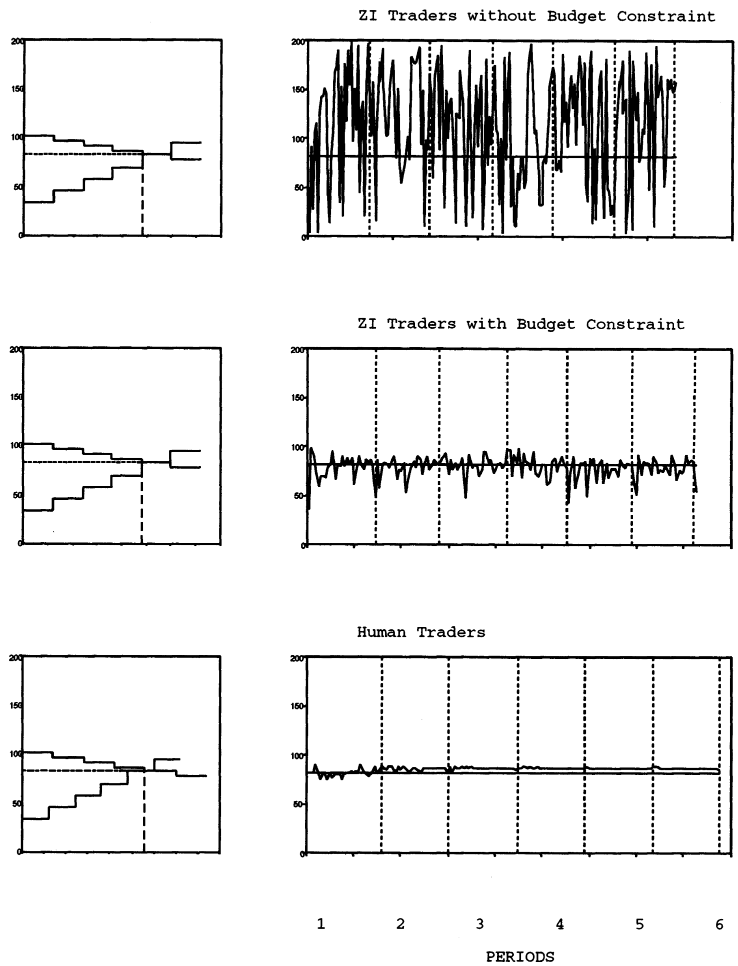

Figure 1: The results (Fig 1) from the first trial in Gode and Sunder (1993). The left part of the figure depicts the implied supply and demand functions, of the market, with the implied competitive equilibrium price and quantity. The right hand side of the figure shows the sequence of transactional prices over a number of trading periods. The horizontal line indicated the theoretical equilibrium price of the market. For more information see *Footnote 9*.

The resulting price series are quite different in appearance, from the erratic prices of the ZI-U trader to the quickly adapting price series of human traders. Gode and Sunder (1993) compares five trials with different demand and supply functions, to examine different market structures (number of sellers, buyers, costs and redemption values).

A regression analysis is also done of the mean square error (MSE) as a function of transactional number. This analysis indicates that the MSE decreases with increasing number of transactions for the ZI-C agents but not the ZI-U agents[10]. We note the quick adaptation to the equilibrium price of the human traders and that the trading periods have little effect on the price level.

Using the extension of the ZI-agents provided by Cliff et al. (1997), we have reproduced their findings together with simulations of the ZI-agents[11]. The supply and demand functions together with the mean transactional prices, are shown on the left hand side of *Figure 2*. The center plots shows the sequence of transactional prices for the three compared agent types. The shape of these processes are quite dissimilar, from the random noise of the ZI-U agents to the more human like pattern (compare with the right hand side of *Figure 1*) of the ZI-P agents. The volatility (variation is transactional prices) of the ZI-U agents is quite high and is much lower for the ZI-P agents, higher in the start of each trading period and then smaller after some transactions.

On the right hand side in *Figure 2* the root mean square errors (RMSE) are compared between the three types of agents and the theoretical equilibrium price (dashed lines are $\pm$ one standard deviation). The downward slope of the RMSE of the ZI-P agents is statistically significant[12].

[10]Gode and Sunder (1993) makes no statement about the number of simulations the regression analysis is based on. As such it is impossible to draw any conclusions on the basis of the statistical significance of the regression analysis (and the following conclusions).

[11]Using the maximum redemption value $\Lambda_v = 2.00$ and maximum production cost $\Lambda_c = 2.00$ with 200 agents ($N_b = 100$ buyers and $N_s = 100$ sellers) over 6 trading periods, each lasting until all agents have traded or the agents have made a total of 10000 shouts during the current trading period. All agents have one tradeable unit of commodity or currency during each trading period.

[12]These results only hold for symmetric supply and demand functions (same maximum cost and redemption value), when asymmetric, the ZI-U and ZI-C preforms poorly as already noted in Cliff et al. (1997).

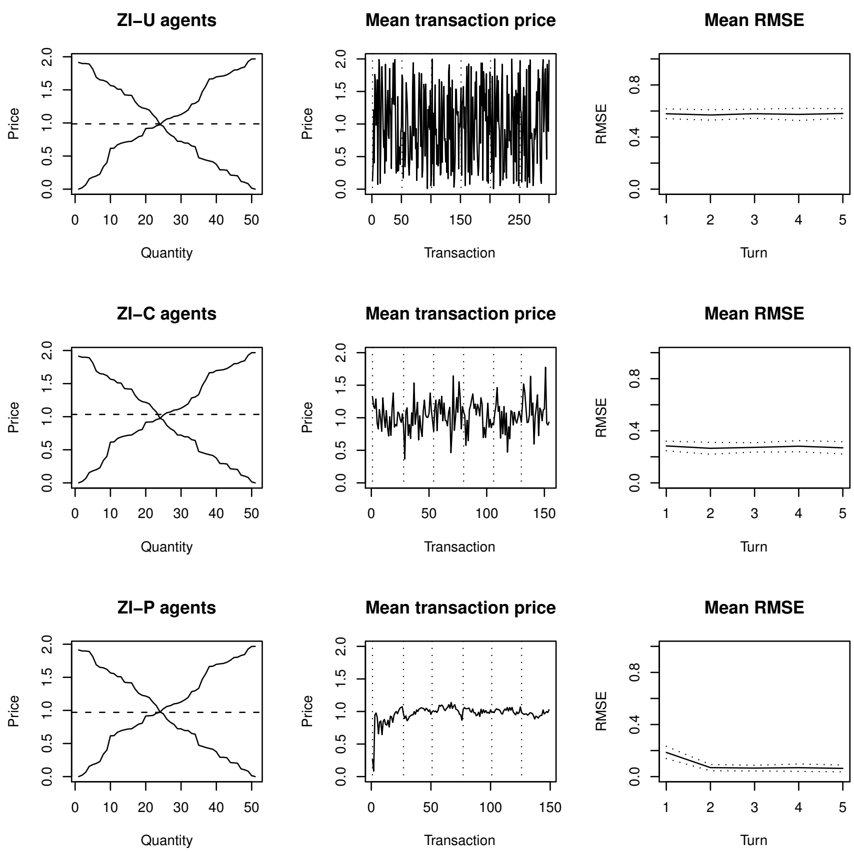

Figure 2: Results for comparison between the effectiveness of the ZI-U and ZI-C agents by Gode and Sunder (1993) and the ZI-P agents proposed by Cliff et al. (1997). The left hand side of the figure shows the supply and demand functions used in these simulations, together with the mean of the transactional prices (dashed line). In the center, the sequence of transactional prices during 6 trade periods. On the right hand side, the mean RMSE is shown calculated as a mean of 50 runs. Dotted lines indicates confidence intervals of length two standard deviations. The vertical dotted lines indicates the starting point of each trading period. All parameters used in the simulation run is presented in *Appendix B* and in *Footnote 11.*

3 Modified agents with production and storage

In this section we modify the ZI-agents and ZI-P agents by Gode and Sunder (1993) and Cliff et al. (1997). The material of this section is supplemented by *Appendix A*, which includes a summary of all the involved variables describing the agents and the algorithms used to simulate their behaviour.

We generalize the agents in five different steps; (i) enabling the agent to trade multiple units of commodity and currency, (ii) storing commodity between periods by paying some storage fee, (iii) enable the sellers to trade with a loss to avoid storage cost, (iv) adding production thereby removing the trading periods and finally by (v) introducing *Zero Intelligence Plus-enhanced* (ZI-Pe) agents with some adaptive behaviour to add the estimated storage cost into the shout generation.

3.1 Agents with storage

We begin by adding the ability to buy and sell multiple units of commodity while allowing for different redemption values and costs for each unit. We assume, as previously stated in *Section 2.1*, that the producers experience a growing production cost and sellers experience diminishing value of owning one more unit of commodity. The vectors of the costs and redemption values are generated in the same manner as previously described but are sorted. With ascending values, we find the *cost vector* $c_i = (c_{i,1}, c_{i,2} \dots c_{i,n})$ for seller $i \in \mathcal{A}_1$ and n commodities. With descending values, we find the *value vector* $r_j = (r_{j,1}, r_{j,2}, \dots, r_{j,m})$ for buyer $j \in \mathcal{A}_2$ and m commodities. Since the units must be traded in sequence by the assumption above, each agent have some maximum number of transferable units during each trading period. The agents are active for as long as they have units left to trade. When deactivated, an agent is no longer able to shout ask/bid prices.

A further generalization is added by allowing the agents to store surplus units until the next trading period. Thus creating a market where different agents have different amounts of units to trade. We allow the agent to store units by increasing its maximum number of transferable units. A consequence of the above, is that it will be more difficult to sell the last unit of commodity then the first. The sellers are also required to pay a storage fee, Γ, for each unit and trading period left in storage. The total storage cost is, $\Gamma N_{i,L}$, where $N_{i,L}$ is the number of transferable units left or seller $i \in \mathcal{A}_1$ after L trading periods. This storage cost is added to the vector of profits made by the agents in the beginning of each trading period[13].

In the original models, the ZI-agents must shout prices that yield some net profits. With the introduction of storage costs, we now relax this assumption for the sellers. Thus allowing the seller to choose the lesser of two losses generated by selling units with a loss and the storage cost. To incorporate this, we modify the requirement of ZI-C and ZI-P agents of positive net profit to allow for a net loss during a trade, if this loss is smaller then the storage cost. To avoid the problem of agents having infinite losses, we only allow this modified pricing limitation if the agent have a positive total net profit. We also introduce bankruptcy, if an agent have some total net profit, π_i, smaller then some arbitrary limit value, Π, such that $\pi_i < \Pi$, the agent is removed from the market for the remainder of the simulation.

3.2 Agents with production

The role of trading periods in the previous models is that an agent may sell or buy a number of units during a trading period. Each period lasts for a certain number of attempted trades or until all agents have traded (or if the mark-up is 0 for all agents). Using the framework of the multiple unit agents above, we introduce the ability of production for the agents. Each agent produces units, that increases the maximum number of transferable units. The production is determined by the production time, ω_i, which is a uniformally distributed random number, $\omega_i \sim \mathcal{U}[0, \Omega]$, for some agent $i \in \mathcal{A}$ and some predetermined maximum production time Ω. An agent $i \in \mathcal{A}$ produces a unit if $0 \equiv l \pmod{\omega_i}$, where l is the number of shouts. An agent starts with some number of allocated units $N_{i,0}$, which increases with production (after a certain number of shouts). If a successful trade occurs then the maximum number of transferable units decreases for both involved agents.

To include a cost of storage without trading periods, we let the sellers pay a fee per shout instead of per period (as used above). This is motivated by assuming that each shout requires some time and therefore may be used as a time step. This new storage fee is significantly smaller then the storage

[13] One could also add some cost of storage for the buyers, much like a form of inflation, where a part of the redemption value is removed, for each unit after each period in *storage*. We however not add this feature in this market, but it remains an interesting area of further improving the model. One could for example study the intertemporal allocation of consumption for buyers in a economy with inflation.

fee per trading period, due the the large number of shouts required to yield a successful trade. This storage cost is added to the vector of profits after each shout. The cost is calculated as $\gamma N_{i,l+1}$, for some seller $i \in \mathcal{A}_1$, shout l and some storage fee γ per unit and shout.

3.3 ZI-P agents with storage cost awareness (ZI-Pe)

Finally, we propose a small enhancement of the ZI-P agents to include the storage cost in pricing. Using ZI-P agents, the sellers are unaware of the fact that a random storage cost is added onto their cost of production. In the real world, this is a factor that sellers needs to include when pricing a commodity. When neglected, this would yield random net losses for the sellers and bankruptcy is likely to follow. To give the ZI-P agents the ability to escape bankruptcy, we introduce some changes in the price calculations. This allows the sellers to add the expected cost of storage into the offers submitted to the market. This is done by adding the expected cost into the seller pricing equation in (3) as,

$$a_i(t) = c_{i,k}(1 + \mu_i(t)) + S_{i,k}$$

where $c_{i,k}$ is the production cost for seller $i \in \mathcal{A}_1$ and unit $k \in \{1, 2, \ldots, n\}$. The expected storage cost $S_{i_k} = \gamma E[T_s]$, with E as the expectation operator, T_s as the random storage time and γ as the storage fee defined above.

We will estimate the expected storage time, $E[T_s]$, by calculating the average storage time previously experienced by the seller. The initial expected storage time will be a predetermined constant (calculated by some pilot simulations), which then will be updated by the average experience storage time of all sellers[14].

[14]In a more realistic setting, we would not expect a seller to know the expected storage time of the other competing sellers. We could justify this simplification by the following argument. Each agent knows the expected number of shouts, from other agents, until itself may present the market with a shout. Therefore the sellers could keep track of each other on the market, noting the time of successful trades and the time of the next shout from the seller. If the sellers had no units in storage during this time, the expected time of production must have been the time of the first shout (after reactivation) minus the expected number of shouts, from other agents, until itself may present the market with a shout. Using this observation each seller could get an approximate value of the average storage time, of all sellers, and thereby use this information to better its own estimated storage cost.

If we did not allow this simplification, the sellers would experience a slow learning

The enhancement above, is however not taking into account that the time of storage is not constant during the simulation. The storage time depends on the rate of production and on the distribution of production times. In the beginning when all agents have a unit to trade the storage time could be higher/smaller then the storage time at later shouts. If agents have many units in storage at a later stage the initial storage time is lower then the later storage time. The opposite occurs if the production is low and many agents have no units in storage.

We could correct this by using a moving average, which only using a last fraction of the recorded storage times, instead of the (complete) average used above. This could however result in a worse estimator of the expected storage time, due to a smaller sample size. Because the number of trades per agent is quite small, only using a fraction of this data could produce a worse estimator of the expected storage time/cost.

---

process, which requires many transactions for each seller to get a good estimate. This could be compensated by using a large value of the maximum number of trades in one simulation, T_{max}. However we will later see that this is impracticable, due to the heavy losses incurred by the sellers during the adaptation process. Therefore this simplification is the only viable solution, conforming to the KISS (Keep-It-Simple-Stupid) principle often used in agent-based modelling, see Axelrod (1997).

4 Results and discussion

This section presents and discusses the three main results of this paper; (i) We examine the RMSE[15] in the model with production and free storage model with the previous results presented by Cliff et al. (1997). (ii) We investigate the efficiency and behaviour of ZI-agents in models with different storage cost. (iii) We compare the performance of the modified ZI-Pe agents with the original ZI-P agents on markets with different storage costs.

All regression analysis in this section have been validated by the use of residual diagnostics and validation of assumptions made by the models. This includes check of irregularities in the residuals, test for heteroscedasticity, verification of global minimization of the risk function etc.[16]

4.1 Agents with production and free storage

We begin by analysing the differences in market performance and agent behaviour of markets with and without trading periods. As previously discussed the trading periods introduces some small fluctuations in the transactional price series in markets populated by machine agents. This fluctuation does not appear in markets populated by human agents. This may be contributed to one of two explanations; human agents accounts for the trading periods which machine agents are unaware of, or some unique mechanism existing only in markets with machine agents. The latter explanation covers the mechanism of inactive agents, a trading period lasts as long as there are active agents. When a new trading period starts, all agents are reactivated and some agents have costs or values largely differing from the main population. This results in an advantage for these agents, who may buy or sell a unit at significantly larger/lower price then the main population of agents.

[15]The Root Mean Square Error is defined as RMSE $= \sqrt{\frac{1}{N}\sum_{i=1}^{N}(x_i - p_i)^2}$, for N estimated values and where p_i is the theoretical equilibrium price and x_i the transactional price. The theoretical equilibrium prices are found as the intercept of the empirical supply and demand functions, using linear interpolation. This solution will introduce an error if there are few active agents on the market and/or due to the assumption of linear and deterministic demand/supply functions. The costs of the sellers are the sum of the deterministic production cost and the stochastic storage cost. Therefore the supply function will be stochastic, which introduces an error into the equilibrium price and therefore the RMSE.

[16]The interested reader may contact the author for more information about the validation and its conclusions.

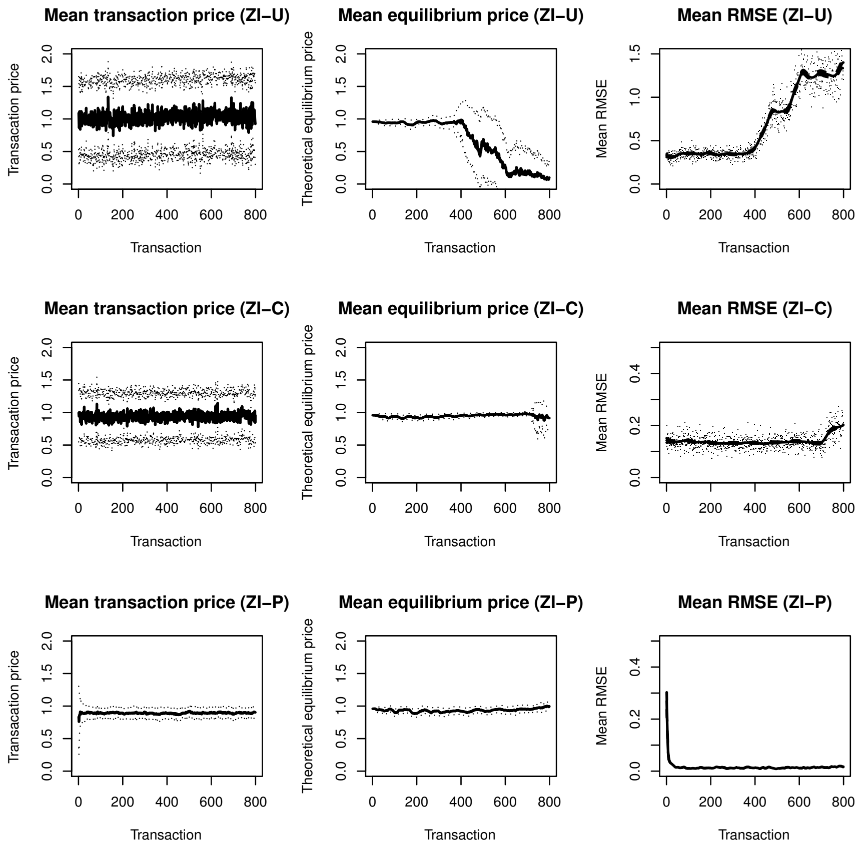

Figure 3: The mean transaction price, equilibrium price and RMSE from 50 simulation runs of a symmetric market populated by agents using production and free storage. Each run lasts a maximum of 10000 shouts or until 800 successful trades. The simulation uses $N_s = 50$ sellers, $N_b = 50$ buyers with maximum cost and redemption value $\Lambda_c = \Lambda_v = 2$. Dotted lines indicates confidence intervals of length two standard deviations. The regression line in the mean RMSE graphs are calculated by non-parametric local linear polynomial regression. Compare with the simulation with trading periods in *Figure 2* in *Section 2.4*.

In *Figure 3* we study a simple case[17] without trading periods on a symmetric market[18]. We will assume constant cost and value for each agent, independent of the number of units sold/bought. We note that the transactional price series lack the small fluctuations, found in previous simulations with trading periods. This result is explained by the previous discussion and therefore also anticipated. When trading periods are removed and the number of transferable units changes uniformly during the simulation. We would therefore not anticipate any larger fluctuations in the transactional prices. This shows that the removal of the trading periods generates a result more similar to the data discussed from markets populated by human agents (previously presented in *Figure 1*).

Studying *Figure 3*, we note the difference in the equilibrium price for the different agent types. The ZI-U agents have a series of equilibrium prices that decreases rapidly after 400 trades. This drop in equilibrium prices is the result of a high number of sellers going bankrupt. The reason for this is the lack of budget constraint, as discussed previously this could generate significant losses for the agent. The position for the sudden drop in the equilibrium price is determined by size of the bankruptcy limit. For the ZI-C and ZI-P agents, the equilibrium price is quite stable during the entire simulation.

We also present the mean RMSE calculated by non-parametric local (linear) polynomial regression[19]. The ZI-P agents have the smallest RMSE and are also the most efficient in trading at the equilibrium prices. The increase in RMSE for the ZI-C and ZI-U agents at the end of the simulation runs, are the result of growing variations of the mean equilibrium price. The ZI-C agents have some larger variations in the equilibrium price at the end of the simulation run. This is explained by an increasing agent bankruptcy, reducing the number of agents and changing the supply and demand functions.

---

[17]Many different market set-ups have been analysed with the same conclusion. Simulation studies have been done on markets varying the market parameters presented in *Appendix A*.

[18]With $N_s = N_b = 50$ sellers and buyers and maximum cost and redemption value $\Lambda_c = \Lambda_v = 2$. All other parameters as described in *Appendix B*.

[19]We have used local (linear) polynomial kernel regression with the use of the Gaussian kernel, to estimate the regression function. The smoothing parameter was chosen by minimizing the risk function, the leave-one-out cross-validation score. The 95 % confidence intervals were estimated by the same method. For an introduction to non-parametric regression, see Wasserman (2010) and for statistically advanced readers, see Haerdle (2004).

4.2 Agents with production and storage with cost

We continue by introducing a fee for sellers storing units instead of trading them. In *Figure 4* we present a symmetric market[20] with no storage fee, $\gamma = 0$. This simulation run is used as a benchmark with other simulation runs where storage fees are introduced. Much of the analysis is covered above but we additional note the large number of active ZI-C agents (the result of difficulties in finding trading partners). Visible by observing that ZI-C agents have 700 successful trades compared to 800 trades for ZI-U and ZI-P agents.

In *Figure 5* we present the results of simulation on a market with a small storage fee, $\gamma = 0.0001$ (per units and shout). In comparison with the benchmark simulation, we note the decrease in the number of active agents and transactions. The most dramatic decrease is found in the number of transactions and active agents of the ZI-U agents. As before, this result is the consequence of a large number of bankrupt ZI-U agents. The ZI-C and ZI-P agents preforms better with more trades and a larger number of active agents on the market.

All agent types experiences a large increase in the mean RMSE and the equilibrium prices at higher number of transactions. This is the result of agents finally going bankrupt and the available number of transferable units on the market decreases rapidly. None of the agent types are therefore capable to handle the storage fee in the long run. We also note that all agents have larger RMSE, compared to the case with storage costs, as predicted by the discussion before this could be contributed to the difficulty of obtaining the theoretical price using the stochastic supply function.

Simulation runs was also conducted with increasing storage fees $\gamma \in \{0.001, 0.01, 0.1\}$ (per units and shout). The analysis is the same as on the market with a smaller storage fee, $\gamma = 0.0001$, above. The increase in the storage fee only decreases the number of trades made, before the rapid decrease of the mean equilibrium prices and the resulting increase in mean RMSE.

[20]With $N_s = N_b = 50$ sellers and buyers, maximum cost and redemption value $\Lambda_c = \Lambda_v = 2$. All other parameters as described in *Appendix B*.

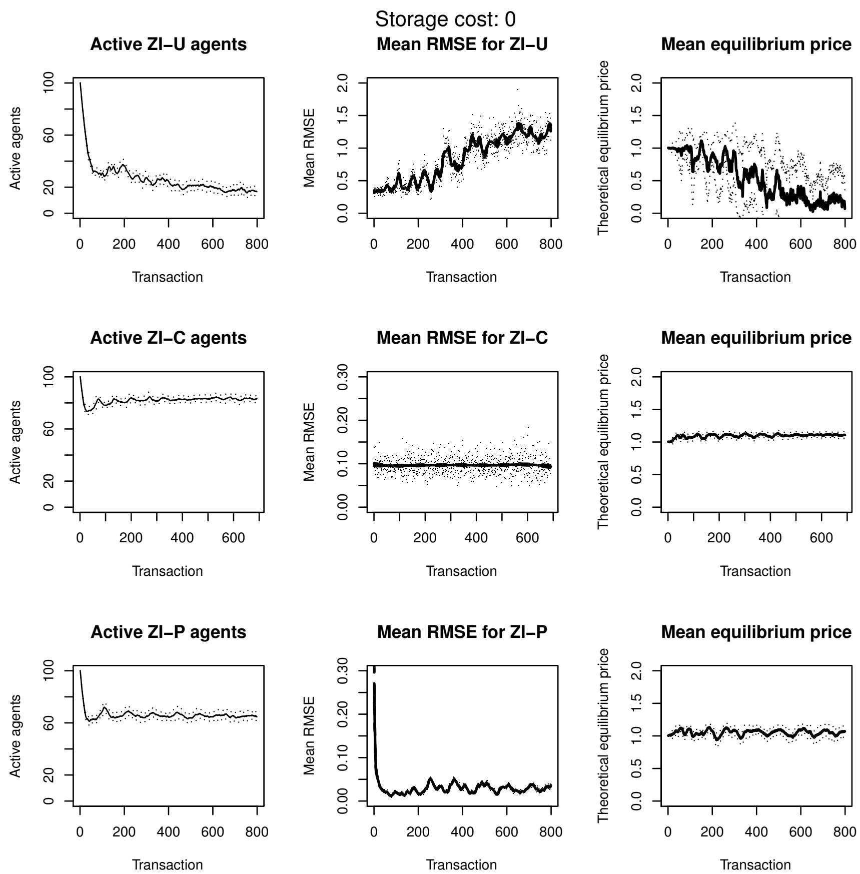

Figure 4: The mean transaction price, equilibrium price and RMSE from 50 simulation runs of a symmetric market populated by agents using production and free storage, $\gamma = 0$. Each run lasts a maximum of 10000 shouts or until 800 successful trades. The simulation uses $N_s = 50$ sellers, $N_b = 50$ buyers with maximum cost and redemption value $\Lambda_c = \Lambda_v = 2$. Dotted lines indicates confidence intervals of length two standard deviations. The regression line in the mean RMSE graphs are calculated by non-parametric local linear polynomial regression.

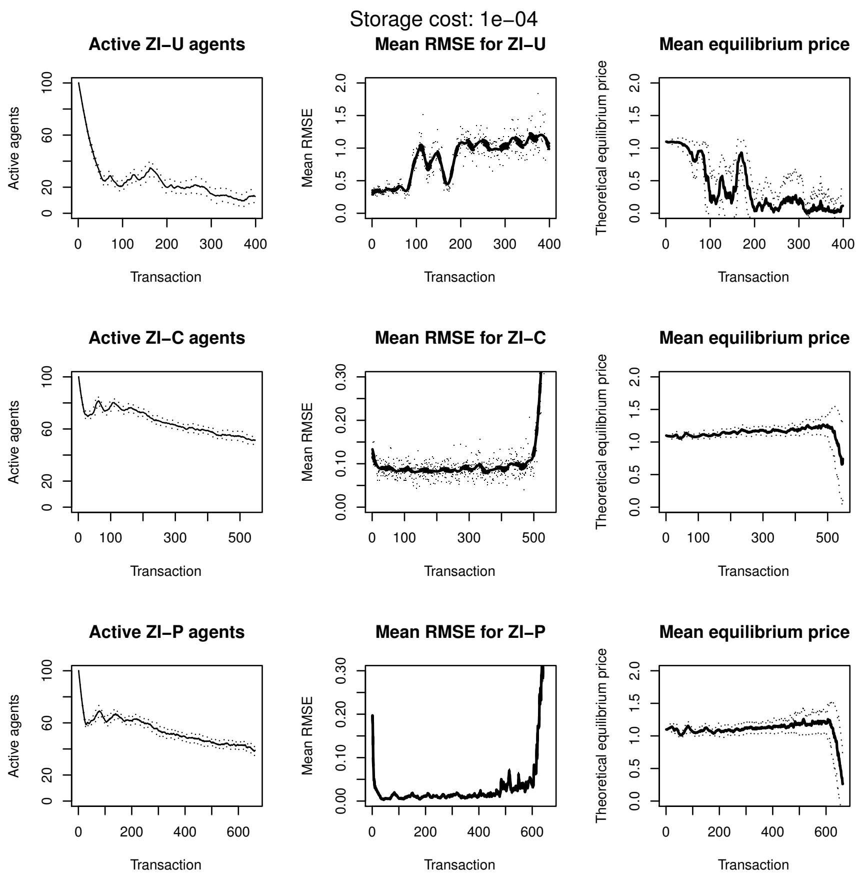

Figure 5: The mean transaction price, equilibrium price and RMSE from 50 simulation runs of a symmetric market populated by agents using production and costly storage, $\gamma = 0.0001$. Each run lasts a maximum of 10000 shouts or until 800 successful trades. The simulation uses $N_s = 50$ sellers, $N_b = 50$ buyers with maximum cost and redemption value $\Lambda_c = \Lambda_v = 2$. Dotted lines indicates confidence intervals of length two standard deviations. The regression line in the mean RMSE graphs are calculated by non-parametric local linear polynomial regression.

4.3 ZI-P agents with storage cost awareness

We propose to resolve the problem with bankruptcy by adding the expected storage cost into the price asked by the ZI-P agents. This allows the sellers to cover some/all of the storage expense by increasing the ask prices. To estimate this expected storage cost, the sellers will use an average of the previous storage costs. On average, the storage costs should be shared by the buyers and seller[21]. These enhanced ZI-P agents will be referred to as ZI-Pe agents.

Simulation studies comprised of 50 runs was conducted to study the efficiency of the original ZI-P agents by Cliff et al. (1997) and the new ZI-Pe agents. In *Figure 6* we study a symmetric market[22] with a small storage fee, $\gamma = 0.0001$, and initial expected storage time as $\bar{t}_{s,0} = 50$ shout attempts.

The main difference between the original ZI-P agents and the ZI-Pe agents, is the number of successful trades during 10000 shouts. The original agents trades 600 units and the modified agents trades 800 units. The ZI-Pe agents do not experience the same large increase in the mean RMSE as the original ZI-P agents. This results in a smoother equilibrium price series without the sudden drop at large number of completed trades. We also note the increase in the mean equilibrium price, which is the result of storage costs being added to the asked prices from the sellers. This results does not generally hold on markets with larger storages fees[23] $\gamma \geq 0.01$, on which the ZI-Pe agents preform almost as bad as the original agents. One important note is that ZI-Pe agents holds onto the market a bit longer, thereby generating more successful trades then ZI-P agents at the same number of attempted shouts. Which in itself is an encouraging result and shows that the ZI-Pe agents work to some extent as intended.

A remaining problem is to increase the pace of which the ZI-Pe agents form the correct storage cost expectations. Preliminary simulations were done where the seller was not allowed to use the storage times of the other sellers. This however created a market where sellers went bankrupt before a

---

[21] From standard economic theory it is known that distribution of an additional cost depends on the buyers price elasticity.

[22] With $N_s = N_b = 50$ sellers and buyers, maximum cost and redemption value $\Lambda_c = \Lambda_v = 2$. All other parameters as described in *Appendix B*.

[23] The size of the largest possible storage fee is also related to the rate of production, Ω, the maximum redemption value, Λ_v, and the maximum production cost, Λ_c. Exactly how these parameters are related to the efficiency of the market populated by ZI-P and ZI-Pe agents remains an area of further study.

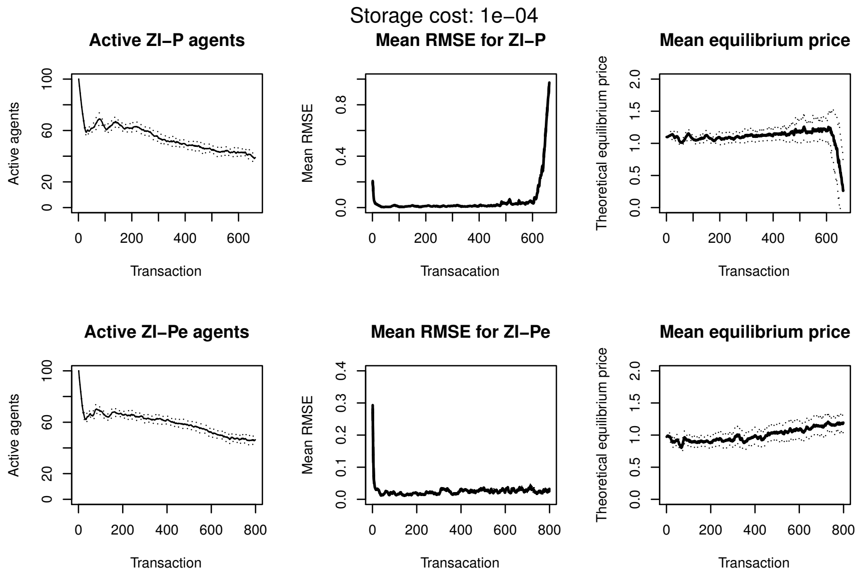

Figure 6: The number of active agents, mean RMSE and mean equilibrium price from 50 simulation runs of a symmetric market populated by agents using production and costly storage, $\gamma = 0.0001$. Each run lasts a maximum of 100000 shouts or until 800 successful trades. The simulation uses $N_s = 50$ sellers, $N_b = 50$ buyers with maximum cost and redemption value $\Lambda_c = \Lambda_v = 2$. Dotted lines indicates confidence intervals of length two standard deviations. The regression line in the mean RMSE graphs are calculated by non-parametric local linear polynomial regression.

good estimate for the storage cost could be found. To resolve this we then allow the seller to know the storage costs of all the other sellers in simulation presented in this paper. This could however create a situation where all agents adds the same expected cost onto the price. This is an undesirable outcome because in real life, both the beliefs and previous experience are used to find the expected storage cost. Therefore the expected storage cost should differ from seller to seller. This could be resolved by adding some random term to the estimated expected storage time. Which would act as some noise thereby giving the sellers heterogeneous estimated storage costs.

5 Concluding remarks

The aim of this paper was to analyse the effects of removing trading periods, in the agent-based models populated by ZI-agent. This was done by introducing storage and production into the market. As discussed in the previous section, this created a price series more like the one observed in experiments of markets with human agents. This new model is therefore a step towards finding the smallest needed intelligence and the simplest possible market, that will mimic the human markets satisfactory.

We have successfully demonstrated that the removal of the trading periods, gave no larger differences in comparison with the previous work done by Gode and Sunder (1993) and Cliff et al. (1997). But small changes created a market model that better corresponds with the real-world market, where buyers and sellers enters at random trading different number of units. The generalization of the market is therefore relevant in the context of trying to create more realistic markets. It is also desirable to keep the agents as simple as possible and it is therefore an interesting result that no more rationality is needed, for the agents to act efficiently on this new market.

The next step was the introduction of a storage fee into the market. Even when a small fee was charged sellers, for storing units instead of trading, the ZI-agents did not create an efficient market. Most of the agents went bankrupt and thereby limiting the market efficiency. We might therefore conclude that ZI-agents presented by Gode and Sunder (1993) and Cliff et al. (1997) are unable to adapt to an environment with storage fees. Some more rational behaviour (intelligence) is needed to enable the agents to survive longer and therefore getting some more revenues to stay off bankruptcy. This requires some foresight and planning capability of the agents, they need to have some belief of how long they have to store the unit, before they may sell it on the market. This result shows that for machine agents to trade on a market with storage and production (random events associated with revenue and cost) require much more intelligence then the ZI-agents, discussed in this paper.

This could have real-world implications when DA-markets are used at stock exchanges etc. where the use of algorithm trading and problems with transaction costs and asymmetric information are complicating factors. Simulations could be used to analyse these situations and thereby conclude if the Walrasian equilibrium is reached using the market rules and ZI-agents.

In the last step, we enhanced the ZI-P agents with the simplest mechanism to form an expected future storage cost. This was based on the average of the mean storage time of all the sellers. This mechanisms enabled the modified ZI-P agents to preform somewhat better then the original agents, in market with small storage fees. This mechanism was not able to handle the markets with higher storage fees and more work needs to be done to find a better (but still simple) mechanism to adapt to storage costs.

There are many areas in which further work may be done to improve the agents and markets presented in this paper. The algorithms that enables the ZIP-agents to adapt their mark-up, needs to be extensively modified to allow the agents to also include the storage cost in the mark-up calculations. This will require the agents to calculate some expected storage time (cost) with the information available, it is possible to use some kind of adaptation[24] using *reinforced learning* or *belief-based learning* to solve this problem. The market model could also be improved upon by letting the buyers experience some inflation in the redemption values, waiting longer would give then a lower net utility. This would also require some modification of the buyers, so that they are aware of this inflation and may act to minimize its effects.

A more complete study of the sensitivity of the model is also needed, to ensure that the results of this paper holds even with small variations in the parameters determining the model. This would include more simulation trials in different market setups, with a range of different demand/supply functions and number of agents/buyers/sellers. The interaction between storage fees, bankruptcy limits and production times should also be studied to find suitable values for these parameters, that matches observed real-world markets.

Further work could also be devoted to study DA-markets with friction populated by machine agents. As discussed above, the real-world applications of DA-markets are often found in the financial industry. It would therefore be justified to study the implications of transaction fees (generalizing the storage fee), introducing assets[25] instead of commodities, generating some dividends (random profits for sellers and buyers). In this model one could study phenomenon related to the recent/current financial crisis, i.e. the role of debt/losses on rational agents and market crashes. Some previous work in this area includes experimental markets in Smith et al. (1988) and markets populated by machine agents in Duffy and Ünver (2006).

[24]These methods are presented in Duffy (2006) and compared in Feltovich (2000) and Salmon (2001).

[25]The use of experiments to analyse asset markets is discussed in Sunder (1995).

References

Axelrod, R. (1997): *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration*, Princeton University Press, 1st printing ed.

Becker, G. S. (1962): "Irrational Behavior and Economic Theory," *The Journal of Political Economy*, 70, 1–13.

Brewer, P., M. Huang, B. Nelson, and C. Plott (2002): "On the Behavioral Foundations of the Law of Supply and Demand: Human Convergence and Robot Randomness," *Experimental Economics*, 5, 179–208.

Cason, T. N. and D. Friedman (1996): "Price formation in double auction markets," *Journal of Economic Dynamics and Control*, 20, 1307 – 1337.

Cliff, D., J. Bruten, and F. Road (1997): "Zero is Not Enough: On The Lower Limit of Agent Intelligence for Continuous Double Auction Markets," Tech. rep., Hewlett-Packard Research Labs Paper HPL-97-157, Bristol, England.

Duffy, J. (2006): "Agent-Based Models and Human Subject Experiments," in *Handbook of Computational Economics: Agent-based Computational Economics*, ed. by L. Tesfatsion and K. L. Judd, North-Holland, no. 13 in Handbooks in Economics, chap. 19, 949–1011.

Duffy, J. and M. U. Ünver (2006): "Asset price bubbles and crashes with near-zero-intelligence traders," *Economic Theory*, 27, 537–563.

Feltovich, N. (2000): "Reinforcement-Based vs. Belief-Based Learning Models in Experimental Asymmetric-Information," *Econometrica*, 68, 605–642.

Friedman, D. (1984): "On the Efficiency of Experimental Double Auction Markets." *American Economic Review*, 74, 60.

——— (1993): *The Double Auction Market: Institutions, Theories, And Evidence (Santa Fe Institute Studies in the Sciences of Complexity Proceedings)*, Westview Press.

Gilbert, N. (2007): *Agent-based Models*, no. 153 in Quantitative Applications in the Social Sciences, Sage Publications.

Gjerstad, S. and J. Dickhaut (1998): "Price Formation in Double Auctions," *Games and Economic Behavior*, 22, 1–29.

Gode, D. K. and S. Sunder (1993): "Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality," *Journal of Political Economy*, 101, 119–137.

Gravelle, H. and R. Rees (2004): *Microeconomics*, Harlow: Prentice Hall, 3 ed.

Haerdle, W. (2004): *Nonparametric and semiparametric models*, Springer Series in Statistics, Springer.

Mas-Colell, A., M. D. Winston, and J. R. Green (1995): *Microeconomic theory*, New York: Oxford Univ. Press.

Salmon, T. C. (2001): "An Evaluation of Econometric Models of Adaptive Learning," *Econometrica*, 69, 1597–1628.

Smith, V. L. (1962): "An Experimental Study of Competitive Market Behavior," *Journal of Political Economy*, 70, 322.

——— (1982): "Microeconomic Systems as an Experimental Science," *American Economic Review*, 72, 923–55.

Smith, V. L., G. L. Suchanek, and A. W. Williams (1988): "Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets," *Econometrica*, 56, 1119–51.

Sunder, S. (1995): *Experimental asset markets: A survey*, Princeton NJ: Princeton University Press, chap. 6, 445–500.

Tesfatsion, L. (2006): "Agent-based computational economics: A constructive approach to economic theory." in *Handbook of Computational Economics: Agent-based Computational Economics*, ed. by L. Tesfatsion and K. L. Judd, North-Holland, no. 13 in Handbooks in Economics, chap. 16, 831—894.

Varian, H. (2006): *Intermediate microeconomics: a modern approach*, Norton.

Wasserman, L. (2010): *All of Non-parametric Statistics*, Springer Texts in Statistics, Springer.

A Algorithms

This appendix presents the algorithms used in the simulations of the ZI-agents, ZI-P agents and the modified agents with production and storage.

A.1 ZI-U and ZI-C agents

The ZI-agents are introduced in *Section 2.2* and does only depend on the sequence of the program presented in *Algorithm 1* together with the calculation of shout prices. Which are random variables, presented in *Section 2.2*, offers as $a_{i,k} \sim \mathcal{U}[0, \Lambda_c]$ or bids $b_{j,k} \sim \mathcal{U}[0, \Lambda_v]$ for ZI-U agents. For ZI-C agents the shouts are generated as, offers $a_{i,k} \sim \mathcal{U}[c_{i,k}, \Lambda_v]$ and bids $b_{j,k} \sim \mathcal{U}[0, r_{j,k}]$ because these shouts must yield some profit (or no loss) for the agent. The redemption value (for agent j and good k) $r_{j,k}$ and the cost for production (by agent i and good k) $c_{i,k}$, was also calculated as an uniformly distributed random variable, as $c_{i,k} \sim \mathcal{U}[0, \Lambda_c]$ and $r_{j,k} \sim \mathcal{U}[0, \Lambda_v]$, for some values of Λ_c and Λ_v. The parameters are summarized in *Table 1* below.

Description	Name
Max. Cost for Seller	Λ_c
Max. Redemption Value for Buyer	Λ_v
No. Sellers	N_s
No. Buyers	N_b

Table 1: The required parameters of the ZI-U and ZI-C agents.

Algorithm 1 Sequence: Gode and Sunder (1993) Double Auction Market with ZI-agents.

Require: Some maximum values of the possible value Λ_v and costs Λ_c for all agents.
  for Each seller $i \in \{1, 2 \ldots, m\}$ and buyer $j \in \{1, 2 \ldots, n\}$. **do**
    $r_{j,k} \sim \mathcal{U}[0, \Lambda_v]$ and $c_{j,k} \sim \mathcal{U}[0, \Lambda_c]$.
  end for
  for A number of turns, T. **do**
    while Not all agents have traded **do**
      Select a random agent l and generate an offer or a bid (depending on agent type). For ZI-U as in (1), and for ZI-C as in (2).
      if The agent is a seller **then**
        Compare the offer to the order book, if there exists no other offers this will be the best offer $a_l \to a^*$. If this offer is smaller then the current best offer, $a_l < a^*$ this offer will replace the best offer $a_l \to a^*$. If this offer is smaller then the current best bid $a_l \leq b^*$, then *trade* at b^* and clear the order book.
      else
        We have a buyer. Compare the bid to the order book, if there exists no other bids this will be the best bid $b_l \to b^*$. If this bid is greater then the current best bid, $b_l > b^*$ this offer will replace the best bid $b_l \to b^*$. If this bid is greater then the current best offer $b_l \geq a*$, then *trade* at a^* and clear the order book.
      end if
    end while
  end for

A.2 ZI-P agents

The ZI-P agents introduced in *Section 2.3* are comprised of a number of algorithms. The sequence of the simulation is shown in *Algorithm 2* below and includes an algorithm for determining mark-up changes in *Algorithm 3* and an adaptive algorithm (that calculates the mark-up changes) in *Algorithm 4*. The program needs a lot of parameters specified, to be able to conduct the simulations. These parameters are specified in connection to the results presented in the report and in *Appendix B*. A summary of the relevant parameters in shown below in *Table 2*.

Description	Name	Value
Max. Cost for Seller	Λ_c	
Max. Redemption Value for Buyer	Λ_v	
No. Sellers	N_s	
No. Buyers	N_b	
Learning coefficient	η	$\eta_i \sim \mathcal{U}[0, \eta_0]$ with $\eta_0 = 0.1$.
Momentum coefficient	β	$\beta_i \sim \mathcal{U}[\beta_{0,1}, \beta_{0,1} + \beta_{0,2}]$
		with $\beta_{0,1} = 0.1$ and $\beta_{0,2} = 0.30$.
Initial Mark-up (sellers)	$\mu_i(0)$	$\mu_i(0) \sim \mathcal{U}[\mu_{0,1}, \mu_{0,1} + \mu_{0,2}]$.
Initial Mark-up (buyers)	$\mu_i(0)$	$\mu_i(0) \sim \mathcal{U}[-(\mu_{0,1} + \mu_{0,2}), -\mu_{0,1}]$.
Seller Cost	c_i	The ordered vector (ascending) c_i,
		where $c_{i,(k)} \sim \mathcal{U}[0, \Lambda_c]$ for $k \in \{1, 2, \ldots, n_s\}$.
Buyer Redemption Value	r_i	The ordered vector (descending) r_i,
		where $r_{j,(k)} \sim \mathcal{U}[0, \Lambda_r]$ for $k \in \{1, 2, \ldots, n_b\}$

Table 2: The parameters used in the simulation of ZI-P agents using the sequence in *Algorithm 2*.

Algorithm 2 Sequence: Cliff, Bruten and Road (1997) ZI-P agents.

Initialize all required variables and set constants to pre-specified values.
Generate the agent characteristics, using randomized values of the constants.
for A number of trading turns **do**
  Reactive all agents and reset some variables.
  while The number of trades $T < T_{limit}$ for some limit value T_{limit} and as long as all agents are not inactive. **do**
    Use *Algorithm 3* to adjust the mark-up.
    Select an active agent l by random and generate a shout, by (3).
    if The agent is a seller **then**
      Compare the offer to the order book, if there exists no other offers this will be the best offer $a_l \to a^*$. If this offer is smaller then the current best offer, $a_l < a^*$, this offer will replace the best offer $a_l \to a^*$. If this offer is smaller then the current best bid $a_l \leq b^*$, then *trade* at b^* and clear the order book and make the trading agents inactive. Let $T \leftarrow T + 1$.
    else
      We have a buyer. Compare the bid to the order book, if there exists no other bids this will be the best bid $b_l \to b^*$. If this bid is greater then the current best bid, $b_l > b^*$, this offer will replace the best bid $b_l \to b^*$. If this bid is greater then the current best offer $b_l \geq a*$, then *trade* at a^* and clear the order book and make the trading agents inactive. Let $T \leftarrow T + 1$.
    end if
  end while
end for

Algorithm 3 Cliff, Bruten and Road (1997) Price adaptation in Double Auction Market with ZI-P agents.

```
Require: The latest shout price q(t) and if it was accepted. The last bid b_i
  or offer a_i for the agent i.
  if The agent is a seller, i then
    if The last shout q was accepted then
      Any seller i for which a_i ≤ q should raise its margin using Algorithm
      4.
      if The last shout was a bid then
        Any active seller i for which a_i ≥ q should lower its margin using
        Algorithm 4.
      end if
    else
      if The last shout was an offer then
        Any active seller i for which a_i ≥ q should lower its margin using
        Algorithm 4.
      end if
    end if
  else
    The agent is a buyer, i
    if The last shout q was accepted then
      Any buyer i for which b_i ≥ q should raise its margin using Algorithm
      4.
      if The last shout was an offer then
        Any active buyer i for which b_i ≤ q should lower its margin using
        Algorithm 4.
      end if
    else
      if The last shout was a bid then
        Any active buyer i for which b_i ≤ q should lower its margin using
        Algorithm 4.
      end if
    end if
  end if
```

Algorithm 4 Cliff, Bruten and Road (1997) Mark-up adaptation in Double Auction Market with ZI-P agents.

Require: The last shout price $q(t)$, the last offer/bid from the agent $p_i(t-1)$.

if The agent would like to increase its mark-up $\mu_i(t)$ **then**

Calculate a new target price $\tau_i(t)$ for the agent using (6) with positive values for $\mathcal{R}_i(t)$ and $\mathcal{B}_i(t)$.

else

(The agent would like to decrease its mark-up $\mu_i(t)$)

Calculate a new target price $\tau_i(t)$ for the agent using (6) with negative values for $\mathcal{R}_i(t)$ and $\mathcal{B}_i(t)$.

end if

Calculate the new value of the mark-up correction, $\Delta_i(t-1)$, using (7).

Calculate the new momentum-based correction $\psi_i(t)$ using (8).

Calculate the new mark-up $\mu_i(t)$ for the agent using (5).

if The new mark-up will generate a loss, $\mu_i(t) < 0$ for sellers or if $\mu_i(t) > 0$ or $\mu_i(t) < -1$ for buyers. **then**

Use the old mark-up, $\mu_i(t) = \mu_i(t-1)$.

end if

return The mark-up $\mu_i(t)$ and the value of the momentum-based correction $\psi_i(t)$.

A.3 Agents with production and storage

The modified agents uses nearly the same price adaptations as for the original ZI och ZI-P agents but with some minor changes to use more then one redemption value and cost. We also need to modify the sequence to incorporate the production of units, bankruptcy, storage costs and agents entering and exiting the market. A general outline of the sequence of the simulation is shown in *Algorithm 5* and is used with some different methods to generate the pricing. For ZI-U agents we use the same price generating process as before i.e.

$$\text{offers: } a_{i,k} \sim \mathcal{U}[0, \Lambda_c] \qquad \text{bids: } b_{j,k} \sim \mathcal{U}[0, \Lambda_v],$$

where Λ_c and Λ_v are predetermined values for the maximum production cost and redemption value. For ZI-C agents we introduce that the agents may sell with some losses if this loss is smaller then the expected total storage cost of holding the unit. Therefore we rewrite the price generating process as i.e.

$$\text{offers: } a_{i,k} \sim \mathcal{U}[c_{i,k} + S_{i_k}, \Lambda_c] \qquad \text{bids: } b_{j,k} \sim \mathcal{U}[0, r_{j,k}]$$

where $S_{i_j} = \gamma \mathbf{E}[T_s]$ is the expected storage cost calculated as discussed in *Section 3*. The redemption value for agent j of good k, $r_{j,k}$ and the cost for production by agent i and good k, $c_{i,k}$, was also calculated as a uniformly distributed random variable as $c_{i,k} \sim \mathcal{U}[0, \Lambda_c]$ and $r_{j,k} \sim \mathcal{U}[0, \Lambda_v]$ for some values of Λ_c and Λ_v. The ZI-P agents who is sellers uses *Algorithm 6* to update their prices with the modification that the cost is comprised of the production cost $c_{i,k}$ and the expected storage cost $S_{i,k}$ as

$$a_i(t) = c_{i,k}(1 + \mu_i(t)) + S_{i,k} \tag{9}$$

These programs requires a lot of parameters specified to be able to conduct the simulations. These parameters are specified in connection to the results presented in the report and in *Appendix B*. A summary of the relevant parameters in shown below in *Table 3*.

Description	Name	Value
Max. Cost for Seller	Λ_c	
Max. Redemption Value for Buyer	Λ_v	
Max. time of Production	Ω	
Limit of Bankruptcy	Π	
Storage Cost (sellers)	γ	
No. Sellers	N_s	
No. Buyers	N_b	
Max. No. Sold Units (per turn)	n_s	
Max. No. Bought Units (per turn)	n_b	
Learning coefficient	η	$\eta_i \sim \mathcal{U}[0, \eta_0]$ with $\eta_0 = 0.1$.
Momentum coefficient	β	$\beta_i \sim \mathcal{U}[\beta_{0,1}, \beta_{0,1} + \beta_{0,2}]$
		with $\beta_{0,1} = 0.1$ and $\beta_{0,2} = 0.30$.
Initial Mark-up (sellers)	$\mu_i(0)$	$\mu_i(0) \sim \mathcal{U}[\mu_{0,1}, \mu_{0,1} + \mu_{0,2}]$
Initial Mark-up (buyers)	$\mu_i(0)$	$\mu_i(0) \sim \mathcal{U}[-(\mu_{0,1} + \mu_{0,2}), -\mu_{0,1}]$.
Seller Cost	c_i	The ordered vector (ascending) c_i,
		where $c_{i,(k)} \sim \mathcal{U}[0, \Lambda_c]$ for $k \in \{1, 2, \ldots, n_s\}$.
Buyer Redemption Value	r_j	The ordered vector (descending) r_j,
		where $r_{j,(k)} \sim \mathcal{U}[0, \Lambda_r]$ for $k \in \{1, 2, \ldots, n_b\}$
Production level	ω_i	$\omega_i \sim \mathcal{U}[0, \Omega]$.

Table 3: The required parameters used with the modified agents with storage and production.

Algorithm 5 Sequence: Agents with production and storage.

Initialize all required variables and set constants to pre-specified values.
Generate the agent characteristics using randomized values of the constants.
while The number of trades $T < T_{limit}$ for some limit value T_{limit} and as long as all agents are not inactive. **do**
  Check if agents are bankrupt (if $\sum_k \pi_{i,k} < \Pi$), add storage fees for producers with surplus units (add loss $\gamma N_{i,n}$, where $N_{i,n}$ is the number of units in storage for agent i at shout attempt n.) and produce units if necessary (if $\mod(n, \omega_i) = 0$ for some agent i at shout attempt n).
  Pricing adaptations for ZI-P agents: Use *Algorithm 6* to adjust the mark-up.
  Select an active agent l by random and generate a shout.
  if The agent is a seller **then**
    Compare the offer to the order book, if there exists no other offers this will be the best offer $a_l \to a^*$. If this offer is smaller then the current best offer, $a_l < a^*$ this offer will replace the best offer $a_l \to a^*$. If this offer is smaller then the current best bid $a_l \leq b^*$, then make a trade at b^* and clear the order book and decrease the maximum number of transferable units for the trading agents. Also check if the agents still are active of if they have no more units to trade. Let $T \leftarrow T + 1$.
  else
    We have a buyer. Compare the bid to the order book, if there exists no other bids this will be the best bid $b_l \to b^*$. If this bid is greater then the current best bid, $b_l > b^*$ this offer will replace the best bid $b_l \to b^*$. If this bid is greater then the current best offer $b_l \geq a*$, then make a trade at a^* and clear the order book and decrease the maximum number of transferable units for the trading agents. Also check if the agents still are active of if they have no more units to trade. Let $T \leftarrow T + 1$.
  end if
  Let $k \leftarrow k + 1$.
end while

Algorithm 6 Mark-up adaptation in Double Auction Market with ZI-P agents, multiple units and storage costs.

Require: The last shout price $q(t)$, the last offer/bid from the agent $p_i(t-1)$.
  if The agent would like to increase its mark-up $\mu_i(t)$ **then**
    Calculate a new target price $\tau_i(t)$ for the agent using (6) with positive values for $\mathcal{R}_i(t)$ and $\mathcal{B}_i(t)$.
  else
    (The agent would like to decrease its mark-up $\mu_i(t)$)
    Calculate a new target price $\tau_i(t)$ for the agent using (6) with negative values for $\mathcal{R}_i(t)$ and $\mathcal{B}_i(t)$.
  end if
  Calculate the new value of the mark-up correction, $\Delta_i(t-1)$, using (7).
  Calculate the new momentum-based correction $\psi_i(t)$ using (8).
  Calculate the new mark-up $\mu_i(t)$ for the agent using (5).
  if The new mark-up, $\mu_i(t) < 0$, for a seller will generate a loss that is greater then the storage cost **then**
    Use the old mark-up, $\mu_i(t) = \mu_i(t-1)$.
  end if
  if The new mark-up will generate a loss for buyers i.e $\mu_i(t) > 0$ or $\mu_i(t) < -1$. **then**
    Use the old mark-up, $\mu_i(t) = \mu_i(t-1)$.
  end if
  return The mark-up $\mu_i(t)$ and the value of the momentum-based correction $\psi_i(t)$.

B Parameters used in simulation runs

This appendix presents the parameters used in the simulation studies.

Description	Name	Value
Max. Cost for Seller	Λ_c	2.00
Max. Redemption Value for Buyer	Λ_v	2.00
No. Sellers	N_s	50
No. Buyers	N_b	50
Learning coefficient	η	$\eta_0 = 0.1$.
Momentum coefficient	β	$\beta_{0,1} = 0.1$ and $\beta_{0,2} = 0.30$.
Initial Mark-up (sellers)	$\mu_i(0)$	$\mu_{0,1} = 0.05$ and $\mu_{0,2} = 0.30$.
Initial Mark-up (buyers)	$\mu_i(0)$	$\mu_{0,1} = 0.05$ and $\mu_{0,2} = 0.30$.
Number of maximum trade periods		6
Number of maximum shouts		10000

Table 4: The parameters used in the simulation run shown in *Figure 2*.

Description	Name	Value
Max. Cost for Seller	Λ_c	2.00
Max. Redemption Value for Buyer	Λ_v	2.00
Max. Production Level	Ω	1000
No. Sellers	N_s	50
No. Buyers	N_b	50
Limit of Bankruptcy	Π	-5
Storage Cost (sellers)	γ	0
Learning coefficient	η	$\eta_0 = 0.1$.
Momentum coefficient	β	$\beta_{0,1} = 0.1$ and $\beta_{0,2} = 0.30$.
Initial Mark-up (sellers)	$\mu_i(0)$	$\mu_{0,1} = 0.05$ and $\mu_{0,2} = 0.30$.
Initial Mark-up (buyers)	$\mu_i(0)$	$\mu_{0,1} = 0.05$ and $\mu_{0,2} = 0.30$.
Number of maximum shouts		10000
Number of maximum trades		800

Table 5: The parameters used in the simulation run shown in *Figure 3*.

Description	Name	Value
Max. Cost for Seller	Λ_c	2.00
Max. Redemption Value for Buyer	Λ_v	2.00
Max. Production Level	Ω	1000
No. Sellers and No. Buyers	N_s, N_b	50
Limit of Bankruptcy	Π	-5
Storage Cost (sellers)	γ	0, 0.0001 and 0.01
Learning coefficient	η	$\eta_0 = 0.1$.
Momentum coefficient	β	$\beta_{0,1} = 0.1$ and $\beta_{0,2} = 0.30$.
Initial Mark-up (sellers)	$\mu_i(0)$	$\mu_{0,1} = 0.05$ and $\mu_{0,2} = 0.30$.
Initial Mark-up (buyers)	$\mu_i(0)$	$\mu_{0,1} = 0.05$ and $\mu_{0,2} = 0.30$.
Number of maximum shouts		10000
Number of maximum trades		800

Table 6: The parameters used in the simulation run shown in *Figure 4* and *Figure 5*.

Description	Name	Value
Max. Cost for Seller	Λ_c	2.00
Max. Redemption Value for Buyer	Λ_v	2.00
Max. Production Level	Ω	1000
No. Sellers and No. Buyers	N_s, N_b	50
Limit of Bankruptcy	Π	-5
Storage Cost (sellers)	γ	0.0001 (and 0.01)
Learning coefficient	η	$\eta_0 = 0.1$.
Momentum coefficient	β	$\beta_{0,1} = 0.1$ and $\beta_{0,2} = 0.30$.
Initial Mark-up (sellers)	$\mu_i(0)$	$\mu_{0,1} = 0.05$ and $\mu_{0,2} = 0.30$.
Initial Mark-up (buyers)	$\mu_i(0)$	$\mu_{0,1} = 0.05$ and $\mu_{0,2} = 0.30$.
Initial expected storage time (sellers)	$\bar{t}_{s,0}$	50
Number of maximum shouts		100000
Number of maximum trades		800

Table 7: The parameters used in the simulation run shown in *Figure 6*.