Experimental Based, Agent Based Stock Market *

Jakob Grazzini

University of Turin

July 20, 2011

Abstract

This paper builds an agent based model to reproduce the results of an experimental stock market that studies how the market aggregates private information. The aim is to contribute to the relationship between experiments and agent-based modeling and to understand the behavior of the agents. Using the experimental environment and results, it is possible to formulate a hypothesis about the behavior of the subjects and thereby formalize (algorithmically) the behavior of the traders. This allows a better understanding of how the market converges toward the equilibrium and the mechanism that allows for the dissemination of private information in the market.

1 Introduction

The aim of this paper is to explore the possible relations between experiment and agent based modeling and to understand how financial markets aggregate traders' private information.

Experiments and agent-based models are "natural allies" (Duffy 2006) and can integrate each other in a prolific way (Contini et al. 2007). In experiments it is possible to control the economic environment and the market structure but not to control the motives and the individual characteristics of the subjects (Chan et al. 2001). A model can use the experimental environment and results to formalize the behavior of the subjects using the generative approach proposed by Epstein & Axtell (1996). The experiment provides the foundation and the results to the model, while the model provides insights into the behavior of the subjects. According to Smith (1982) it is possible to define a

*The paper was written while the author was visiting the *Center for Nonlinear Dynamics in Economics and Finance - University of Amsterdam.* The kind hospitality and the stimulating working environment is gratefully acknowledged. The author would like to thank also Martin Berner, Francesco Feri and Charles R. Plott (and in particular Charles R. Plott as corresponding author) for sending the data of the experimental stock market described in Berner, Feri & Plott (2005). E-mail: jakob.grazzini@unito.it

microeconomic system, $S = (e, I)$, as a microeconomic environment (e) and a microeconomic institution (I). The microeconomic environment is defined by Smith (1982) as a list of N economic agents, $K + 1$ commodities and the characteristic of each agent i, $e = (e^1, ..., e^i, ..., e^N)$. The environment is thus a set of initial conditions that cannot be altered by the agents. The microeconomic institution defines the rules under which the agents act and interact. In a stock market the institution is represented for example by the continuous double action trading mechanism that imposes a set of rules on how the agents can act, how they issue orders and how contracts are made. The behavior of the agent i is a function $\beta^i(e^i|I)$. Given the institution, the environment and the behavior of the agents, the macro-behavior of the system is defined. The aim of the experimentally based model is to use the environment, the institution and the outcome of the experiments to formalize a plausible set of micro-behavior. It is not possible to produce any sufficient and necessary result, but by reproducing the experiment it is possible at least to propose a set of sufficient conditions on the behavior of the agents. As stated by Gode & Sunder (1993, p.120): "It is not possible to control the trading behavior of individuals. Human traders differ in their expectations, attitudes toward risk, preferences for money versus enjoyment of trading as a game, and many other respects. The problem of separating the joint effects of these variations, unobservable to the researcher, can be mitigated by studying market outcomes with participants who follow specified rules of behavior. We therefore replaced human traders by computer programs." Once the behavior is defined, it is possible to test the effect of different hypotheses and isolate the influence of different factors on the outcome of the system. In the literature there are several examples of modeling experiments using heterogeneous agent models (Hommes 2011, Hommes et al. 2005, Anufriev et al. 2010) or agent based models (Chan et al. 2001, Gode & Sunder 1993, Gjerstad & Dickhaut 1998, Boero, Bravo & Squazzoni 2010, Boero, Bravo, Castellani & Squazzoni 2010, Duffy & Ünver 2006). This paper will contribute to the literature by building an agent based model of an experimental stock market analyzing the plausible behaviors of the agents. In particular the interest concerns how the market aggregates the information, where the market is simply defined by a set of bounded rational agents motivated by profits with asymmetric information operating in a continuous double auction.

How the information is incorporated in the prices is a crucial question in financial economics. The problem has been dealt with experimentally in Plott & Sunder (1982, 1988) and in the following literature (Camerer & Weigelt 1991, Copeland & Friedman 1986, 1991, Forsythe & Lundholm 1990, Berner et al. 2005). The experimental results show that markets can aggregate information but sometimes they make mistakes (Plott 2000). Using the experimental data published by Berner et al. (2005), this paper builds an agent-based model of the experimental market and tries to understand how the market aggregates the information and what happens when the mechanism fails. The agent based model tool was chosen since the main properties of the stock market are emergent properties [1]. Starting from a the simple market described in Smith (1962), an agent

[1]Emergent properties are system-level properties that cannot be explained directly by the micro-level properties (Gilbert 2001). The price convergence for example is an emergent property of the system,

based model was developed following the model proposed by Cliff & Bruten (1997). Particularly section 2 shows how the results in Smith (1962) can be reproduced using a very simple heuristic representing two basic incentives of the traders: maximizing the profits for each trade and maximizing the probability of trading. Considering this behavior as a fundamental behavior, the model is extended to a more general type of market (section 3), and then using the experiment in Berner et al. (2005), described in section 4, extended to a more complicated cognitive situation (section 5) where the agents have to evaluate a random variable (representing the value of the traded asset) and learn the information held by the insiders in the market (section 6). The aim is to find the conditions under which the information is correctly disseminated into the market. The result is that markets tend to disseminate correctly the information due to the incentives of the traders and the transparency of the continuous double auction. The problem is that also the failures seem to depend on idiosyncratic features of the traders and are therefore difficult to prevent. Section 7 will conclude and propose some future research.

2 The simple market

The aim is to build an agent based stock market starting from experimental economics results. In order to isolate the basic market behavior of the subjects it is useful to start by building a model reproducing the behavior of a simple stock market. In Smith (1962) an experimental market is described where the traders are divided into buyers and sellers. Each subject receives a card containing an induced value for the fictitious commodity. The trade is conducted through a continuous double auction over a sequence of periods. Each trader is free to place an order at any time (constrained only by the role and the value), once a contract is concluded, the interested traders drop out of the market for the given period. The induced value for each trader is the same in all periods. The aim in Smith (1962) was to study the behavior of the price in a controlled situation where demand and supply schedule were well defined over a unit of time. Figure 1 shows the market environment, determining the supply and demand schedules and the theoretical equilibrium as well as the transaction prices in a sequence of periods. To measure the convergence of the price over periods Smith (1962) uses a "coefficient of convergence" α, shown in the figure 1:

$$\alpha = \frac{100\sigma}{P_0} \tag{1}$$

where σ is the standard deviation of the transaction price from the theoretical equilibrium P_0 in the given period. A lower α means smaller deviation of the transaction prices around the theoretical equilibrium.

The experiment shows how a small number of inexperienced traders converge (α is decreasing) rapidly to a competitive equilibrium under the double auction mechanism (Smith 1962, p.157). The result is extremely interesting because it shows how the

since its definition makes sense only at the system level

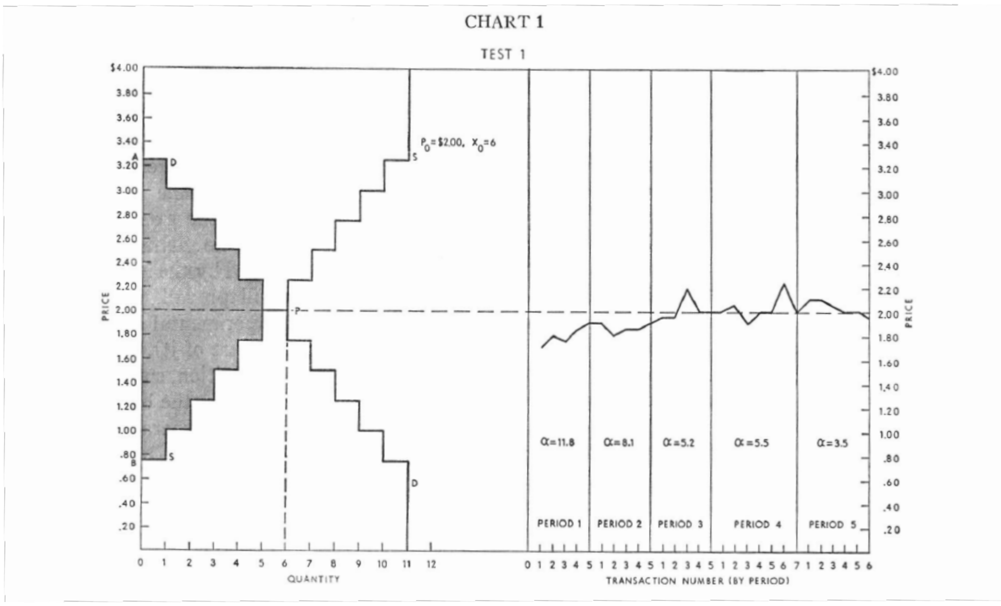

Figure 1: The figure represents the induced market environment on the right and the transaction price for a sequence of periods

interaction between the traders allows the emergence of the equilibrium price and how this equilibrium price is predictable by the classical microeconomic theory (Cliff & Bruten 1997). The private profit-seeking incentives allow the market to reach the equilibrium and even if the market environment (demand and supply schedule) is unknown the traders learn it in a few periods [2]. The aim of reproducing the experiment is to formalize the behavior of the traders and understand how the convergence occurs. The equilibrium price can be considered as an emergent property since it is a property of the system, not explainable directly by the properties of its elements: it is the outcome of individual incentives, the set of induced values and the interaction between the agents through the continuous double auction. Given the experimental environment, institution and the results of the experiments, it is possible to use the generative approach (Epstein & Axtell 1996) to specify a behavior for the agents consistent with the outcome. The interaction between the agents, thus the orders and their arrival, is fundamental to the understanding of the behavior of the system. A representation of the market with a purely mathematical model would result in a very complex dynamic system, difficult to handle and impossible to resolve. The reason is that the agents might be heterogeneous in the time dimension (a continuous double auction typically involves asynchronous bidding) and this influences the order arrival and therefore the information introduced into the market. An agent based model made by a set of individual autonomous agents

[2] In Smith (1962), note 5: “It is only through some learning mechanism of this kind that I can imagine the possibility of equilibrium being approached in any real market.”

implemented as software objects, each agent with private variables, states and rule of behavior will be used to reproduce the experimental market. The starting point of the analysis of the experimental continuous double auction is the computational model proposed by Gode & Sunder (1993). The behavior of the market with human traders is compared with the outcome of a computational market. The agents in the computational market generate independent random bids or offers depending on the assigned role (buyer or seller). The agents do not try to maximize profit, do not react to market conditions, do not use any information from the market to choose the price - they simply bid or offer at a random price uniformly distributed between 0 and 200; therefore they have been called "zero-intelligence" traders (ZI). In order to isolate the effect of market rules and profit-seeking behavior Gode & Sunder (1993) introduce also an augmented ZI trader called "zero-intelligence with budget constraint" trader (ZI-C). This latter computational trader behaves exactly as the ZI trader but the support of the random variable used to determine the price is constrained by its private value. A buyer with value v will generate random bids with the price $p \in [0, v]$; a seller with cost c will generate random offers with price $p \in [c, 0]$. Gode & Sunder (1993) interpret the constraint as imposed by the market rules: "the market forbade traders to buy or sell at a loss because then they would not have been able to settle their accounts" (Gode & Sunder 1993, p.123). The difference between the performance of the ZI-C traders and human traders can be attributed to the features of human subjects. The differences between the markets with ZI-C and ZI traders can be interpreted as the contribution of market rules to the behavior of the price. The results of the human experiment and computational experiments are shown in figure 2.

The astonishing result claimed by Gode & Sunder (1993) is that the market rules are the most important element in the observed convergence of the price toward the theoretical equilibrium, while the rationality and profit-seeking behavior of human traders can account only for a small fraction of the efficiency of the market. ZI-C traders achieve a high level of efficiency trading very near to the theoretical equilibrium and a decreasing standard deviation in every period. The standard deviation is simply computed as the root mean square deviation of the prices from the equilibrium *within* each period (the measure is similar to Smith's α, but note that in Smith the convergence is inter-period, while here the convergence is intra-period). The efficiency measure is defined as the total profits actually earned by the traders during a given period divided by the maximum possible profit (i.e. the sum of buyer and seller surplus). The interpretation of the results has been criticized by several papers. In particular Gjerstad & Sachat (2007) show that the claim of convergence in Gode & Sunder (1993) is not correct. Moreover, the (reported) convergence is intra-period not inter-period, and it depends on particular features of the model: the traders with higher value (and lower costs) will trade before and the traders with lower values (and higher costs, thus closer to the equilibrium price) will trade later in the period producing the intra-period convergence (see Gode & Sunder 1993, p.129). The convergence in Smith (1962) is due to a learning mechanism that allows the convergence of the price period by period. Also the interpretation of the constraint as a budget constraint imposed by the market is problematic. There is no

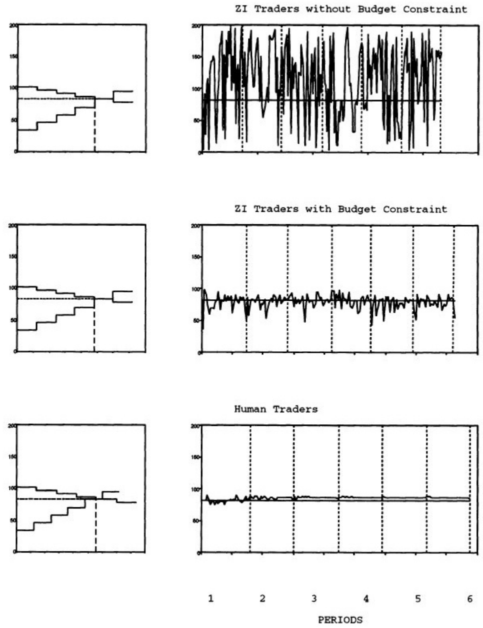

Figure 2: Demand and supply functions and the transaction price time series, from Gode and Sunder (1993) p. 124

necessary link between the private value of a trader and the money that the same trader brings to the market. The market institution has neither the capacity nor the interest in knowing the private value of the traders and imposing a budget constraint on them based on that value. The constraint is not an institutional constraint but a behavioral constraint, the traders decide to adopt a non-loss behavior, it is a constraint due to a (bounded) rational behavior. Following such an argument Gjerstad & Sachat (2007) conclude that the increase in market efficiency observed in Gode & Sunder (1993) with ZI-C cannot be interpreted as a direct consequence of the double auction rules. Cliff & Bruten (1997) show that the increase in efficiency and the convergence of the price in Gode & Sunder (1993) is a statistical artifact. The reduction of volatility and the apparent convergence are due to the random structure of the model, which in turn depends on the private value of the traders. The Gode & Sunder (1993) model is very important since it is the first computational model that replicates a realistic continuous double auction and shows that the market mechanism coupled with extremely simple strategies can lead to the equilibrium in particular situations but it is not enough to model the experimental market. Two different models of a continuous double auction can be found in Gjerstad & Dickhaut (1998) and Cliff & Bruten (1997). In these models the agents use simple heuristics to understand the market environment. The double auction is an important feature of the market since it allows the traders to gather information from bids and offers, but for convergence there is the need of more sophisticated traders that are able to use the information made available by the market institution. In the following the Cliff & Bruten (1997) stock market model will be described and reproduced. The market is formed by a set of agents divided between buyers and sellers, each with a private value for the traded asset. The market mechanism is a continuous double auction and the trade takes place in a sequence of periods. The aim is to reproduce exactly Smith's experimental environment and institution. Therefore the agents are free to bid and offer at any time and they withdraw from the market for the given period when they successfully close a deal. The price proposed by agent i at time t can be written as:

$$p_i(t) = v_i(1 + \mu_i(t)) \tag{2}$$

where v_i is the certain induced value and $\mu_i(t)$ is the profit margin, positive for sellers and negative for buyers. The budget constraint is still active: no bid or offer can be made with a loss. The profit margin evolves over time following a very simple heuristic. The pseudo algorithms for a seller and a buyer are Algorithm 1 and Algorithm 2 [3]. The agents observe the book and use the information about the last proposal to understand the market.

Algorithm 1 The basic behavior of a Seller: adapting the profit margin

Seller
if the last shout was accepted at price q **then**
  1. any seller s_i for which $p_i \leq q$ should raise the profit margin
  2. if last shout was a bid and $p_i \geq q$, any active seller s_i should lower its margin
else
  if the last shout was an offer and $p_i \geq q$ any active seller s_i should lower its margin
end if

Algorithm 2 The basic behavior of a Buyer: adapting the profit margin

Buyer
if the last shout was accepted at price q **then**
  1. any buyer b_i for which $p_i \geq q$ should raise its profit margin
  2. if last shout was an offer and $p_i \leq q$, any active buyer b_i should lower its margin
else
  if the last shout was a bid and $p_i \leq q$ any active buyer b_i should lower its margin
end if

The algorithm is well explained in Cliff & Bruten (1997). As described above, in the Smith (1962) experiment the time was divided into periods and each trader had the opportunity to trade only once [4] during each period. The traders start the period as

[3] On http://www.jakob.altervista.org/Python-model1.rar it is possible to download the python files of the model

[4] In some experimental sessions the traders were able to trade more than once, but always for a given number of times. This procedure is useful as it provides a definition of demand and supply schedule.

active traders and become "non-active" after having agreed on a contract. The aim of the traders is to trade at the best possible condition, i.e. with the maximum possible profit margin. The seller s_i might for example start offering at a given price $\bar{p}_i(t)$. From equation 2 we know that the offered price depends on the private (constant) value and on the profit margin. If s_i observes that the last order was accepted at a price q greater than $\bar{p}_i(t)$, the incentive to maximize the profits will induce the seller to increase the profit margin. The observed order tells seller i that there are buyers willing to buy at a higher price. If on the contrary the last accepted order was an offer with a price q lower than $\bar{p}_i(t)$, the incentive to trade induces s_i to lower the profit margin (if it is greater than zero). The seller s_i will use the information contained in the last order to infer on the behavior of the other sellers. To be able to trade she must reduce the selling price (reducing the profit margin) to undercut the competition. It is important to note that the reduction of the profit margin by a seller is triggered only by offers: the traders undercut their competitors. If on the other hand the sellers reacted also to very low bids, the buyers could coordinate and reduce artificially the price. For the buyers the algorithm works symmetrically. The crucial point is to understand how the traders adapt and that only some bids and offers influence the market. Extra-marginal traders and exceptional bids and offers (very low bids and very high offers) have no effect on the market. This simple algorithm allows the traders to understand the optimal pricing strategy by adapting the profit margin. In order to adapt there is the need for some form of updating rule. Cliff & Bruten (1997) propose the Windrow-Hoff "delta rule":

$$A_{t+1} = A_t + \Delta_t \tag{3}$$

where A_{t+1} is the output after the update, A_t is the current output and Δ_t is the change in output in time t and depends on the difference between actual A_t and the desired output D_t and a learning rate coefficient β:

$$\Delta_t = \beta(D_t - A_t) \tag{4}$$

The traders want to update the proposed price by updating the profit margin. Given $p_i(t)$ and a target price $\tau_i(t)$ it is possible to compute $\Delta(t)$ from equation 4,

$$\Delta_i(t) = \beta_i(\tau_i(t) - p_i(t)) \tag{5}$$

and the new profit margin rearranging equation 2:

$$\mu_i(t+1) = \frac{p_i(t) + \Delta_i(t)}{v_i} - 1 \tag{6}$$

The target price is defined using the price of the last shout $q(t)$ in the following way:

$$\tau_i(t) = R_i(t)q(t) + e_i(t) \tag{7}$$

where R_i is a random coefficient and e_i is a random perturbation. If the aim is to increase the last shout, $R_i > 1$ and $e_i > 0$, if the aim is to decrease the last shout

$0 < R_i < 1$ and $e_i < 0$. Finally, to give weight to past events, Cliff & Bruten (1997) introduce a "momentum coefficient" denoted by $\gamma_i \in [0, 1]$. Given $\Delta_i(t)$, computed as above, the momentum based update is given by:

$$\Gamma_i(t+1) = \gamma_i \Gamma_i(t) + (1 - \gamma_i)\Delta_i(t) \tag{8}$$

Defining $\Gamma_i(0) = 0$, and substituting $\Gamma_i(0)$ in place of $\Delta_i(t)$ in equation 6, the updating rule used by the ZIP traders is the following:

$$\mu_i(t+1) = \frac{p_i(t) + \Gamma_i(t)}{v_i} - 1 \tag{9}$$

This simple heuristic learning produces a complex behavior at the system level, each intra-marginal order is going to influence the profit margins, and thereby the following orders. From this formalization it is also possible to understand the mechanism that allows the market to converge toward the equilibrium. The agents have two different and opposite economic incentives: they try to increase the profit margin to maximize the profit per trade and decrease the profit margin to undercut the competitors and increase the probability of trading. This double incentive allows the price to converge stochastically to the theoretical equilibrium defined by the supply and demand schedule (a simulation with 11 buyers and 11 sellers is shown in figure 3). Note that the computational nature of the agent based model impedes an analytical demonstration of the equilibrium properties, but it is possible to infer about the behavior of the system using multiple simulation (Axtell 2000). In order to have a quantitative evaluation of the behavior of the model the stationarity and ergodicity tests proposed in Grazzini (2011*b*) were used. Testing the first moment of the price time series means to test whether the model converges toward a stochastic equilibrium (the price time series has a mean and stays around it). To understand whether the system has a unique equilibrium we use the ergodicity test. If the process is stationary and ergodic it means that the model has an equilibrium and that the equilibrium is always the same regardless of the initial conditions (regardless of the random seed). The outcome of the tests is that, given the private values of the traders, the system has a unique stable equilibrium (and this is true also for the Gode & Sunder (1993) model)[5]. In order to characterize the model as a good representation of the experiment it is possible to analyze the behavior of the transaction price computing α (as computed in Smith (1962), see equation 1) and the allocative efficiency of the market in each period, where the allocative efficiency is defined as the total profit actually earned by the traders divided by the maximum total profit that could have been earned (i.e., the sum of producer and consumer surplus) (Gode & Sunder 1993, p.131). As shown in figure 4a the former decreases, showing that learning is occurring from period to period, and in figure 4b the latter increases. In both figures the average per period value of 100 simulations is shown.

[5] See the appendix for more details. From the stationarity test it is also possible to infer that the model is not strictly stationary, this is because the profit margins are influenced by the contracts and therefore there is a dependence between successive transaction prices.

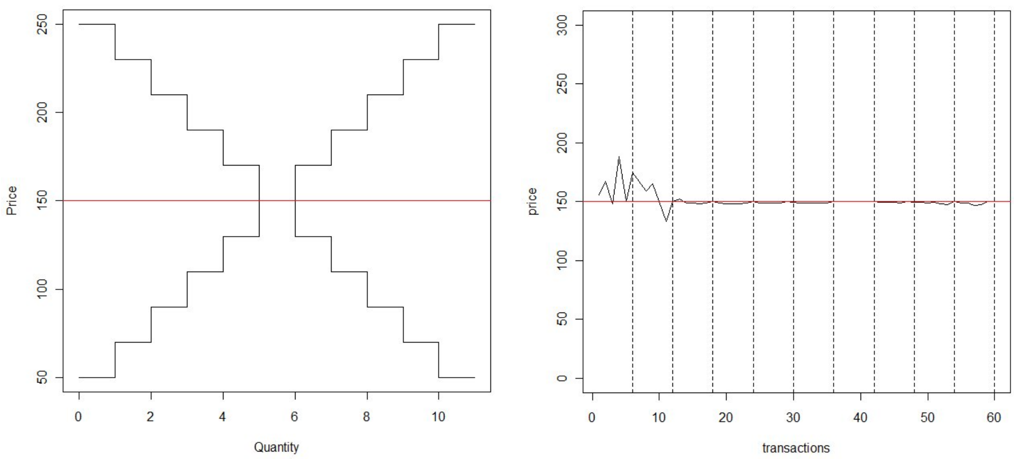

(a) Demand and supply schedules

(b) Prices per transaction. The vertical lines divide the periods

Figure 3: The simulation. The model is run for ten periods.

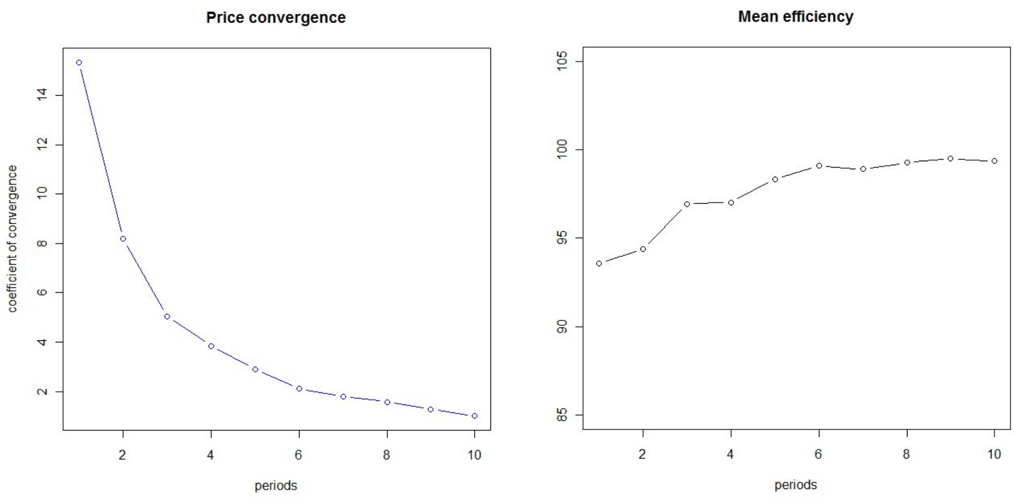

(a) Mean α per period

(b) Mean efficiency per period

Figure 4: Price convergence and efficiency of the market. Average values over 100 simulations.

During each simulated period every agent will issue an order - if active - on average every 20 seconds. One period lasts 500 seconds and trading normally takes place in the first part of the period [6]. The timing has some relevance, as the traders use the last proposal to gain information about the market, if the agents acted simultaneously they would use less orders and the learning mechanism would be more instable and slow. This is one example of how important it is to explicitly model the price formation mechanism since it shapes the market behavior. The asynchronicity of actions is crucial as it allows the traders to understand the market environment.

3 Extension of the simple market

Accepting the behavior of the agents as the basic behavior and the incentives at work as the basic incentives, it is possible to try to generalize the model step by step. In order to have a more realistic stock market the first step is to allow traders to be both buyer and seller, and let them decide whether to buy or sell with regard to the market situation. This means that the traders will buy when the profit margin is negative (when they think that the market evaluates the asset less than they do) and sell when the profit margin is positive. The number of buyers and sellers change with respect to the market situation, and the outcome of the system is the equilibrium price in this case simply defined as the price at which buyers and sellers are in equal number (note that the quantity traded by each trader in each period is one, the equilibrium price is the price at which demand and supply pressure are balanced). Starting from the particular case of symmetric demand and supply (as in the example in figure 3) the equilibrium price does not change, but the quantity traded in equilibrium changes, since all traders can potentially trade in this case, and the extra-marginal traders do not exist anymore. The second extension to the model is the elimination of the period structure. The time structure will therefore be continuous, the traders will use a random variable to act and have an endowment of money and assets. The definition of demand and supply is now average demand and average supply and depends on the level of the price. The equilibrium is the price at which there is the same number of buyers and sellers. Random perturbation of the equilibrium price due to the agents' random activity rate will tend to create a difference between buying and selling activity that brings the price back toward the equilibrium. The trade activity of agents with a fixed private value produces a unique stable global equilibrium [7]. The traders have an induced value for the asset and their activity shapes the demand and supply per given time unit. Different traders might act at different rates, for example in Gjerstad & Dickhaut (1998) the agents' activity depends on the expected profit. In the present model a constant and homogeneous activity rate among the agents has been imposed. The reason is that the aim of the paper is to understand how the system reaches the equilibrium value. With constant private values different activity

[6] In Smith's experiment the traders could freely offer and bid; within the time limit of a trading period this procedure was continued until bids and offers were no longer leading to contracts.

[7] The unique stable equilibrium of the model is inferred using the tests used also in section 2. See the Appendix.

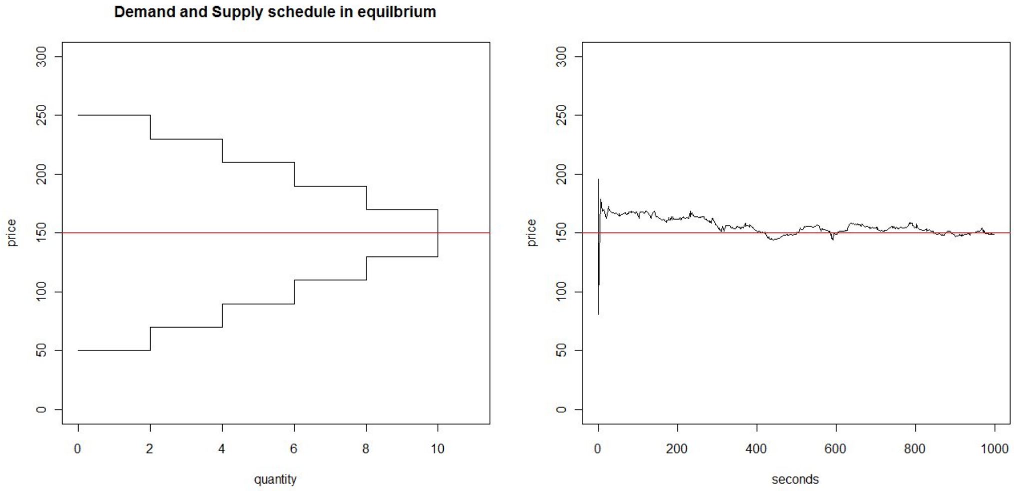

(a) Average supply and demand for 20 seconds time interval in equilibrium

(b) Transaction price over 1000 seconds

Figure 5: Average supply and demand schedules and the transaction price behavior.

rates among the agents would only change the equilibrium (which is already exogenous since it depends on the induced values) and have very little effect on the convergence process. If on the contrary the asset evaluation is directly influenced by the market (i.e. positive feedback between the market and the evaluation) then the equilibrium will be influenced by the convergence process and by the activity of the agents. The market with one type of trader is shown in figure 5 [8]. The convergence measure is shown in figure 6, the efficiency is in this case difficult to compute since the periods have been eliminated. In allocative terms there is a transfer of assets from the traders with low value to the traders with high value.

The periods are in this case only an arbitrary division of time. The trader trades for 1000 seconds with the only possible constraint of the endowments (here not binding). Each "period" in this latter model consists of 100 seconds and has the only role of observing the convergence of the price and orders. The model predicts that the orders and transaction prices will converge toward the expected equilibrium. The traders start from their own private value using a random profit margin. The competition among the agents releases information about the environment, the agents "learn" how to behave in the market and the system converges toward the equilibrium. The comparison with the "basic model" shows how the convergence is fast but the price is less stable around the equilibrium. As shown in figure 6 the convergence coefficient is strongly reduced in the first two periods but it never goes below 2. The high convergence speed is due to the fact that all traders are introducing useful information in the market, the market pressure

[8]The market consists of 22 traders. The code is available on http://www.jakob.altervista.org/Python-OneTrader.rar

Figure 6: The price and order convergence measured using the average α coefficient over 100 simulations.

toward equilibrium is high since there are no extra marginal traders. Once the market reaches the equilibrium it stays around it, but the equilibrium is less stable here than in the previous model. This latter feature is due to the random nature of the activity of the agents. On average the equilibrium is the theoretical equilibrium, and the actual equilibrium depends on the actual activity of the agents. When the price departs from the equilibrium, the tendency is to move back toward the equilibrium since (on average) there will be a difference between buyers and sellers. The actual difference between buyer ans sellers depends on which traders actually act.

4 Berner, Feri and Plott Experiment

The extension made in the previous section is compared with experimental data. The experiment used to extend the model is published in Berner et al. (2005),[9] and in the following the experiment is quickly described. There are n subjects (the number changes in the different sessions, from 9 to 17) with an endowment of 10,000 francs [10] and 10 assets. Each of the 5 sessions is divided into 5 years, and each year into three periods during which the trade can take place. The market is for a single asset that pays a single dividend depending on the state of nature. There are four possible states of nature corresponding to four different dividends of 800,600,400,200. The actual state of nature depends on the draw of a random variable that eliminates in each period one of the two extreme dividend values with a probability (1/2) known by the subjects. Figure 7 repre-

[9] I would like to thank the authors of Berner et al. (2005) for providing the data

[10] Francs are the experimental currency which will be converted into $ at a fixed rate (1000 francs = 1$) at the end of the experiment.

sents an experimental year. At the end of each period the draw (up or down) is publicly announced and one of the two extreme dividend values is eliminated (either 800 or 200). Since the information is public, all agents move to the next node that corresponds to a different expected value for the final dividend value. At the end of the third period the last draw is made and the subjects owning assets are paid accordingly. Two series of experiments were conducted: the "Symmetric Information" and the "Asymmetric Information". For each session the first year represents the Symmetric Information series in which all subjects have the same information and know that no insiders are present in the market.

In particular they know the dividend time structure (figure 7) and the probability distribution over the dividend value in each period, thus all agents are informed about the expected value of the dividend. The data from the first year will be used for the first extension of the model: the agents are free to buy and sell (constrained by the endowments) and have to evaluate a random dividend structure to assign a value to the owned asset. In the next section the first years will be analyzed to understand the behavior of the agents using as magnifying glass the model described in section 3. In the Asymmetric Information series the experimenters draw the random variable that determine the elimination of one of the extreme dividends at the beginning of each period and provide such information to a subset (one-third) of the subjects. The expected value of the asset for informed agents and non-informed agents is different. Non informed agents know that insiders are present in the market. The analysis of the Asymmetric Information series are very interesting as it allows us to understand how the non-informed traders look for the information they know is present in the market. The aim is to understand the signal the uniformed traders use to infer the information present in the market, design a plausible "learning mechanism" that uses that signal and to understand under which conditions the mechanism works (the price converges toward the "informed equilibrium"), it works only partially or even does not work at all (the price converges toward a wrong equilibrium).

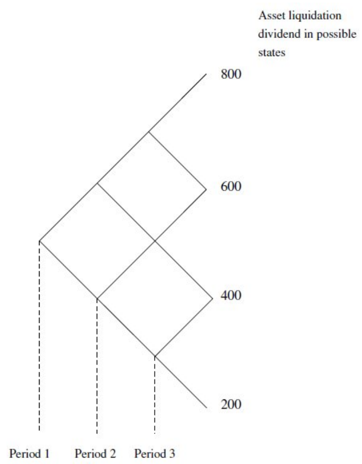

Figure 7: The time structure of the dividend.

5 The first year

The result of the Symmetric Information case is that "prices converge to fully revealing Rational Expectations equilibrium (the risk neutral expected value given all the information available to the market)" Berner et al. (2005) p. 92. From the vast literature about bubbles in experimental markets (see for example Smith et al. 1988, King et al. 2001, Smith et al. 2000, Noussair et al. 2001, Lei et al. 2001) it is known that speculative motives and confusion can produce positive feedbacks between the market behavior and individual evaluation of the asset. The first important cognitive effort made by the subjects is to attach a monetary value to the expected dividend, i.e. the money they are willing to pay for one asset and the money they are willing to accept when selling one asset. The subjects' trade implies that the evaluation of the common expected dividend is different and this probably depends on the individual characteristics of the traders such as their risk aversion. The basic assumption in the following analysis is that the value for each subject is constant during a period. Note that it is possible to observe only the apparent demand and supply (the proposed price that contains both the value and the profit margin). The constancy of the value can be supported by the absence of any evident positive feedback between the market behavior and the individual behavior. A positive feedback would create bubbles and instability in the market. Assuming that there is no positive feedback implies the assumption that the subjects understand the random structure of the dividend and trade consequently without pursuing speculative profits.Traders who evaluate the asset highly will buy,and traders who evaluate the asset lowly, will sell. There is evidence that traders who buy in a given period might sell in another period. This suggests that the relative evaluation of the asset is not constant; traders who seem to be relatively risk averse in a given period (they sell) might become relatively risk lovers in another period (buyers). To implement such a behavior the choice is to avoid the complexity of the individual choice and simply model the individual evaluation of the asset with a draw at the beginning of each period of a normal random variable with a mean value equal to the theoretical expected dividend and given variance. The constancy of the evaluation during a given period allows us to use the model described above and assume that traders use the orders they observe in the market to adjust the profit margin.

5.1 Analysis of convergence

The main prediction of the stock market model is that there is convergence of the transaction prices and of the orders toward the theoretical equilibrium. The convergence toward the equilibrium means also an increase in the efficiency of the market, which in the present case cannot be computed since it is neither possible to know the theoretical efficient profit nor the actual profits of the agents, since the individual evaluation of the asset is not observable. To compute the convergence of price and orders, the α coefficient in equation 1 will be used over the whole set of periods available in the Symmetric Information Series (15 periods). Since periods last 5 minutes it is possible to divide every period into minutes. The orders and contracts are then divided into five sets depending on the minute they were made and regardless of the period. The standard deviation of the transaction prices is computed using the expected value of the dividend. The "theoretical equilibrium price" is the price at which demand and supply are equal, therefore it depends on the actual distribution of private values among the traders. The actual individual values are not observable. The assumption is that the actual individual values are distributed as a Normal around the theoretical expected value of the dividend, this assumption implies that the expected value of the theoretical equilibrium is indeed the expected value of the dividend [11]. The best estimate of the actual theoretical equilibrium in the market is therefore the expected value of the dividend. Computing α as in equation 1 for all transaction prices in the first year sorted by minute it is possible to see that also in the experimental market there is a significant convergence. To be clear, all contracts were grouped with respect to the time of execution (regardless of the period) and α were computed using the theoretical equilibrium valid in the period in which the contract was made. The first minute group therefore includes all the contracts made in the first minute of period one, two and three of the first year. In figure 8a the transaction price convergence is shown. The absence of transaction costs has to be considered during the analysis of the convergence of bid and offer prices. While the contract prices can be considered as the result of faithful actions, since they actually imply a transaction and therefore a profit or a loss, the orders do not have the same limitation. For example, if a subject has only 200 francs left and wants to buy, he could try to issue orders with 200 as a proposed price even if the market price is much higher. The trader knows that the market price is higher but since making orders is free nothing disincentives such an action. If the convergent coefficient is computed without taking this problem into account, we have very high values for the last periods [12]. To avoid these outliers, the implausible orders [13] are eliminated from the sample and the result is shown in figure 8b.

It is interesting to note that the deviation of the orders in the last minute increases significantly. This phenomenon can be seen as an "end of the experiment" effect and explained using the model: the weakening of the trade incentive (reduces the undercutting) and the constancy of the profit incentive increases the profit margins and therefore the deviation of the order prices from the theoretical equilibrium. Note that the effect on the transaction price is almost null; this means that the agents identify the equilibrium price and contracts are made only in the neighborhood of the equilibrium. The model seems to explain the qualitative results of the experiment in a satisfactory way and it can therefore be used as a good approximation of the experimental market: the traders observing the market understand the equilibrium price and converge toward it by adjusting their profit margin.

The convergence result of the transaction prices divided by periods (and minutes) is similar, except for the fact the there seems to be also an inter-period convergence

[11] The theoretical equilibrium in a market with many traders would be the expected value of the dividend.

[12] Resulting α from the first to the fifth minute: 14.6, 18.1, 12.8, 1414213552.47, 1055280.66.

[13] All orders below 200 and above 1000. Over 1221 orders, only 74 are eliminated. For example there is a sell order of 100 billion francs (this is probably an error) and many buy orders of almost 0.

(a) Transaction price convergence coefficient (b) Order price convergence coefficient

Figure 8: Average α coefficient per minute over all 15 periods.

and that in the third period, the transaction prices start very near to the expected theoretical equilibrium. Grouping the transaction prices only by period (obtaining three sets of transaction prices) it is possible to observe how, as the periods progress, average α is reducing [14]. The inter-period reduction of the α coefficient can be due to the reduction of the variance in the dividend and therefore to a smaller degree of heterogeneity in the evaluation of the agents (which in turn simply means that less learning time is needed to converge toward the equilibrium).

Figure 9a shows an example of experimental market (session 1, year 1) with sell orders, buy orders and contracts made. In figure 9b a simulated market is shown, with the same dividend draws[15]. The main difference is the volatility of the orders: in both markets the orders converge toward the equilibrium, but in the real market the agents make limit orders with high profit margin (this behavior was reported also in Smith (1962)). The high volatility of the orders can be interpreted as the result of the absence of order costs; the effort of issuing an order is almost zero in the experimental market, since no other activity is available and there is no "fundamental" risk (since all possible states are known). In the model the agents have always the double incentive of trading and maximizing the profit margin. The time variable and the period structure are not considered in the model. The subjects probably choose the trading quantity not only with respect to the expected profits of the single trade but also with respect to the time variable. This behavior might be explained by considering some sort of precautionary behavior suggesting that the whole endowment should not be consumed in one period. Therefore they stop trading and start issuing orders with a high profit margin in order

[14] period 1: 32.14; period 2: 20.77; period 3: 9.63

[15] The code is available on http://www.jakob.altervista.org/Year1.rar

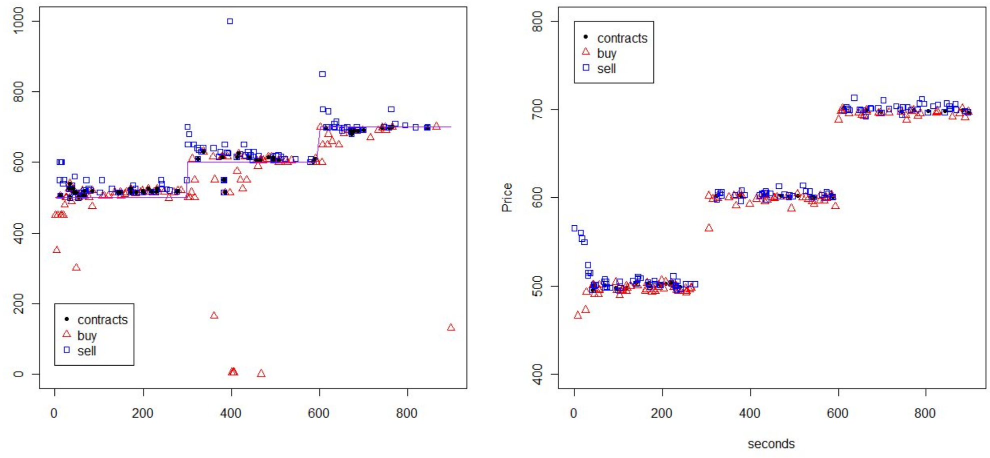

(a) Experimental market (session1) (b) Simulated market

Figure 9: Comparing the Experimental market (a) and a simulated market (b).

to trade only with no risk (for example sell only at the maximum possible value or even more). Since the "extra-marginal orders", i.e. the orders that are very distant from the equilibrium, have no effect on the traders' beliefs this behavior has no effect on the price convergence. The agents adjust their beliefs about the market only in response to contracts and to competitive (undercutting) orders. To be clear, the model does not consider a behavior that leads a subject to bid with a price equal to zero, but on the other hand this behavior can be considered simply as "noise" since such a bid will not affect the trading behavior of the other subjects and will therefore not affect the emergent properties of the system.

6 Asymmetric Information

The aim of the experiment made in Berner et al. (2005) is to understand how the market gathers the information. After the first year the experimenters introduce asymmetric information in the market by telling 1/3 of the agents which of the two extreme possible dividends will be eliminated at the end of the period. The insiders thus have a different expected value in the given period. The research question is whether the non-informed agents are able to "read" the information from the market and consequently whether the price is going to converge toward the informed equilibrium (defined as "fully revealing rational expectation equilibrium" by Berner et al. (2005)). The results in Berner et al. (2005) is that in most cases the price converges toward the rational expectation equilibrium.

If the informed traders did not use the information in the market, the equilibrium

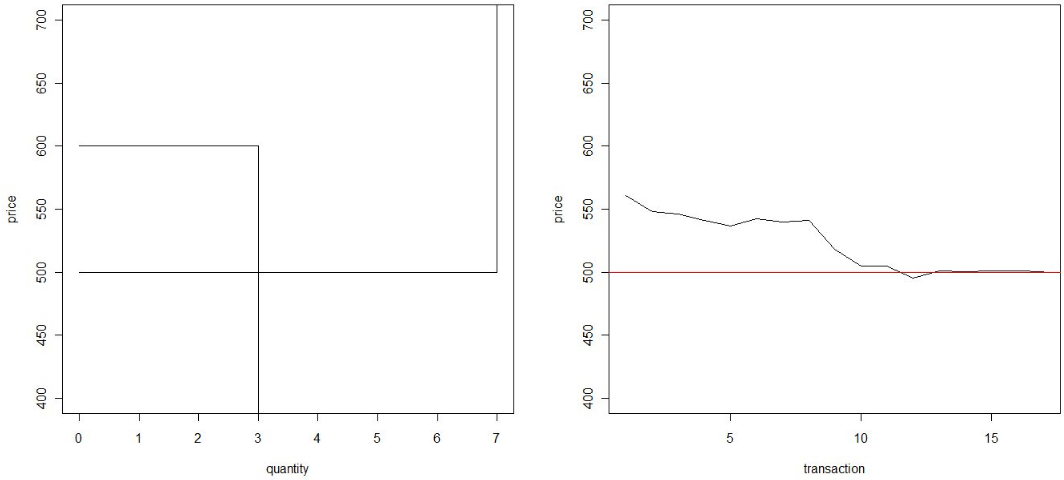

Figure 10: On the left the average supply and demand schedules without learning, on the right the transaction price in the market.

price would depend on the private values of the traders and on the average activity of the traders in the market. By defining an average activity rate for the agents it is possible to simulate the case in which the agents do not use the market information in the trade and to define an average theoretical equilibrium. Suppose for example that all traders have the same activity rate and they all have the possibility to buy or sell (they have not met their endowment constraint), that the non-informed agents evaluate the asset 500 and the informed traders evaluate the asset 600. By considering a time interval in which all traders make on average one order (the length of the interval is not relevant since it is the same for every trader) the average supply and demand schedules are as in figure 10 and the equilibrium price would be 500. At any price above 500 all the non-informed traders would sell and all the non-informed traders would buy. Given the activity rate, it is possible to determine the average competitive market equilibrium. Above 600 and below 500 no market is possible since no trade is possible. Indeed, the competition among traders who want to buy at a price above 500 is such that the actual equilibrium will be at the price of 500 with excess supply and the average number of transaction per time interval will be 3. In the simulation shown in figure 10 the traders acted on average every 40 seconds and the trading period lasted 300 seconds [16].

The “Result 2” in Berner et al. (2005) states that “in most periods, prices either move in the direction of or converge close to fully revealing Rational Expectation equilibrium” (Berner et al. (2005) p. 93). The “fully revealing rational expectation equilibrium” implies that the non-informed agents learn about the information and that the informed agents reveal their information through trading. The experiment supports the

[16]The code is available on http://www.jakob.altervista.org/nolearning.rar

hypothesis of competition among informed traders [17], Berner et al. (2005) note that the non-informed traders are exploiting this competition using limit orders and learning the informed value of the asset. The dissemination of information in financial markets with insider trading is supported also in the literature. Kyle (1985) for example studies a continuous auction with one risk-neutral insider, random noise traders and a market maker. The insider acts as a monopolist and maximizes the profits taking into account the effect his actions have on the market-maker's price setting. The result is that the information is gradually incorporated into the prices through the trades made by the informed trader. The conclusion is that the market converges toward the informed equilibrium price even in presence of a monopolistic insider. Holden & Subrahmanyam (1992) extends the Kyle (1985) model by introducing multiple informed traders. Their basic result is that the competition makes the informed traders trade very aggressively. Even with only two insiders the common private information is incorporated into the price very quickly. In the agent-based modeling literature, asymmetric information is analyzed by Chan et al. (2001) where the agents are endowed with different learning capabilities and are under some conditions able to converge toward the Rational Expectation Equilibrium. The difference between the present model and the "traditional" literature is the wish to simulate a realistic market. The belief is that the reproduction of specific features of a continuous double auction, like for example the asynchronous arrivals of orders, the price formation and the information arrival, is important for the understanding of the experimental market. The difference with the computational model developed by Chan et al. (2001) is that in the present case the analysis starts from the experiment in order to arrive at the model and then back to the experiment trying to replicate some fundamental features of the experimental market. The "experiment-based market" is a method that considers plausible strategies of the agents and tries to gain insight into the experimental events.

In order to introduce the learning into the model it is necessary to understand where the non-informed traders find the information. Starting from the literature above and from the results obtained by Berner et al. (2005) it is possible to find the "signal" used by the non-informed traders to learn the informed equilibrium and the actual behavior that favors the convergence toward the informed equilibrium or the non-convergence. In particular, among all the periods in the experiments, the ones with full convergence and bubbles will be selected and analyzed to understand the difference between the two situations (the definition of convergence and bubble will be given in next section)

6.1 The Signal

The ability to infer information from the behavior of others is normal in everyday life (Plott 2000), and this happens also in stock markets. In this section the aim is to find the behavior that signals the information to the non-informed traders. It can be done by analyzing the behavior of informed and non-informed traders. The non-informed trader knows the structure of the dividend and its expected value, knows about the presence of

[17] Berner et al. (2005) call it ITRRH: Informed Traders Rat Race Hypothesis

insider traders and knows what kind of information is owned by the informed traders. In particular they know that the expected value of the informed traders can be one of two known possible values (i.e. "up" or "down"). The non-informed traders will look for behaviors that are distinguishable from the behavior of other non-informed traders; the signal has to carry information about the state of the world. The hypothesis made here is that the signals are what will be called "non-compatible orders". Supposing that the expected value of the asset for the non-informed traders is 500 (e.g. in period 1 of every year) and taking into account the risk, the non-informed trader will evaluate the asset around 500 francs. The same trader knows that also all the other non-informed traders will do the same. This means that trader i, using as threshold her own evaluation of the asset (v_i), will consider a buy below v_i or a sell above v_i as compatible with the non-informed "rational" behavior. A non-compatible order for trader i is a buy order above v_i or a sell order below v_i. The sign of the signal tells the non-informed traders the direction of the information. Supposing that the informed traders have "up" information, the expected value of the asset conditioned to their information set is 600, and they will evaluate the asset around 600. For the informed trader j - with an evaluation (for example) $v_j = 600$ - it is "rational" to buy below 600 and sell above 600. The informed traders know the expected value of the non-informed and of the informed, thus they know that selling above 600 will be very difficult. The informed trader therefore starts buying at prices that might start below 500 (not a signal), and the competition among the informed traders will bring the bids up to above 500 starting to release information to the non-informed. If the information is "down", the reversal situation occurs, with offers below 500. The non-compatible orders are the signals; they are released due to competition among informed traders and are simple to be read by the non-informed traders. In Berner et al. (2005) the period is divided between "convergent", "mirage" and "bubble" periods using two different models: the Ashenfelter-El Gamal model and the AR1 model. Using the two models Berner et al. (2005) evaluate statistically the behavior of the market. In particular they estimate the final price in each period and define the period as convergent if the theoretical price is within the 95% confidence interval; as a mirage if the movement of the price is in the right direction and the confidence interval of the estimated price contains values that differ at least 100 from the theoretical equilibrium; as bubbles the periods in which the price moves in the wrong direction and the estimated price confidence interval contains values that differ at least 100. The models can give different estimation of the price and therefore classify some periods in different categories. It is interesting to note that most of the periods do converge, that all bubbles are bubbles using both models and that all strict mirages ("strict" in the sense that both model estimations result in a mirage) come after a bubble (see the table at p. 87 of Berner et al. (2005)). The two extreme cases are thus bubbles and strict convergence (periods that result convergent using both models), and these will be analyzed to understand how the learning mechanism works and why sometimes it does not work. The focus on strict convergence and on bubbles is due to the desire to have well defined results and to understand why a mechanism that normally works (the learning mechanism) does sometimes not work. In order to get an idea of what

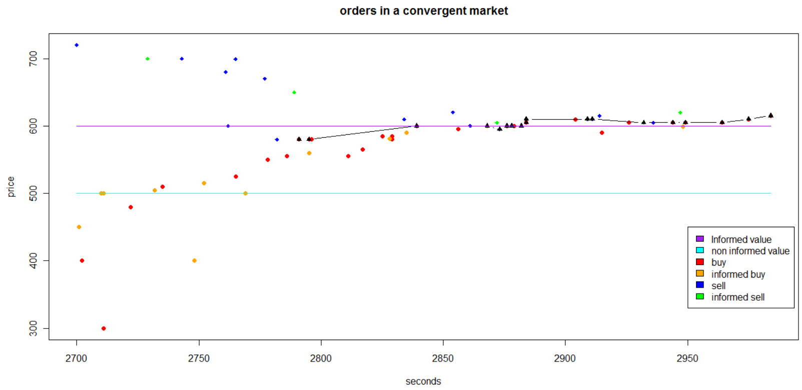

Figure 11: Orders and contracts in a strictly convergent period in the experimental stock market.

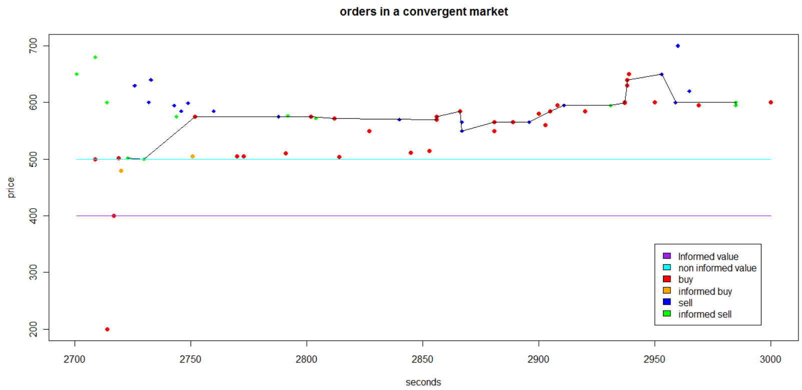

Figure 12: Orders and contracts in a bubble period in the experimental stock market.

happens during a strict convergence and during a bubble, figures 11 and 12 show a strict convergence period and a bubble period.

Consider figure 11. The non-informed expected value is the cyan line and the informed expected value is the purple line. Note that the first orders are compatible orders. There are two non-compatible orders made by informed traders; they seem to trigger the competition among informed and non-informed who understand the state of the world. The first contract is made at a price very near to the Rational Expectation equilibrium. The non-informed traders understand the information present in the market and they trade according to that information. The bubble case is illustrated in figure 12, the price converges toward the wrong informed equilibrium: even if the informed expected value is 400, for some reason the non-informed traders understand that the informed expected value is 600. In this case the competition among informed traders does not release any information in the market, the non-informed are convinced that they have understood the information and trade according to it. In 37 of 45 periods, the direction

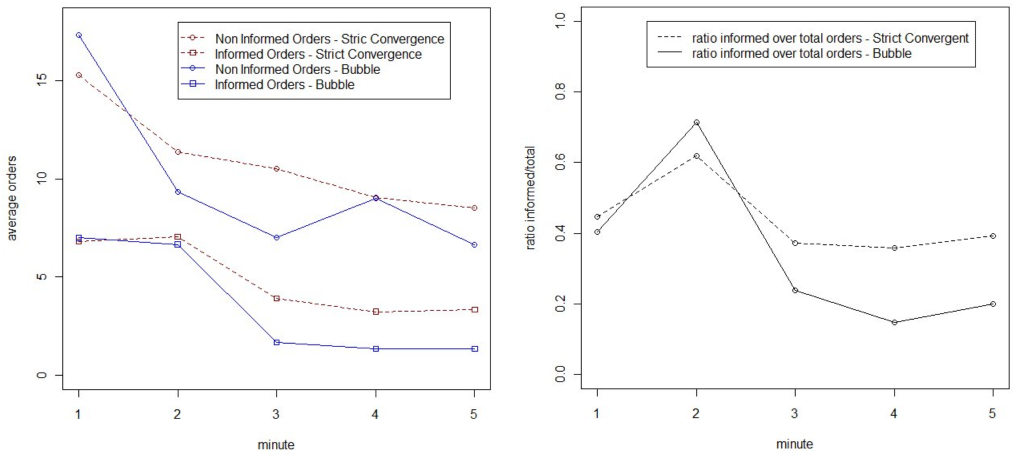

(a) The average number of orders per minute

(b) Order ratio per minute

Figure 13: Average number of orders made by informed traders and non informed traders per minute in bubble and strict convergent periods and their ratio.

of the first non-compatible order is the same as the final direction of the price, in 42 of 45 periods the first minute of non-compatible orders (the direction in the first minute is defined simply as the direction of the majority of non-compatible orders) is the same as the final direction of the price. The non-compatible orders thus seem to have a close connection with the price, the assumption is that this close connection is due to the role of information carrier the non-compatible orders have in the market. It is interesting to investigate the differences between the periods, paying particular attention to the signal behavior. In particular it is interesting to understand how informed and non-informed traders behave in the different periods. Figure 13a shows the average number of orders made every minute by informed agents and non-informed agents in bubble periods and strict convergence periods and figure 13b shows the ratio between the number of informed orders and total orders. A hypothesis to explain the emergence of bubbles is that for some reason the informed agents act less during the first minute leaving the stage to the non-informed agents, however this hypothesis is rejected by the data. Figure 13a shows that the mean number of informed and non-informed orders is similar in the two cases. The first minutes, which are relevant for learning, seem to have a very similar behavior, with a reduction of the informed orders in the last minutes of a bubble.

Since also the average ratio between non-compatible and compatible orders is similar between bubble and strict convergent periods, the next analysis will investigate directly the signal in the different situations. The behavior of the non-compatible orders per minute is shown in figure 14. Since the non-compatible orders are the signal used by non-informed traders, it is important to understand the origin of the signal: informed or non-informed traders. Figure 14a shows the average number of non-compatible orders in strict

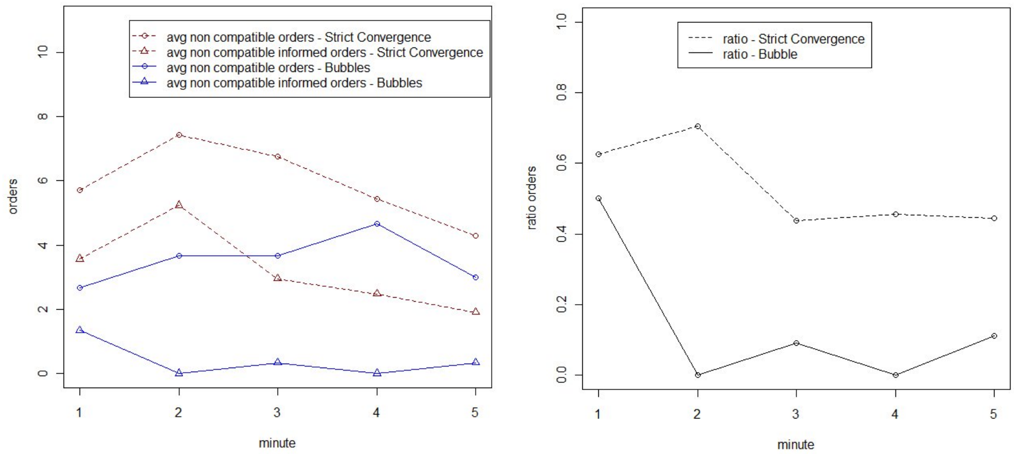

(a) Average number of non compatible orders (b) Ratio informed non compatible orders over total non compatible orders

Figure 14: Source of the signal. The figure shows the number of non compatible orders made by informed traders and total number of non-compatible orders and their ratio both in strict convergent periods and in bubbles periods.

convergence and bubble periods, figure 14b shows the ratio between non-compatible orders made by informed traders and the total number of non-compatible orders. The average number of non-compatible orders is smaller in bubble periods than in strict convergence periods, and the ratio of informed signals is higher in strict convergence periods. A non-informed trader can make a non-compatible order in the right or wrong direction depending on his evaluation of the asset, and the non-compatible orders issued by informed traders should always be in the right direction. Therefore the ratio between the average number of non-compatible orders made by informed traders and the average total number of non-compatible orders (figure 14b) can be considered as evidence of the reliability of the signal introduced into the market. The consequence of the emergence of a bubble is that informed traders stop making non-compatible orders since the market is going in the wrong direction. For example, if the non-informed expected value is 500, the informed expected value is 400 and the market is converging to 600, the non-informed traders will make sell orders following the market price (which are compatible orders), while the non-informed who have learned the wrong information are sustaining the bubble by buying (making incompatible orders).

The last analysis regards the direction of the non-compatible orders. As seen above there is a difference between the source of signals in bubble periods and convergent periods. While a non-compatible order made by an informed agent is in the right direction, the non-compatible orders of non-informed traders can be both in the right and in the wrong direction, there is no *a priori* reason to believe that any direction is predominant

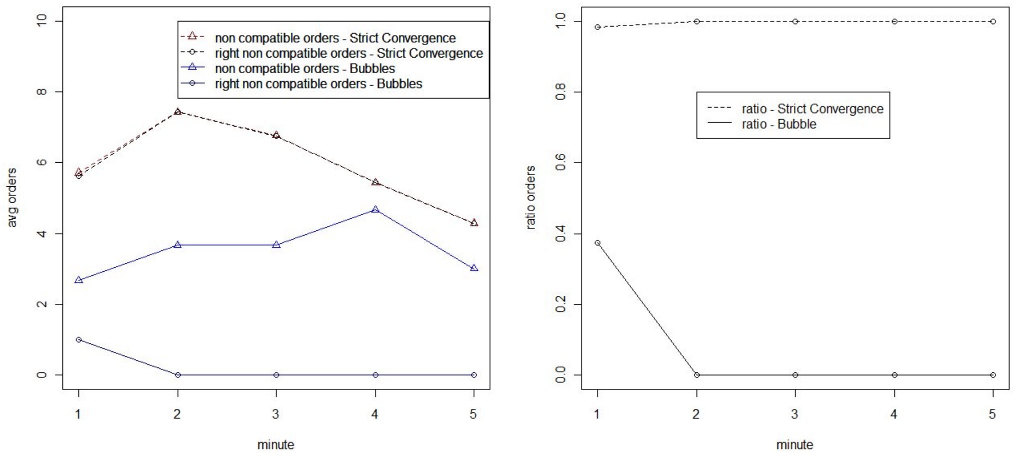

(a) Average number of non compatible orders (b) Ratio right non compatible orders over total non compatible orders

Figure 15: The Signal. The figures show how the signal is the right signal in strict convergent periods and the wrong signal in the bubble periods.

(they should not even make incompatible orders, and why they do it will be investigated in the next sections). The point is therefore that a high percentage of non-informed, non-compatible orders is necessary for the emergence of a bubble (since the informed traders are the only traders that will ever make a non-compatible order in the wrong direction) however this is not sufficient: the non-compatible orders must also be in the wrong direction. Figure 15 shows the average number of non-compatible orders in the right and in the wrong direction in strict convergence periods and bubble periods.

The behavior of agents does not seem to change much in terms of activity, but it changes the quality of the activity. In particular in the bubble, the non-compatible orders are made mostly by non-informed traders and they take the wrong direction. In the first minute the learning occurs after the signal (right or wrong) has been introduced into the market. The analysis above suggests that the non-compatible order is the signal used by the traders to understand the information in the possession of the insiders in the market. In the next section a learning mechanism is defined using Bayesian Learning. Both in strict convergence periods and in bubble periods learning takes place, but the problem is whether the learned equilibrium is the right one or not. Having a learning model will allow us to understand under which conditions the market produces bubbles or convergence.

6.2 Bayesian Learning

The agents have information about the possible states of nature. They know that *a priori* the probability of "up" and "down" are 0.5. Moreover, the agents know that there are insiders in the market, thus they observe the market looking for evidence that can reveal the information possessed by the informed agents. In the previous paragraph a possible signal (evidence) was identified. The structure of the problem fits a Bayesian learning approach; for a good introduction to Bayesian Machine Learning see Mitchell (1997). Given the information possessed by the non-informed agents they have a *a priori* distribution of the dividend. When the non-informed agents observe a non-compatible order they update their beliefs about the state of nature. We can therefore build a learning mechanism that dynamically updates the beliefs of the agents. Every time a signal (an order that in the agents' interpretation carries information) is observed, the (non-informed) agents use that signal to improve their understanding of the environment. The update of the belief that "up" is the state of nature after the observation of an "up" event (E_{up} is a non-compatible buy order) is:

$$P(up|E_{up}) = \frac{P(E_{up}|up)P(up)}{P(E_{up}|up)P(up) + P(E_{up}|down)P(down)} \tag{10}$$

and the update of the belief that "up" is the state of nature after a "down" event (E_{down} is a non compatible sell order) is:

$$P(up|E_{down}) = \frac{P(E_{down}|up)P(up)}{P(E_{down}|up)P(up) + P(E_{down}|down)P(down)} \tag{11}$$

The agent may observe three types of orders: E_{up}, a non-compatible buy order, E_{down} a non-compatible sell order and $\overline{E}$ a compatible order. The first two are considered signals by the agents, while the third carries no information. Suppose that $P(up)$ is the probability assigned by an agent to the event that the information of the insiders is "up" prior to the evidence, and suppose that the same agent observes a non-compatible buy order. By assumption a non-compatible buy order is a signal that is interpreted as evidence in favor of the "up" event. The confidence the agents have in that evidence is measured by $P(E_{up}|up)$: the probability (which may be subjective) that the non-compatible buy order actually comes from a situation in which the state is "up". If the agent is rational and risk-neutral and believes that the other agents are rational and risk-neutral, then $P(E_{up}|up) = 1$. The only possible case in which a non-compatible buy order is made by a risk-neutral rational agent is when he has "up" information. In such a case the non-informed agent should not trade and update the value to the inferred level. Given the behavior of the market we know that the agents make different evaluations of the same random variable, which can be interpreted as different evaluations of the risk. The agents know that they make different evaluations of the asset simply by observing the trade. This behavior introduces noise in the signal since a non-compatible order can be either from an informed agent trying to exploit the information or by a non-informed agent who makes a different evaluation of the random dividend. Therefore we assume that $P(E_{up}|up) < 1$. The assumption that the non-compatible buy order is a signal in favor of the "up" state implies that $P(E_{up}|up) > P(E_{up}|down)$. A reasonable assumption is that the weight the agents give to "up" evidence is equal to the weight given to "down" evidence. As noted above, the weight given to the signals is an evaluation of the market rationality, and since the "up" and "down" evidence comes from the same market, they should be evaluated in same way:

$$P(E_{up}|up) = P(E_{down}|down) \tag{12}$$

A further reasonable assumption is that the probability of observing the wrong signal is the same both in the "up" and "down" state:

$$P(E_{up}|down) = P(E_{down}|up) \tag{13}$$

This last assumption is equivalent to the hypothesis that the probability of observing a compatible order is the same in both possible states. Indeed the probability of observing E_{up} when the state of nature is "up" is:

$$P(E_{up}|up) = 1 - P(E_{down}|up) - P(\overline{E}|up) \tag{14}$$

where $\overline{E}$ is a compatible order. The probability of observing E_{down} when the state is "down":

$$P(E_{down}|down) = 1 - P(E_{up}|down) - P(\overline{E}|down) \tag{15}$$

Since $P(E_{down}|down) = P(E_{up}|up)$ by assumption, it is possible to rearrange equations 14 and 15 to get

$$P(E_{up}|down) = 1 - P(E_{up}|up) - P(\overline{E}|down) \tag{16}$$

Finally, since the possible state is either up or down $P(up|E_{up}) = 1 - P(down|E_{up})$. The belief of state "up" is complementary to state "down", and this is true also for the *a priori* probability: , $P(up) = 1 - P(down)$. These arguments are necessary to simplify the formalization of the learning process. To completely define the Bayesian learning mechanism it is necessary to give a value only to $P(E_{up}|up)$ and $P(E_{up}|down)$ which in behavioral terms represent the way the agents interpret the behavior of the other agents. If for example $P(E_{up}|up) = P(E_{up}|down)$ we would be in a situation in which no information is carried by non-compatible orders (a non-compatible order comes with the same probability from an "up" state or from a "down" state). Assuming $P(E_{up}|up) > P(E_{up}|down)$ means to assume that the agents are interpreting the non-compatible orders as evidence and that the relative magnitude indicates the reliability of the evidence. Since the agents collect more than one evidence, it is possible to rewrite the equation 10 as a dynamical equation where the timing t follows the signal arrival process. Suppose that in t a signal E_{up} arrives. The *a priori* probability of state "up" of the agent (resulting from the evidence collected until t) can be defined p_t. The posterior probability resulting from the update required by the arrival of E_{up} is p_{t+1}. By defining $\omega^{up} = P(E_{up}|up)$ and $\omega^{down} = P(E_{down}|up)$ it is possible to rewrite equation 10 as:

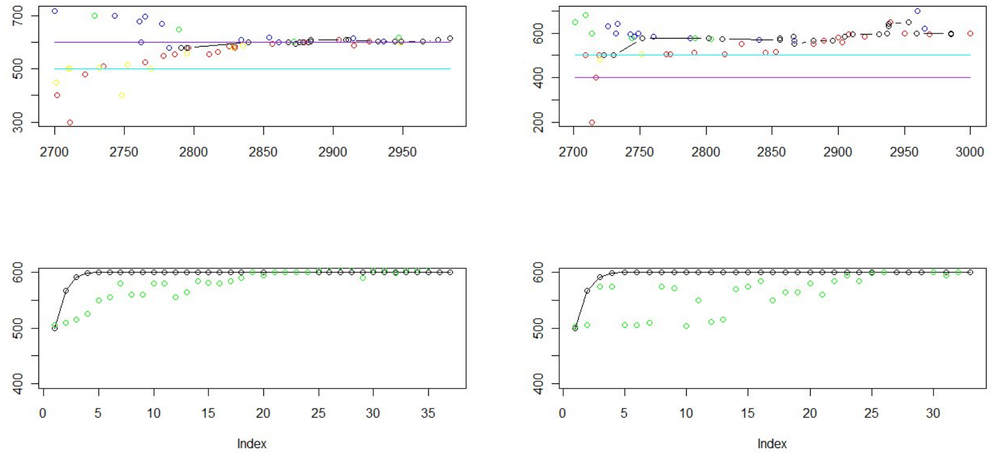

(a) Learning in a convergent period

(b) Learning in a bubble period

Figure 16: The Bayesian learning mechanism was tested in the experimental market.

$$p_{t+1} = \frac{\omega^{up} p_t}{\omega^{up} p_t + \omega^{down}(1 - p_t)} \tag{17}$$

and equation 11 as:

$$p_{t+1} = \frac{\omega^{down} p_t}{\omega^{down} p_t + \omega^{up}(1 - p_t)} \tag{18}$$

By assumption $\omega^{up} > \omega^{down}$. The relative magnitude influences the rate of learning. Equation 17 describes the change in p_t if a non-compatible buy order arrives, while equation 18 describes the change in the beliefs if a non-compatible sell order arrives. $p_0 = 0.5$ is defined by the information possessed by non-informed traders at the start of the period.

In order to understand how the learning mechanism works, it is possible to introduce an artificial agent in the experimental market. The task of the agent is only to learn using the Bayesian learning mechanism just described. The idea is to check whether the artificial agents learn what the human agents learn (i.e. the final equilibrium in the experimental market). In figure 16 two examples are shown with $\omega^{up} = 0.4$ and $\omega^{down} = 0.1$, in both cases the starting expected value is 500, and the trader needs to understand the state of nature; the informed expected value can be either 400 or 600. The artificial agent is learning the equilibrium that is actually reached. Similar behavior of the learning algorithm can be observed in all considered periods.

6.3 Simulation

Given the behavioral model described in the previous sections, it is possible to simulate a market with asymmetric information. The informed traders will trade trying to profit from their information; the non-informed traders will instead trade very safely, with high profit margins while waiting for the information to get out in the market. The main behavioral difference, as noted also in Berner et al. (2005), is the hurry with which the two types of traders trade. The informed traders know that other informed traders are in the market and they start competing with each other. As shown in figure 11 also the non-informed traders learn the information circulating the market and start competing with informed traders. From the model, it is possible to understand when the learning mechanism does not work. We know that the traders make different evaluations of the same expected value and that they will trade depending on the value they assign to the asset and their profit margin. In a normal situation the competition among informed traders will introduce information into the market. In some cases, even if the mechanism is present, the market converges toward the wrong equilibrium. What is happening? Given the market model and the learning model we know that the reached equilibrium depends on the type of orders that are made at the start of the period. As soon as the non-informed traders collect information they use it to trade, and with their trading they are signaling the information they are using to trade. The first part of the period is therefore the part in which the beliefs are formed. The only way in which it is possible to converge toward the wrong equilibrium is that the non-informed traders observe non-compatible orders in the wrong direction. These non-compatible orders cannot be made by informed traders unless we assume a very risky and difficult coordination among informed traders or irrationality. Note that the informed traders do not know who the other informed traders are and coordination is therefore very difficult. Trading to mislead the non-informed is risky and expensive. The other option is that the non-informed traders are making the non-compatible orders. This is plausible, if a non-informed trader has not the patience to wait for the information, she will try to trade very near to the non-informed expected value where she is sure to be able to trade. A non-informed trader with a very low profit margin can influence the other non-informed traders' beliefs by issuing non-compatible order [18]. If at the same time the informed traders use high profit margins, the convergence is not assured anymore. The behavioral model therefore helps us to understand what happens when a bubble arises. The informed traders start with high profit margins that allow them to avoid introducing information into the market. In a situation where also the non-informed traders use high profit margins, the stronger competition among informed traders would allow the right convergence of the market. If instead at least one non-informed trader starts with a low profit margin she might make an order which is considered non-compatible by another non-informed trader and such an event may start the learning process. The direction of a non-compatible order made by a non-informed trader depends on her private value that is

[18] Non compatibility has a subjective definition since it derives from the comparison between the private value and the price of the order.

supposed to be a random variable around the theoretical expected value. The result can be either a convergence or a bubble. Different initial profit margins therefore radically change the properties of the market. In the “normal” case, the market converges always to the right equilibrium; it is therefore stationary and ergodic. In the “not normal” case the convergence of the market depends on the first traders who make orders, and it therefore depends on the initial conditions: the market is stationary and not ergodic with two possible equilibriums. Figure 17 shows a simulated market using the Bayesian Learning with two different profit margin settings. The market is always convergent when the profit margins are the same for both the informed and the non-informed agents (in the simulation the initial profit margin is drawn from a uniform distribution $U(-0.25, 0.25)$). Figure 17b shows the simulation of a bubble market. In this case the learning mechanism is triggered by early non-informed non-compatible orders. With the setting used in the simulation shown in figure 17b (profit margin distributed as $U(-0.01, 0.01)$ for the non-informed and $U(0.55, 0.75)$ or $U(-0.75, -0.55)$ for the informed traders) about 30% of the simulations (over a total of 100 simulations) converged toward the wrong equilibrium [19]. The conclusion is that if the informed traders compete among each other, the market will converge 100% toward the informed equilibrium, but if the informed traders are not competing, the market may go in the right or wrong direction depending on the first orders. In the model the non-informed traders have evaluations symmetrically distributed around the theoretical non-informed expected value and the informed traders evaluate the asset symmetrically around the informed expected value, so the market will tend to go toward the right informed equilibrium (simply because there are more traders on that side of the market). What is crucial for the bubble is that the informed traders do not compete and this is due to the behavior of the non-informed traders who use low profit margins. The informed traders can trade with profit without signaling their information when the non-informed traders use low profit margins. The market will start trading around the non-informed expected equilibrium, and the direction of the price will depend on how the orders arrive and how the contracts are made.

To verify this last conclusion, it is possible to compute the mean profit margin used during the first ten seconds by informed and non-informed traders. There are several problems in this test. The first is that the agents start learning as soon as some signal arrives in the market; it is therefore difficult to estimate the initial profit margin. By considering only the first ten seconds the learning problem is limited but the number of observations is small. The second problem is that, as noted above, it is possible to observe only the price proposed by the trader, and this price is the combination of profit margin and individual evaluation of the asset. The only option for computing the profit margin on the experimental data is to suppose a value for the personal evaluation. Such evaluation rarely coincides with the theoretical expected value. However, it is known from previous analyses that the theoretical expected value is the best estimation of the individual evaluation of the asset. Such an approximation would be disastrous in

[19] In both simulations the learning parameters are $\omega^{up} = 0.44$ and $\omega^{down} = 0.1$. The code for the converging market is available on http://www.jakob.altervista.org/converging.rar. The code for the bubble market is available on http://www.jakob.altervista.org/noconverging.rar

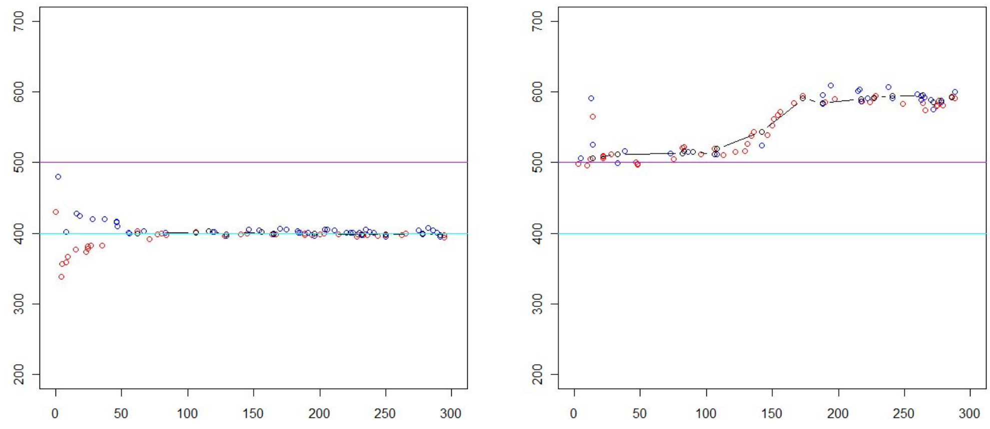

(a) Simulation of a convergent period (b) Simulation of a bubble period

Figure 17: Simulation of a market with non-informed traders using Bayesian learning. The cyan line is the informed expected value, the purple line is the non-informed expected value. The blue dots are sell orders, the red dots are buy orders.

the market simulation, since it would imply no trade, but it can be used in this case were the aim is to compute the mean profit margin used by informed and non-informed traders. The computation of the mean difference between proposed price and theoretical expected value in bubble periods is 0.079 for non-informed traders and 0.457 for informed traders. In the strict convergence periods the results are 0.20 for non-informed traders and 0.233 for informed traders. The results seem to support the mechanism described above. The bubbles arise when the non-informed traders use small profit margins while the informed traders use large profit margin. On the contrary the price converges when the non-informed traders use profit margin high enough to avoid that their orders are considered as signals by other non-informed traders [20] and informed traders use so small profit margins that their orders are considered as signals [21].

7 Conclusions

The aim of this paper is to develop an agent-based model using the information of the experimental market. Even if the replication of the traders' behavior is not perfect, the model seems to mimic the fundamental characteristics of the experimental market. In particular it has been possible to understand how the market converges toward the

[20] The size of the profit margin depends on the variance of the distribution of the individual evaluations of the assets, the higher the variance the higher the profit margin has to be to avoid misunderstandings.

[21] It should be kept in mind that the informed traders might sell above 600 and buy below 500 if their margins were large enough, not creating any signal.

theoretical equilibrium and how the information is disseminated in the market. The behavior of the system depends heavily on the behavior of the elements of the system through interaction. It is evident that a change in one parameter in the behavior of the agents can change the basic properties of the model. The good news is that most of the time the information is correctly understood by the traders, the bad news is that the possibility of bubbles depends on the behavior of the single agents and they are therefore difficult to prevent. Another conclusion is that the transparency of continuous double auction is useful for the traders to understand the environment in order to converge toward the theoretical equilibrium and to understand the insider information present in the market.

It should be underlined that the experiment and the model reflect a very particular situation. The traders have to evaluate a random variable completely known and relatively simple, the insider has certain knowledge about the future and the non-informed traders know about the presence of insiders in the market. This condition has important implications on the outcomes of the market. The agents are able to understand the value of the asset and they normally trade at prices close to it. The certainty of the knowledge in the possession of the insider is important for an aggressive competition among the traders. It has been shown by Plott & Sunder (1988) that if the information is not certain, some problems in the dissemination of the information may arise. Also the informed traders' knowledge is important for inducing the non-informed to look actively for the information. The relatively simple random variable representing the asset value is important for avoiding speculative (or even non-speculative) bubbles in most cases (see for example Smith et al. 1988, King et al. 2001, Smith et al. 2000, Noussair et al. 2001, Lei et al. 2001). The model built in this paper has shown that it is possible to transfer the basic behavior of the agents to different experiments. It would be interesting to build a model able to explain the bubble phenomenon in experimental markets. Bubbles occur for speculative motives and, as shown by Lei et al. (2001), also to some non-speculative motives. It is probable that the non speculative motives derive from confusion in the evaluation of the random variable thereby creating the incentives for the traders to search for information in the market. Gneezy (1996) for example shows that subjects use anchoring (Tversky & Kahneman 1974) to judge compound probabilities. If the information is not present, this mechanism can easily produce positive feedbacks. It is also important to understand the simplification proper of an experimental setting. In a real market it is difficult to extract the information from the market, and what is more important the random variable over the asset value is not defined. The experimental literature combined with agent-based models can give interesting insight into the behavior of the traders, and using the experiment-based agent-based model it might be possible to draw some general conclusions also related to more realistic settings that cannot be replicated in the laboratory.

Appendix

The appendix describes shortly the nonparametric stationarity and ergodicity tests. An agent-based model is a very powerful tool for the study of complex systems: it is possible to build a realistic model, understand emergent properties and experiment different settings. The main drawback is that it is difficult to generalize from the properties of the model. As stated by Axtell (2000, p.3), each run of the model yields a sufficiency theorem but a single run does not provide any information on the robustness of such theorem. In order to understand the behavior of the model using different parameters and different random seeds it is necessary to compute multiple runs to assess the robustness of the results. In economics the property of a unique stable equilibrium is fundamental to the understanding of the behavior of the model and thereby the behavior of the system under analysis. In a "traditional" mathematical model it is usually possible to formally prove the existence (or not) of a unique stable equilibrium. In an agent-based model such formal proof is not possible by construction since the agent-based model "consists of individual agents, commonly implemented in software as objects" Axtell (2000, p.2). The way followed in this paper to assure the presence of a unique stable equilibrium is to use non-parametric statistical tests over the artificial data. Stationarity of the artificial data assure that the model reaches an equilibrium state. Ergodicity assures that the equilibrium is unique regardless of the initial condition[22]. The tests used are described in Grazzini (2011*b*). The tests are performed over the model described in section 2 which in the following will be called "model 1" and on the model in section 3 which will be called "model 2".

The stationarity test tests whether a given moment is constant. The test is very simple; the first step is to produce a time series with n observations using the model. The time series has to be divided into windows with a length of n_w observations. The non-centered moment of order m on each window, w_t, has to be computed to create a new time series $\{w_t\}$. If the non-centered moment of order m is constant it means that there is no systematic variation in w_t and that$\{w_t\}$ is well-fitted by the non-centered moment of order m computed over the whole time series. The non parametric test used is the Wald-Wolfotwitz test (Wald & Wolfowitz 1940), for details see Grazzini (2011*b*). The choice of n and n_w (which implies in turn the choice of the number of windows) is important for the power and the size of the tests. In particular, in order to have full power the number of windows has to be higher than 50. The size of the test depends on the properties of the time series. If the time series is strictly stationary, the test can detect the stationarity even using windows of just one observation. If the time series is weakly stationary, the test needs longer windows to have a good estimate of the window moments. In the stock market model described in sections 2 and 3, the number of observations depends on the number of trade periods and on the traders' activity. The time series produced by model 1 was tested using a number of periods corresponding to an average of 6000 observations (transactions) and n_w = 100. Model 2 was tested

[22]Provided that the model is well specified and that a sufficient amount of data is available, ergodicity and stationarity assure also the possibility of a consistent estimation of the model (Grazzini 2011*a*).

using a number of periods that corresponds to an average of 90,000 observations and $n_w = 1000$. As noted also in section 3, model 2 is less stable around the equilibrium, therefore longer windows are needed to detect the stationarity. The tests cannot reject the null hypothesis; the models have a stable equilibrium. Over 100 simulations of model 1, the rejection rate is of 4% and the rejection rate for model 2 is 6%; both are coherent with the chosen α-error (5%). From the stationarity test it is also possible to infer that the models are not strictly stationary since we need long windows to detect the stationarity. This is because the profit margins are influenced by the contracts and therefore there is dependence between successive transaction prices.

The ergodicity test tests whether a given moment is constant in different runs of the model. If the time series is stationary (tested above), the time series is ergodic if the sample moments computed over a set of k_1 windows with a length of n_w observations from one long time series come from the same distribution as the sample moments computed over k_2 different processes with a length n_w. Stationarity means that the given moment is constant *within* one run while ergodicity means that the given moment is constant *between* different runs. If the moment is constant during one run we infer that the model has a stable equilibrium, if the moment is constant in different runs we infer that the model has only one equilibrium regardless of the starting conditions. The number of periods chosen for the long time series and for the short time series for the model in section 2 are such that $k_1 = 50$, $k_2 = 50$ and $n_w = 100$, for the model in section 3 are such that $k_1 = 50$, $k_2 = 50$ and average $n_w = 1000$. Note that these are average values since the number of contracts (and therefore of observations) is not deterministic as it depends on the activity of the traders. The power of the test depends again on the number of observations in the samples (k_1 and k_2) and the size of the test depends on n_w. The windows have to be long enough to detect the stationarity of the process. The test cannot reject the null hypothesis of ergodicity; the models have a unique stable equilibrium. Over 100 tests the rejection rate for model 1 is 4% and the rejection rate for model 2 is 6%, both coherent with the chosen α-error (5%). The Python files used for the tests of model 1 can be downloaded from:
http://www.jakob.altervista.org/Python-Test-model1.rar.
The Python files used for the tests of model 2 can be downloaded from:
http://www.jakob.altervista.org/Python-Test-model2.rar.

Bibliography

Anufriev, M., Hommes, C. & Philipse, R. H. (2010), 'Evolutionary selection of expectations in positive and negative feedback markets', *CeNDEF Working paper 10-05 University of Amsterdam* .

Axtell, R. (2000), 'Why agents? on the varied motivations for agent computing in the social sciences', *Center on Social and Economic Dynamics Working Paper* (17).

Berner, M., Feri, F. & Plott, C. R. (2005), 'On the microstructure of price determination and information aggregation with sequential and asymmetric information arrival in an experimental asset market', *Annals of Finance* **1**, 73–107.

Boero, R., Bravo, G., Castellani, M. & Squazzoni, F. (2010), 'Why bother with what others tell you? an experimental data-driven agent-based model', *Journal of Artificial Societies and Social Simulation* **13**(3), 6.
URL: *http://jasss.soc.surrey.ac.uk/13/3/6.html*

Boero, R., Bravo, G. & Squazzoni, F. (2010), 'Trust and Partner Selection in Social Networks: An Experimentally Grounded Model', *ArXiv e-prints* .

Camerer, C. & Weigelt, K. (1991), 'Information mirages in experimental asset markets', *Journal of Business* **64**(4), 463–493.

Chan, N. T., LeBaron, B., Lo, A. W. & Poggio, T. (2001), 'Agent-based models of financial markets: A comparison with experimental markets', *MIT Sloan Working Paper* (4195-01).

Cliff, D. & Bruten, J. (1997), 'Minimal-intelligence agents for bargaining behaviors in market based environments', *HP Laboratories Bristol* (HPL-97-91).

Contini, B., Leombruni, R. & Richiardi, M. (2007), 'Exploring a new expace.the complementarities between experimental economics and agent-based computational economics', *Journal of Social Complexity* **3**(1), 13–22.

Copeland, T. E. & Friedman, D. (1986), 'The effects of sequential information arrival on asset prices: An experimental study', *Journal of Finance* **42**(3), 763–797.

Copeland, T. E. & Friedman, D. (1991), 'Partial revelation of information in experimental asset markets', *Journal of Finance* **46**(1), 265–295.

Duffy, J. (2006), Chapter 19 agent-based models and human subject experiments, Vol. 2 of *Handbook of Computational Economics*, Elsevier, pp. 949 – 1011.

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

Epstein, J. M. & Axtell, R. (1996), *Growing Artificial Societies: Social Science from the Bottom Up*, The MIT Press.

Forsythe, R. & Lundholm, R. (1990), 'Information aggregation in an experimental market', *Econometrica* **58**(2), pp. 309–347.

Gilbert, N. (2001), Holism, individualism and emergent properties. an approach from the perspective of simulation, *in* J. J. Heckman & E. E. Leamer, eds, 'Modelling and Simulation in the Social Sciences from the Philosophy of Science Point of View', Elsevier.

Gjerstad, S. & Dickhaut, J. (1998), 'Price formation in double auctions', *Games and Economic Behavior* **22**(1), 1–29.

Gjerstad, S. & Sachat, J. M. (2007), 'Individual rationality and market efficiency', *Institute for Research in the Behavioral, Economic and Management Science* (1204).

Gneezy, U. (1996), 'Probability judgments in multi-stage problems: Experimental evidence of systematic biases', *Acta Psychologica* **93**, 59–68.

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

Grazzini, J. (2011*a*), Consistent estimation of agent based models, LABORatorio R. Revelli Working Papers Series 110, LABORatorio R. Revelli, Centre for Employment Studies.
URL: *http://ideas.repec.org/p/cca/wplabo/110.html*

Grazzini, J. (2011*b*), 'Emergent properties: Stationarity and ergodicity', *MIMEO - Department of Economics "Cognetti de Martiis", University of Turin* .

Holden, C. W. & Subrahmanyam, A. (1992), 'Long-lived private information and imperfect competition', *Journal of Finance* **47**(1), 247–270.

Hommes, C. (2011), 'The heterogeneous expectations hypothesis: Some evidence from the lab', *Journal of Economic Dynamic and Control* **35**, 1–24.

Hommes, C., Sonnemans, J., Tunistra, J. & van de Velde, H. (2005), 'Coordination of expectations in asset pricing experiments', *The Review of Financial Studies* **18**(3), 955–980.

King, R. R., Smith, V. L., Williams, A. W. & van Boening, M. (2001), 'The robustness of bubbles and crashes in experimental stock markets', *Econometrica* **69**(4), 831–859.

Kyle, A. S. (1985), 'Continuous auction and insider trading', *Econometrica* **53**(6), 1315–1335.

Lei, V., Noussair, C. N. & Plott, C. R. (2001), 'Nonspeculative bubles in experiemntal asset markets: Lack of common knowledge of rationality vs actual irrationality', *Econometrica* **69**(4), 831–859.

Mitchell, T. M. (1997), *Machine Learning*, McGraw-Hill.

Noussair, C., Robin, S. & Ruffieux, B. (2001), 'Price bubbles in laboratory asset markets with constant fundamental values', *Experimental Economics* **4**, 87–105.

Plott, C. R. (2000), 'Market as information gathering tools', *Southern Economic Journal* **67**(1), 1–15.

Plott, C. R. & Sunder, S. (1982), 'Efficiency of experimental security markets with insider information: An application of rational-expectations models', *The Journal of Political Economy* **90**(4), pp. 663–698.

Plott, C. R. & Sunder, S. (1988), 'Rational expectations and the aggregation of diverse information in laboratory security markets', *Econometrica* **56**(5), pp. 1085–1118.

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

Smith, V. L. (1982), 'Microeconomic systems as an experimental science', *American Economic Review* **72**, 923–955.

Smith, V. L., Suchanek, G. L. & Williams, A. W. (1988), 'Bubbles, crashes, and endogenous expectations in experimental spot asset markets', *Econometrica* **56**(5), 1119–1151.

Smith, V. L., van Boening, M. & Wellford, C. P. (2000), 'Dividend timing and behavior in laboratory asset markets', *Economic Theory* **16**(3), 567–583.

Tversky, A. & Kahneman, D. (1974), 'Judgment under uncertainty: Heuristics and biases', *Science* **185**(4157), 1124–1131.

Wald, A. & Wolfowitz, J. (1940), 'On a test whether two samples are from the same population', *The Annals of MAthematical Statistics* **11**(2).