SIMPLE AGENT-BASED COMPUTATIONAL MODEL OF COSTLESS MARKET WITHOUT INTERMEDIATION

Michal Kvasnička, Masaryk University, Faculty of Economics, qasar@econ.muni.cz

progress report, January 12, 2009

1 Inspiration and Goal

Smith (1962) has experimentaly shown that markets could be efficient (reach the equilibrium price predicted by the competitive market theory) even though the conditions usually regarded as necessary (many traders, perfect knowledge etc.) are not satisfied. His results can be contrasted with results of the previous experiments done by Chamberlin (1948). Since their experiments differred in the institutional setting (Smith used the continuous double auction markets,⋆ CDA, with many "trading days" (repetitions) while Chamberlin's market was unorganized and carried out only once), the difference might suggest that it is the *market institutions* that cause the efficiency.

This judgment might be strengthened by the results of the simulations by Gode and Sunder (1993). They have shown that artificial CDA markets were efficient (measured with total traders' surplus) even when they replaced human traders by 'zero-intelligence' software robots (ZI-C traders) that submitted random bids and offers.

Later Cliff (1997) has demonstrated some deficiencies of Gode and Sunder's approach, and designed a simple trading algorithm for the CDA markets (ZI-P traders) that does not suffer from the problems affecting Gode and Sunder's traders, and secured much higher efficiency of the CDA market.

Both Gode and Sunder's and Cliff's simulations and most of the Smith's experiments cited above were done on the CDA market. The question is how would the market efficiency change when we change the market institutions.

The goal of this paper is to explore this question: to assess to what extent does the market efficiency depend on the market institutions, or to put it in other words, how would the market efficiency deteriorate when we depart from the CDA market. Our initial hypothesis is that the market efficiency (measured as Smith's α) depends on market institutions—it decreases (α rises) as we depart from the CDA market.

This question could be possibly answered experimentally. However, we have chosen another approach: agent-based computational model⋆⋆ similar to Cliff's one. In this report, we present the first results of our simulations in the hope to gain a feedback. The results of our simulation do support the initial hypothesis only partially—the results are more complex than expected, and some of them surprised us.

2 Description of the Model

The model is an agent-based computational model. It consists of the traders—ZI-P software robots (we plan to add other trading algorithms later), and the trade world described by market institutions. The markets is simulated many times, and the results are analyzed statistically.

2.1 Traders

Traders are Cliff's ZI-P software robots.⋆⋆⋆ They are either buyers, or sellers. Each trader has his own private reservation price, and one unit of a commodity to trade in each trading day. They trade directly (i.e. without intermediation) for money. The monetary endowment of the buyers is unlimited so that the only constraint are the reservation prices.

The traders have (beside others) these state variables:

- `non-satisfied-quantity`
- `reservation-price`
- `haggling-price`

⋆ The continuous double auction (CDA) market is a market where any trader can in any moment bid or ask, and receive the last bid or offer in which case he and the last bidding (offering) trader make a binding contract. All traders are gathered together so that any trader can hear all bids and offers on the market, and can possibly accept them.

⋆⋆ To learn more about the agent-based computational economics, see L. Tesfatsion's web http://www.econ.iastate.edu/tesfatsi/ace.htm.

⋆⋆⋆ Since the traders are robots, not human beings, we will refer to them as to "males". There are no "she-robots" in our artificial world.

The non-satisfied-quantity says how many units is the trader still about to buy or sell in this trading day (presently one, or zero). It is reset to one at the beginning of each trading day. When the trader makes a deal, the value of this variable is decreased by one. When the non-satisfied-quantity is positive, the trader is said to be "active"; otherwise he is "inactive".

The reservation-price is the minimum price for which a seller can sell, or the maximum price for which a buyer can buy. A seller (buyer) is allowed to sell (buy) at his reservation-price (having no gain) but he is not allowed to sell (buy) cheaper (dearer). If a seller (buyer) can sell (buy) for more (less), he gains a gain from trade equal to the absolute value of the difference between the actual price and his reservation-price. Each trader tries to maximize his gain from trade; however, when the gain from trade is zero, he is still willing to trade. The reservation-price of each trader is his own private knowledge—no other trader can know it. Each trader's reservation-price is constant in all trading days of one simulation.

Each trader tries to maximize his gain from trade. For this reason he has his haggling-price: a seller (buyer) accepts only those bids (offers) that are higher (lower) or equal to his haggling-price. Initially, the haggling-price of a seller (buyer) is set to his reservation-price plus (minus) some random margin, and it is adjusted later as the trader gains more information about the "market conditions" from the actual transactions, see below.

The haggling-price initialization based on a random margin seems to be intuitive—it seems natural to add (subtract) a margin (a given percentage) to (from) the one's reservation price when a trader has no information what the market price could be. However, it has an interesting outcome: we can expect that the actual prices would in average be below the theoretic equilibrium price in the first trading day, at least if the supply and demand curves are symmetric. The reason is obvious from the figure 1: since the intra-marginal buyers are the buyers with the highest reservation prices, they lower their "haggling" prices in comparison to their reservation prices most; on the other hand, the intra-marginal sellers are those with the lowest reservation prices, so they increase their "haggling" prices very little above their reservation prices. The space between the "haggling" supply and demand is biased downward. Therefore, it seems likely that few first deals would be below the theoretic equilibrium. Moreover, since the traders can learn from the market via ZI-P algorithm (described below), they herd toward the last bids, which can further strengthen this tendency for a while.

It seems that this tendency might be more than just a property of our model. For instance, Chamberlin (1948) reported the average prices to be below the theoretic equilibrium in his experiments. He tried to explain it on psychologic grounds. However, it seems that the "haggling" prices initialization is a more likely culprit.*

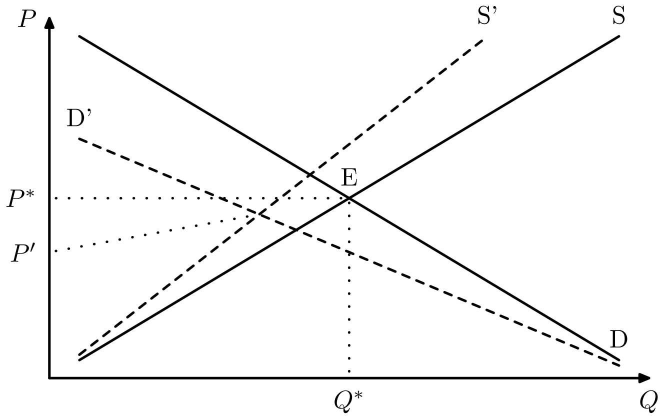

Figure 1 Initial deals are more likely to be below the theoretic equilibrium price than above it if the supply and demand curves are symmetric. The "true" supply and demand derived from the traders' reservation prices are the solid lines labeled as S and D. The theoretic equilibrium price is P^*. Let us suppose that all buyers subtract from their reservation prices a 30 % margin, and all sellers add it. The resulting "haggling" supply and demand curves are the dashed lines labeled S', and D'. The space between these "haggling" curves is downward biased, and their intersection is below the theoretic equilibrium price—compare the horizontal dotted line ending in P^* to the declining dotted line ending in P', which halves the space between the "haggling" curves.

The trade is possible between two traders if they are active, can see each other (see below), one of them is a seller, and the other is a buyer, and if the seller's haggling-price is lower or equal to the buyer's one.

2.2 Market

The traders live in a rectangular toroid world, which defines natural "distance" between them. Initially, traders are randomly placed in the world, see figure 2.

Market institutions are given by five parameters (they are the same for all traders):

- vision% (0–100 %)
- who-offers? (buyers / sellers / both)
- public-offers? (true / false)
- public-hearing? (true / false)
- moving-type (moving / not-moving)

The vision% parameter determines how much of the world can each trader see. The vision% = 30 % means that each trader can see all other traders positioned within a circle in center of which he stands; the size of the circle covers 30 % of the world, so he can see 30 % of all traders in average, see figure 2. If the vision% = 100 %, each

* The judgment that my model behaves like this is based on the above stated theoretic reasoning, and on a casual observations; I haven't done the proper statistical analysis of the price evolution yet.

trader can see all other traders. A trader can trade only with the traders he can see. The `vision%` parameter can be interpreted as the degree of a formal market integration. If the `vision%` = 100 %, then the market is fully integrated; if the `vision%` is lower, then the whole market is broken into many overlapping sub-markets (the sub-markets overlap because the vision-circles of the traders overlap).

The `who-offers?` parameter determines what type of traders (buyers, sellers, or both) are able to initiate an exchange (i.e. bid, or offer).

The `public-offers?` parameter determines whether the trader initiating the deal can address *all* active traders of the opposite type within his vision area (true value), or *only one* randomly selected active trader of the opposite type within his vision area (false value).

The `public-hearing?` parameter determines which traders can hear the outcome of the bids and offers (i.e. wheather it was a bid, or offer, what what the bid (offer) price, and whether there was a deal, or not). If the value of the `public-hearing?` is false, then only the initiating trader, and his selected partner know the type of the proposal (bid, or ask), the proposed price, and whether the proposal was accepted, or not. If the value of the `public-hearing?` is true, then this information is communicated to all traders that the proposing agent can see.⋆ In the first case, each agent can learn only from his own experience; in the later one, he can learn also from the public market information (i.e. from the experience of the others).

The `moving-type` parameter can have two values: "moving", and "not-moving". If the agents are "moving", they move randomly in each step of the simulation; otherwise they stay on the same spot all the time. If the traders can move, the market is integrated informally in this way. (Later we plan to implement other types of "moving" to see whether it would change our results.)

These parameters create a *parametric space* of the possible market institutions. We can "mix" many types of markets, for instance:

- CDA market: `vision%` = 100 %, `who-offers?` = both, `public-offers?` = true, `public-hearing?` = true; `moving-type` is indifferent.
- Chamberlin's market: `vision%` < 100 %, `who-offers?` = both, `public-offers?` and `public-hearing?` uncertain, `moving-type` = moving.
- Retail market: `vision%` < 100 %, `who-offers?` = sellers, `public-offers?` = true, `public-hearing?` uncertain, `moving-type` = moving.

⋆ If the `public-hearing?` = true, only the traders that can see the proposing trader can "hear" the bid (offer), and the potential deal. In the real world, perhaps all traders could hear it who can see *either* the proposing trader, *or* his partner. Such a definition is easy to implement, but it would make the interpretation of the simulation results more difficult because the "hearing-range" would be of a variable size, and different from the size of the "vision-range". That's why we keep this unintuitive definition.

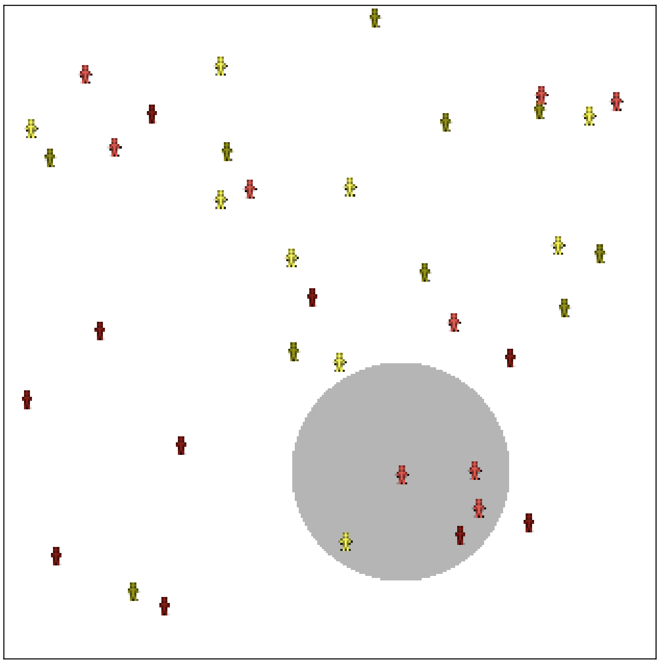

Figure 2 The trading world: the trader in the middle of the gray circle can see all traders within the circle, and no traders outside it.

2.3 Initialization And Resetting

At the beginning of each market simulation:

- traders are created, and randomly positioned in the world,
- they are assigned their `reservation-price` derived from a given supply and demand curves,
- traders' `haggling-prices` are assigned their initial values,
- the parameters of the ZI-P adjusting algorithm β and γ (see below) are set up—every trader has his own randomly chosen individual values, the same for all trading days, and
- other traders' state variables are reset to zero.

At the beginning of each trading day:

- each trader's `non-satisfied-quantity` is reset to 1 unit, and
- other traders' state variables are reset.

2.4 Trading

Each market simulation is (like Smith's experiments) divided into trading days. Each trading day lasts until there is no room for more trading, or no deal takes place for a specified time. Within each trading day, these steps are repeated:

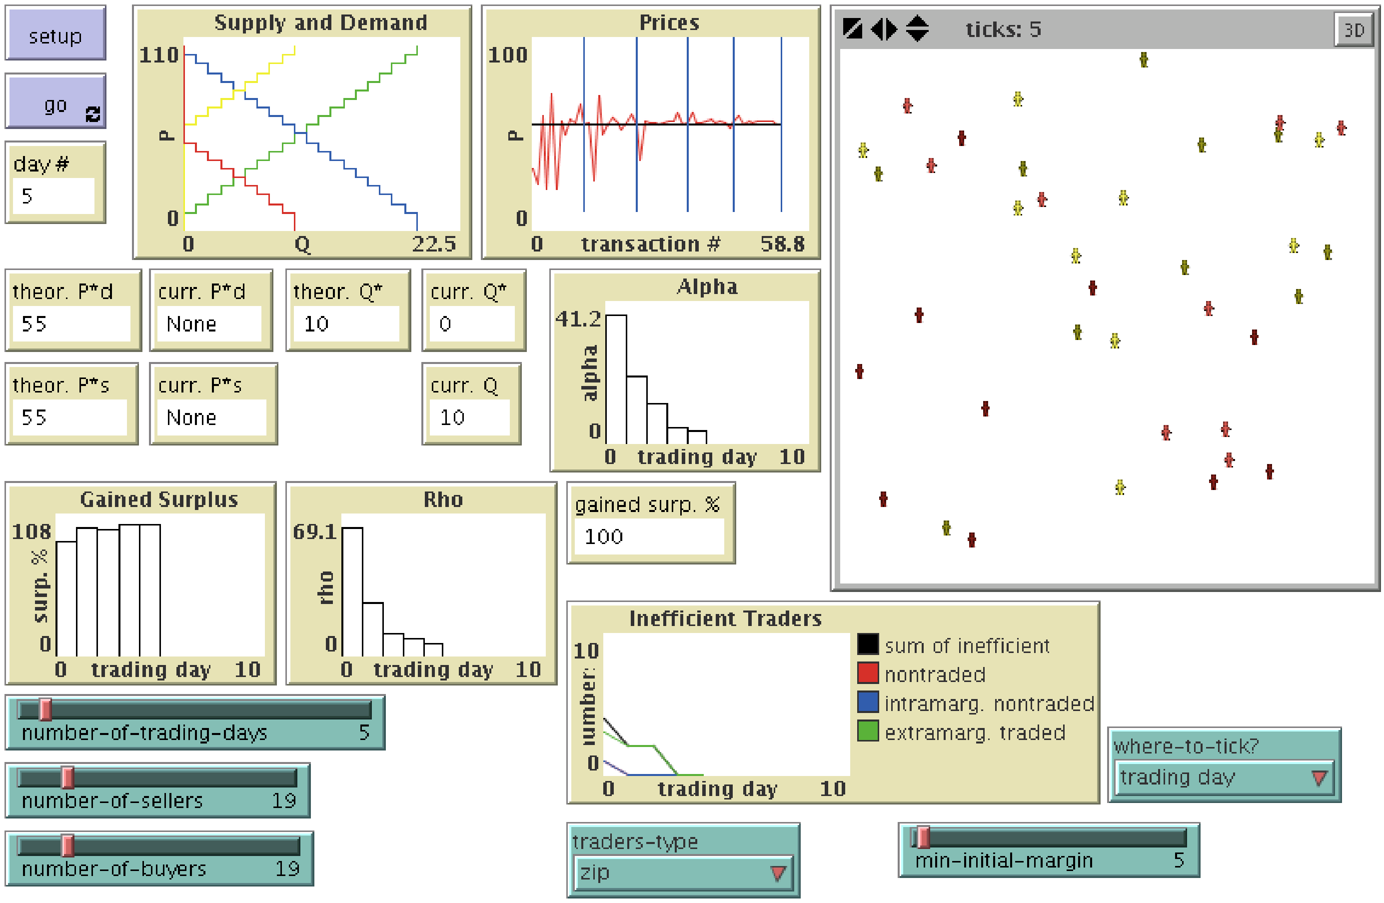

Figure 3 The user interface of the model web-based application http://www.econ.muni.cz/~qasar/marketmodel/.

1. If `moving-type` = move, then all traders move.
2. One active trader (let's call him X) of the `who-offers?` type is randomly selected.
3. If `public-offers?` = false, X randomly selects one active trader of the opposite type within his vision range as his partner. Otherwise he randomly selects as his partner one active trader of the opposite type within his vision range with whom the trade is possible (which is the same as if he tried all active traders in a random sequence, and made the deal with the first one with which the trade is possible).
4. If there is a partner and the trade is possible between X and his partner, they trade, and drop out of the market (they become "inactive").
5. If there was a partner, traders adjust their `haggling-prices`: if `public-hearing?` = true then all traders that X can see adjust them, otherwise only X and his partner adjust them.

2.5 Haggling Price Adjustment

After every ask and bid, all traders that could hear it adjust their `haggling-price`. They use Cliff's ZI-P algorithm for this (see Cliff, 1997).

The following algorithm says a seller when and in what direction he should change his `haggling-price` (the algorithm for a buyer is similar). Let the seller's `haggling-price` be p_t, and the last quotation price be q_t.

If the last shout was accepted at price q_t

- then if $(p_t \leq q_t)$
  - then raise p_{t+1} *(could have asked more)*
  - else if ((the last shout was a bid) and (I'm active)) then lower p_{t+1} *(otherwise undercut by a competitive seller)*
- else if (the last shout was an offer)
  - if ((I'm active) and $(p_t \geq q_t)$) then lower p_{t+1} *(wouldn't get the deal either)*

When a trader adjusts his `haggling-price`, he uses the following algorithm to decide how much to change it:

First, he sets the target price τ_{t+1} to

$$\tau_{t+1} = R_{t+1} \cdot q_t + A_{t+1}$$

where R_{t+1} and A_{t+1} are random numbers. If the trader is raising his haggling price, then R_{t+1} and A_{t+1} are uniformly distributed over the interval of positive numbers, if he is lowering the price, they are uniformly distributed from intervals of negative numbers. This way the trader "continuously tests" the market.

There is inertia in the `haggling-price` movements Γ_{t+1}:

$$\Gamma_{t+1} = (1 - \gamma) \cdot \beta \cdot (\tau_{t+1} - p_t) + \gamma \cdot \Gamma_t$$

where $\gamma \in \langle 0, 1\rangle$, and $\beta \in \langle 0, 1\rangle$ are the trader's parameters, and the initial momentum $\Gamma_0 = 0$. The new value of the `haggling-price` p_{t+1} is then

$$p_{t+1} = p_t + \Gamma_{t+1}.$$

If the new $p_{t+1} \geq$ `reservation-price` (for a seller, the opposite for a buyer), the `haggling-price` is assigned p_{t+1}, and is rounded with a given precision; otherwise it is reset to its old value p_t.

2.6 Measuring of Efficiency

There are many (imperfect) ways how to measure the market efficiency. We measure several of them.

The simplest one is the relative amount of the total gains from trade which is a sum of all traders' actual gains divided by sum of all traders' potential gains on a hypothetical "perfect" market. This value lies between zero (for a completely failed market) and unity (for a "perfect" market). However, this measure is very weak because it drops from its "perfect" value 1 only if an intra-marginal trader fails to trade, or if an extra-marginal trader manages to trade. The actual prices do not affect it. For instance, if all intra-marginal traders, and only them traded, the ratio would always be unity with no respect to at what prices they traded. In most simulations described below this ratio was very close to unity.

A better measure of the market efficiency is Smith's α, measured in each trading day as

$$\alpha = \frac{\sqrt{\dfrac{\sum_{i=1}^{n}(P_i - P^*)^2}{n-1}}}{P^*} \cdot 100\ \%$$

where P_i is the actual price of i-th deal, P^* is a theoretic equilibrium price, and n is the number of contracts made within the trading day. The α measures the average deviation of actual prices from the theoretic equilibrium price.⋆ This measure is weak too: for instance, if only two contracts were made out of ten possible but they were made at the theoretic equilibrium price, then α would be zero indicating "perfect market"—even though the market almost completely failed.

For this reason we also count in every trading day the number of intra-marginal traders that failed to trade, the number of extra-marginal traders that managed to trade (both of it lowers the market efficiency), the number of traded units, and the number of units that could have been traded but were not.

We have also defined another measure of the deviation from the "perfect market"—it is the most complex of them. It measures the average redistribution of gains from trade among the traders. Let us call it ρ:

$$\rho = \frac{\sum_{j=1}^{n} |g_j - \gamma_j|}{\sum_{j=1}^{n} \gamma_j}$$

where γ_j is the potential theoretic gain from trade of the j-th trader on a perfect market, while g_j is his actual gain. The ρ is zero on the perfect market, and it rises if the market efficiency decreases.

However, all these measures measure the market efficiency only at a given point of time—how much do actual prices, traded quantities, or gains from trade deviate in a given period from their potential values on the hypothetic perfect market. But the term "efficiency" could also mean (as in the financial theory) a *speed of convergence* to the prices etc. efficient in the first sense. The market is than declared efficient if and only if it converges quickly to its equilibrium after a shock (or if it quickly integrates a new piece of information to the price). We have implemented no such a measure yet, though our simulations show that it could be very important, see below.

2.7 Implementation

The model was implemented and simulated in NetLogo 4.0.2. Data were processed and analyzed in Matlab.

The model has a fancy user interface, see figure 3. You can play with it in your web browser.

3 Preliminary Data Analysis

In this section we present the results of the first model simulations. The simulations were carried out with 19 buyers and 19 sellers. Their reservation prices were generated from symmetric linear demand and supply curves with reservation prices 5, 10, 15, …, 100. Thus the theoretic equilibrium price was 55, and there were 10 intra-marginal pairs of traders.

Each simulation was run 10 times for each combination of parameters (`vision%` being 10, 20, …, 100 %). Each simulation consisted of 50 trading days. We present only results of selected parameter constellations.

The adjustment parameters of ZI-P algorithm β and γ were taken from Cliff (1997) and were not fine-tuned for our markets.

So far, the impact of the market institutions on the market efficiency was explored only with the Smith's α as the measure of the efficiency. The robustness of the results has not been tested yet.

The first question is how much does the formal integration of the market (`vision%`) affects its efficiency. We will start with the CDA market, and step-by-step decrease

⋆ The α is similar to the standard deviation; the only difference is that the equilibrium value is subtracted from the actual values instead of their actual average.

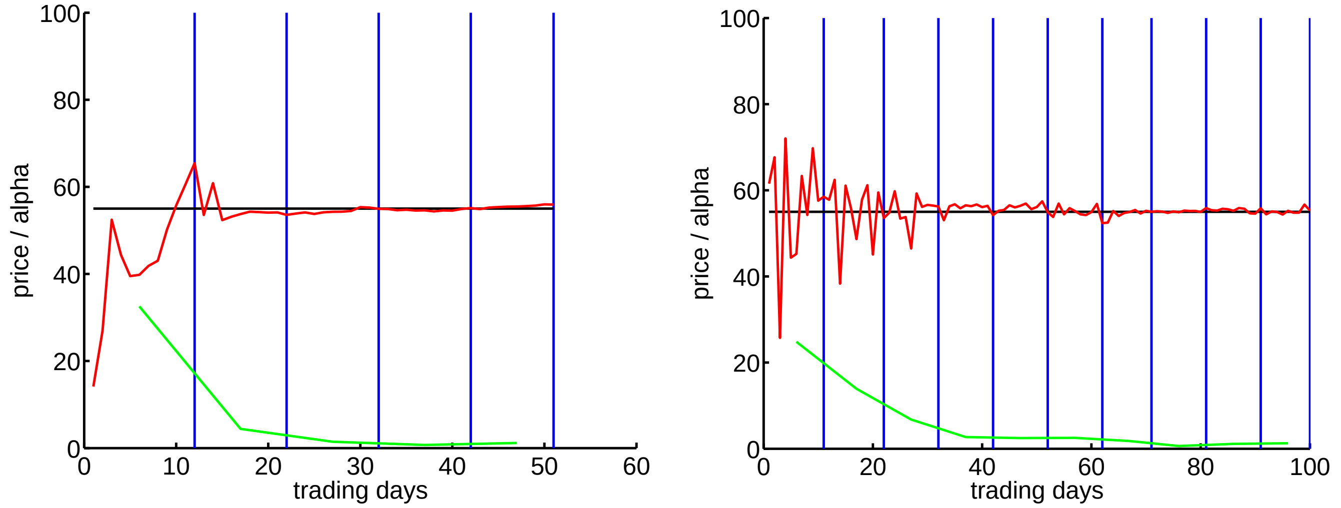

(a) `vision%` = 100 %, `who-offers?` = both, `public-offers?` = true, `public-hearing?` = true, 5 trading days.

(b) `vision%` = 30 %, `who-offers?` = both, `public-offers?` = true, `public-hearing?` = true, `moving-type` = moving, 10 trading days.

Figure 4 Typical evolution of prices on a CDA-like market with full and limited vision. Red line are actual prices, green one is the average α (calculated once a trading day). The horizontal black line depicts the theoretic equilibrium price. Trading days are delimited with vertical blue lines.

the `vision%`. The results are presented in the figure 5 which shows the evolution of the average α in time for various levels of the `vision%` parameter, and in table 1 which shows the average eventual values of α for different settings of parameters. The results of the simulations should be compared to the experimental results: Smith (1962, p. 117) obtained eventual α between 0.6 % and 9.4 % with median 3.5 % in his CDA experiments (experiments 1–7 before the supply or demand curve shifts). However, the small number of trading days might not allowed his experimental markets to converge.

A typical evolution of prices on the CDA market could be seen in the panel (a) of the figure 4: the actual prices converge very soon to the theoretic equilibrium price. If we lower the `vision%`, the convergence is less smooth and takes more time but it is achieved eventually (for an example see panel b) unless the `vision%` is too low. Actually, the eventual α is practically the same (lower than 1.5 %) for the `vision%` between 30 and 100 %, only the speed of convergence is slower with the lower `vision%`. The α is much higher only for very low `vision%` < 30 %: it seems that the vision limited so that a trader can see only one or two intramarginal partners in average is too low to secure the same level of efficiency as the fully integrated market. However, even with the `vision%` =

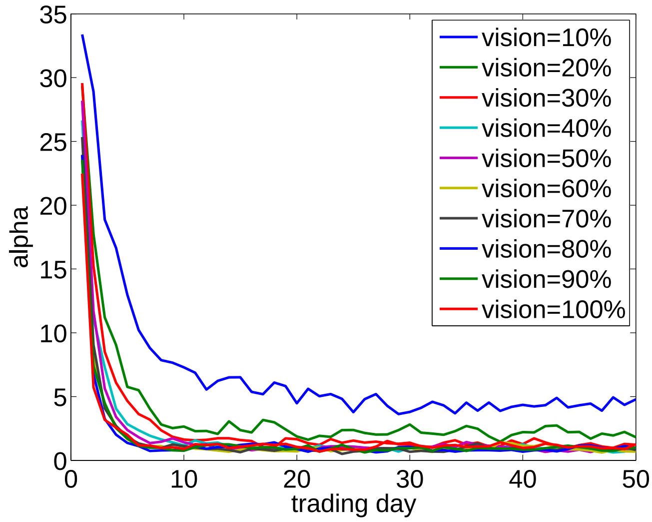

Figure 5 Impact of "vision" parameter on α: `vision%` variable, `who-offers?` = both, `public-offers?` = true, `public-hearing?` = true, `moving-type` = moving.

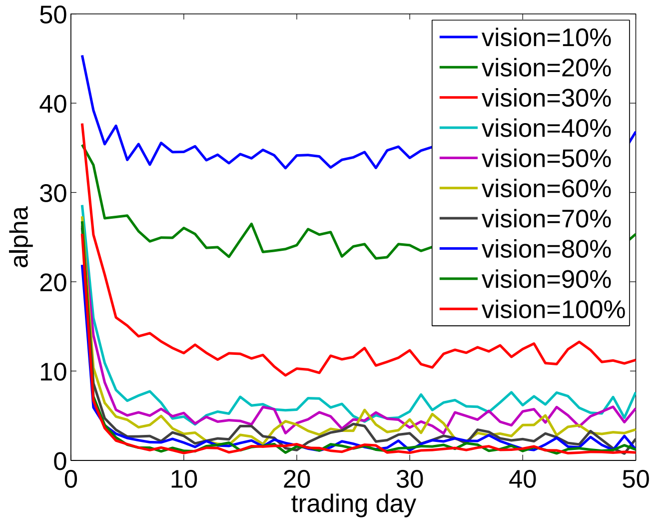

Figure 6 Impact of "non-moving" on α: `vision%` variable, `who-offers?` = both, `public-offers?` = true, `public-hearing?` = true, `moving-type` = not-moving.

vision%	10 %	20 %	30 %	40 %	50 %	60 %	70 %	80 %	90 %	100 %
true/true/moving	4.43	2.19	1.24	0.87	**0.78**	0.88	0.92	0.96	0.94	1.08
true/false/moving	6.57	2.88	1.58	1.15	1.06	0.76	0.67	0.60	0.63	**0.55**
false/true/moving	6.50	4.28	3.70	2.96	3.47	2.86	2.64	2.60	2.60	**2.59**
false/false/moving	7.94	6.63	5.27	5.22	5.69	4.89	4.82	5.03	**4.70**	5.10
true/true/not-moving	35.19	24.54	11.71	6.42	5.13	3.52	2.16	1.80	1.24	**1.01**
true/false/not-moving	35.43	30.06	20.88	13.25	9.84	7.32	4.93	1.88	1.23	**0.53**
false/true/not-moving	37.26	20.84	18.37	10.64	8.11	5.69	4.59	3.13	2.95	**2.84**
false/false/not-moving	35.45	29.52	24.21	20.29	13.85	9.83	9.63	6.23	5.33	**4.97**

Table 1 Summary statistics of the eventual α when both sellers and buyers could ask or bid. The average α is calculated from the last 10 trading days of 10 runs with the parameter `who-offers?` = both. The first column consists of three pieces of information: the first one is the `public-offers?`, the second one is `public-hearing?`, and the third one is `moving-type`. Thus the combination "false/false/not-moving" in the last row means `public-offers?` =false, `public-hearing?` =false, and `moving-type` = not-moving. The minimal α in a row is typeset in bold.

vision%	10 %	20 %	30 %	40 %	50 %	60 %	70 %	80 %	90 %	100 %
true/true/moving	4.33	2.11	1.28	0.99	0.93	**0.84**	0.85	0.88	0.90	1.00
true/false/moving	6.30	2.94	1.78	1.22	0.97	0.84	0.74	0.64	0.63	**0.60**

Table 2 The average α calculated from the last 10 trading days of another fifty runs with the parameter `who-offers?` = both. The table interpretation is as in the table 1. This table corresponds to its first two rows.

vision%	10 %	20 %	30 %	40 %	50 %	60 %	70 %	80 %	90 %	100 %
true/true/moving	4.23	2.50	1.21	1.02	0.90	0.88	**0.81**	0.93	0.93	0.99
true/false/moving	6.41	2.93	1.74	1.14	1.00	0.90	0.70	0.67	0.65	**0.63**

Table 3 The average α calculated from the very last trading day of the same fifty runs as in the table 2 with the parameter `who-offers?` = both. The table interpretation is as in the table 1. This table corresponds to its first two rows.

10 %, the average α was much lower than 10 %.

Remarkable is the fact that the eventual α *does not* monotonicly decrease with the rise of the `vision%`. Actually, the CDA-like market is the most efficient (the eventual α is minimal) with the vision limited to 50 %, see the first row of the table 1. To be certain that this is the model property and not just a random error, we carried out more simulations with a different random seed, and calculated the average α in two different ways. The result remained, though the `vision%` minimizing the eventual average α was different in every case, see the tables 2 and 3.

So far, it seems that the formal market integration (`vision%`) is not crucial for the market efficiency. It might be because it is partially substituted by an informal market integration either through the market overlapping, or through agents' "walking". To discriminate between these two hypotheses, we forbade the traders to move. The results are presented in the figure 6 and in the fifth row of the table 1. Since the market efficiency is much lower when the agents are "not-moving", we can conclude that it is the traders' ability to move, which is the major substitute to the formal market integration. However, is seems that the overlapping of the markets alone can generate some efficiency: for `vision%` $\geq$ 40 %, the eventual α was still much lower than 10 %. This result is strong: it means that an ability to see in average as many as *four* intra-marginal traders is sufficient to secure quite a good market efficiency. However, this result should be tested by more simulation runs—it seems obvious that when traders are "not-moving", the traders' initial positions are important, and the effect might be biased with only ten runs.

Another interesting question is the impact of the ability or inability to ask or bid publicly. The figure 7 shows what happens when we depart from the CDA market in two ways: we limit the vision and all offers could be only private (`public-offers?` = false). If the traders can move, the market is still pretty efficient but the α is slightly (but significantly) higher than with `public-offers?` = true (compare the figures 5 and 7, and the first and third row of the table 1). This result surprised me—I expected that this setting would only slow down the convergence of the prices, but would not lower the eventual efficiency. What happened was the opposite.

Next we tested the impact of the ability or inability

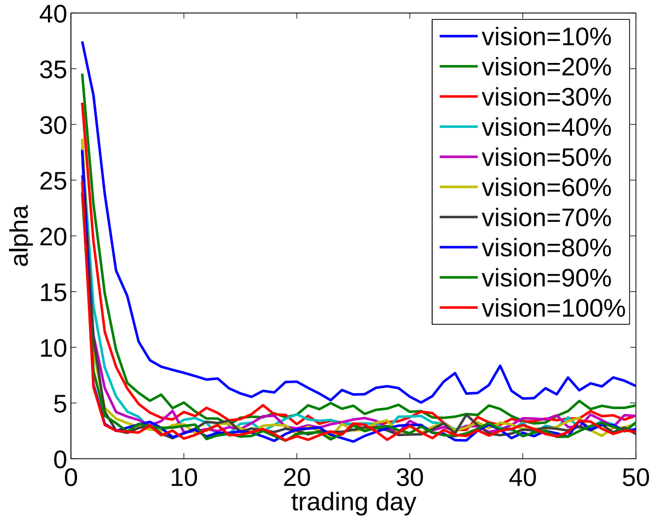

Figure 7 Impact of "private offers" on α: `vision%` variable, `who-offers?` = both, `public-offers?` = false, `public-hearing?` = true, `moving-type` = moving.

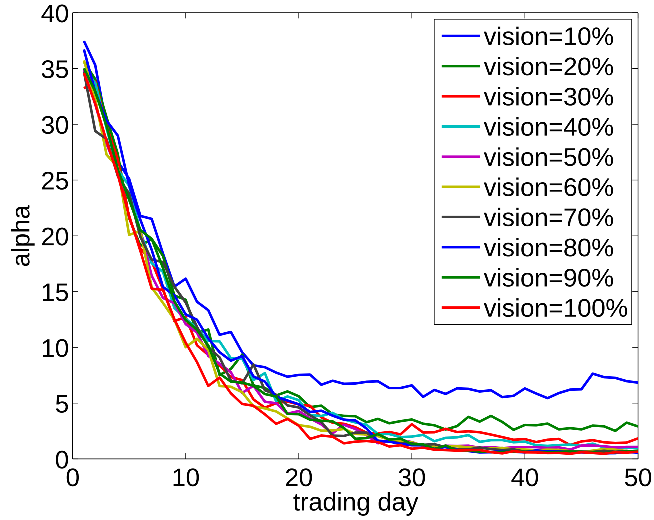

Figure 8 Impact of "private hearing" on α: `vision%` variable, `who-offers?` = both, `public-offers?` = true, `public-hearing?` = false, `moving-type` = moving.

to learn from experience of the others. We forbade the traders to learn from the others (setting `public-hearing?` = false), and combined it with various levels of the `vision%`. I expected that this setting would decrease the eventual market efficiency, and again, I was surprised. The process of price convergence was (not surprisingly) much slower but the eventual efficiency was the same or *higher* than with `public-hearing?` = true, compare the figures 5 and 8, and the first and second row of the table 1.

Surprisingly, the most efficient market institution was not the CDA, but this fully integrated market with public offering but *without* ability to observe the behavior of the others (it has the lowest α in the table 1). The result is paradoxical: it seems that a lower amount of information is better than a higher one. This result could correspond with the observation that the CDA-like market (compare the first and second rows in the table 1) were the most efficient when the `vision%` was limited. One possible explanation for this result is that it is specific to the ZI-P algorithm. With it, the traders adjust their haggling prices in such a way that they continuously "test" the market, and learn from each other (if not prohibited to). Perhaps this could lead to "bubbles" in the market prices: when one trader randomly departs from the equilibrium price, and when he is successful (i.e. there is a deal), others might follow him for a while. The bubble bursts when some intra-marginal traders are not able to trade, and they "test" the market in the opposite direction. I haven't tested this hypothesis yes—doing it would mean to explore the model behavior on the agents' level, which I haven't done yet.

If this behavior is really caused by the ZI-P algorithm, it would rise a question: under what conditions this happens? Could agents learn from each other, and not cause such bubbles? What types of haggling price adjusting algorithms cause this effect? Is this effect present in the real and experimental markets where traders continuously test the market and learn from the public data too?

The composite impact of the inability to offer publicly, and learn from the experience of the others (`public-offers?` = false, `public-hearing?` = false) was both slow convergence, and significantly lower eventual efficiency, compare the figure 9 to the figures 5, 7 and 8, and the fourth row of the table 1 to its first three rows. But the market efficiency measured by α is still surprisingly high (quite lower than 10 %). However, more testing is needed—we cannot be certain that the process had already converged within 50 trading days in this case.

It is interesting that the impact of the `vision%` was not

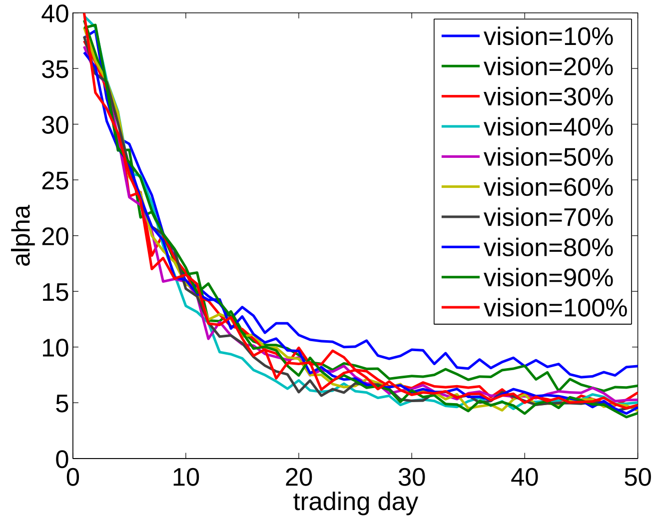

Figure 9 Impact of "private hearing" and "private offers" on α: `vision%` variable, `who-offers?` = both, `public-offers?` = false, `public-hearing?` = false, `moving-type` = moving.

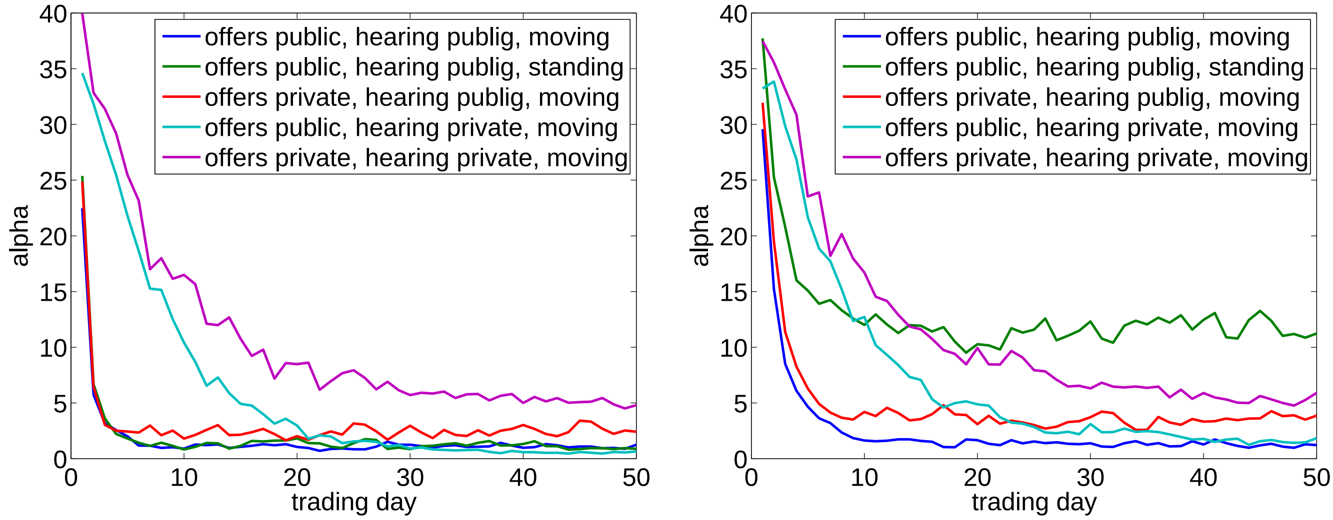

(a) vision% = 100 %, public-offers? = true, rest variable.

(b) vision% = 30 %, public-offers? = true, rest variable.

Figure 10 The impact of selected combinations of parameters on the evolution of the α.

very important in all these cases—it was substituted by the ability of the traders to move across the markets, and to some extent also by the overlapping of the markets.

A more general picture could be derived from the figure 10. It presents impacts of selected combinations of the parameters on the evolution of the α—the panel (a) with the full vision (`vision%` = 100 %), the panel (b) with a quite a limited vision (`vision%` = 30 %). First, we can see that the `vision%`, and the ability to walk across the markets are substitutes: with the full ability to “see”, the impact of walking is nil; with very limited `vision%`, the ability to walk is crucial—the ability to “see” as such is not that important if the traders can move. Second, the inability to hear the outcomes of other’s trading slows the convergence down but has ambiguous impact on the market efficiency. Third, the inability to place public offers lowers the market efficiency significantly, especially when accompanied with other departures from the CDA market institution as a low vision or the inability to hear the others. However, if the `vision%` is high or the traders can move, the α is still lower than 10 %.

Another interesting question is what would happen if only sellers could offer. Smith (1962) reported that the eventual market prices lied *below* the theoretic equilibrium in such a case. Our model results are somewhat more complex: the average market prices lie below the theoretic equilibrium for many trading days, however, if the market is close to the CDA and the `vision%` is high, the prices

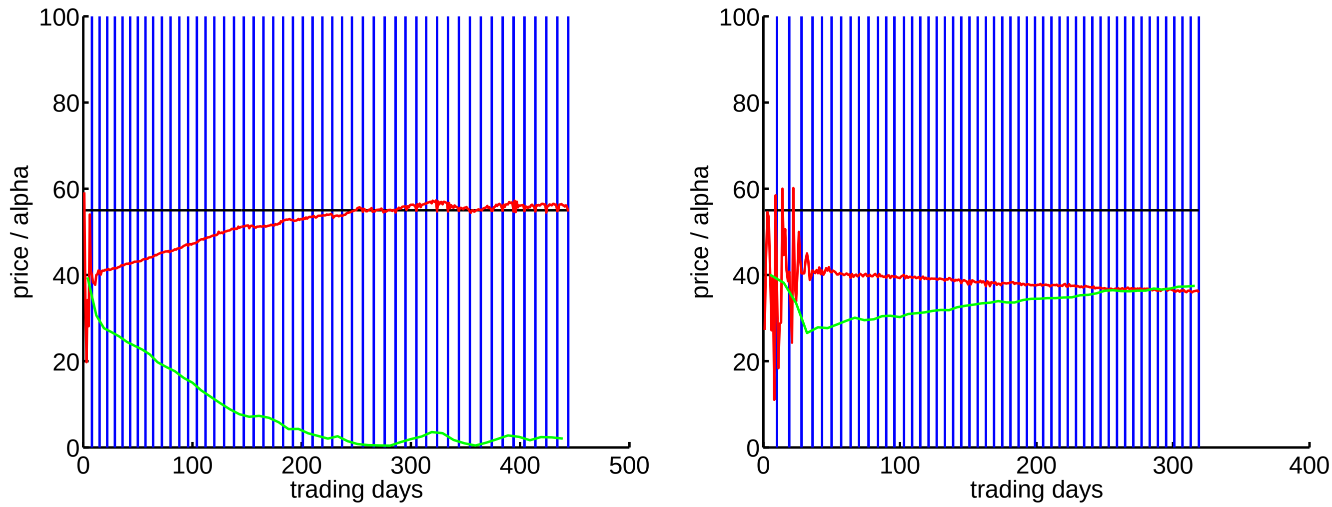

(a) `vision%` = 100 %, `who-offers?` = sellers, `public-offers?` = true, `public-hearing?` = true, `moving-type` = moving.

(b) `vision%` = 30 %, `who-offers?` = sellers, `public-offers?` = true, `public-hearing?` = true, `moving-type` = moving.

Figure 11 Examples of the evolution of prices when only sellers can quote.

`vision%`	10 %	20 %	30 %	40 %	50 %	60 %	70 %	80 %	90 %	100 %
true/true/moving	53.69	42.08	32.70	25.44	21.57	15.07	8.62	4.04	**1.22**	1.72
true/false/moving	83.66	77.79	76.79	77.10	75.72	75.04	77.83	78.12	**73.18**	74.77
false/true/moving	64.86	**56.21**	57.62	62.74	63.51	67.55	72.82	74.72	77.82	85.36
false/false/moving	94.97	88.05	**82.11**	86.26	84.22	83.76	86.65	89.20	86.72	87.40
true/true/not-moving	58.04	38.62	31.80	28.28	17.32	12.21	7.64	4.43	2.84	**1.81**
true/false/not-moving	–.–	–.–	78.72	**73.42**	73.68	75.22	74.07	75.84	76.18	77.52
false/true/not-moving	–.–	62.67	61.99	**61.36**	64.96	66.66	69.54	71.09	79.88	82.29
false/false/not-moving	–.–	–.–	–.–	–.–	**84.98**	87.95	89.58	88.23	86.15	87.10

Table 4 Summary statistics of the eventual α when only sellers offered: the average α calculated from the last 10 trading days of 10 runs with the parameter `who-offers?` = sellers. The rest is as in the table 1. The "–.–" symbol means that less than two trades took place so it was not possible to calculate the α.

eventually converge to the theoretic equilibrium. On the other hand, if the `vision%` is low, the market prices either *diverge* or converge to a *wrong equilibrium*. For examples of individual price evolutions see figure 11, for the summary see the table 4. The turning point in the `vision%` seems to lie somewhere between 50 and 60 %—but this result is only preliminary because most prices haven't converged within our 50 trading days, and the turning point depends most likely on the number of the traders. I haven't explored other parameter constellations yet.

4 Discussion

We presented some results of our simulations together with a preliminary analysis above. There are two major questions however: first, how to interpret the results, and second, in what relation they might be with the real-world markets.

The interpretation is not straightforward because the results are highly non-linear and there are many substitution effects between the parameters. The overall result seems to be a surprisingly strong efficiency of our artificial markets if both sellers and buyers could offer or bid—such markets must have been seriously "crippled" (the departure from the CDA market must have been strong along at least two lines) before the eventual α was higher than 5 % (it was hardly ever higher then 10 %). Surprisingly, the CDA market was not the most efficient market—a departure from it had an ambiguous effect on the market efficiency. Generally, it seems that it is rather the *intelligence* of the traders than a particular form of market institutions what makes the market efficient (measured by α).

As for the sellers' markets, the results are less certain. It seems that if the market institutions are in other dimensions close to the CDA, the market could eventually converge to the theoretic equilibrium. What happens otherwise is uncertain because the number of trading days in our simulations was too short for the α to converge. Further research is needed here.

It may seem that the market institutions are not very important for the market efficiency measured by α; however, they do affect the market efficiency in another sense: they influence the *speed* of the convergence. It may have two implications: First, only dynamic markets (with swift changes of underlying data) need to be integrated and have well designed institutions that secure a swift convergence to a new equilibrium. If the changes are small and slow (as they were e.g. in the Middle Ages) then people could trade quite well without such institutions. Second, Chamberlin (1948) must have gave up too early—most likely it were not poor institutions that made his markets inefficient, but short time given to his traders to learn. The same might hold true for Smith (1962) when he concluded that if only sellers can offer, the market prices lie below the theoretic equilibrium.

The correspondence of the model to the real-world markets is even tougher problem. We employed Cliff's ZI-P traders that are quite simple-minded. However, they performed well under various institutions (they were able to bring the actual market prices quite close to the theoretic equilibrium). Thus we can expect that the real-world markets consisting of much cleverer traders should perform even *better*—our simulation results should be the *lower limit* of the real-world markets performance. However, it seems that there might be a place for "bubbles" when traders learn from each other.

However, one reservation is needed here: our model has shown the ability of markets to communicate the *given* private knowledge (the reservation prices) well under many different institutions: given this "data", our artificial markets were usually able to reach eventually the equilibrium determined by them. However, if the reservation prices are *not given*, but *endogenous* (as are for instance on the financial markets where analysts have to assess the "true fundamental" values of financial assets), it is not guaranteed that the market would miraculously reach the "true values"—it would at best converge to the equilibrium values given the actual traders' assessments. If the analysts' assessments are influenced by the market conditions themselves (as it seems to be true) than various "bubbles" are possible—bubbles not only in the adjusting process (discussed above), but in the determination of the "theoretic equilibrium" too.

5 Known Issues and Possible Enhancements

The biggest know problem is the poor definition of the traders' "walking". It seems obvious that the speed of their movements could influence the market efficiency if the vision is limited. However, I have found no natural way how to define the speed yet. It seem obvious that there should be such a definition, and the speed should be a parameter of the model rather than an arbitrary constant as is now.

Moreover, there could be other "walking algorithms" than the "wiggle walk" that is implemented now. I plan to implement a "circular walk" where the traders would walk in circles, which would integrate some but not all markets together.

The second big issue is the explicit trading, moving, and searching costs. There are no such costs now. This is very good for keeping the model simple but it might seriously lower the correspondence of the model to the real world.

Third, with the limited vision, it seems there might be a profit (at least a temporary one) from a wise searching for a partner. Perhaps the mindless traveling through the marketplace (implemented now) should be replaced with a partner searching algorithm.

Fourth, our model assumes that the traders have given reservation prices, and sell and buy for money. It is very good for a simple interpretation but it means the world is not "closed". It would be nice to implement a market with producers of goods exchanging their products with each others. Other trading algorithms than Cliff's ZI-P would be needed in such a case.

Fifth, the robustness of the model must be tested. More simulations must be run for all parameter configurations. The model should be tested with different shapes of the supply and demand curves, different numbers of traders, more units to trade than one, and with other trading algorithms than Cliff's ZI-P. The efficiency of the market institutions should be measure with more measures. To explore behavior of the markets where only sellers can offer, the simulations must be repeated with more trading days.

Sixth, beside the eventual market efficiency measured by α and the like, the speed of convergence should be measured rigorously too. For this I have to find a good general measure of convergence suitable for our needs.

Seventh, we should explore how does the efficiency (both the eventual one, and the speed of convergence) depend on the number of traders, shapes of the demand and supply curves, adjustment algorithm and its parameters (β and γ above) etc.

Eight, the model should be also explored at the level of individual prices and the agents' level to explore why the α is not minimized on the CDA market, to test the "bubble" hypothesis stated at the page 8, to learn what precisely does a big α mean (whether a big departure from the equilibrium value, or a big oscillation around it), and to learn why the sellers' market is so different from the "double" markets, and why it sometimes converges to the theoretic equilibrium value, and sometimes not.

Ninth, the model should be more rigorously tested against outcomes of experimental markets. I have to find papers that describe experimental markets similar to some of my non-CDA parameter constellations. This could allow me both to test the model, and perhaps to calibrate it to the experimental data.

6 References

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

Cliff, D. (1997). Minimal-intelligence agents for bargaining behaviors in market-based environments. Technical Report, School of Cognitive and Computing Sciences, University of Sussex.

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

Smith, V. L. (1962). An experimental study of competitive market behavior. *The Journal of Political Economy*, 70(3):322–323.

Tesfatsion, L. Agent-based computational economics: Growing economies from the bottom up. http://www.econ.iastate.edu/tesfatsi/ace.htm.