Working Paper Series

Adaptive Information Systems and Modelling in Economics and Management Science

Heterogeneous Trade Intervals in an Agent Based Financial Market

Alexander Pfister

Working Paper No. 99
October 2003

Working Paper Series

October 2003

SFB
'Adaptive Information Systems and Modelling in Economics and Management Science'

Vienna University of Economics
and Business Administration
Augasse 2–6, 1090 Wien, Austria

in cooperation with
University of Vienna
Vienna University of Technology

http://www.wu-wien.ac.at/am

This piece of research was supported by the Austrian Science Foundation (FWF) under grant SFB#010 ('Adaptive Information Systems and Modelling in Economics and Management Science').

Heterogeneous Trade Intervals in an Agent Based Financial Market

Alexander Pfister
Department of Economics
Vienna University of Economics and Business Administration

October 7, 2003

Abstract

This paper studies the dynamics of an asset pricing model based on simple deterministic agents. Traders are heterogeneous with respect to their time horizon, prediction function and trade interval. Concerning the trade interval we distinguish between intraday traders and end-of-day traders. Intraday traders update their portfolio every period, whereas end-of-day traders adjust their positions only at the closing price of each trading day. The parameter values of the model were partially determined by an adapted Markov chain Monte Carlo sampling method. We analyse the properties of the time series and find that they exhibit low autocorrelation of the returns, volatility clustering and fat tails. Particularly heterogeneous trade intervals seem to be an important factor for generating time series showing "stylized facts".

Keywords: agent-based models, heterogeneous agents, endogenous stock price fluctuations, artificial markets

1 Introduction

One important paradigm of modern finance is the efficient market hypothesis, which was developed by Fama, 1965. It is assumed that the actual market price contains all available information or in Malkiel's words (Malkiel, 1989): 'Thus, investors cannot devise an investment strategy to yield abnormal profits on the basis of an analysis of past price patterns'. Especially past prices cannot help in forecasting future price changes (see also Fama, 1970 and Fama, 1991). Much effort has been invested in proving or disproving the efficient market hypothesis. Current opinion of the financial community is that there is much evidence that financial markets are not efficient (see Malkiel, 1996; Lo and MacKinlay, 1999; Savit, 1988; Savit, 1989; Hinrich and Patterson, 1985; Scheinkman and LeBaron, 1989). This is also affirmed by the empirical trading behaviour of many investors. Chartists or technical traders believe that prices may be predicted by extrapolation of trends, technical trading rules or other patterns generated by past prices.

Another paradigm is the rational expectation equilibrium (REE), introduced by Muth, 1961. The rational expectations model makes two assumptions (see also Sargent, 1986). First agents are rational and are able to optimize an objective function. Second the same information is available to all agents. Therefore agents are expected to be fully informed and to know all equations of the economic model. Perfectly rational agents maximize their utility function and are able to solve complicated optimization problems. This seems to be highly demanding, and therefore bounded rationality models have been proposed. Note that models with agents being bounded rational and having access to different information sets may still converge to the rational expectation equilibrium, e.g. Sargent, 1993. Therefore these two markets may be indistinguishable on a macro-level.

Examples of bounded rational models are agent-based simulations of financial markets (e.g. Chen et al., 2001; Farmer, 2000; Cont and Bouchaud, 2000), especially the Santa Fe artificial stock market (Arthur, et al., 1997; LeBaron et al., 1999; LeBaron, 2001a and the review LeBaron, 2000) with various modifications (e.g. Ehrentreich, 2002 correcting a faulty mutation operator or

Tay and Linn, 2001 introducing expectation formation based on fuzzy inductive reasoning) or evolutionary models (e.g. Bullard and Duffy, 1999; Brock and Hommes, 1997; Brock and Hommes, 1998; Beltratti and Margarita, 1992; Chen et al., 2001).

Noe et al. (2003) examine corporate security choice by simulating an agent based economy. Adaptive agents learn about the structure of security returns and prices. A firm needs external financing in order to undertake a project paying out random future payoff. The firm can choose from different securities, e.g. debt, equity, convertible debt. These securities are issued in a perfect market. The cash flow of the project is then split between the firm and outside investors according to the terms specified by the security issued. Security issuance and pricing strategies are learned using a genetic algorithm. Their results show that strong preferences over capital structure choices emerge even in perfect and competitive markets. These preferences are consistent with observed empirical regularities.

In a stationary world all past observations are relevant and should be included in a prediction function. But if the data generating mechanism is changing over time, it might be better to concentrate only on recent observations. Several papers have reported evidence of breaks in economic time series (see Bai and Perron, 1998; Garcia and Perron, 1996 or regarding stock price time series Timmermann, 1998). Therefore the length of the history of past prices and dividends is of importance.

In LeBaron (2001) investors with different time horizons are modeled in an evolutionary setting. LeBaron finds that agents with a short horizon increase the volatility of the stock price significantly, which is far beyond that driven by fundamental variability. Furthermore patterns in volatility and trading volume are similar to empirical markets.

Thurner et al. (2002) present a deterministic model of traders, who are all fundamentalists. Investors are heterogeneous in the sense that they work with a different time frequency of the stock price. The artificial time series exhibits fat tails of returns distributions, volatility clustering and power laws.

This paper presents a standard capital market model. Investors are heterogeneous with respect to their prediction function, their time horizon and their trade interval. There are two types of prediction functions which lead to fundamentalists and chartists. Fundamentalists base their prediction on past dividends whereas chartists use past prices to forecast the future price. In our model the time horizon is modeled as length of history of past prices and dividends used for the prediction. In other words time horizon means how far back agents look into the past to predict the future price. Agents trade in two different time intervals which lead to two types of investors: Intraday traders and end-of-day traders. Whereas intraday traders update their position every time period, the latter only trade at the closing price of each day.

The validation technique of the model is based on an adapted Markov chain Monte Carlo sampling. By using this technique the influence of model parameters can be investigated. We can explicitly investigate the correlation of model parameters and under which conditions the artificial price series exhibits empirical stylized facts. This new validation technique was introduced by Sallans et al., 2003.

The paper is organised as follows. In section 2, a standard asset pricing model is presented. Moreover we calculate the rational expectation equilibrium, which acts as a benchmark for our simulation. Section 3 focuses on the formation of expectation of our agents. In section 4 we present the results of the simulation and section 5 provides a conclusion and an outlook on further research.

2 Model

Within this section we will present a standard capital market model (see e.g. Arthur, et al., 1997; Brock and Hommes, 1998; Dangl et al., 2001). Myopic investors maximize their next period's utility subject to a budget restriction. At time t agents invest their wealth in a risky asset (a stock or index of stocks) with price p_t and in bonds, which are assumed to be risk free. There are S stocks paying a stochastic dividend d_t. The risk free asset is perfectly elastically supplied and earns the risk free and constant interest rate r. There are no restrictions on short selling of the stock or the bond. The model considered in the following is a discrete time model. Investors are allowed to change their portfolio in every time step[1]. The wealth of investor i at time $t + 1$ is given by

[1]See section 3 for the different length of a time step for an intraday trader and an end-of-day trader.

$$W_{i,t+1} = (1+r)\,W_{i,t} + (p_{t+1} + d_{t+1} - (1+r)\,p_t)\,s_{i,t} \tag{1}$$

where $W_{i,t}$ is the wealth at time t and $s_{i,t}$ the number of stocks of the risky asset held at time t. As in Brock and Hommes, 1998 , Levy and Levy, 1996 , Chiarella and He, 2001 or Chiarella and He, 2002 the demand function of the following model is derived from a Walrasian scenario. This means that each agent is viewed as a price taker (see Brock and Hommes, 1997 and Grossman, 1989)

Investors have a constant absolute risk aversion (CARA) utility. Therefore agents determine their investment in the risky asset independent of their wealth. This seems to be unrealistic because it is clear that in actual markets wealthier investors on average hold more stocks and therefore have a greater impact on prices. On the other hand if the wealth of investors is more or less evenly distributed, there should be little difference between investors with a constant absolute risk aversion versus agents with a relative absolute risk aversion.

Denote

p_t:	Price (ex dividend) per share of the risky asset at time t
d_t:	Dividend at time t
r:	Risk free rate
S:	Total number of shares of the risky asset
N:	Total Number of investors
$s_{i,t}$:	Number of shares investor i holds at time t
$W_{i,t}$:	Wealth of investor i at time t
ζ_i	Risk aversion of investor i

Let an investor i with wealth W_i maximize his/her utility of the form

$$u\,(W_i) = -e^{-\zeta_i W_i} \tag{2}$$

with ζ_i as constant absolute risk aversion. Denote by $F_t = \{p_{t-1}, p_{t-2}, ...p_0, d_t, d_{t-1}...d_0\}$ the information set available at time t[2]. Let $E_{i,t}$ and $V_{i,t}$ be the conditional expectation and conditional variance predicted by investor i at time t based on F_t. Then the demand for the risky asset $s_{i,t}$ solves

$$\max_{s_{i,t}} \left\{ E_{i,t}\,(W_{i,t+1}) - \frac{\zeta_i}{2} V_{i,t}\,(W_{i,t+1}) \right\} \tag{3}$$

i.e.,

$$s_{i,t} = \frac{E_{i,t}\,(p_{t+1} + d_{t+1}) - p_t\,(1+r)}{\zeta_i V_{i,t}\,(p_{t+1} + d_{t+1})} \tag{4}$$

Let S be the total number of shares. The market clearing price p_t is implicitly given by the equilibrium equation

$$S = \sum_{i=1}^{N} s_{i,t} \tag{5}$$

In the following section we will calculate the rational expectations equilibrium (REE) as a benchmark for our simulations. For determining the REE we assume homogeneous expectations. This means that all investors hold the same view about conditional return and variance i.e. $E_{i,t} \equiv E_t$ and $V_{i,t} \equiv V_t$.

[2]Note that at time t price p_t is not given.

2.1 Rational Expectation Equilibrium

Let the dividend follow an autoregressive process of order one,

$$d_t = \bar{d} + \alpha \left(d_{t-1} - \bar{d}\right) + \varepsilon_t \text{ with } \varepsilon_t = N\left(0, \sigma_d^2\right) \tag{6}$$

where the independent and identically distributed noise ε_t is drawn from a normal distribution with mean zero and variance σ_d^2.

We assume that agents conjecture that the price of the risky asset is a linear function of the dividend, i.e.

$$p_t = f d_t + g \tag{7}$$

Now rewrite the conditional expectation by replacing p_t from equation (7) :

$$\begin{aligned} E_{i,t}\left(p_{t+1} + d_{t+1}\right) &= E_{i,t}\left((1+f)\, d_{t+1} + g\right) \\ &= E_{i,t}\left((1+f)\left(\bar{d} + \alpha\left(d_t - \bar{d}\right) + \varepsilon_{t+1}\right) + g\right) \\ &= (1+f)\left(\bar{d} + \alpha\left(d_t - \bar{d}\right)\right) + g \end{aligned} \tag{8}$$

For the conditional variance we get

$$\begin{aligned} V_{i,t}\left(p_{t+1} + d_{t+1}\right) &= V_{i,t}\left((1+f)\, d_{t+1} + g\right) \\ &= V_{i,t}\left((1+f)\left(\bar{d} + \alpha\left(d_t - \bar{d}\right) + \varepsilon_{t+1}\right) + g\right) \\ &= V_{i,t}\left((1+f)\, \varepsilon_{t+1}\right) \\ &= (1+f)^2\, \sigma_d^2 \end{aligned} \tag{9}$$

Let all investors be equally risk averse and have homogeneous expectations on both mean and variance, i.e.

$$\begin{aligned} &\zeta_i \equiv \zeta, \\ &E_{i,t}\left(p_{t+1} + d_{t+1}\right) \equiv E_t\left(p_{t+1} + d_{t+1}\right) \text{ and} \\ &V_{i,t}\left(p_{t+1} + d_{t+1}\right) \equiv V_t\left(p_{t+1} + d_{t+1}\right) \text{ for } i = 1, \cdots, N \end{aligned} \tag{10}$$

Now we substitute the equation for the conditional expectation (8) and conditional variance (9) into the demand equation (4). Under homogeneous investors all agents have to hold S/N. There is no reason why one investor should hold more stocks than other investors (see the theorem on constant portfolios Kenneth et al.,2000) .

$$\frac{S}{N} = \frac{(1+f)\left(\bar{d} + \alpha\left(d_t - \bar{d}\right)\right) + g - \left(f d_t + g\right)(1+r)}{\zeta\,(1+f)^2\, \sigma_d^2} \tag{11}$$

The left hand side of the equation is a constant, so the right hand side must be a constant too. Therefore the time dependent variable d_t has to vanish.

$$\begin{aligned} &0 = d_t\left((1+f)\,\alpha - (1+r)\, f\right) \\ &f = \frac{\alpha}{(1+r-\alpha)} \end{aligned} \tag{12}$$

Substituting 12 into 11 and solving for g leads to

$$g = \frac{(f+1)\left(\bar{d}\,(1-\alpha)\right) - \zeta\,(1+f)^2\, \sigma_d^2 \frac{S}{N}}{r} \tag{13}$$

Therefore the solution for the REE-price p^* is given as

$$p^* = \frac{\alpha}{(1+r-\alpha)} d_t + \frac{(f+1)\left(\bar{d}(1-\alpha)\right) - \zeta(1+f)^2 \sigma_d^2 \frac{S}{N}}{r} \tag{14}$$

2.2 Sequence of Events

Let us have a look at the timing of the events within the model. First the dividend d_t of the current period is announced and paid. The next step is the formation of expectations. Based on past prices and dividends, including d_t an investor i forms his/her expectation about the distribution of the next period's price and dividend. According to equation (8) and (9) the investor calculates $E_{i,t}(p_{t+1}+d_{t+1})$ and $V_{m,t}(p_{t+1}+d_{t+1})$. Plugging the expectations into Equation (4) the agent is able to determine the demand function, which is submitted to the stock market via limit buy orders and limit sell orders[3]. After the orders of all agents are collected, the stock market calculates this period's equilibrium price p_t.

The market uses a sealed-bid auction, where the clearance mechanism chooses the price at which trading volume is maximized. The constructed supply and demand curves are based on the transaction requests. Example supply and demand curves are shown in 1

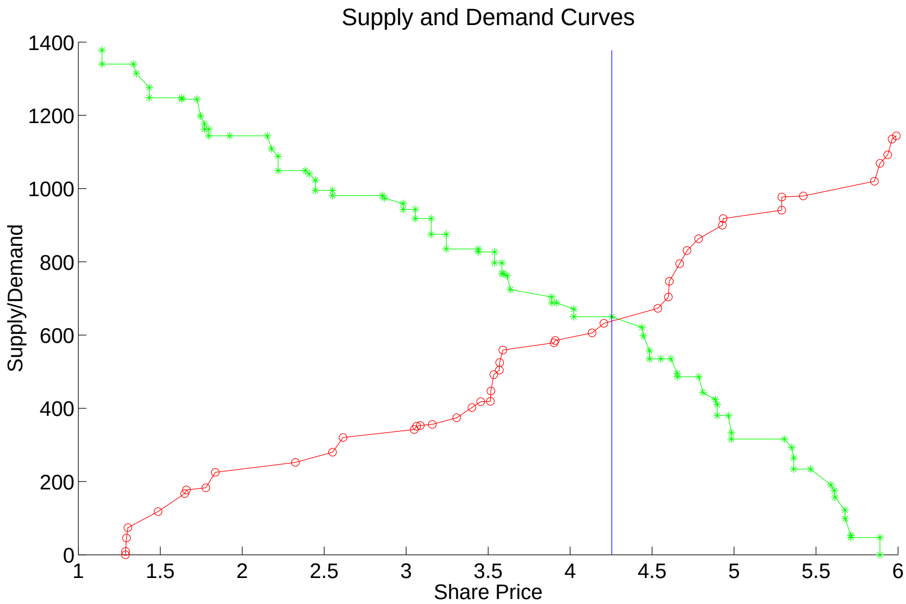

Figure 1: Supply and demand curves. Supply is marked with "O" and increases with price. Demand, marked with "*" decreases with price. The market price (vertical line) is set to a price which maximizes the volume traded. In this case, the market price is 4.25.

3 Formation of expectations

It is well known that expectations play a key role in modeling dynamic phenomena in economics. Heterogeneous expectations are introduced in the following way:

$$\begin{aligned} E_{i,t}(p_{t+1}+d_{t+1}) &= F_i(p_{t-1}, \cdots p_{t-h_i}, d_{t-1}, \cdots, d_{t-h_i}) \\ V_{i,t}(p_{t+1}+d_{t+1}) &= G_i(p_{t-1}, \cdots p_{t-h_i}, d_{t-1}, \cdots, d_{t-h_i}) \end{aligned} \tag{15}$$

[3] A limit order is an instruction stating the maximum price the buyer is willing to pay when buying shares (a limit buy order), or the minimum the seller will accept when selling (a limit sell order).

Heterogeneity arises from different information sets and different prediction functions. Agents have three characteristics:

- Type of prediction function: Fundamentalist or chartist
- Time horizon: Length of history of past prices and dividends used for prediction
- Trade interval: Intraday trader or end-of-day trader

These characteristics are initialized at the beginning of the simulation and are held fixed thereafter. In order to keep the simulation simple, there are only two types of prediction functions and the trade interval can only take one out of two different values. The above three characteristics can be combined in any way, e.g. chartists trading intraday or fundamentalists, who are end-of-day traders.

First let's have a closer look at the type of prediction function. As in many other heterogeneous agents models we assume that two kinds of investors exist: Fundamentalists and chartists. Chartists use past prices to predict the future price whereas fundamentalists calculate a "fair value" based on past dividends. Therefore for a fundamentalist $E_{i,t}(p_{t+1})$ is a function of past dividends and for a chartist $E_{i,t}(p_{t+1})$ is a function of past prices.

The second characteristic of agents is their time horizon. The time horizon is the length of history h_i of past prices and dividends used for their predictions. Time horizon means how far back agents look into past to predict next period's price and dividend. A long time horizon, i.e. high h_i means that many past observations are used as a forecast. The time horizon is drawn from an uniform distribution. The range of this uniform distribution depends on the agents type of prediction function. It was found that for fundamentalists and chartists different time horizons were appropriate. In particular chartists have to identify trends. The more observations they use, the further they look into the past which also means that longer trends can be identified. For example, if prices followed an upward trend, thereafter a downward trend, a short term chartist, i.e. a chartist with a short time horizon, would extrapolate the downward trend. This trader would predict a further drop in the stock price. For a long term chartist the last two trends might cancel each other and this agent would predict no change in the stock price. However the situation for fundamentalists may be different. Fundamentalists use past dividends for their price prediction. Depending on the type of the dividend process, it might be favorable to base the prediction only on the last observation. The minimum time horizon for a fundamentalist is 1, for a chartist it is 2 (the chartist needs at least 2 prices to calculate a trend, see also equation (23)). We allow two different maximum values for the time horizon for fundamentalists and chartists. For more details please see the results (section 4).

The third characteristic of an agent is the trade interval ϕ_i. Investors update their portfolio every ϕ periods. In our model we assume five price fixings per day with the last price setting being the closing price of the day. Therefore every 5^{th}, 10^{th}, 15^{th} etc. price is a closing price. We distinguish between two types of agents concerning their trading interval: Intraday traders with $\phi_i = 1$ and end-of-day traders with $\phi_i = 5$. Intraday traders update their position at every time step, whereas end-of-day traders only trade at the closing price of each day. This means that in the first four price settings per day only intraday traders rebalance their portfolios. In the fifth price fixing (the closing price) all agents trade. End-of-day traders ignore intraday prices, i.e. prices between their updates. The dividends of the five previous periods are accumulated. In order to keep consistent notation between intraday and end-of-day traders, we will define prices and dividends as used by the intraday traders as they are related to the highest frequency information available. Given the high-frequency price and dividend information $\widehat{p_t}$ and $\widehat{d_t}$, the price p_{t-1} and dividend d_t used by a trader with interval ϕ_i is given by:

$$\begin{aligned} p_{t-1} &= \widehat{p}_{t-\phi_i} \\ d_t &= \sum_{j=0}^{\phi_i - 1} \widehat{d}_{t-j} \end{aligned} \tag{16}$$

Therefore the information set used by trader i with interval ϕ_i is given by

$$F_{h_i,\phi_i,t}:\left\{\widehat{p}_{t-\phi_i},\widehat{p}_{t-2\phi_i},...,\widehat{p}_{t-h_i\phi_i},\sum_{j=0}^{\phi_i-1}\widehat{d}_{t-j},\sum_{j=0}^{\phi_i-1}\widehat{d}_{t-\phi_i-j},...,\sum_{j=0}^{\phi_i-1}\widehat{d}_{t-h_i\phi_i-j}\right\} \\ \rightarrow \{p_{t-1},p_{t-2},...,p_{t-h_i},d_t,d_{t-1},...,d_{t-h_i}\} \tag{17}$$

Now let's have a closer look at the formation of expectation. Let's begin with the expectation for the variance $V_{i,t}$. Agents determine $V_{i,t}$ in the following way

$$V_{i,t}\left(p_{t+1}+d_{t+1}\right)=\frac{1}{h_i}\sum_{j=1}^{h_i}\left(p_{t-j}+d_{t-j}-M\right)^2 \\ \text{with } M=\frac{1}{h_i}\sum_{k=1}^{h_i}\left(p_{t-k}+d_{t-k}\right) \tag{18}$$

Fundamentalists and chartists have the same prediction function for the variance. Their information set depend on their history h_i of past prices and dividends and their trade interval ϕ_i.
Now we take a look at the expectation of next period's price and dividend. First we split $E_{i,t}\left(p_{t+1}+d_{t+1}\right)$ into $E_{i,t}\left(p_{t+1}\right)$ and $E_{i,t}\left(d_{t+1}\right)$. Investors form their expectations over the next periods dividend d_{t+1} in the following way:

$$E_{i,t}\left(d_{t+1}\right)=\frac{1}{h_i+1}\sum_{j=0}^{h_i}d_{t-j} \tag{19}$$

Fundamentalists determine their price expectation according to a model based on fundamental information, which in our model are past dividends. A fundamentalist i assumes that the fair price $p_{i,t}^{\text{Fair price}}$ is a linear function of past dividends, i.e.

$$p_{i,t}^{\text{Fair price}}=F_i\left(d_t,\cdots,d_{t-h_i}\right) \\ =f\frac{1}{h_i+1}\sum_{j=0}^{h_i}d_{t-j}+g \tag{20}$$

This leads to the following price and dividend expectation for the fundamentalist

$$E_{i,t}\left(p_{t+1}+d_{t+1}\right)=\left(f+1\right)\frac{1}{h_i+1}\sum_{j=0}^{h_i}d_{t-j}+g \tag{21}$$

Chartists use the past history of the stock prices in order to form their expectations. They assume that the future price change per period equals the average price change during the last h_i periods[4].

$$E_{i,t}\left(p_{t+1}\right)=p_{t-1}+\frac{p_{t-1}-p_{t-h_i}}{h_i-1} \tag{22}$$

Note that at time t, p_t is not included in the information set F_t. This leads to the following price and dividend expectation of the chartists

$$E_{i,t}\left(p_{t+1}+d_{t+1}\right)=p_{t-1}+\frac{p_{t-1}-p_{t-h_i}}{h_i-1}+f\frac{1}{h_i+1}\sum_{j=0}^{h_i}d_{t-j}+g \tag{23}$$

[4]Note that for chartists $h_i >= 2$ for computing an average price change

4 Results

The aim of our model is to reproduce stylized facts with simple agents. Our agents are completely deterministic. The only source of randomness in the simulation are the dividends. The dividends follow an AR(1) process. We are interested in whether our simple agents are able to reproduce time series which exhibit stylized facts present in real world market data (see for example the survey by Cont, 2001).

This section is organized as follows. First we will recall the most important stylized facts in which we are interested. These are the stylized facts we want to reproduce with our model. Thereafter we will present the parameters of our model and how we determined the values for them. The main part of this section will be devoted to the results of our simulation.

4.1 Stylized Facts

The most important characteristic of a financial time series is the absence of significant autocorrelation of the returns. This is a necessary condition for the weak form of the efficient market hypothesis. The efficient markets hypothesis states that stock prices reflect all available information. If markets are efficient then new information is reflected quickly into market prices. To test the efficient markets hypothesis, the null hypothesis is that the financial market is a ”fair game”. That is, the difference between actual and expected returns is unpredictable. Fama, 1970 defines three forms of the efficient markets hypothesis: Weak form, semi-strong form and strong form. If stock prices are weak form efficient, past prices contain no information about future changes and price changes are random. Kendall, 1953 found that stock and commodity prices follow a random walk. A random walk implies zero correlation between price changes at t and price changes at t+1. Hence we derive the autocorrelation function of our artificial return series and check for significance.

The next important characteristic of a typical financial time series is persistent, partially predictable heteroskedasticity (”volatility clustering”). Since the work of Mandelbrot, 1963 the significant autocorrelation of the volatility is well known. Volatility clustering means that large or small price changes tend to follow one another. The clustering in periods of high and low volatility is not due to the clustered arrival of news, but the result of market dynamics. After new information is announced the investors need some time to deal with the shock. First there might be an exaggeration caused by technical traders, who amplify the short-term trend. After some time the price calms down again. As a volatility measure we take the squared returns. We therefore derive the autocorrelation function of the squared returns and check for significance.

Another feature of the returns of the stock price is a non-Gaussian distribution. Stock returns exhibit fat tails. This feature can be measured with the kurtosis. Hence we calculate the excess kurtosis of our artificial time series.

To sum up, in order to check for these stylized facts we derive the autocorrelation function of the returns, the autocorrelation of the squared returns and the kurtosis of the distribution of the returns.

4.2 Model Parameters

Although we use simple agents in a standard capital market model, there are still some parameters which have to be set. These parameters were described in the earlier sections. The aim of this section is to find parameter combinations which generate price series exhibiting empirical stylized facts.

First we run several preliminary simulations and divided the parameters into two groups. The first group consists of parameters which have a small influence on the outcome of the simulation. A variation of their values within a reasonable range has a small effect on the qualitative results of the simulations. These parameters, their values and their references are summarized in table 1. There are 100 agents, each given 100 stocks (S/N) as an initial endowment. The cash of an investor is not tracked. This is due to their CARA utility which means that an investment decision is made independent of an agent's wealth or cash.

The fundamentalist's prediction parameters f and g for the intraday trader were set according to equation (13) and equation (12) respectively. The risk free interest rate equals 0.4% and the mean dividend $\bar{d}$ is set to 3. Note, that for end-of-day traders the dividend d_t is the accumulated dividend of the trading day which corresponds to 5 periods. For end-of-day traders the mean dividend $\bar{d}$ equals 15 and the risk free interest rate is set to 2%. For end-of-day traders the fundamentalist's parameters f and g can not be computed from equation (13) and equation (12) because the dividends do not follow an AR(1) process (see section 3). For simplicity reasons f was set to $1/\phi^{th}$ of the value of intraday traders and g was set to the intraday trader's value.

Parameter	Description	Value	Reference
N	Number of agents	100	Eq. (5)
S	Number of stocks	10000	Eq. (5)
ζ	Risk aversion	0.00001	Eq. (2)
r	Risk-free interest rate per period	0.02 (0.004)	Eq. (4)
$\bar{d}$	Mean of dividend	15 (3)	Eq. (6)
α	Autoregressive parameter of dividend process	0.998	Eq. (6)
σ_d^2	Noise of dividend process	0.01	Eq. (6)
g	Parameter of prediction function	181	Eq. (21)
f	Parameter of prediction function	33.3 (166.3)	Eq. (21)
ϕ	trade interval of end-of-day trader	5	Eq. (17)

Table 1: Parameters, which were held fixed for the MCMC simulation. If parameter values for intraday traders and end-of-day traders are different, the parameter values for intraday traders are written in brackets.

The second group of parameters consists of those parameters having a big influence on the outcome of the simulation. We identified three important parameters: The proportion of fundamentalists, the time horizon of the agents and the proportion of intraday traders. We have to search for appropriate values for these parameters so that the generated artificial price series exhibits empirical stylized facts. In order to find "useful" parameter combinations we could "grid" the space of parameters and then run a simulation for each grid point. The disadvantage of this method is the high number of simulation to be conducted. For example three parameters each taking twenty different values would lead to 20^3 simulations. From these 8000 samples many would be wasted in uninteresting areas of the parameter space. In order to speed up the search process, a novel validation technique based on adapted Markov chain Monte Carlo (MCMC) sampling is used to determine the value of the parameters of the second group. Using the Metropolis algorithm (see Metropolis et al., 1953), samples are more likely to be drawn in parameter regions generating good results. We therefore need fewer samples to find "good" parameter combinations. This method allows for the efficient exploration of large parameter spaces. We can investigate how different model parameters are correlated, and under what conditions the model reproduces empirical stylized facts.

For this directed random search we need to define an energy function, which evaluates a simulation. In other words, we need to define what we mean by a "good" parameter combination. A low energy level corresponds to a better reproduction of stylized facts. A good parameter combination should generate a price series with no significant autocorrelation of the returns, significant autocorrelation of the squared returns and significant excess kurtosis. Preliminary simulations showed that most of our artificial price series exhibited significant excess kurtosis. Therefore we focused on the autocorrelation functions of the returns and squared returns. We define the energy function as follows:

$$E = \sum_{j=1}^{5} \left(\rho_j - \rho_j^*\right)^2 + \sum_{j=1}^{5} \left(\gamma_j - \gamma_j^*\right)^2 \tag{24}$$

where ρ_j denotes the autocorrelation coefficient of returns at lag j, γ_j denotes the autocorrelation coefficient of squared returns at lag j. The variables $\rho^* = \{0, 0, 0, 0, 0\}$ and $\gamma^* = \{0.15, 0.14, 0.13, 0.12, 0.11\}$ are the idealized autocorrelation coefficients, which correspond to typical empirical values. Thus the energy function is the sum of squared differences between actual autocorrelation coefficients and the idealized autocorrelation coefficients up to lag five for both, returns and squared returns. The autocorrelation coefficients were calculated over the last 2000 time steps. Using the adapted MCMC sampling method we searched the parameter space for regions minimizing the energy function E defined in equation (24). We held the parameters of the first group fixed (see table 1). The dividend sequence was generated by an AR(1) process at the beginning of the MCMC sampling. Thereafter each drawn sample of parameter combination was run with the same dividend sequence. The reason for using the same dividend sequence for all samples is to reduce randomness. Therefore given the dividend sequence, each parameter combination generates a deterministic price path and energy level. Figure 2 presents a shortened window of the energy level across iterations. In this example the parameter combination at the beginning is best, i.e. the energy function is minimized at the first few iterations. For more details about the Metropolis algorithm and the Markov chain Monte Carlo method see Sallans et al., 2003 .

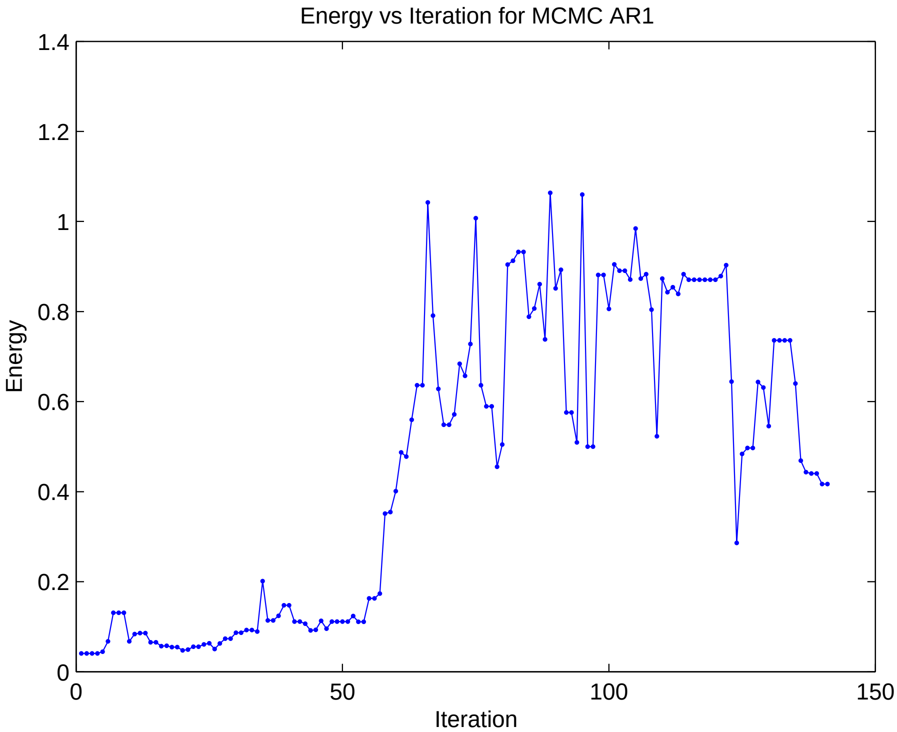

Figure 2: The energy level across iterations.

The values of the most influential parameters (see table 2) were determined by the best sample, i.e. the sample, which minimized the energy function. Let us have a closer look at the ”optimal” values.

The proportion of fundamentalists γ was set to 0.97 which means that 97% of all investors in our model are fundamentalists and only 3% are chartists. A possible explanation for the high proportion of fundamentalists is the following: The danger of the emergence of regular patterns and therefore significant autocorrelation of returns is high, because the agents in our model are completely deterministic. There are no noise traders and there is no learning, i.e. an agent always uses the same prediction function. In order to prevent regular patterns the price path has to stick closely to the fundamental value. Simulations showed, that reducing the proportion of fundamentalists increases the autocorrelation function for both returns and squared returns. Interestingly the proportion of fundamentalists was not set to 1, therefore it seems that a small proportion of chartists is an important factor for nonlinearity and volatility clustering.

Another important parameter of our model is the maximum time horizon. The time horizon of each investor is initialized at the beginning of the simulation and it is drawn from an uniform distribution between the minimum time horizon and the maximum time horizon. We allow for different value ranges for fundamentalists and chartists. The minimum time horizon of a fundamentalist is set to 1 and the minimum time horizon of a chartist is set to 2 (the chartist needs at least 2 prices to calculate a trend, see also equation (23)). The energy function was minimized by setting the maximum time horizon of the fundamentalists to 1. Thus the time horizon of a fundamentalist is always 1. This means that the price prediction of a fundamentalist is based only on the last dividend. The maximum time horizon of the chartists is set to 6.

The third important parameter of the model is the proportion of intraday traders. This variable was set to 0.31. This parameter is very important in order to generate price series exhibiting volatility clustering. Simulations without heterogeneous trade intervals did not lead to volatility clustering.

Parameter	Description	Value	Reference
γ	Proportion of fundamentalists	0.97	Section 3
h	Time horizon of fundamentalists	1	Section 3
	Time horizon of chartists (uniform distribution)	2-6	Section 3
δ	proportion of intraday traders	0.31	Section 3

Table 2: Parameters, which were varied in the MCMC simulation. Values were found to be optimal in the sense of reproducing the desired stylized facts best.

An interesting feature of the adapted MCMC method is the possibility to get some insights into the correlation of the parameters. The density of the samples in the parameter space is highest in regions of low energy because the Metropolis algorithm assigns a higher probability of samples to be drawn from low energy regions than from high energy regions. Figure 3 shows the correlation between the proportion of intraday traders and the maximum time horizon of chartists.[5] It can be seen that there are two regions (dark-coloured) of low energy, i.e. good reproduction of stylized facts. The second local minimum is found at a proportion of intraday traders of approximately 45% and the maximum time horizon of chartist being 17. However the adapted MCMC method found out that a proportion of intraday traders of 31% and a maximum time horizon of chartist of 6 were minimizing the energy function.

4.3 Simulation Results

Given the values of the parameters, we conducted 40 simulations with different AR1 dividend sequences lasting 3000 periods each. We applied our statistical tests to the last 2000 periods of each simulation.

Figure 4 shows a typical price path of a shortened window. The window was shortened so that the visual relation between the stock price and the REE price is not obscured by the compression necessary when graphing the entire history. The fundamental value calculated from equation (14) is plotted as dashed line. As one can see the price roughly reflects the underlying fundamental price. The reason for this phenomenon is the high proportion of fundamentalists, which forces the stock price to follow its fundamental value very closely (see also section 4.2).

There are many empirical studies about the distribution of returns. As already mentioned in section 4.1 the property of fat tails is an important stylized fact. In contrast to the normal distribution one finds more mass in the mean and in the tails of the distribution. Fat tails are revealed by the significant excess kurtosis. Figure 5 shows the returns of the simulated prices. The returns were calculated according to

$$ret_t = log\left(\frac{p_t}{p_{t-1}}\right) \tag{25}$$

for the last 2000 periods of the simulation run.

The histogram of returns is presented in Figure 6. As a point of reference a normal distribution with the same mean and variance as the distribution of returns is plotted. It can be seen, that the return distribution is not Gaussian. There is more mass in the mean and in the tails.

Table 3 shows some summary statistics of the returns calculated for the last 2000 periods of each run and averaged over 40 runs. Under a Gaussian distribution of returns the kurtosis is given with 3, whereas real world financial time series exhibit significant excess kurtosis, i.e. kurtosis larger than 3. Our model also generates time series showing excess kurtosis. The mean, standard deviation and skewness of the returns are within a reasonable range.

[5]The density plots were generated using the kernel density estimator for Matlab provided by C.C. Beardah at http://science.ntu.ac.uk/msor/ccb/densest.html [Beardah and Baxter 1996].

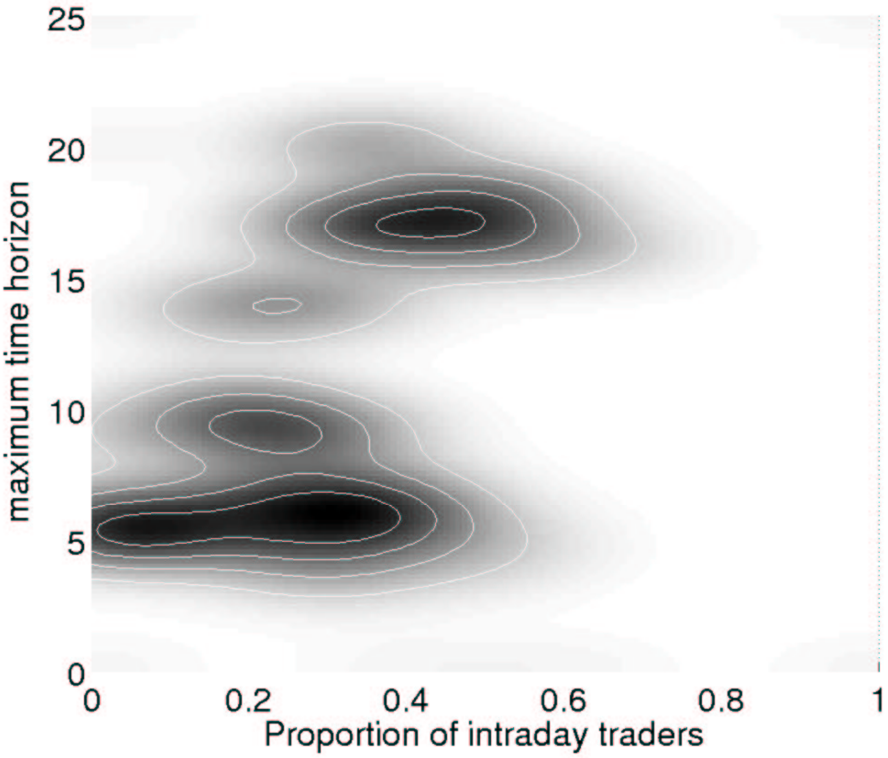

Figure 3: Correlation between proportion of intraday traders (x-axis) and maximum time horizon of chartists (y-axis) . There are two regions (dark-coloured) of good reproduction of stylized facts.

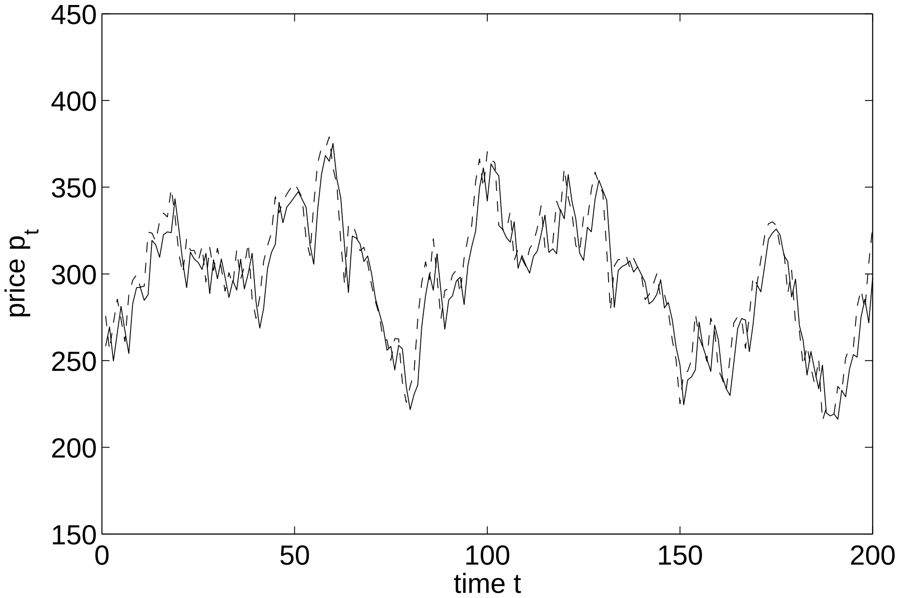

Figure 4: Price series and the REE price over a typical window. Solid line is the market price, dashed line is the REE price.

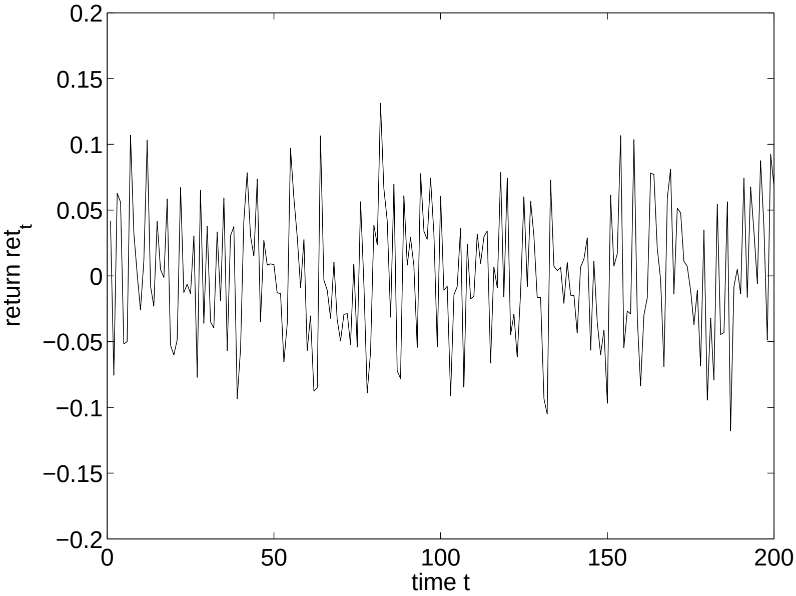

Figure 5: Return series over a a typical window.

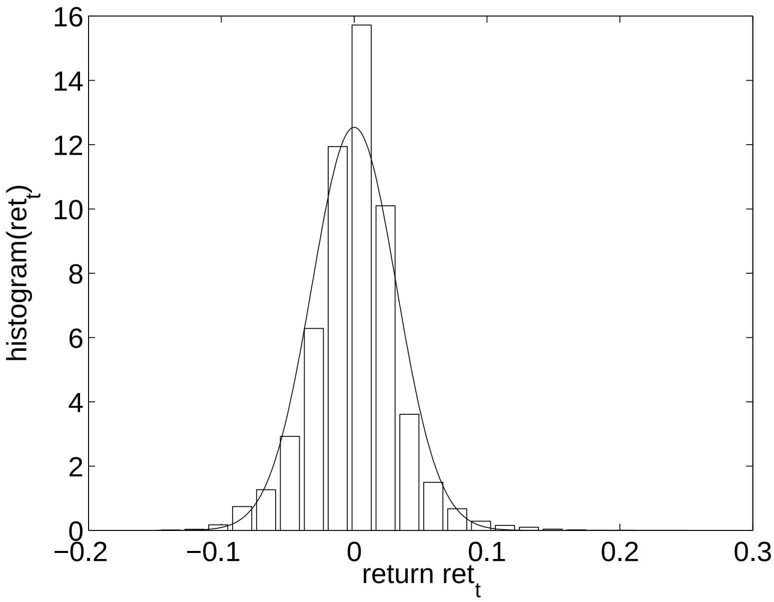

Figure 6: Histogram of returns.

Figure 7 and figure 8 present the autocorrelation coefficients of returns and squared returns for different time lags. They were again computed for the last 2000 periods of each run and averaged over 40 runs. The standard deviation of the autocorrelation coefficients across the 40 runs are plotted as error bars. As in real world markets no significance at a 95% confidence interval, given as $\pm 2/\sqrt{T}$, was found for the mean of autocorrelation of returns, although at lag 1 few runs exhibited significant autocorrelation. As already mentioned, the reason for the absence of significant autocorrelation of returns is the high proportion of fundamentalists. Therefore the stock price sticks close to its fundamental value. Moreover the time horizon

Variables	Mean	Standard deviation
Mean	1.7661e-05	3.2495e-04
Standard deviation	0.0309	0.0077
Kurtosis	4.1135	0.9341
Skewness	0.1461	0.1729

Table 3: Summary statistics of returns calculated for the last 2000 periods of each run and averaged over 40 runs.

of a fundamentalist is set to 1 and therefore the change of the fundamental value corresponds to the change in dividends.

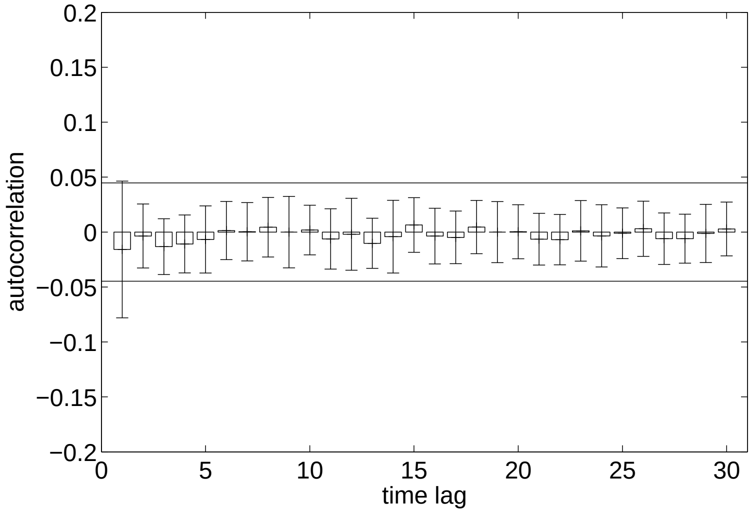

Figure 7: Autocorrelation of returns: Mean and standard deviation across 40 runs. Each simulation run lasted 2000 periods.

This is in contrast to the autocorrelation coefficients of the squared returns. They show significant positiv autocorrelation over long periods, i.e. the simulated prices exhibit volatility clustering. There are two reasons in our model for volatility clustering: Chartists and heterogeneous trade intervals. Although only 3% of all investors are chartists, they are able to introduce some positive autocorrelation of squared returns. However, the driving force for volatility clustering is the existence of heterogeneous trade intervals. Approximately $\frac{2}{3}$ of the agents only trade at the closing price of each day. This causes higher trading volume at closing prices compared to intraday prices (see also figure 9). As already mentioned, simulations with homogenous trade intervals did not reveal any volatility clustering.

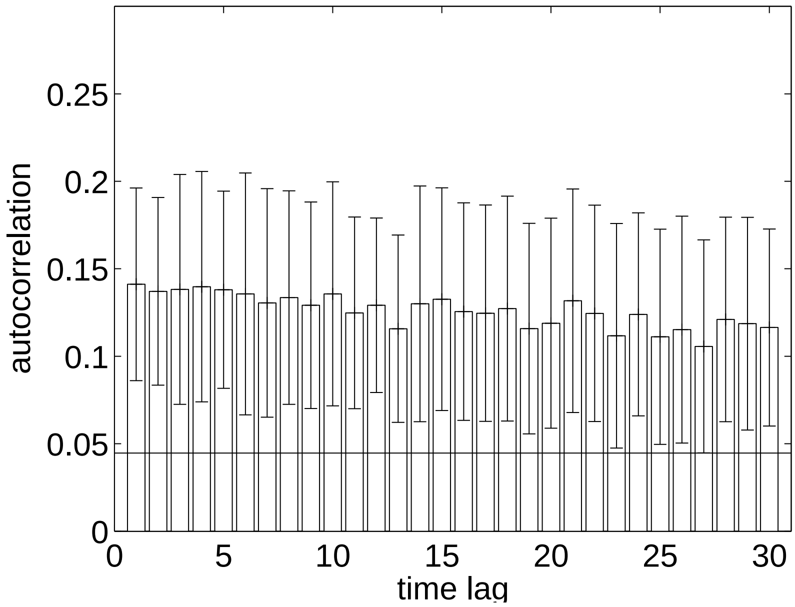

Figure 8: Autocorrelation of squared returns.

Figure 9 presents a snapshot of the trading volume over a typical window. Notice the increased trading activity at the closing price of each trading day. This is also found in empirical data of trading volume see Vieru, 2000. There are five price fixings per day, the last being the closing price. At the closing price all agents update their position, whereas at all other periods only intraday trader place their orders.

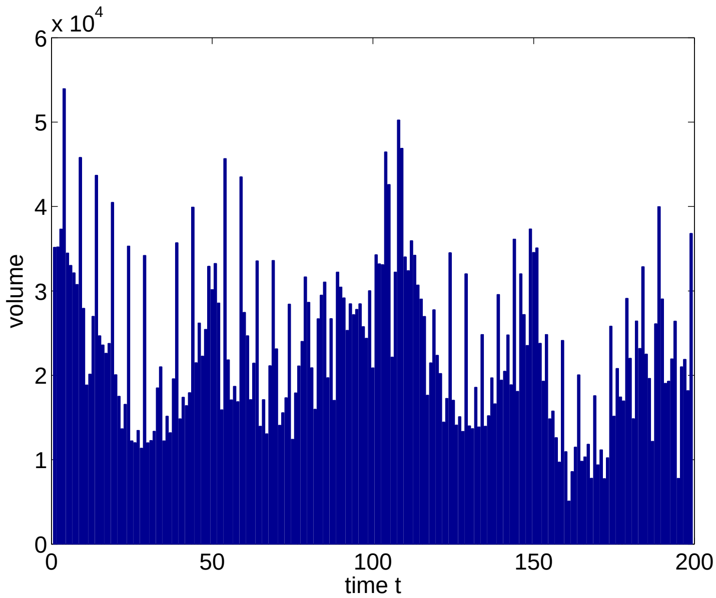

Figure 9: Trading volume in number of stocks over a a typical window.

Consistent with real markets, volume is significantly autocorrelated. This can be seen in figure 10 presenting the autocorrelation coefficients of the trading volume.

To sum up, the results show that our simple agents are capable of generating price series, which exhibit

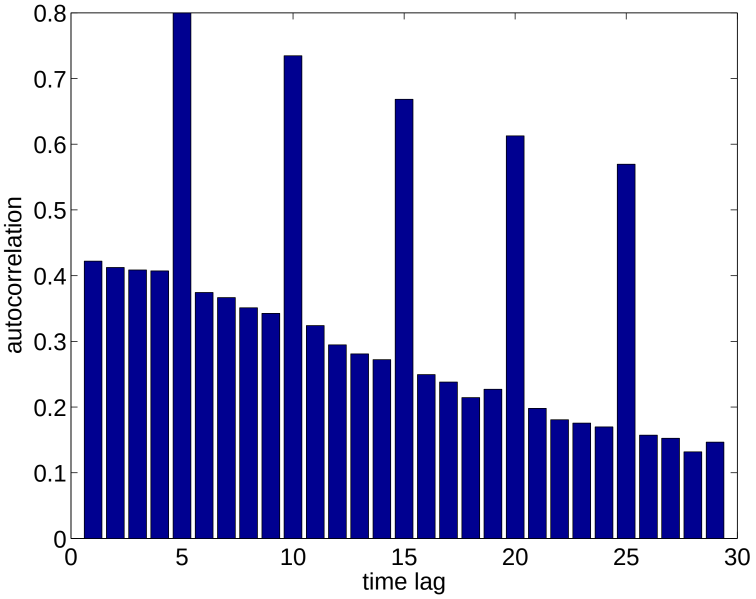

Figure 10: Autocorrelation of trading volume.

real world stylized facts. The returns are non-Gaussian distributed and excess kurtosis exists. Prices are efficient in that there is no significant autocorrelation of returns. Significant autocorrelation of squared returns shows that the artificial price series exhibits volatility clustering.

5 Conclusions

This paper presents a standard asset pricing model (see e.g. Arthur, et al., 1997; Brock and Hommes, 1998; Dangl et al., 2001) with simple deterministic agents. These agents are heterogeneous in three ways: Their prediction function, their time horizon and their trade interval. There are two different prediction functions, one based on fundamentals and the other one based on past prices. Fundamentalists assume that the fair price is a linear function of past dividends. Chartists interpolate past price trends into the future. The time horizon is modeled via the length of the history of past prices and dividends which is plugged into the expectation. Further we distinguish between intraday traders, who update their position every period and end-of-day traders. The latter only trade at the closing price of each day. Using the Metropolis algorithm in an adapted Markov chain Monte Carlo simulation we found the parameter combination which best generates price series exhibiting typical financial stylized facts. Interestingly, introducing heterogeneous trade intervals is an important factor in order to generate volatility clustering. The results are promising in that simple deterministic agents are able to reproduce efficient stock price returns with insignificant autocorrelation of returns at all lags, volatility clustering and fat tails.

References

Arthur,W. B., Holland, J., LeBaron, B., Palmer, R. and Tayler, P. (1997), Asset pricing under endogenous expectations in an artificial stock market, in W. B. Arthur, S. Durlauf and D. Lane, eds, 'The Economy as an Evolving Complex System II', Addison-Wesley, Reading, MA, pp. 15-44.

Bai, J. and Perron, P. (1998), 'Estimating and Testing Linear Models with Multiple Structural Changes', Econometrica 66, 47-78.

Beardah, C. and Baxter, M. (1996), Matlab routines for kernel density estimation and the graphical presentation of archaeological data, in H. Kammermans and K. Fennema, eds, 'Interfacing the Past: Computer Applications and Quantitative Methods in Archaeology 1995', Analecta Prehistorica Leidensia 28(1), Leiden.

Beltratti, A. and Margarita, S. (1992), Evolution of trading strategies among heterogeneous artificial economic agents, in J. A. Meyer, H. L. Roitblat and S. W. Wilson, eds, 'From Animals to Animats 2', MIT Press, Cambridge, MA.

Brock, W. and Hommes, C. (1997), 'A rational route to randomness', Econometrica 65, 1059-1095.

Brock,W. and Hommes, C. (1998), 'Heterogeneous beliefs and routes to chaos in a simple asset pricing model', Journal of Economic Dynamics and Control 22, 1235-1274.

Bullard, J. and Duffy, J. (1999), 'Using Genetic Algorithms to Model the Evolution of Heterogeneous Beliefs', Computational Economics 13, 41-60.

Chen, S., Lux, T. and Marchesi, M. (2001), 'Testing for non-linear structure in an artificial financial market', Journal of Economic Behavior and Organization 46 (3), 327-342.

Chiarella, C. and He, X. (2001), 'Asset pricing and wealth dynamics under heterogeneous expectations', Quantitative Finance 1, 509-526.

Chiarella, C. and He, X. (2002), 'Heterogeneous beliefs, risk and learning in a simple asset pricing model', Computational Economics 19, 95-132.

Cont, R. and Bouchaud, J.-P. (2000), 'Herd behavior and aggregate fluctuations in financial markets', Macroeconomic dynamics Vol 4, 170 - 196.

Cont, R. (2001), 'Empirical properties of asset returns: stylized facts and statistical issues', Quantitative Finance 1, 223-236.

Dangl, T., Dockner, E., Gaunersdorfer, A., Pfister, A., Soegner, A. and Strobl, G. (2001), 'Adaptive Erwartungsbildung und Finanzmarktdynamik', Zeitschrift für betriebswirtschaftliche Forschung, Vol. 53, 339-365.

Ehrentreich, N. (2002), 'The Santa Fe Artificial Stock Market Re-Examined - Suggested Corrections', Economics Working Paper Archive at WUSTL, http://ideas.repec.org/p/wpa/wuwpco/0209001.html

Fama, E. (1965), 'The behaviour of stock market prices', Journal of Business, January, 34-105.

Fama, E. (1970), 'Efficient capital markets: a review of theory and empirical work', Journal of Finance 25, 383-417.

Fama, E. (1991), 'Efficient capital markets: II', Journal of Finance 46, 1575-1617.

Farmer, J. (2000), 'A simple model for the nonequilibrium dynamics and evolution of a financial market', International Journal of Theoretical and Applied Finance 3 (3), 425-441.

Garcia, R. and Perron, P. (1996), 'An Analysis of the Real Interest Rate under Regime Shifts', Review of Economics and Statistics 78, 111-125.

Grossman, S. (1989), 'The Informational Role of Prices', MIT Press, Cambridge, MA.

Hinrich, M. and Patterson, D. (1985), 'Evidence of Nonlinearity in Daily Stock Returns', Journal of Business and Economic Statistics Vol. 3, No. 1, 69-77.

Kendall, M (1953), 'The Analysis of Economic Time Series. Part I: Prices', Journal of the Royal Statistical Society, 96.

Kenneth, J., Kubler, F. and Schmedders, K. (2000), 'Asset Trading Volume with Dynamically Complete Markets and Heterogeneous Agents', Discussion Paper 1294, Northwestern University, Center for Mathematical Studies in Economics and Management Science.

LeBaron, B. (2000), 'Agent based computational finance: Suggested readings and early research', Journal of Economic Dynamics and Control 24(5-7), 679-702.

LeBaron, B. (2001a), 'A builder's guide to agent-based financial markets', Quantitative Finance 1(2), 254-261.

LeBaron, B. (2001b), 'Evolution and time horizons in an agent based stock market', Macroeconomic Dynamics 5(2), 225-254.

LeBaron, B., Arthur, W. B. and Palmer, R. (1999), 'Time series properties of an artificial stock market', Journal of Economic Dynamics and Control 23, 1487-1516.

Levy, M. and Levy, H. (1996), 'The danger of assuming homogeneous expectations', Financial Analysts Journal 52(3), 65-70.

Lo, A. and MacKinlay, C. (1999), 'A Non-Random Walk Down Wall Street', Princeton University Press.

Malkiel, B. (1989), 'Efficient Market Hypothesis', Norton, New York, 1989. in The New Palgrave Finance, Eds. Eatwell, J., Milgate, M., Newman, P.

Malkiel, B. (1996), 'A Random Walk Down Wall Street'. Norton, New York.

Mandelbrot, B. (1963), 'The Variation of Certain Speculative Prices', Journal of Business 36, 394-419.

Metropolis, N., Rosenbluth, A., Rosenbluth, M., Teller, A. and Teller, E. (1953) 'Equation of state calculations by fast computing machines', Journal of Chemical Physics 21, 1087-1092.

Muth, J. F. (1961), 'Rational Expectations and the Theory of Price Movements', Econometrica 29, 315-335.

Noe, T., Rebello, M. and Wang, J. (2003), 'Corporate Financing: An Artificial Agent-based Analysis', The Journal of Finance 58, 943-973.

Sallans, B., Pfister, A., Karatzoglou, A. and Dorffner, G. (2003), 'Simulation and Validation of an Integrated Markets Model', forthcoming The Journal of Artifical Societies and Social Simulation.

Sargent, T. (1986), 'Rational Expectations and Inflation', Harper and Row, New York.

Sargent, T. (1993), Bounded Rationality in Macroeconomics, Oxford University Press, Oxford, UK.

Savit, R. (1988), 'When Random Is Not Random: An Introduction to Chaos in Market Prices', Journal of Futures Markets 8, 271-290.

Savit, R. (1989), 'Nonlinearities and Chaotic Effects in Option Prices', Journal of Futures Markets 9, 507-518.

Scheinkman, J. and LeBaron, B. (1989), 'Nonlinear Dynamics and Stock Returns', Journal of Business 3, 311-337.

Tay, N. and Linn, S. (2001), 'Fuzzy inductive reasoning, expectation formation and the behavior of security prices', Journal of Economic Dynamics and Control 25, 321-362.

Timmermann, A. (1998) 'Structural Breaks, Incomplete Information and Stock Prices'. Manuscript, Financial Markets Group, LSE.

Thurner, S., Dockner, E. and Gaunersdorfer, A. (2002), 'Asset Price Dynamics in a Model of Investors Operating on Different Time Horizons', Working Paper 93, http://www.wu-wien.ac.at/am/Download/wp93.pdf

Vieru, M. (2000), ’Intraday Trading Behavior around Interim Earnings Announcements on the Helsinki Stock Exchange’, Finnish Journal of Business Economics 49:1,109-137.