Heterogeneous Trading Agents

L. Neuberg
University of Namur
Rempart de la Vierge 8,5000 Namur, Belgium
K. Bertels
Delft University of Technology
Mekelweg 4, 2628 CD Delft, the Netherlands
koen@dutepp0.et.tudelft.nl

ABSTRACT
In this paper, we present a multi agent system (MAS) simulation of a financial market and investigate the requirements to obtain realistic data. The model consists of autonomous, interactive agents that buy stock on a financial market. Transaction decisions are based on a number of individual and collective elements. The former being risk aversion and a set of decision rules reflecting their anticipation of the future evolution of prices and dividends. The latter is the information arriving on the market influencing the decision making process of each trader. We specifically look at this process and the following observations hold : The market behaviour is determined by the information arriving at the market and agent heterogeneity is required in order to obtain the right statistical properties of the price and return time series. The observed results are not sensitive to changes in the parameter values.

KEYWORDS
multi agent systems, financial markets, agent based simulation, classifier systems

I. Introduction

In this paper, we present a model that simulates the behaviour of a heterogenous collection of financial traders on a market.[1] Each trader is modelled as an autonomous, interactive agent and the agregation of their behavior results in market behaviour. We emphasize that the main goal of the paper is not to predict the future evolution of any stock, but rather to gain a deeper understanding of the phenomena observed in financial markets and to study the conditions under which realistic data is obtained.
The main contributions of the paper are the following :

- The simulations suggest that the information arriving at the market determines to a high degree how the market will behave.
- In function of the information arriving at the market, crashes or speculative bubbles appear.
- Furthermore, it appears that in introducing heterogeneity, the overall market dynamics changes. An even stronger claim is that only by introducing heterogeneity does the model reproduce a market dynamics similar to real world financial price dynamics.

[1] Some of the results of this paper are also described in [13].

The paper is organized as follows. We first introduce the model and explain how each agent is modelled and how their interaction results in the overall market behavior. We then present the results of three simulation runs, one for a crash situation, one in which a speculative bubble appears and one representing a ’normal’ situation.

II. Related Research

The Santa Fe Artificial Stock Market, as described in [1], served as a starting point for the model described in this paper. In their model, financial agents are recursive in nature as they form beliefs and expectations about the market on the basis of what they believe will be the other agents’ expectations. They claim to provide an answer to an old debate in which practitioners claim that there are speculative opportunities in the market, whereas academics believe in its efficiency. Their model shows that both views are correct, given the degree of explorative capabilities of the agents. In short, when agents are only marginally allowed to explore alternative expectational models, the market price converges to the rational expectations equilibrium price. However, when agents can explore alternative models, a complex price pattern emerges allowing the emergence of bubbles or crashes.

Another model of a stock market using a similar approach is described in [4]. Each agent is described by a mathematical function and he uses a set of rules to form expectations about the future prices of a stock. This approach is different from the one used in [1], or in this paper, as the learning is implemented as a modification of the parameters of the mathematical function describing each agent. The main findings are that the initial wealth held by an agent and the method used to predict future prices largely influence the success of that agent on the market.

In [5], volatility clustering is explained in terms of certain proportions between chartist and fundamentalist trading strategies present on the market. Their model, using also the interactive agents approach, shows that when a certain threshold is surpassed, a sudden outbreak of volatility occurs. They see similarities with the on-off intermittency behavior in physics. They furthermore verify that the artificially generated time series have the same properties as real world financial data.

III. Description of the Model

A. The financial Agent

We distinguish between different kinds of traders on the market, each having their own rationality and knowledge. As any financial trader, the agent must be able to evaluate an action and form an expectation with respect to its future price. On the basis of this expectation, he will propose a transaction price and quanity. This offer can then be evaluated by other traders on the market. The expectations are the result of some kind of reasoning and decision making. Depending on the success of the proposed transaction, measured in terms of financial profit, the agent will modify his decision rules and thus learn.

Decision making and expectations formation : as explained above, each agent needs to be able to decide whether he wants to buy or sell a particular stock, and at what price. He therefore needs to have decision rules that allow him to make some kind of expectation as to the future evolution of the price. He will do so on the basis of information at his disposal. In our model, we have chosen to implement a classifier system where different decision model are represented as if-then rules. At a given moment, if a condition of his set satisfies the present situation in the environment, the agent will take the corresponding action. The condition of each rule is a chain of characters(" 0 ", " 1 ", or " # ") determining whether the rule is equivalent to the market situation. This equivalence is achieved if the characters along the chain of the condition are similar to the characters along the chain of the market situation. In the case of character " # ", there is always equivalence to the extent that it expresses the indifference between the characters " 1 " and " 0 ". As for the action, it is a chain of characters representing the value of two parameters *a* and *b* in binary fashion. These parameters allow to compute the expected future prices and dividends in the following way :

$$E[P_{t+1} + d_{t+1}] = a(P_t + d_t) + b \tag{1}$$

where P_t is the stock price at time t and $d_t$is the dividend of the risky asset at time t. For each agent, a set of rules allowing to calculate these expected prices and dividends will be generated using genetic algorithms. Initially, 900 rules are generated. This number will be reduced during the learning process. Risk aversion is expressed in terms of the CARA utility function which, for the sake of comparability, is taken from [1].

$$U(w) = -e^{-\lambda w} \tag{2}$$

w represents the wealth of the trader and λ indicates the degree of risk aversion and is defined in the following way :

$$\lambda = -\frac{U''(w)}{U'} \tag{3}$$

The higher the value of λ, the higher the risk aversion of the agent.

Learning : in this original set of 900 rules, some may be more efficient than others. Those rules yielding more accurate expected prices and therefore a higher financial gain will have a higher reproduction rate and a higher probability to survive. The frequency of the re-actualization of the rules will depend on each agent's ability to learn. Before starting the actual trading sessions, we introduced an initialisation phase which makes the 900 initial rules less random and already in some way tuned for the trading agents. Where other models also had similar initialisations, e.g. [1], requiring over 200000 iterations before some kind of homogenous rational behavior emerged, we decided to introduce an initial learning phase based on real historical data resulting in more realistic decision rules. The effectiveness of the decision rules is defined in function of the error generated by the rule and is computed as follows :

$$Error = (E[P_{t+1} + d_{t+1}] - P_{t+1})^2 \tag{4}$$

using equation (1) results in

$$Error = (a(P_t + d_t) + b - P_{t+1})^2 \tag{5}$$

A perfect rule will compute an expected value equal to the price and the future dividends and the error will be zero. These rules will have a maximal evaluation value. If we represent this maximal value by C, we obtain a rule evaluation function that is defined as follows :

$$EVAL(rule) = C - (a(P_t + d_t) + b - P_{t+1})^2 \tag{6}$$

also called the strength of the rule.

B. The market

Fig. 1. The Information Frequency Distribution : Normal Situation

Information : as in real life, expectations with respect to prices and dividends are largely influenced by information arriving on the market. In our model, information arrives at the market at regular intervals of time. This information may vary from 'very negative (-3)' over 'neutral (0)' to 'very positive (+3)'. Figure 1 represents the distribution of the different kinds of news flashes for a normal situation. We emphasize that not every agent may interpret the same piece of information in the same way.

Price formation and market clearing : Intersecting orders to buy and sell are going to create the dynamics of asset prices (see Figure 2). The market clearing mechanism is similar to the one used in [1] in which bids are continuously resubmitted until a price is formed that clears the market. At each period of time, the agents try to optimize the allocation of risky and non-risky assets. Initially, the price and dividend previsions made by agent i at time t are normally distributed with an average of $E_{i,t}[p_{t+1}+d_{t+1}]$ and a variance $\sigma^2_{t,i,p+d}$. Demand (or supply) by agent i at time t is given by :

$$x_{i,t} = \frac{E_{i,t}(p_{t+1}+d_{t+1}-(1+r)p_t)}{\lambda\sigma^2_{i,t,p+d}} \quad (7)$$

where p_t is the price of the asset at time t and λ is the degree of risk aversion.

In order to close the system, total demand must be equal to the number of available goods on the market :

$$\sum_{i=1}^{N} x_{i,t} = N \quad (8)$$

Fig. 2. Aggregated Market Behaviour

IV. Simulation Results

It is important to emphasize that we are focusing on simulation and not on prediction. We therefore do not look at real world markets and this for two reasons. First, even though there are large data series with stock market data available, they all lack one important aspect : a link between the price evolution and the information arriving on the market, such as news flashes from Bloomberg or Reuters. We therefore generate an artificial time series in which there is a direct relation between the information and the price evolution. Secondly, the main objective of this research, as explained above, is to establish whether real world stock market phenomena can be reproduced. The ultimate validation of the model with respect to this objective is done in terms of looking at the statistical properites. Real world financial data have particular statistical properties.[3][2][8] [7] The first one is that the returns on the market have zero or positive skewness, which is the degree of asymmetry of the distribution with respect to its mean. A possible cause of this property could be that the the underlying process is highly unlinear. [11] The second property is that the returns have a positive excess kurtosis. One explanation for this observation is the presence of different kinds of actors on the market inducing a higher frequency of extreme events (high losses or profits) and thus fatter tails. [8]. These statistics are easily computed and serve as benchmark with respect to the validity of the approach.

Sim.–RAT	10	7	5	4	3	2
Normal	0.94	0.96	0.91	0.88	0.32	0.14
Crash	0.99	0.95	0.92	0.88	0.72	0.70
Bubble	0.99	0.99	0.98	0.98	0.96	0.96

TABLE I
Correlation Coefficient

Sim–RAT.	10	7	5	4	3	2
Normal	1217	894	587	519	234	258
Crash	4561	2505	1598	1216	362	356
Bubble	5120	3638	2813	2641	1733	1254

TABLE II
Volatility (Standard Deviation)

The time series, representing the price evolution in function of a particular information vector was generated using the following equation :

$$P_t = (1+\alpha I_{t-1})P_{t-1} \quad (9)$$

where I_{t-1} represents the information, P is the price and α is the sensititivy to the news arriving on the market. Before the actual simulations, we had the traders learn this particular relation. The goal of this ”mode setting” or initial learning phase is twofold : First, we can reduce considerably the 900 rules to a more manageable couple of hundreds. And secondly, most of the rules obtained as the result of the learning process, will make more sense than the original ones who were generated randomly.

In the remainder of the paper we will use the following terms :

- Normal Agents : are those agents that will have learned this mechanism of how to use the information to compute a future price.
- Perturbating Agents : are those agents who will deviate from this mechanism.
- Reference Time Series : this is a time series computed using equation 9 on the basis of a new information vector.
- Generated Time Series : these are the ones generated by the interacting agents.

In each table, the heading Sim.-RAT indicates the number of rational agents used in the simulation. For each of the simulations, we compute the following statistics which are summarized in Tables I to IV and which are based on the data plotted respectively in figures 3,5 and 7 :

- correlation coefficient between the reference time series of the prices and the generated one.
- The standard deviation measures price volatility and indicates the riskiness of a particular product.

Sim.–RAT	10	7	5	4	3	2
Normal	0.07	-0.43	-0.29	-0.21	0.65	0.70
Crash	-0.99	-1.22	-0.89	-0.48	0.28	0.29
Bubble	1.26	1.62	2.83	3.36	4.53	4.11

TABLE III
SKEWNESS

Sim.-RAT	10	7	5	4	3	2
Normal	1.68	2.01	4.09	4.48	14.89	14.50
Crash	1.01	2.79	4.46	6.04	13.33	12.69
Bubble	4.17	5.85	13.19	16.26	23.39	24.08

TABLE IV
EXCESS KURTOSIS

- Skewness and Kurtosis are computed on the returns. As explained above, a positive skewness and positive excess kurtosis are characteristic for real world financial data.

For each set of simulations, we introduce heterogeneity by giving, in consecutive runs, a certain number of agents different kinds of decision making behavior. Besides the agents that will use Equation 9, we will have crazy agents who behave totally random, agents who will always be the opposite of the ’normal’ ones (inverse) and the third category of heterogeneous agents are the ones that attach less importance to extreme values of the information arriving at the market (filter). Table V summarizes the proportion of each type for each of the runs.

V. NORMAL REGIME

We first discuss the results of the different simulations using a normal distribution of the news flashes arriving at the market. In three consecutive simulations, we increased the number of deviating agents as indicated by the table headings Sim.-RAT that goes from 10 to 2. The statistical information can be found in the different tables having as line heading ’Normal’.

Simulation 1

In this simulation, we modeled 10 agents having different sets of parameters (a and b in Equation 9) and each having his learned set of decision rules. The information vector used for this simulation is different than the one used during initial learning. If the agents have learned well the price dynamics during the initial learning phase, we expect that they should be able to reproduce similar (but differerent) dynamics. The differences could then be primarily due to the differences in the exogeneous variable I_t. The price dynamics is given in Figure 3. Where all the agents are rational,the correlation coefficient between the reference time series and the generated time series is 0.94 which shows a great similarity between the two. The skewness is very small but positive(0.07), indicating that the distribution of the returns is quasi normal. A positive excess kurtosis implies that the distribution is peaked. This seems to imply the following :

- the agents reproduce the correct dynamics. This claim is supported by the correlation coefficient of 0.94.
- The interaction on the market does not introduce a higher (positive) skewness even though the returns have a peaked distribution.
- We might also advance that the agents are apparently applying the decision mechanism they were taught.

Type	7 RAT	5 RAT	4 RAT	3 RAT	2 RAT
inverse	1	2	2	3	3
Filter	1	2	3	3	3
Crazy	1	1	1	1	2

TABLE V
INTRODUCING HETEROGENEITY IN THE SIMULATIONS

Fig. 3. Price Dynamics of Simulations 1 and 2

Simulation 2

We now investigate whether or not the presence of perturbating agents can influence the market in such a way that the dynamics change. This boils down to introducing heterogeneity in the agents. To this purpose, we introduce, in consecutive simulations,from 1 to 8 perturbating agents that will systematically react differently than the others. Their interpretation of the information arriving on the market will be different, pushing them to make a different decision. The simulation counts the same number of periods and the same information vector has been used. This way, we can better compare the resulting prices with the time series of the previous simulation. The correlation coefficient goes down from 0.94 to 0.14, as the number of perturbating agents increases. Volatility decreases from 1217 to 258, as measured in terms of standard deviation. We also observe a systematic increase of the excess Kurtosis (see Table IV, Normal Regime) indicating the occurence of a fat tail in the returns. Conform with current explanations (see above), this seems to be the direct consequence of introducing heterogeneity among the trading agents. In addition, skewness goes from 0.07 to 0.7 which again is conform realistic market data. Several similar runs of the model seem to indicate that due to differences in initial states, the proportion at which a positive skewness occurs differs. From the above, we can conclude that introducing heterogeneity results in the appropriate statistical properties of the timeseries generated.

Fig. 4. The Information Frequency Distribution : Crash

Fig. 5. Price Dynamics of Simulation 3 and 4

VI. Market Crash

Simulation 3

We now introduce a new information vector as the basis for the market dynamics. Rather than looking at a normal market situation, where there is no dominant trend in the information, we now simulate the situation in which bad news arrives at the market in a more or less constant way. The distribution of the information is given in Figure 4. As we can observe from Figure 5, there is a clear negative trend in the market. We also see from the computed standard deviations, that the volatility has increased drastically (from 1217 to 4755) which is in concordance with reality. Markets in crisis behave always more nervously than markets in a normal state. From Table I, we can see that the correlation coefficient is still very high.[2] This leads us to suppose that still the same underlying decision taking mechanism is applied. Skewness and excess Kurtosis give contradictory indications as the skewness is negative, which is not conform real world markets and the excess kurtosis has the appropriate sign.

Simulation 4

We again introduce, in consecutive runs, a number of perturbating agents. However, these agents are different than the perturbating ones in simulation 2. The heterogeneity is introduced by imposing these agents to attach less importance to very negative information. The information arriving at the market is the same as in the previous simulation. As we can see from Figure 5 and from Tables I and II, the market trend is still downward in all cases even though the downward trend diminishes as the number of perturbating agents increases. This is logical given the 'rationality' imposed on the perturbating agents. The correlation coefficients remain high (from 0.99 to 0.71) and the volatility diminishes as the number of perturbating agents increases. The volatility remains systematically higher than the volatility in the normal regime. Again, a positive skewness occurs when more heterogeneity is introduced and in all of the simulations, an excess kurtosis is found. The following conclusions can be advanced :

- Analogously to simulation 2, we observe that in introducing heterogeneity in the agents, the generated time series have properties similar to those of real world financial data.
- We furthermore see that the constant inflow of bad news, causes the market to crash. The price dropped 50% and volatility, compared to the volatility of simulation 2, has risen with a factor of 4.

Fig. 6. The Information Frequency Distribution : Bubble

Fig. 7. Price Dynamics of Simulation 5 and 6

[2] We now use as reference time series one that uses the same 'bad news' information vector as a point of comparison.

λ	10	7	5	4	3	2
0.5	7.529	8.239	9.188	9.546	15.171	23.937
1	7.512	8.130	8.798	9.401	14.708	22.416
2	7.514	7.908	8.726	8.895	15.017	21.614

TABLE VI
KURTOSIS-STATIONARY REGIME

λ	10	7	5	4	3	2
0.5	0.002	-0.012	0.099	0.210	1.094	3.605
1	-0.047	-0.015	-0.240	0.145	0.864	2.951
2	-0.138	-0.249	-0.785	-0.772	-0.248	2.733

TABLE VII
SKEWNESS-STATIONARY REGIME

α	10	7	5	4	3	2
0.025	3.021	3.295	4.565	4.562	4.562	4.975
0.01	7.689	8.018	10.094	10.480	20.897	20.504
0.005	2.983	3.428	4.680	5.021	7.584	7.730

TABLE VIII
KURTOSIS-STATIONARY REGIME

α	10	7	5	4	3	2
0.5	-0.217	0.108	0.497	0.542	0.542	0.672
1	0.079	-0.43	-0.29	-0.21	0.650	0.706
2	-0.1168	0.111	0.365	0.449	0.913	1.004

TABLE IX
SKEWNESS-STATIONARY REGIME

VII. SPECULATIVE BUBBLE

Simulation 5

In a third series of simulations, we introduce a new information vector (see Figure 6) where there is systematically good news arriving at the market. The price evolution is given in Figure 7. In the first run in which all agents are similar, we see a clear upward trend of the market. The volatility is very high, compared to the normal regime, and we immediately have a positive skewness and an excess kurtosis.

Simulation 6

We introduce heterogeneity in consecutive runs to see how the market behaves. Again, we observe that the bubble is less pronounced and even seems to disappear when heterogeneity is increased. The correlation coefficient remains high (from 0.99 to 0.96). Volatility goes down from 5120 to 1254 but, similar to the crash situation, remains systematically higher than the volatility in the normal case. Skewness increases with the number of heterogeneous agents, just as the excess kurtosis increases from 4 to 23.

VIII. SENSITIVITY ANALYSIS

Even though an exhaustive exploration of the parameter space is impossible, we did investigate how sensitive the results are for differences in certain parameter. We more specifically look at λ from equation 2 and α from equation 9. We also investigate the influence of the news distribution on the volatility of the stock prices in each of the different regimes. Tables VI and VII contain the computed values for the skewness and kurtosis for different values of λ. We can see from these tables that the values of both statistics increase as λ decreases. This implies that as far as the statistical characteristics are concerned, the findings still hold and the data still has the correct (real world) properties.

As far as α is concerned, a similar observation can be made on the basis of the data in tables VIII and IX. Irrespective of the value of α, the skewness and kurtosis increase as the heterogeneity of the traders increases. For all of the other market regimes, a similar behaviour is observed. Full details can be found in [12].

Finally, we turn to the influence of the distribution of the information vectors on the price volatility. In real markets, volatility is higher during a crash than in the stationary or speculative bubble market regime. Consequently, it is also expected to find a similar observation. Table X represents the volatilities for the 3 different market regimes. Conform to real world markets, the volatility is higher when the market crashes than in any other regime. This observation holds for both the fully homogeneous and heterogeneous population of trading agents.

IX. CONCLUSION AND FURTHER RESEARCH

In this paper, we have presented the main results of a simulation experiment implementing an artificial financial market. The main goal of the research was not stock market prediction but to study the methodological issues involved in reproducing the major behavioural phenomena of financial markets such as crashes and speculative bubbles. The main conclusions are the following. Information plays a crucial role in the way the market behaves. Each set of simulations clearly shows a different behavior of the market when different information sets are used. The second main conclusion is that the model generates realistic market dynamics only when introducing heterogeneity among the trading agents. A third conclusion that can be drawn from the above results is that, as far as the volatility is concerned, the increase of this statistic is observed whenever the market tends towards a crash or a speculative bubble. Further research is needed to confirm the above results. One of the things one might look at is what the influence is of different proportions of normal and perturbating agents and to introduce a richer psychological profile for each of the traders.

Regime	10	7	5	4	3	2
Stationary	1237	1051	1052	847	743	999
Crash	3692	3738	3209	3111	2282	2166
Bubble	5086	3216	1611	1466	1137	2343

TABLE X
PRICE VOLATILITY

References

[1] Arthur W.B., Holland J.H., LeBaron B., Palmer R., Tayler P.; *Asset Pricing Under Endogenous Expectations in an Artificial Stock Market*, in [6], p.15-44.

[2] Bouchaud JP, Potters M., Theory of Financial Risks, Cambridge University Press, Cambridge, 2001, 218 p.

[3] Campbell J.,Lo A.W.,MacKinlay A.C.,*The econometrics of financial markets*, Princeton University Press, Princeton, New Jersey, 1997.

[4] Wan H.A., Hunter A.;*On Artificial Adaptive Agents Models of Stock Markets*, Simulation 68:5, pp. 279-289.

[5] Lux, T., Marchesi M.;*Volatility Clustering in Financial Markets : A Micro- Simulation of Interactive Agents*, Proceedings of the 3rd Workshop on Economics and Interacting Agents, Ancona, 1998, http://www.econ.unian.it/dipartimento/siec/HIA98/papers/program.htm

[6] Arthur W.B., Durlauf S.N., Lane D.A.; *The Economy as an Evolving Complex System II*, Santa Fe Institute, Addison-Wesley, Volume 27, 1997, 583 p.

[7] De Grauwe P., Dewachter H., Embrechts M.; *Exchange Rate Theory. Chaotic Models of Foreign Exchange Markets*, Blackwell, Oxford, 1993, 273 p.

[8] Guillaume D. et.al.;*From the bird's eye to the microscope : A survey of new stylized facts of the intra-daily foreign exchange markets*, Finance and Stochastics 1, pp. 95-129, 1997.

[9] Holland J.H., Miller J.H.; *Artificial Adaptive Agents in Economic Theory*, American Economic Review, Papers and Proceedings 81, May 1991, p.365-370.

[10] Kirman A.P.; *Economies with Interacting Agents*, Working Paper, EHESS and Universite d'Aix-Marseille III, 1995.

[11] Mignon V.,Marchés financiers et Modélisation des Rentabilités Boursiéres, Economica 98.

[12] Neuberg L.; *Marché Financier Artificiel*, PhD-thesis, University of Namur, 1999, 161 p.

[13] Neuberg L.; Bertels K.;An Artificial Stock Market, Proc. of the IASTED International Conference, 2002,ACTA Press, Anaheim, pp.308-313.

[14] Sanglier M., Romain M., Flament F.;*A Behavioral Approach of the Dynamics of Financial Markets*, Decision Support Systems 12, pp. 405-413, 1994.