Financial Market in the Laboratory[†]

Andrea Morone*

Abstract

This paper investigates experimentally a market inspired by two separate strands of economic literature. The first strand is that of herd behaviour in non-market situations and the second that of the aggregation of private information in markets. The first suggests that socially undesirable herd behaviour may result when information is private; the second suggests that in a market context the private information may be aggregated efficiently through the price mechanism. The latter literature therefore suggests that socially undesirable behaviour may be eliminated through the market mechanism. We tested this hypothesis experimentally, in a very simple extension of a herd model into a market context, and found that many of the stylised facts of financial markets (i.e. fat tails of the distribution of returns and autoregressive dependence in volatility) can be reproduced in our experimental market.

Keywords: herd behaviour, fat tail volatility clustering.
JEL:

[†] I would like to thank the European Commission in the form of the TMR grant ENDEAR (ERB FMR XCT 98 0238) for funding the experiments reported in this paper. The experiments were run at the laboratory of EXEC at the University of York. Pilot experiments were also run at the laboratory of ESSE at the University of Bari. For their support with discussions and hints I also thank John Hey, Thomas Lux and Eleni Samanidou.
* Dr. Andrea Morone, Dipartimento di Scienze Economiche, Universita' degli Studi di Bari, via Camillo Rosalba n. 53, 70124 Bari, Italy.

I. Introduction

It is now well established (Mandelbrot [32] and Fama [19]) that changes in asset prices (returns) do not have a Normal distribution. In fact, if we assume that returns are normally distributed we have to accept an unrealistically high number of 'outliers'. This leads us to reject the normality assumption.

There is abundant literature studying the empirical features of financial markets. Pagan [33] provided an authoritative survey of these stylized facts and of the econometric techniques to treat them.

There are also several empirical works that analyze the empirical regularity of those markets, i.e. de Vries [9], Guillaume *et al.* [17] and Lux and Ausloos [27].

In the last couple of years, the study of behavioral models of dynamic interaction in financial markets – Beja and Goldman [3], Day and Huang [8], Lux and Marchesi [28, 29], Chen et al. [5, 7], Iori [19], Farmer [11], LeBaron [22], Gaunersdorfer and Hommes [13, 14], Arifovic et al. [1] and Georges [15] – has brought about a better understanding of some of the key stylized features of financial data, namely the *fat tails* of the distribution of returns and the *autoregressive dependence* in volatility. Some possible general explanations seem to emerge from this literature: first, volatility clustering and fat tails may emerge from an indeterminacy in the equilibrium of the dynamics. In particular, with different strategies performing equally well in some kind of steady state, stochastic disturbances lead to continuously changing strategy configurations which at some point generate a burst of activity. This type of dynamics can be found already in Youssefmir and Huberman [18] in the context of a resource exploitation model and can be identified in both the papers by Lux and Marchesi[28, 29].

The purpose of this paper is to verify if the above stylized facts can be reproduced in the laboratory.

In the next section we shall describe our data and in section 3 we shall present some elementary statistics. In sections 4, 5, 6 and 7 we shall compare the features our experimental data with the stylized facts. In particular, we shall check for Unit Root, Fat Tails, Volatility Clusters and Autocorrelation, respectively. Finally in section 8 we draw some conclusions.

II. Data description

The data set analysed in this paper is 'artificially generated' in the laboratory. The experiment is based on at least two important strands of literature. The first of these strands is that of *herd behaviour* in a non-market context. The key references are Banerjee (1992) and Bickchandani, Hirshleifer and Welch (1992), both of which showed that herd behaviour may result from private information not publicly shared. More specifically, both of these papers showed that individuals, acting sequentially on the basis of private information and public knowledge about the behaviour of others, may end up choosing the socially undesirable option. The second strand of literature motivating this paper is that of *information aggregation in market* contexts. A very early reference is the classic paper by Grossman and Stiglitz (1966) which showed that uninformed traders in a market context can become informed through the price in such a way that private information is aggregated correctly and efficiently. A summary of the progress of this strand of literature can be found in the paper by Plott (2000). A third, though less directly relevant, strand is that of the experimental economics literature, which suggests that the market may act as a sort of disciplining device on 'irrational' behaviour in individual contexts. This third strand reconciles, in a sense, the first two strands.

The experiment was programmed using the z-Tree software of Urs Fischbacher [12]. It was piloted at the laboratory of ESSE at the University of Bari, and the main experiment, reported in this paper, was run at the laboratory of EXEC at the University of York.

We have n agents in our market. Each is endowed with a quantity of experimental money and m units of some asset. This asset pays a dividend at the end of the trading period. This dividend is uncertain. There are two possible 'states of the world'- each with equal probability – either the dividend is some positive number *d*, or it is zero. At the beginning of each trading period the true state is determined by the experimenter – but not revealed to the agents. They can, however, buy signals – which are partially but not totally informative as to the true state of the world. These signals take either the value 0 or the value 1. Agents are informed of all the relevant parameters – the positive dividend *d,* the cost of buying a signal *c*, and the two probabilities *p* and *q*[1].

We should now describe the trading process. We use a single-unit double-auction mechanism in which agents are free at any time to make bids and asks and to accept existing asks or bids. We adopted this trading procedure as it is well-known from countless experiments (in simpler contexts) that this mechanism reaches the competitive equilibrium quickly and efficiently.

We run this experiment for four different parameter sets, each one corresponding to a different quadrant in the following diagram:

Treatment 1 low cost/low quality	**Treatment 2** high cost/low quality
Treatment 3 low cost/high quality	**Treatment 4** high cost/high quality

The payment mechanism is the obvious one: agents start with some experimental money and with *m* units of the asset. During the trading process they can increase or decrease the number of units of the asset that they own and, depending upon the prices at which they trade, their stock of experimental money will increase or decrease during each period. At the end of each market period the true dividend for that period is announced and the appropriate dividend is paid in experimental money to the asset owners. Accordingly agents will end up with a stock of experimental money at the end of each trading period – which may be more or less than that with which they started that period. An agent's trading profit for any trading period is the difference between the final stock of

[1] More precisely, *p* is the probability of getting a signal of 1 if the true state of the world is that the dividend is 10; *q* is

experimental money and the initial stock. For the experiment as a whole the total profit to an agent is simply the sum of the profits over all trading periods of the experiment. There was a fixed rate of exchange between experimental money and real money.

Note that agents can make losses. To avoid some of the problems associated with subjects making real losses in experiments, we endowed all agents with a participation fee, which could be used (if the subject agreed) to offset losses. Once this participation fee was exhausted, any further losses had to be covered by the subjects themselves – some subjects chose this option, others chose to leave the experiment once they had exhausted their participation fee[2].

We analysed also daily changes of the German share price index DAX. This will be our benchmark to compare the experimental data to real data.

III. Some elementary statistics

We have already noticed that empirical data in financial markets are not normally distributed. Table 1 reports some elementary statistics for the returns of the DAX and our four experimental treatments.

It is clear that all five distributions exhibit excessive kurtosis. This implies that the experimental financial market, like real markets, exhibits more probability mass in the tails and in the center compared to a Normal distribution. Additionally the Bera –Jarque test for normality leads to a rejection of its null hypothesis.

	DAX	Treatment 1	Treatment 2	Treatment 3	Treatment 4
Mean	0	0	-0,001	-0,005	0
S.D.	0,005	0,252	0,213	0,418	0,261
Skewness	-0.3216	-0.0826	0.0437	0.0602	-0.0185
Kurtosis	10.4481	3.3944	8.8983	6.4489	14.0119
Bera-Jarque test	44908.312 (0.000)	626.579 (0.000)	5094.409 (0.000)	1407.571 (0.000)	11223.906 (0.000)

Table 1

the probability of getting a signal of 1 if the true state of the world is that the dividend is zero.

[2] A more detailed presentation of the experiment can be found in Hey and Morone [18].

Now we will take a closer look at the statistical characteristics of our experimental data sets. More precisely, we will investigate whether and how the experimental market compares with the stylised facts observed in real financial markets: Unit Root, Fat Tail, Cluster Volatility and Autocorrelation.

IV. Unit root property

Levels (or logarithms of levels) of prices have a unit root, which implies that returns (or the difference of logarithms of levels) are stationary. This implies that the price follows, in the simplest case, an autoregressive process of the first order:

$$p(t) = c + \lambda p(t-1) + \varepsilon(t), \tag{1}$$

where $p(t)$ is the price at time t, $p(t-1)$ is the price at time 1-t and $\varepsilon(t)$ is the unanticipated element. Equation (1) is clearly a "no arbitrage condition" (LeRoy [23]).

The usual finding in financial markets is non-stationarity for the price time series and stationarity for its first difference, i.e. the returns.

Table 2 reports the outcome of the Augmented Dickey-Fuller test. For each case the time series have been divided into 10 sub-samples and the test has been run on each sub-sample.

Time Series	Range of λ Min	Max	No. of rejections for one-sides test at 95% level	No. of rejections for two-sides test at 95% level
DAX	0,4148	1,0038	2 out of 10	2 out of 10
Treatment 1	-0,2788	0,2674	10 out of 10	10 out of 10
Treatment 2	-0,0891	0,3301	10 out of 10	10 out of 10
Treatment 3	-0,1763	0,9719	3 out of 10	3 out of 10
Treatment 4	-0,0724	0,9635	8 out of 10	7 out of 10

Table 2

We can, now, try to categorise the four treatments. It seems that we can divide them into two groups. Treatment 1 and Treatment 2 completely fail to exhibit nonstationarity in the price series. On the other hand Treatment 4 has a unit root in two out of ten periods, and Treatment 3 has a unit root in seven out of ten periods. Note that the DAX exhibits a unit root in eight out of ten subsamples. This is an interesting result, since in Treatment 1 and 2 the quality of the signals is very poor (a signal is informative with probability 0.6 and is misleading with probability 0.4) and thus the aggregated information is not very informative. For this reason price fluctuates around the 'un-informed' expected price. On the other hand in Treatments 3 and 4 the quality of the information is higher (a signal is informative with probability 0.8 and is misleading with probability 0.2) thus the price does not fluctuate around the 'un-informed' expected price but converges (in average) to the correct price.

V. Fat tail phenomenon

In section 3 we reported that our four experimental financial markets (as well as the DAX) exhibit excessive kurtosis and we noticed that return distributions exhibit more probability mass in the center and in the tail of the distribution. In the following figures it is clear what we meant by fat tails. In fact it is possible to see how the distributions of the returns are leptokurtotic.

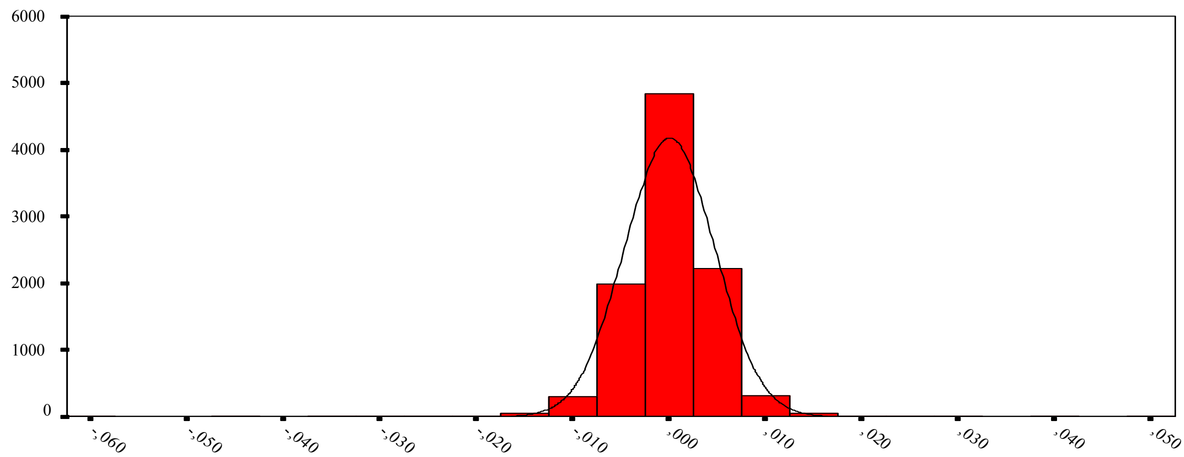

Returns' distribution in the DAX

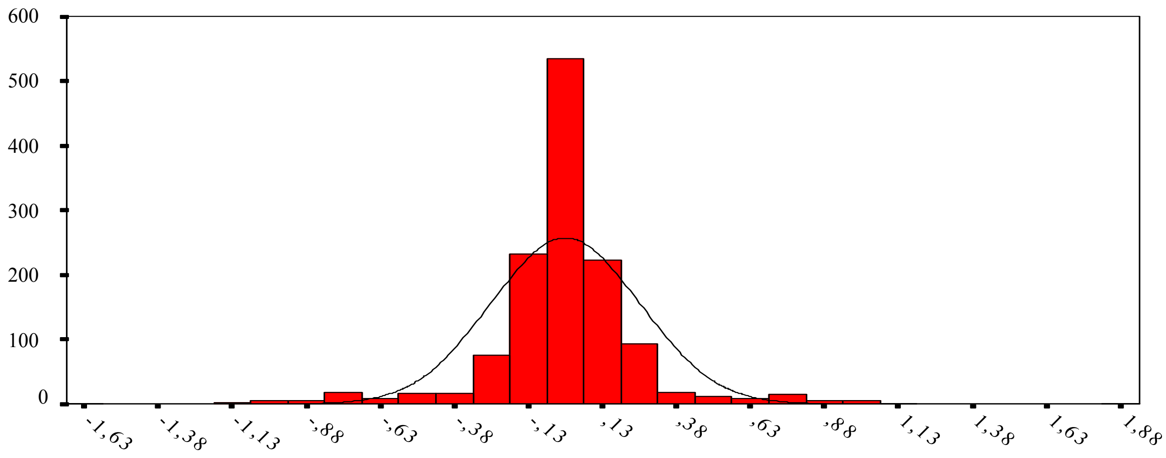

Returns' distribution in treatment 1

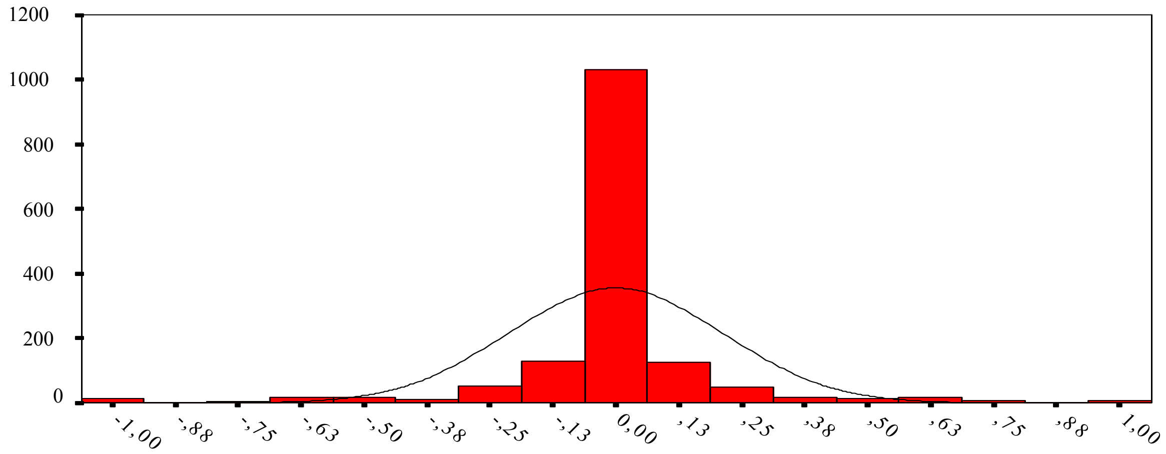

Returns' distribution in treatment 2

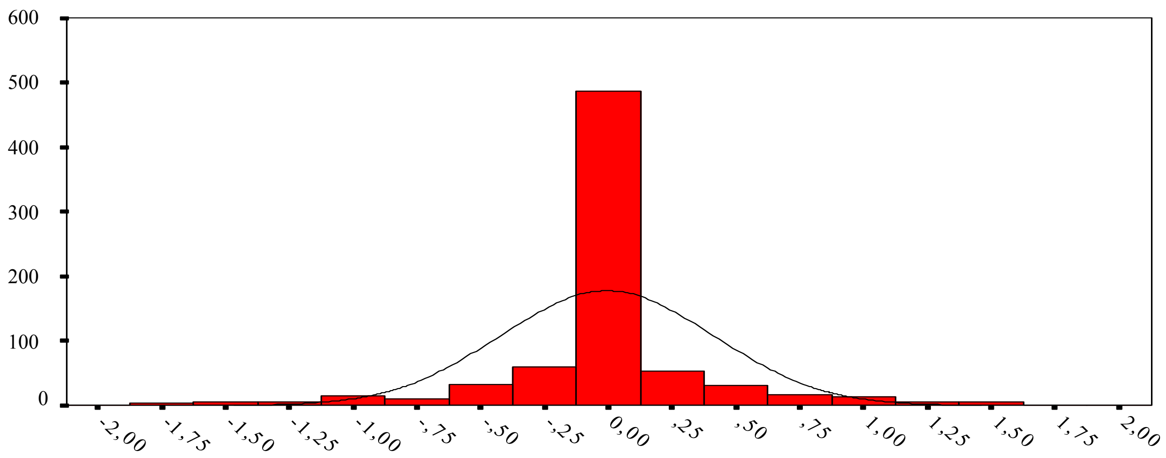

Returns' distribution ot treatment 3

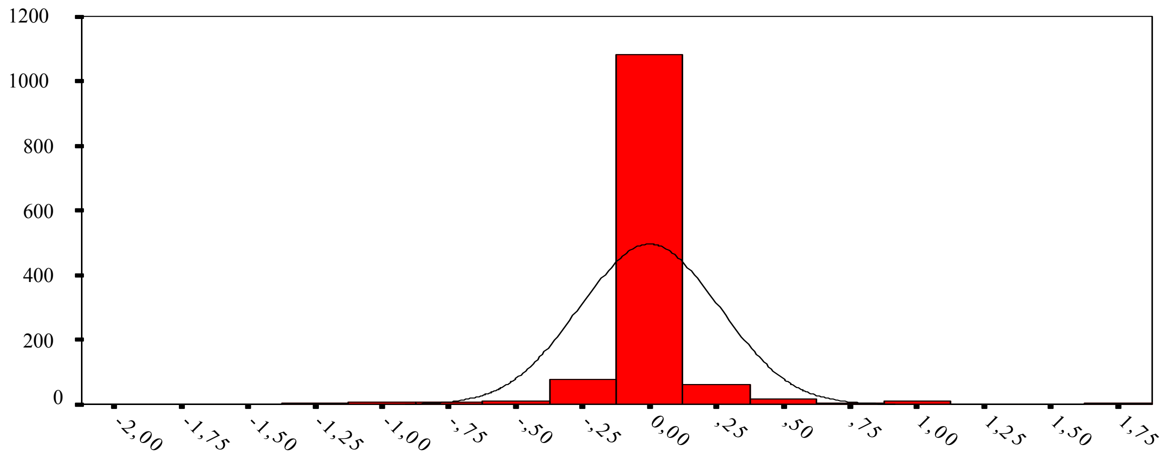

Returns' distribution of treatment 4

We have to note that the kurtosis is, in a certain sense, a poor measure of deviation from normality. For this reason, we need to refer to a sharper characterization of the empirical distribution. It is now well known that the distribution of returns belongs to the class of the 'fat tail' distributions. These distributions exhibit a hyperbolic decline of probability mass[3].

	Treatment 1	Treatment 2	Treatment 3	Treatment 4	DAX[4]
Hill 10%	2,527	1,898	2,320	1,013	2,945
Hill 5%	6,035	2,905	3,422	1,723	2,974
Hill 2.5%	7,684	6,177	4,930	5,937	3,105

Table 3

The visual impression of fat tails is also confirmed by the above table, which reports the Hill estimators for the five distributions (i.e. the four experimental treatments and the DAX). The tail index gives us information about the "fatness" of the tails of the distributions. In fact, given a tail index, the biggest integer number smaller than the tail index is the number of finite moments of the distribution.

[3] For a more detailed analysis see Lux[27]
[4] The New York Stock Exchange Composite Index, the US Dollar – DM exchange rate and the price of gold exhibits similar figures, cf. Lux [26].

Empirical studies show that the Hill estimators usually lie in the range of [2.5, 5]. Examples include Koedijk, Schafgans and de Vries [21], Jansen and de Vries [20], Loretan and Phillips [25], Longin [24], Lux [31] and Lux and Ausloos [27].

From table 3 we can argue that all our four treatments look like a real financial market, even though for a tail size of 2.5% the tails seems not to be very fat. It is interesting to note that in Treatment 2 and Treatment 4 the bursts are so strong that even tail indices below 2 were found. Remembering that in Treatment 2 and 4 the cost of information was high compared to Treatment 1 and 3, it seems that markets in which the information is more expensive have larger price changes.

We can try to rank our four treatments according to their tail index. Treatment 3 and Treatment 2 seem to be very good approximations of a real financial market, whereas Treatment 4 exhibits too fat tails at both 10% and 5% tail size and too thin tails at 2.5%. Treatment 1 exhibits realistic tails at 10% level but too thin tails at both 5% and 2.5% level.

These results seem to be quite encouraging and the "rejection" of the fat tail hypothesis at 2.5% level could be related to the sample size (1303, 1545, 813, 1373 in treatment 1, 2, 3 and 4 respectively). It could be of some interest to note that Lux and Sornette [26] demonstrated that the prevalence of a rational bubble component would lead to an Hill tail index estimator smaller than 1, which would imply non existence of the mean and variance of the data.

VI. Volatility clustering

Plotting the time series of returns it is immediately evident that the results of our experiment are dramatically different from previous experiments on financial markets. In fact we do not have the usual fast convergence to equilibrium, but we see periods of tranquillity interrupted by periods of turbulence. The time series plotted in the figures below are astonishingly similar to the empirical ones. Periods of quiescence and turbulence tend to cluster together.

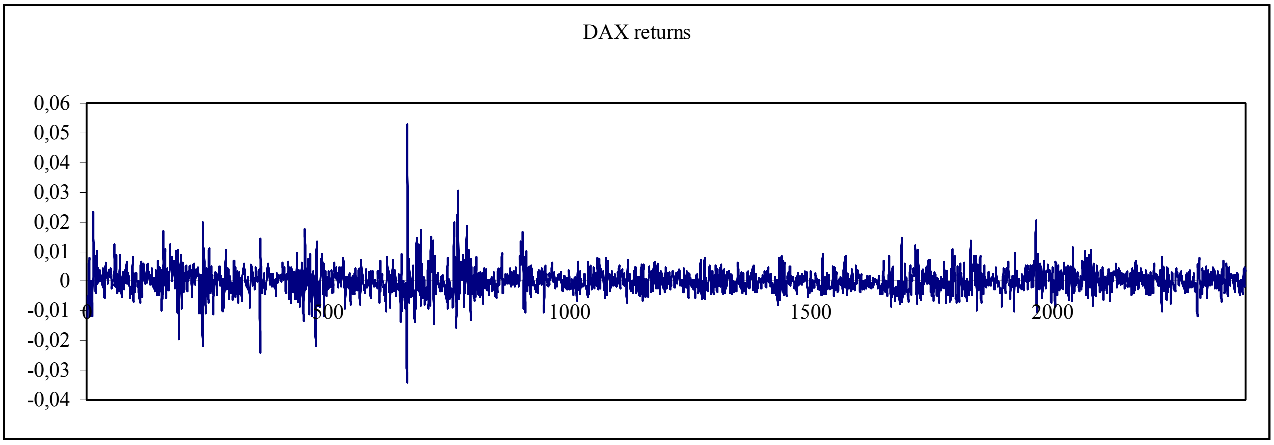

Figure 6

This fact was already pointed out by Mandeldrot [18], but it was by and large neglected until recently. The volatility cluster regularity (which is particularly clear in figures 6, 7, 8, 9 and 10) suggests that there is autocorrelation in the scale of the process i.e. in the second moments.

Also the figure below exhibits clustered volatility in the returns. Treatment 1 and Treatment 2 seem to capture this phenomenon pretty well.

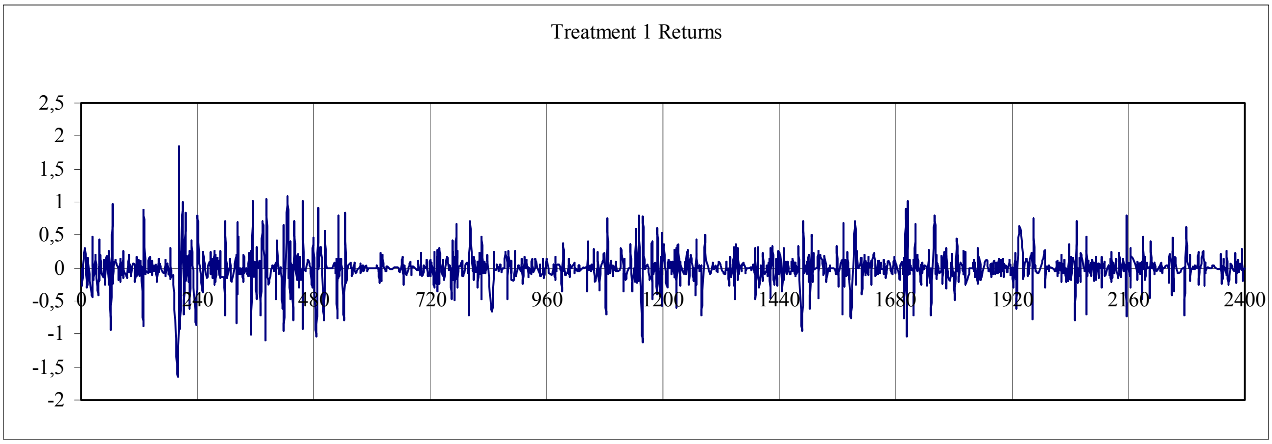

Figure 7

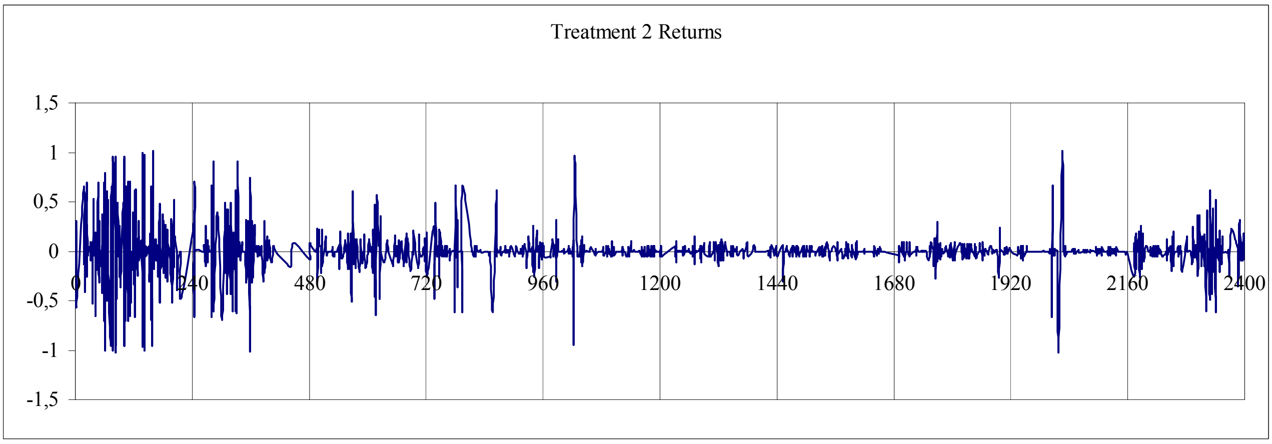

Figure 8

On the other hand Treatment 3 and Treatment 4, even though they exhibit clustered volatility, seem to be different from a real financial market. A simple explanation could be that, because the quality of information in the market is higher compared to Treatment 1 and Treatment 2 the *invisible hand is less trembling.*

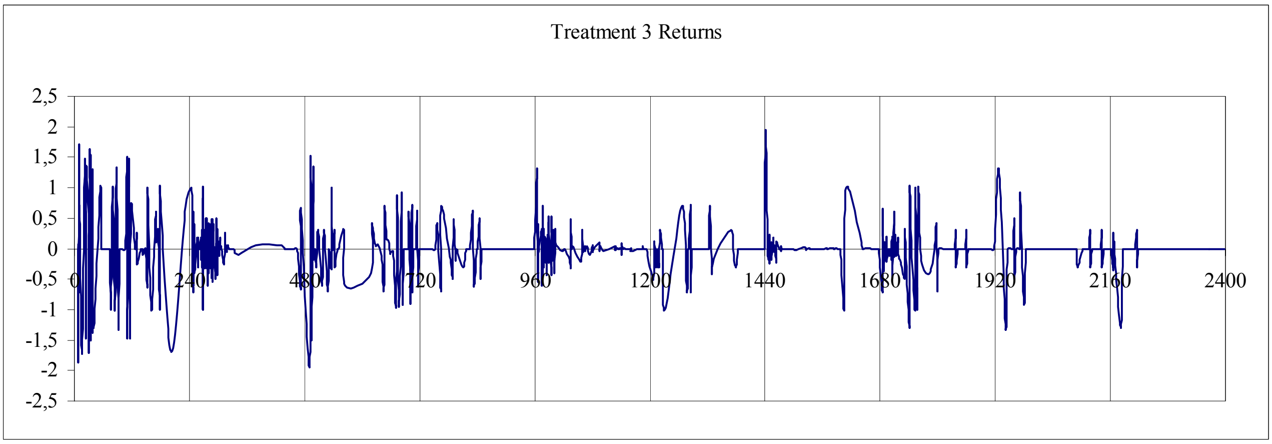

Figure 9

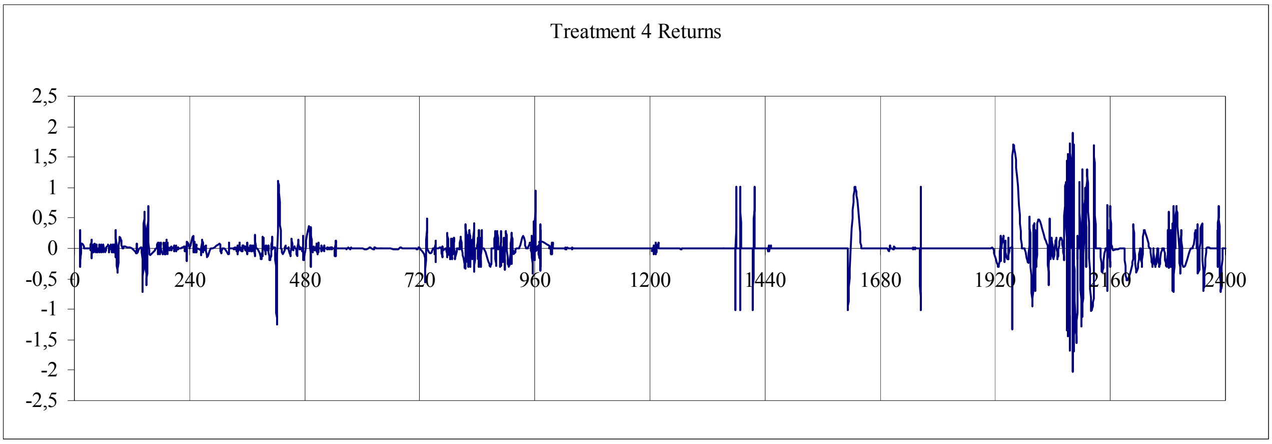

Figure 10

There is abundant literature that studies this phenomenon. Gaunersdorfer and Hommes [15] proposed that clustered volatility emerges from stochastic dynamics with multiple attractors. Some amount of noise added to a deterministic dynamic with two or more attractive states can lead to recurrent switches between one state and another. This kind of explanation seems to be confirmed by our experimental data. In fact we have more volatility clustering in the first two treatments, where the two "attractors" (i.e. 0 and 10) have almost the same power, but in the last two treatments in which the information is more precise both volatility and volatility clustering seem to be weaker.

VII. Absence of autocorrelation

Autocorrelation is often insignificant in raw returns, but highly significant in the volatility

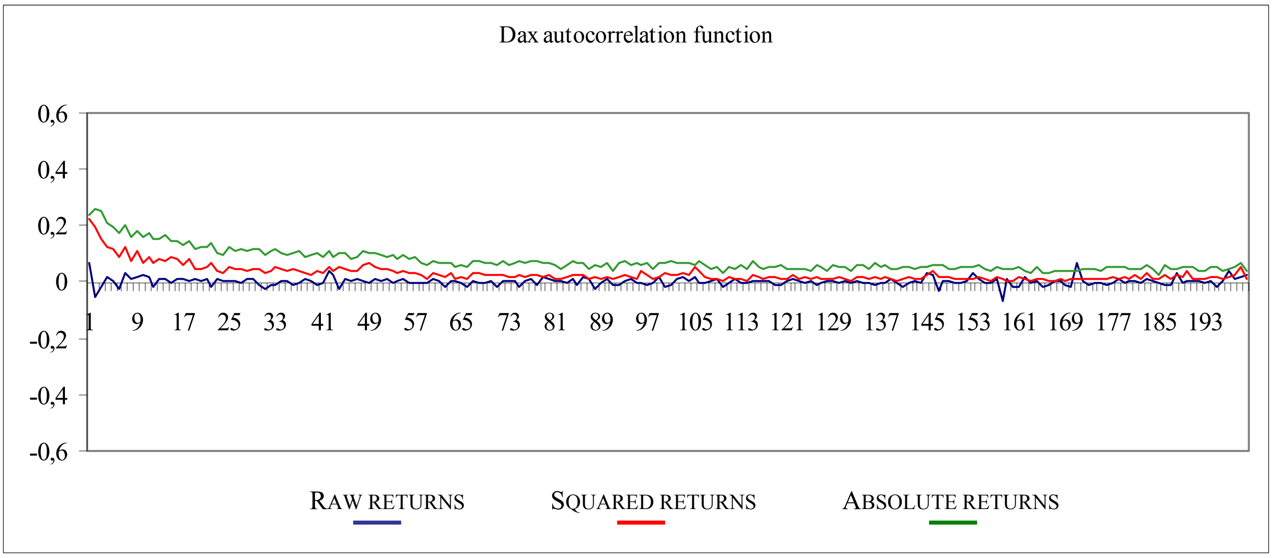

Figure 11

measures, i.e. squared returns and absolute returns. The absence of autocorrelation is a very well-known fact in financial data.

In figure 11 we plotted the autocorrelation functions of the DAX raw returns, squared returns and absolute returns in the period 1959-1999 (daily observations). For each time series we computed the autocorrelation functions for 200 lags. In the first plot it is evident that the autocorrelation of raw returns is not significantly different from zero. On the other hand, considering the squared returns, we can observe a very long autocorrelation, and it is even larger in the case of absolute returns. For squared and absolute values the temporal independence is strongly rejected. These results are common to all financial markets.

In figure 12 we report the autocorrelation of returns for Treatment 1 of our experiment. It is evident that returns, squared returns and absolute returns exhibit temporal independence contrary to financial market empirical evidence. We note that another unconventional feature characterizes figure 12: the presence of negative correlation among absolute and squared returns. This suggests that in Treatment 1 randomness overwrites the financial market features.

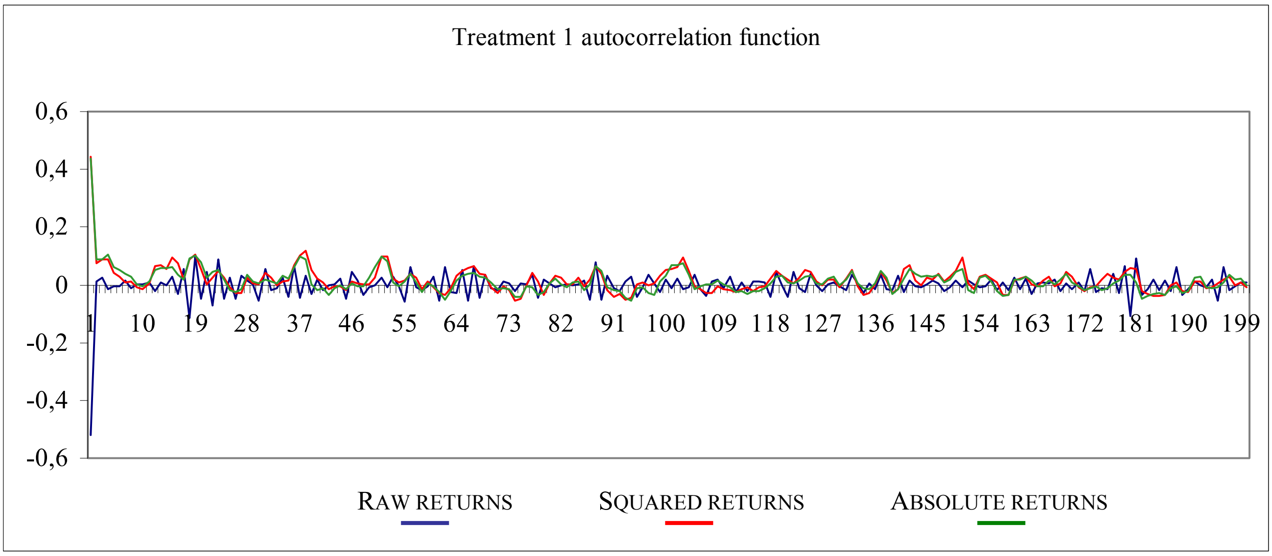

Figure 12

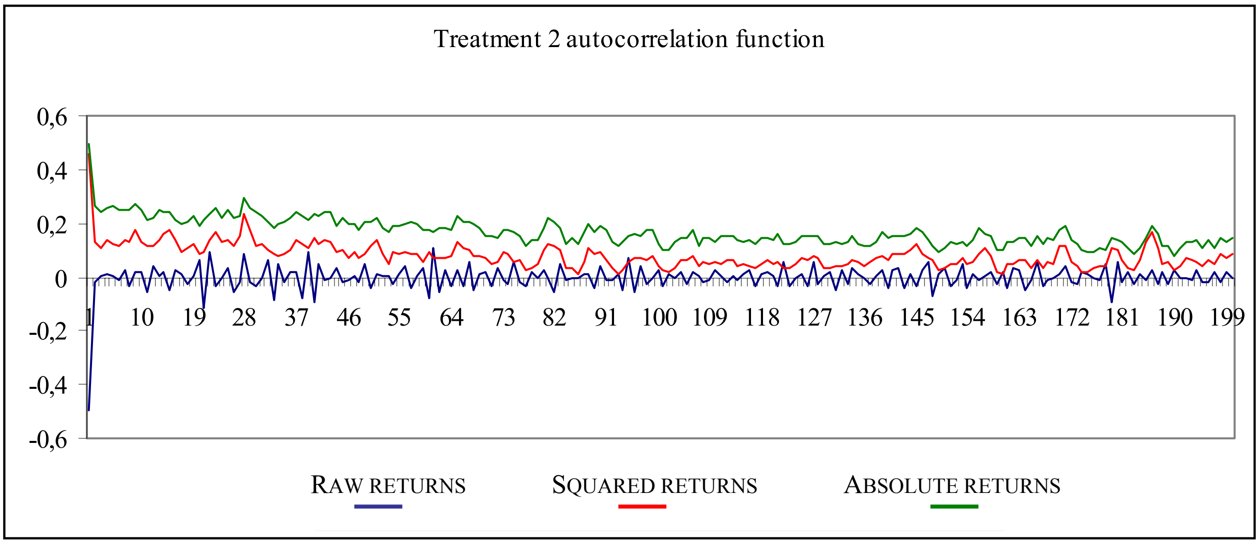

Figure 13

Treatment 2 exhibits autocorrelation functions with features typically characterizing financial markets. In fact, the raw data are completely uncorrelated, the squared returns have very long correlation and the absolute returns exhibit longer and higher autocorrelation.

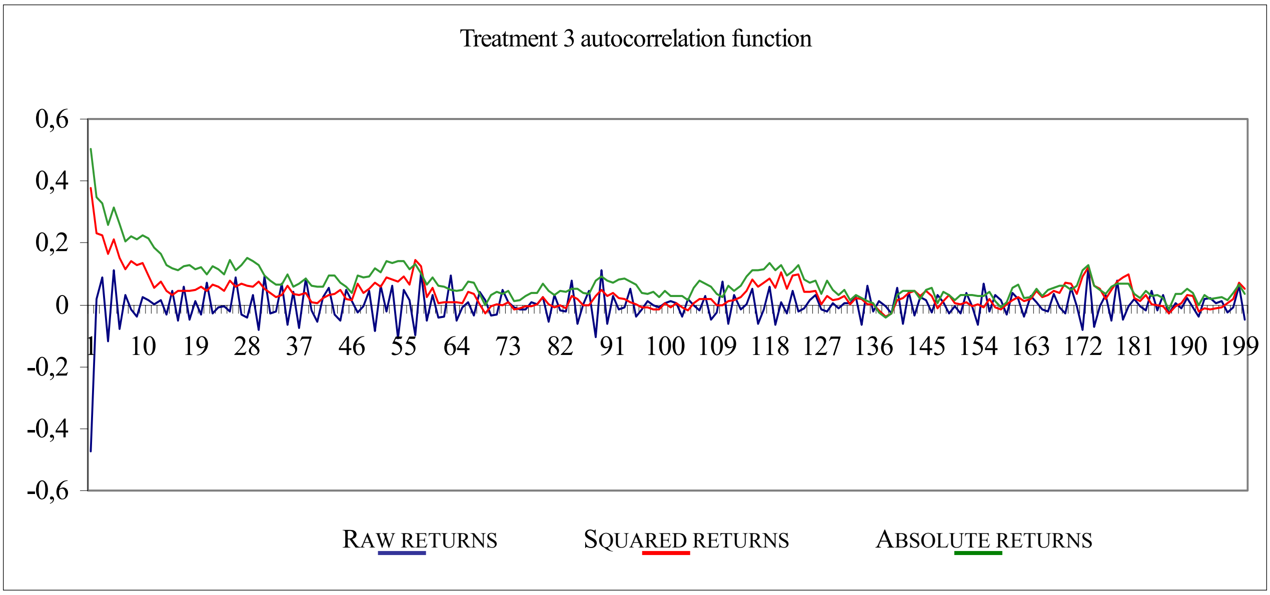

Figure 14

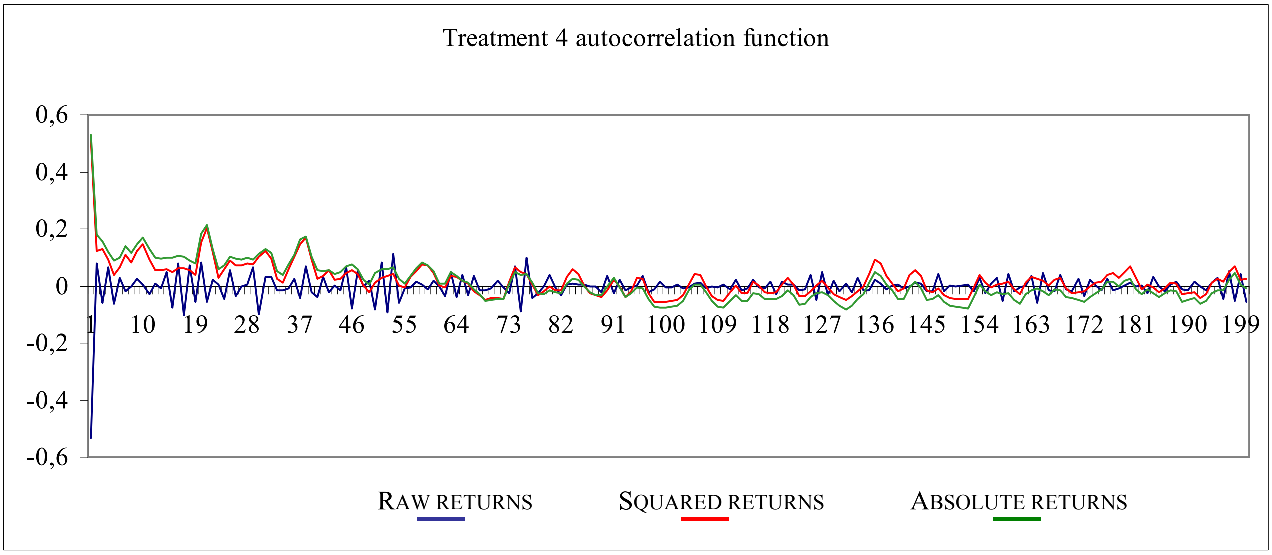

Figure 15

Also Treatment 3 and 4 (figure 14 and figure 15 respectively) have this feature but it is weaker compared to Treatment 2.

To investigate better the autocorrelation structure, we applied the Box-Ljung test (Table 4) to the autocorrelations up to lags 8, 12, 16 for the raw data as well as the squared and the absolute returns.

	Treatment 1			Treatment 2			Treatment 3			Treatment 4			DAX		
	8 lags	12 lags	16 lags	8 lags	12 lags	16 lags	8 lags	12 lags	16 lags	8 lags	12 lags	16 lags	8 lags	12 lags	16 lags
raw returns	353.56	287.26	295.73	386.51	504.60	1073.42	216.72	305.89	658.69	418.19	444.48	554.08	102.46	1641.57	3590.50
significance	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Squared returns	354.32	292.99	299.50	395.44	623.38	1431.08	218.70	344.46	803.71	420.43	510.93	659.40	117.41	1926.47	4682.58
Significance	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Absolute returns	356.73	322.26	315.29	401.30	769.32	1785.33	218.70	344.46	803.71	440.36	528.58	714.63	121.41	2204.23	5592.22
Significance	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Table 4: Box-Ljung test

It is clear that we have to reject the null hypothesis that the absolute values of the autocorrelation coefficients are not significantly different from zero at both the 0.05 and 0.01 significance level (for raw, and for squared and absolute returns as well) .

VIII. Conclusion

In this paper we investigate the characteristics of an experimantal financial market, and we compare it to a real one.

Our market is inspired by two separate strands of economic literature – the first that on herd behaviour in non-market situations and the second that on the aggregation of private information in markets. The first of these two strands of literature suggests that socially undesirable herd behaviour may result when information is private; the second suggests that in a market context the private information may be aggregated efficiently through the price mechanism. The latter strand therefore suggests that socially undesirable behaviour may be eliminated through the market.

The first result we obtain concerne the relationship between information quality and market efficency:

- If the quality of the information is low, market seams to fail to aggregate information and the price fluctuate around the un-informated price.
- If the quality of the information is high, the invisible hand seams to work 'properly'.

The second result we obtain relate the cost of information with the leptokurtosis of the returns distribution:

- The more expensive information is the more leptokurtic returns distribution are.

We finally have evidence that dissemination and aggregation of information through the trading mechanism is possible, but it is no longer defensible to argue that rational expectations can be achieved instantaneously, or precisely.

References

Ariofovic, J., and R. Gencay, Statistical properties of genetic learning in a model of exchange rate, Journal of Economic Dynamics and Control (2000) 24:981-1006.

Banerjee, A. V., A simple model of herd behaviour, (1992) Quorterly Journal of Economics, 107, pp. 797-817.

Beja, A., and M. B. Goldman, On the dynamic behaviour of prices in disequilibrium, Journal of Finance (1980) 35:235-248.

Bickchandani, S., D. Hirshleifer and I. Welch, A theory of fads, fashion, customs and cultural change as informational cascades, (1992) Journal of Political Economy, 100, pp 992-1026.

Chen, S. H., and C. H. Yen, On the Emergent Properties of Artificial Stock Markets: Some Initial Evidence, Journal of Economic Behaviour and Organisation (in press).

Chen, S. H., T. Lux and M. Marchesi, Testing for Non-Linear Structure in an Artificial Financial Market,(2002) mimeo.

Chen, S. H., T. Lux and M. Marchesi, Testing for Nonlinearity Structure in an Artificial Financial Market, Journal of Economic Behaviour and Organisation (in press).

Day, R., and W. Huang, Bulls, bears and market sheep, Journal of Economic Behaviour and Organisation (1990) 14:299-329.

de Vries, C. G., Stylized facts of nominal exchange rate returns, pp. 348-389, in ed. F. van der Ploeg, The Handbook of International Macroeconomics. Blackwell: Oxford (1994).

Fama, E., Mandelbrot and the Stable Paretian Hypothesis, Journal of Business (1963) 35:420-429.

Farmer, D., and S. Joshi, The price dynamics of common trading strategies, Journal of Economic Behaviour and Organisation (in press).

Fischbacher, U., z-Tree – Zurich Toolbox for Readymade Economic Experiment – Experimenter's Manual,

Gaunersdorfer, A., and C. Hommes, A non linear structured model for volatility clustering, Mimeo University of Vienna.

Gaunersdorfer, A., C. Hommes and F. Wagner, Bifurcation routes to volatility clustering, Mimeo University of Vienna.

Georges, C., Learning dynamics in an artificial currency market, Mimeo Hamilton's Collage.

Grossman and Stiglitz, Information and Competitive Price Systems, (1966) American Economic Review, 246-253.

Guillaume, D. M., M. M. Dacorogna, R. R. Dave', U. A. Muller, R. B. Olsen and O. V. Pictet, From the bird's of eye to the microscope: A survey of new stylized facts in the intra-daily foreign exchange markets, Finance and Stochastics (1997) 1:95-129

Hey, J. D., and A. Morone, Do Markets Drive Out Lemmings – or Vice Versa? (2002) mimeo.

Iori, G., A micro-simulation traders' activity in the stock market: the rule of heterogeneity, agents' interactions and trade friction, Journal of Economic Behaviour and Organisation (in press).

Jansen, D. W., and C. G. de Vries, On the Frequency of Large Stock Returns: Putting Booms and Busts into Perspective, Review of Economics & Statistics (1991) 73:18-24.

Koedijk, K. G., M. M. A. Scafgans and C. G. de Vries, The Tail Index of Exchange Rare Returns, Journal of International Economics (1990) 29:93-108.

LeBaron, B., Agent based computational finance: suggested readings and early research, Journal of Economic Dynamics and Control (2000) 24:679-702.

LeRoy, S. F., Efficient Capital Markets and Martingales, Journal of Economic Literature (1989) 27:1583-1621.

Longin, F. M., The Asymptotic Distribution of Extreme Stock Market Returns, Journal of Business (1996) 69:383-408.

Loretan, M., and P. C. B. Phillips, Testing the Covariance Stationarity of Heavy-tailed Time Series, Journal of Empirical Finance (1994) 1:211-248.

Lux, T., and D. Sornette, On Rational Bubbles and Fat Tail, (2002) mimeo.

Lux, T., and M. Ausloos, Market fluctuations I: Scaling, Multi-Scaling and Their Possible Origins, mimeo.

Lux, T., and M. Marchesi, Scaling and Criticality in a Stochastic Multi-Agent Model of Financial Market, Nature (1999) 397:498-500.

Lux, T., and M. Marchesi, Volatility clustering in financial market: a micro-simulation of interacting agents, International Journal of Theoretical and Applied Finance (2000) 3:675-702.

Lux, T., The Limiting Extremal Behaviour of Speculative Returns: An Analysis of Intra-Daily Data from the Frankfurt Exchange, (2002) mimeo.

Lux, T., The Stable Paretian Hypothesis and the Frequency of Large Returns: An Examination of Major German Stocks, Applied Financial Economics (1996) 6:463-477.

Mandelbrot, B., The Variation of Certain Speculative Prices, Journal of Business (1963) 35:394-419.

Pagan, A., The econometrics of financial markets, Journal of Empirical Finance (1996) 3:15-102.

Plott, C. R., Markets as Information Gathering Tools, (2002) Southern Economic Journal, 67, 1-15.