Working Paper Series

Adaptive Information Systems and Modelling in Economics and Management Science

Asset Price Dynamics in a Model of Investors Operating on Different Time Horizons

Stefan Thurner
Engelbert J. Dockner
Andrea Gaunersdorfer

Working Paper No. 93
November 2002

Working Paper Series

November 2002

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').

Asset Price Dynamics in a Model of Investors Operating on Different Time Horizons

Stefan Thurner
Departments of Mathematics NuHAG and ENT
University of Vienna
Währinger Gürtel 18-20
1090 Vienna, Austria

Engelbert J. Dockner*) and Andrea Gaunersdorfer
Department of Business Studies
University of Vienna
Brünner Straße 72
1210 Vienna, Austria

November 7, 2002

Abstract

We present a dynamic asset pricing model based on a heterogenous class of traders. These traders are homogenous in the sense that they are fundamentalists who base their investment decisions on an exogenoulsy given fundamental value. They are heterogenous in the sense that each trader is working with a different frequency of the underlying price data. As a result we have a system of interacting investors who together influence the market price. We derive a system that characterizes out-of-equilibrium dynamics of prices in this market which is structurally equivalent to the Nosé-Hoover thermostat equation in non-equilibrium thermodynamics. We explore the time series properties of these prices and find that they exhibit fat tails of returns distributions, volatility clustering and power laws.

Key Words: dynamic asset pricing, interacting agents and speculative markets, heterogenous agents

JEL classification: G12, D84, E32

*) Corresponding author.

1 Introduction

The random walk hypothesis for modeling asset prices has been critizised by numerous authors and studies (see e.g. Lo and MacKinlay, 2000). It is argued that prices and returns follow predictable patterns, that can be the outcome of herding effects, boundedly rational investment decisions or heterogenous beliefs. From a theoretical perspective the random walk assumption is closely linked to the efficient market hypothesis. In an efficient market prices fully reflect all available information at any time and hence any price change can only be the outcome of unanticipated shocks (Fama, 1970, 1991).

Theoretical studies of security prices on the other hand overemphasize fundamentals, i.e., cash flows and risk preferences of investors with homogenous beliefs. While these are important factors for equilibrium asset prices, both, the organization and the institutional details of asset markets (i.e., the market microstructure) and interacting agents with heterogenous beliefs must have significant impacts on the prices of financial assets. Furthermore, prices do not instantenously react to current changes in the environment. As Beja and Goldman (1980) argue, "operating markets are therefore better described by general processes that admit a *finite* adjustment speed and that allow transactions at the current (non-equilibrium) market prices."

The desire to build models deviating from the rational expectations equilibrium concept has led to active research during the last years, including psychological approaches (see, for example, the survey of Hirshleifer, 2001), agent-based computational models like the Santa Fe artificial stock market (Arthur et al., 1997 and LeBaron et al., 2000; see also the review by LeBaron, 2000) and evolutionary models (e.g. Farmer, 2000; Brock and Hommes 1997, 1998, see also Brock, 1997, and the review of Hommes, 2000). These models show that the interaction among traders and their heterogeneous nature by themselves can give rise to interesting non-linear effects and generate endogenous fluctuations. They may help explaining so-called stylized facts in financial data, like fat tailed distributions of returns, volatility clustering and power laws (e.g. Gaunersdorfer and Hommes 2002, Lux and Marchesi, 1999, 2000). Recently also methods of statistical physics of disordered systems were applied to study models of interacting agents (see the survey by Farmer, 1999).[1]

LeBaron (2001) studies the impact of different time horizons of investors in an evolutionary setting. He demonstrates that a heterogeneous learning market can amplify the variability in a relatively stable fundamental series and generate fat tails, persistent volatility, and trading volume that is persistent and moves with volatility.

Dacorogna et al. (2001) introduce the notion of a heterogeneous market hypothesis. Short-term traders evaluate the market at a higher frequency and have a shorter memory than long-term traders. These authors consider different time scales of the market participants

[1] See also the contributions by the participants of the workshop "Beyond Equilibrium and Efficientcy" held at the Santa Fe Institute in May 2000, which appeared in the first two issues of *Quantititative Finance* (2001).

as the key characteristic of the market. They argue that "the weak form of efficiency coupled with the rational expectation model cannot be attained. Because of the presence of different time components with heterogeneous expectations, current market prices cannot reflect all available information. The price discovery mechanism follows rather a 'dynamic error correction model' where the successive reactions to an event unfold in the price." (p. 354)

In this paper we present a dynamic asset pricing model that introduces the importance of the individual investment horizon for investment decisions. We assume that investors are heterogenous with respect to the time horizon and explore the consequences of this on price dynamics. In particular, we start out that each trader is in a sense a fundamentalist (or in terms of Farmer, 2000, a value investor) who makes his investment decision on the basis of the difference between the actual price and an exogenous given fundamental value. But traders differ with respect to the type of data they use. For example, a day trader uses intraday data to make his investment decisions. A trader who is more interested in the long run might use monthly or quarterly data to rebalance his portfolio. Hence the same types of traders are using data with different frequencies to make their investment decisions. Day traders immediately react to changes of the actual price in relation to its fundamental value, while long run investors only slowly react to these changes since they are operating on a monthly or even quarterly time scale. These differences introduce a set of interacting agents who together determine the current price. Hence, we can argue that due to different investment horizons agents react with different speeds to mispricings in the market. Price dynamics is driven by these interactions.

We derive a system of non-linear differential equations that characterizes the out-of-equilibrium price dynamics. The system has the same structure as the Nosé-Hoover thermostat equation in non-equilibrium thermodynamics (Nosé, 1984a,b, 1991, Hoover, 1985). Like a thermostat investors drive prices into the direction of the fundamental value. We find that for certain parameter values the system exhibits complex behavior. For these values we derive price paths and explore their time series properties. We find that the returns generated by our simple *deterministic* system exhibit many stylized facts.

Our paper is organized as follows. In the next section we present our model and derive its reduced form. In Section 3 we discuss the properties of returns series, Section 4 concludes.

2 A Simple Model of Agents with Different Investment Horizons

Consider an actively traded security for which the current price is given by $P(t)$. Assume that total supply of the asset at any point in time is fixed so that price changes are only driven by excess demand in the market. The relationship between price changes and excess demand D_t in the market at time t is modeled by

$$\frac{\dot{P}(t)}{P(t)} = D_t\left(P(t), F(t)\right), \tag{1}$$

where $P(t)$ is the actual security price and $F(t)$ is its fundamental value at time t. We assume that $F(t)$ is exogenously given and do not derive the fundamental value of the asset through its cash flow profile. We do not assume an equilibrium market clearing condition for price formation, transactions are not made when prices are in equilibrium where supply equals demand (like in a classical Walrasian auction). Instead, price movements are forced by supply and demand imbalances. Thus, we are primarily interested in out-of-equilibrium price dynamics.

Investors believe that the fundamental value is not fully reflected in the price and try to make profit out of this mispricing. They take positive positions when they think that the market is underpriced and negative positions when they think it is overpriced (cf. Beja and Goldman, 1980, and Farmer, 2000).

Since $F(t)$ is the fundamental price of the asset at time t, excess demand for the short run investors D_t^0 evaluated at $F(t)$ is zero, i.e.

$$D_t^0(F(t), F(t)) = 0.$$

We assume a linear demand function for the short run investors (cf. Beja and Goldman, 1980, Chiarella, 1992, Farmer, 2000)

$$D_t^0\left(P(t), F(t)\right) \equiv \alpha[F(t) - P(t)], \tag{2}$$

where α is some positive constant. The interpretation of short run demand is closely related to the behavior of the (short run) fundamentalists. They react immediately to mispricings in the market and hence cause the actual price to move towards the equilibrium price.

Other traders are also reacting to mispricings in the market but they are operating on different time scales, i.e. they use data with different frequencies than the short run investors. In particular, a long run investor is not so much interested in short run deviations of the actual price from its fundamental value but in long run deviations. Hence, he reacts to changes in the difference between the actual price and the fundamental value with some time delay. For example, it could take him a period of $\tau > 0$ until he reacts to price deviations from the fundamental value, i.e.,

$$D_t^\tau\left(P(t-\tau), F(t-\tau)\right) \equiv \alpha(\tau)\left[F(t-\tau) - P(t-\tau)\right].$$

$\alpha(\tau)$ is a density function which specifies the weight (or number) of investors which are characterized by reaction time τ. For simplicity we set $\alpha(\tau) \equiv \alpha$. We assume that there is a continuum of investors with different investment horizons who are characterized by different values of τ. Total excess demand is the sum of individual demands and therefore given by

$$D_t\left(P(t), F(t)\right) \equiv \alpha \int_0^t \left[F(t-\tau) - P(t-\tau)\right] d\tau. \tag{3}$$

Introducing a new variable

$$\xi(t) \equiv \gamma \int_0^t \left[F(t-\tau) - P(t-\tau)\right] d\tau, \tag{4}$$

we get from (1)

$$\frac{\dot{P}(t)}{P(t)} = D_t(P(t), F(t)) = \beta\xi(t),$$

where $\beta = \alpha/\gamma$ and $\xi(t)$ satisfies the differential equation

$$\dot{\xi}(t) = \gamma[F(t) - P(t)]. \tag{5}$$

Investors in our model are aware that they can misjudge the situation about mispricings and as a consequence they make losses. Since they are not entirely sure if their investment will be profitable or not, they will invest only limited amounts. The model should thus be able to restrict the price changes ξ to a region around zero with controllable variance. Technically this can be done by introducing a new dynamical variable ζ and by adding a term $-\xi\zeta$ to the left hand side of Eq. (5),

$$\dot{\xi}(t) = \gamma[F(t) - P(t)] - \xi\zeta. \tag{6}$$

As long as the asset is overpriced, $\dot{\xi}$ is negative. Thus, ξ will decline and eventually become negative and prices will fall. The dynamics of ζ, which keeps the variance of ξ at a value of V, is given by

$$\dot{\zeta}(t) \;\; = \;\; \delta\left[\xi^2(t) - V(t)\right], \tag{7}$$

where δ is another positive constant. This mechanism is inspired by a well known approach in non-equilibrium thermodynamics, the dynamical thermostat equations introduced by Nosé (1984a,b, 1991) and Hoover (1985). $V(t)$ is a factor (maybe time dependent) that controls the variance and therefore can be interpreted as a measure of aggregate risk aversion of the involved agents. If V is small, agents are not willing to take risk and only small price changes will occur. Taking all the arguments together our model results in the following deterministic dynamical system (omitting the time variable),

$$\begin{aligned} \dot{P} &= \beta P \xi \\ \dot{\xi} &= \gamma(F - P) - \xi\,\zeta \\ \dot{\zeta} &= \delta(\xi^2 - V). \end{aligned} \tag{8}$$

Symbol	Meaning in Model	Relates to in physics
$p(t)$	price	momentum
β	primary coupling of the asset to the market	coupling constant
ξ	relative price-change variable	thermostat variable
ζ	control variable for ξ	thermostat variable
$F(t)$	"fundamental" value of asset	temperature
$V(t)$	term of collective risk aversion	temperature
γ, δ	response onset times	response onset times

Table 1: Model parameters and variables and its counterparts in statistical physics.

The parameters γ and δ control the rigidity of the system, i.e. they represent an inverse "onset-time" for the correctional price movements to happen.[2] Table 1 compares the model parameters and variables and its counterparts in statistical physics.

We will show that this simple deterministic model is able to generate some of the stylized facts observed in financial data such as fat tails in the distribution of returns, volatility clustering, long memory in the second moments of returns, and scaling laws in the moments of the return process.

3 Time Series Properties of Simulated Price Data

As was pointed out the model is driven by two forces: the interaction of value traders operating on different time scales and the variance of the prices induced by an average level of risk aversion.

Our primary interest focusses on whether or not the model generates price series that exhibit stylized facts present in real world asset markets. We therefore analyze the properties of the price series generated by the dynamical system (9).

For that purpose let us recall which are the most relevant stylized facts we are interested in.[3] Firstly, the returns of asset prices are characterized by the absence of autocorrelations which is consistent with the efficient markets hypothesis. Therefore we derive the (linear) autocorrelation function of the returns and check for significant autocorrelation coefficients. Secondly, in many asset markets volatility measures like squared returns show significant positive autocorrelations. This is a strong indication for non-linearity,

[2]The first two equations have the form of the Nosé-Hoover thermostat equations. If one considers p^2 instead of p, the first equation stays the same up to a rescaling factor 2. The second equation can be seen as a dynamical thermostat, which keeps the kinetic energy (p^2) at a temperature $F(t)$. In this case γ is the thermostat onset time.

[3]A comprehensive discussion of stylized facts (empirical properties) of asset returns can be found, for example, in Cont (2001). He lists eleven regularities, but it must be expected that our simple stylized model cannot generate all these effects.

volatility clustering, and long memory. To identify this stylized fact we therefore derive the autocorrelation function for the squared returns and check for significant correlations. We show that they decline with a power law. Finally, we analyze if returns exhibit fat tails. This can simply be done by looking at the implied probability distributions of the returns generated by our model.

It should be emphasized here, that our simulated data are entirely driven by a deterministic system, the model does not contain any source of randomness. Thus, if some of the above stylized facts are present, they are the consequences of the non-linear structure of the model and not driven by any exogenous forces.

To derive price paths for our model we apply a Runge Kutta 4^{th} order solver with dynamic time increments Δt ranging from 10^{-3} to 10^{-4}, according to preset relative and absolute error tolerances (10^{-3} and 10^{-6} respectively). Prices are recorded at integer times $t = 1, 2, \ldots, N$, where N is typically a few hundred to a few thousand. We have carefully checked by numerical simulation that for constant V the third model equation in (9) indeed keeps the variance of ξ at the desired value V. Distributions of ξ show fat tails but the sample variance is V, with very little error. If V changes over time the actual variance tracks the functional form of $V(t)$.

As pointed out, the fundamental value $F(t)$ is exogenous in our model. In principle, we could choose any process (also a stochastic one) for the fundamental value. Here we choose a simple time profile of the asset price with exponential growth so that $F(t) = S_0 \exp(rt)$, where $r = 0.005$ and $S_0 = 10$. Note that by this specification (9) is a non-stationary system. We assume that the average risk aversion of agents decreases as their wealth increases. This is taken into account by setting $V(t) = \varepsilon F(t)$. This specification implies that the larger the value of the asset the bigger price changes the average agent is willing to accept.

Figures 1 and 2 present some typical simulation results for the following parameter specifications: $\xi(0) = -1$, $\zeta(0) = 1$; Figure 1: $P(0) = 10$, $\beta = 0.2$, $\gamma = \delta = 10$, $\varepsilon = 0.5$; Figure 2: $P(0) = 1$, $\beta = 0.2$, $\gamma = \delta = 12$, $\varepsilon = 0.6$. It should again be pointed out that the dynamics shown in the price paths and the returns are entirely driven by the feedback mechanism that characterizes the model, but recorded only at integer tines.

Figures 1(a) and 2(a) present the price paths. Prices fluctuate around the exponentially growing fundamental value. Based on these prices continuoulsy compounded returns are calculated according to

$$r(t) = \log \frac{P(t)}{P(t-1)}.$$

The derived returns (Figures 1(b) and 2(b)) are mean stationary, but exhibit heteroscedastic variances for which large changes are mostly followed by large changes, i.e. volatility clustering is present in the data. The histograms for the returns are shown in Figures 1(c) and 2(c). It clearly can be seen that the return distributions are not Gaussian. In particular, they are leptokurtic, exhibit excess kurtosis and fat tails. As a point of reference the solid lines show the densities for the normal distributions with mean zero and sample variance V. Figures 1(d), (e) and 2(d), (e) show the (linear) autocorrelation coefficients of returns and squared returns for different lags. Here we find that the returns have significant negative autocorrelations of first order but small or insignificant correlation coefficients for higher lags. In some real world markets (like in Austria for example) returns show significant positive autocorrelations at a level around 0.3. This, however, indicates thin trading and slowly adjusting prices. The negative significant autocorrelation found in our model is mainly driven by the strong mean reverting property of actual prices. Squared returns show significant positive autocorrelations over long periods of time indicating volatility clustering and long memory. The first 100 lags die out slowly and decay with a power law (straight line) with an exponent of -0.201 resp. -0.048 (Figures 1(f), 2(f)). We found that these values are robust with respect to the length of the returns series, but change significantly with different parameters.

The observation that the first lag autocorrelations of returns are strongly negative seems to be due to the fact that all investors believe that sooner or later prices will return to the fundamental value. Thus, prices are characterized by a mean-reverting process. There is empirical evidence that in real markets traders using very different trading strategies are present. Introducing speculators who extrapolate price trends might destroy these significant autocorrelation (cf. Farmer, 2000).[4]

To sum up, both simulated price series and their continuously compounded returns exhibit many properties that are present in real world markets. As a point of reference we present price paths, continuously compounded returns, empirical distributions, autocorrelations of returns and squared returns, and power laws for the decay of the squared returns of two individual stocks MicroSoft and Coca Cola for which we matched the sample lengths with the examples of the simulated data (Figures 3 and 4). By simply comparing the artificial with the real data it is obvious that there are many similarities. This indicates that the model captures features about investors and their trading strategies that can be used to explain some of the empirical facts. Therefore we explore the mechanisms that drive the results in more detail.

Overall it needs to be stressed that the results presented in Figures 1 and 2 are striking. It is surprising that a simple deterministic system like (9), where the fundamental value follows a smooth exponentially growing proccess, can generate volatility clustering, leptokurtic distributions and fat tails. Let us first concentrate on the fluctuations of the prices around the fundamental value. These fluctuations are entirely driven by (6) and (7). The first part of equation (6) shows that every time the price $P(t)$ is different from its fundamental value $F(t)$ it moves towards it by a linear adjustment mechanism, thus we have a strong mean reverting property. This adjustment is perturbed by total excess demand $\xi(t)$ multiplied by $\zeta(t)$ which controls for the change in the price. Even if the current price is equal to its exogenously given fundamental value it does not stay there since $P(t) = F(t)$ is not a fixed point for the dynamical system (9).

[4] Decreasing the onset times γ and δ reduces also the first autocorrelation coefficient. However, for such parameter values the system tends to unveil regular dynamics.

The property of volatility clustering can also be rationalized easily. From the definition of the continously compounded returns we get

$$\dot{r}(t) = \frac{\dot{P}(t)}{P(t)} = \beta\xi(t) = \beta \int_0^\infty [F(t-\tau) - P(t-\tau)]d\tau.$$

Hence, the returns vary with excess demand. Excess demand, however, is driven by current as well as historical (lagged) mispricings in the market. These lagged reactions together with the interaction of investors of different types cause an overshooting effect and hence result in volatility clustering. This seems to be an interesting explanation for volatility clustering. Heterogenous investors with different time horizons and different reaction speeds to mispricings in the market amplify changes in the returns.

The analysis above has demonstrated that our model generates price fluctuations that are in line with empirical observations. But it should be pointed out that so far we have delt with behavior along a single scale (e.g. when looking at the autocorrelations we analyze price and return changes along a single time scale). As Farmer (1999) points out, a good price-fluctuations model should connect the behavior on multiple timescales. A natural test for this is to look at the behavior of moments $E_t|r_\tau|^q$ for different q and τ. There are strong empirical indications that there is approximate power-law scaling with respect to τ but with different slopes for different values of q. A property like this suggests a fractal random process.

To analyze this property with our simulated data we compute higher moments q for different lags τ, i.e. $E|r_\tau|^q$ for the continuously compounded returns. Figure 6 presents the results for the simulated series presented in Figure 1 as well as for the returns of the MicroSoft stock. In the top panells of Figure 6 we plot the moments of the absolute returns as a function of τ for different values of q, where q ranges from 1.5 to 3.0, for the MicroSoft returns and for the simulated returns. We observe that $E|r_\tau|^q$ scales with respect to τ for a given q, the corresponding log-log slopes are very small (close to horizontal lines). Moreover, the variation of the slopes with respect to q seems to follow a non-linear pattern, which suggests a multi-fractal behavior. To test for this multi-fractal behavior we calculate the ratio

$$\frac{E|r_\tau|^q}{(E|r_\tau|)^q}.$$

This ratio must be constant for a simple fractal but changes for a multi-fractal with changing q. The lower panells of Figure 6 present these ratios again for the MicroSoft and the simulated returns. Here we can observe two things. Firstly, it seems that the simulated as well as the real returns are generated by a multi-fractal process. Secondly, again we observe very strong simularities between the real and the simulated returns. Therefore it is legitimate to argue that our simple model of heterogenous investors with different time horizons is rich enough to mimic many of the power law properties that can be found in real world securities markets.

4 Conclusions

In this paper we present a simple low dimensional deterministic dynamical system that – over a wide range of parameter settings – generates prices and returns series that exhibit some of the most important stylized facts present in real world asset markets. Of course, such a simple, purely deterministic model can never represent all aspects of a financial market, but it is intriguing that it is possible to explain many stylized facts of its resulting time series. The main feature of our model is the assumption that there are heterogenous (value) investors who act on different time scales. This difference in the time horizons of the investors introduces enough heterogeneity and interactions so that complicated dynamics result from this behavior. Our model generates fat tailed distributions and non-significant autocorrelations of returns (except for order one) but significant positive autocorrelations for squared returns that decay hyperbolically. What is even more surprising is the property that the simulated data are characterized by similar power laws as real asset returns. We even found a strong indication that the returns of our model are generated by a multi-fractal process. Empirical capital market research has found that asset returns also seem to be generated by multi-fractal random processes. One of the shortcomings that still exist in our model are the price paths. In particular the long run behavior of actual prices is entirely driven by the exponentially growing fundamental value. Modelling the fundamental value by a stochastic process or introducing further groups of traders, like trend followers, would add additional degrees of real world market behavior to our model. Preliminary simulations with a stochastic fundamental value indicate, for example, that the strong negative first autocorrelation in the returns can be drastically reduced. We leave such extensions for future research.

5 References

Arthur, W.B., Holland, J.H., LeBaron, B., Palmer, R., Taylor, P., 1997, Asset pricing under endogenous expectations in an artificial stock market. In: Arthur, W.B., Durlauf, S.N., Lane, D.A., (Eds.), The economy as an evolving complex system II, Redwood City, Addison-Wesley.

Beja, A., Goldman, M.B., 1980, On The Dynamic Behavior of Prices in Disequilibrium, The Journal of Finance 35, 235–248.

Brock, W.A., 1997, Asset Price Behavior in Complex Environments, In: Arthur, W.B., Durlauf, S.N., Lane, D.A., (Eds.), The economy as an evolving complex system II, Redwood City, Addison-Wesley.

Brock, W.A., Hommes, C.H., 1997, A rational route to randomness, Econometrica 65, 1059–1095.

Brock, W.A., Hommes, C.H., 1998, Heterogeneous beliefs and routes to chaos in a simple asset pricing model, Journal of Economic Dynamics and Control 22, 1235–1274.

Chiarella, C., 1992, The Dynamics od Speculative Behavior, Annals of Operations Research 37, 101–123.

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

Dacorogna, M.M., Gençay, Müller, U.A., Olsen, R.B., Pictet, O.V., 2001. An Introduction to High-Frequency Finance. Academic Press, San Diego, California.

Fama, E.F., 1970, Efficient capital markets: a review of theory and empirical work, The Journal of Finance 25, 383–423.

Fama, E.F., 1970, Efficient capital markets: II, The Journal of Finance 46, 1575–1617.

Farmer, J.D., 1999, Physicists Attempt to Scale the Ivory Tower of Finance, Computing in Science & Engineering (IEEE), November/December 1999, 26–39.

Farmer, J.D., 2000, Market force, ecology, and evolution, SFI Working Paper 98-12-117, http://arXiv.org/abs/adap-org/9812005.

Gaunersdorfer, A., Hommes, C.H., 2000, A Nonlinear Structural Model for Volatility Clustering, SFB Working Paper No. 63, University of Vienna and CeNDEF Working Paper 00-02, University of Amsterdam.

Hirshleifer, D., 2001, Investor Psychology and Asset Pricing, The Journal of Finance 56, 1533–1597.

Hommes, C.H., 2001 Financial markets as nonlinear adaptive evolutionary systems, Quantitative Finance 1, 149–167.

Hoover, W.G., 1985, Canonical Dynamics: Equilibrium Phase-Space Distributions, Physical Review A 31, 1695–1697.

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.

Le Baron, B., 2000, Agent-based computational finance: Suggested readings and early research, Journal of Economic Dynamics and Control 24, 679–702.

Le Baron, B., 2001, Evolution and Time Horizons in an Agent Based Stock Market, Macroeconomic Dynamics 5, 225–254.

Lo, MacKinlay, 2000. A Non-Random Walk Down Wall Street. Princeton University Press, Princeton, New Jersey.

Lux, T., Marchesi, M., 1999a, Scaling and criticality in a stochastic multi-agent model of a financial market, Nature 397, February 1999, 498–500.

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

Müller, U.A., Dacorogna, M.M., Davé, R.D., Olsen, R.B., Pictet, O.V., von Weizsäcker, J.E., 1997, Volatilities of different time resolutions – Analyzing the dynamics of market components, Journal of Empirical Finance 4, 213–239.

Nosé, S., 1984a, A Unified Formulation of the Constant Temperature Molecular Dynamics Method, The Journal of Chemical Physics 81, 511–519.

Nosé, S., 1984b, A Molecular Dynamics Method for Simulation in the Canonical Ensemble, Molecular Physics 52, 255–268.

Nosé, S., 1991, Molecular Dynamics Simulations, Progress of Theoretical Physics Supplement 103, 1–46.

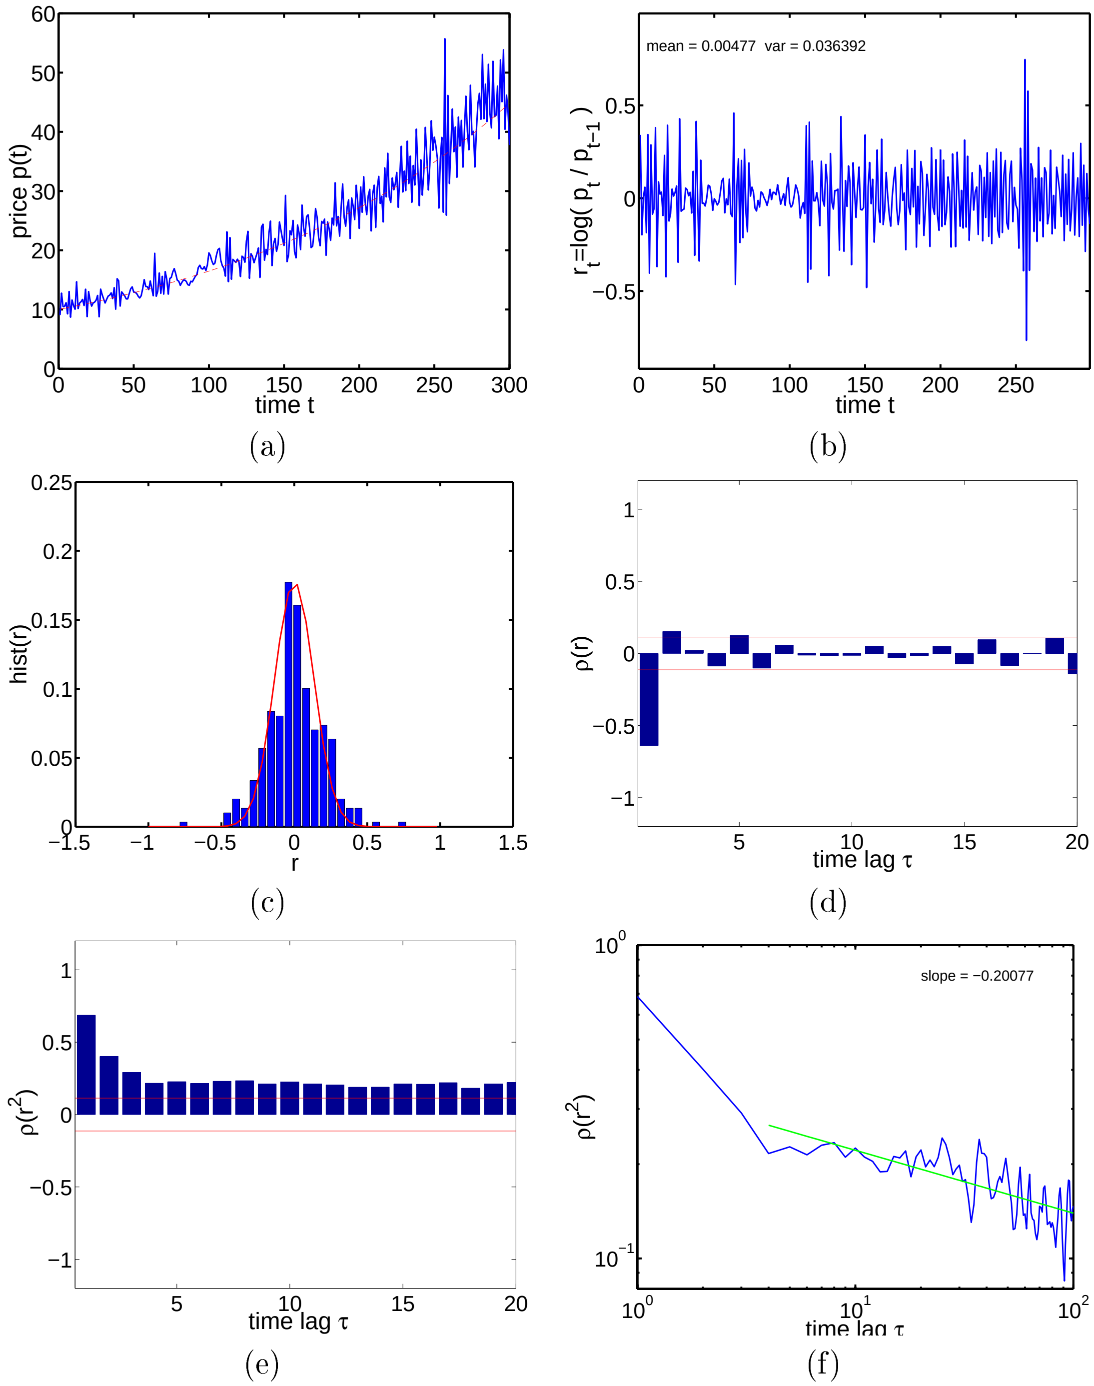

Figure 1: (a) Model price series for 300 time steps, following the fundamental value $F(t) = S_0 \exp(rt)$, (dashed line). Parameter values: $S_0 = 10$, $r = 0.005$, $P(0) = 10$, $\xi(0) = -1$, $\zeta(0) = 1$, $\beta = 0.2$, $\gamma = \delta = 10$, $\varepsilon = 0.5$. (b) Corresponding log-returns $r(t)$, (c) normalized histogram of the return process $r(t)$, (d) autocorrelation function of log-returns (on a 5% significance level, coefficients above 0.11 are significant). (e), (f) Autocorrelation function of squared log-returns $\rho(r^2)$ in linear and log-log scaling. As in real time series it shows a power law decay in about the right order.

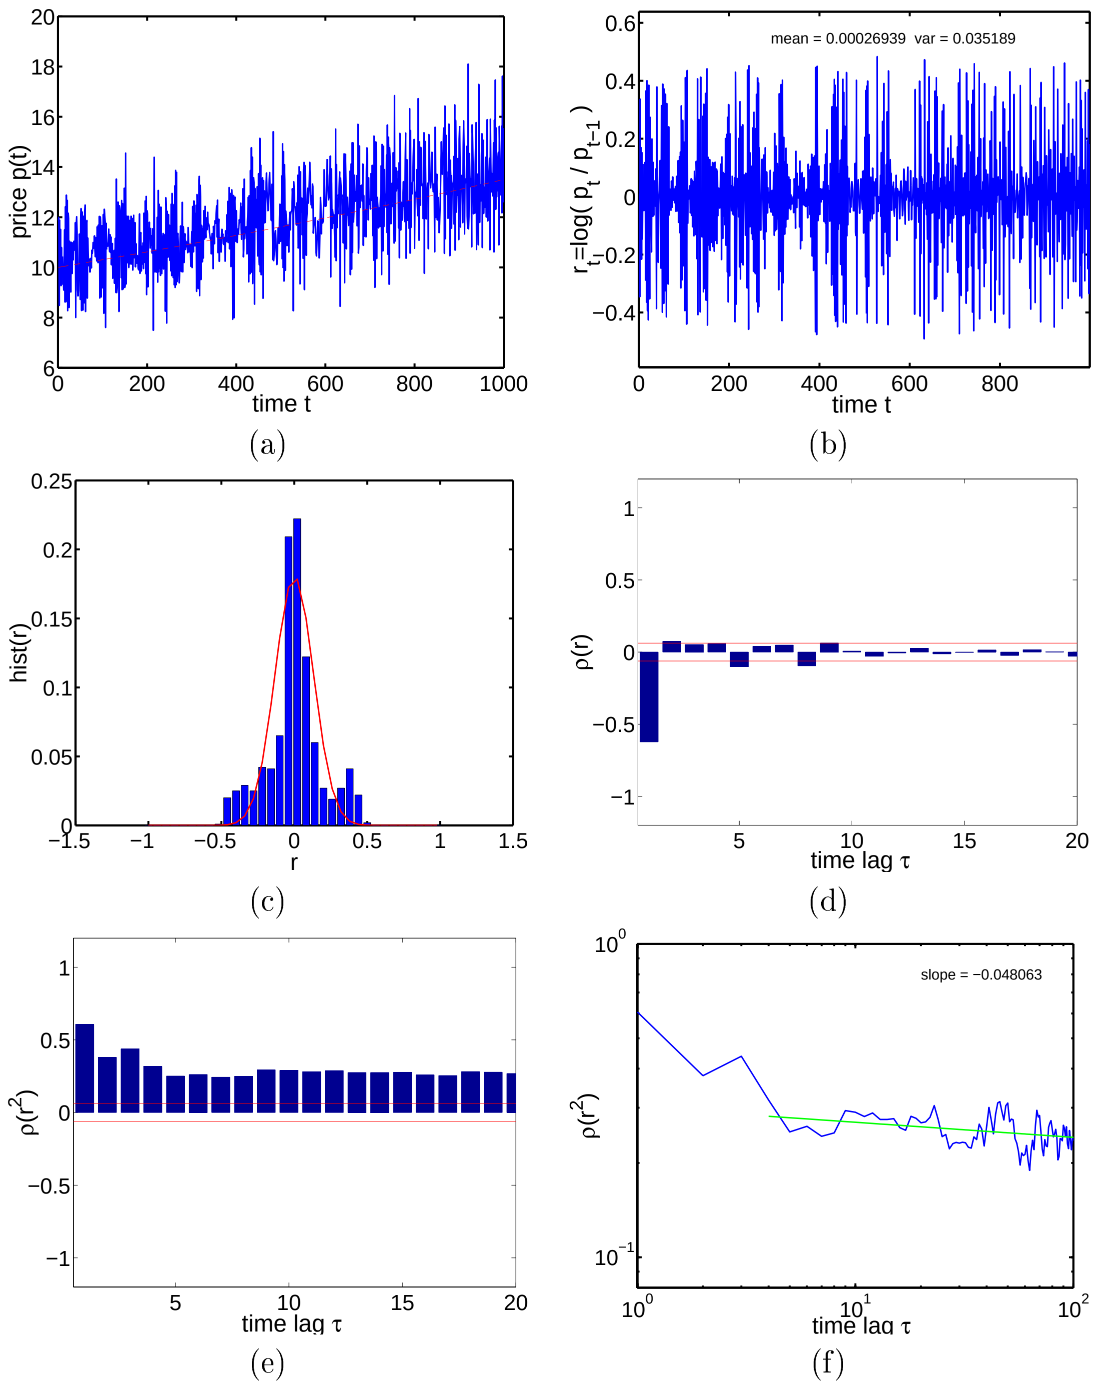

Figure 2: (a) Model price series for for 1000 time steps and parameters: $S_0 = 10$, $r = 0.005$, $P(0) = 1$, $\xi(0) = -1$, $\zeta(0) = 1$, $\beta = 0.20$, $\gamma = 12$, $\delta = 12$, $\varepsilon = 0.6$, (b) corresponding log-returns $r(t)$, (c) normalized histogram of the return process $r(t)$, (d) autocorrelation function of log-returns, (e), (f) autocorrelation function of squared log-returns.

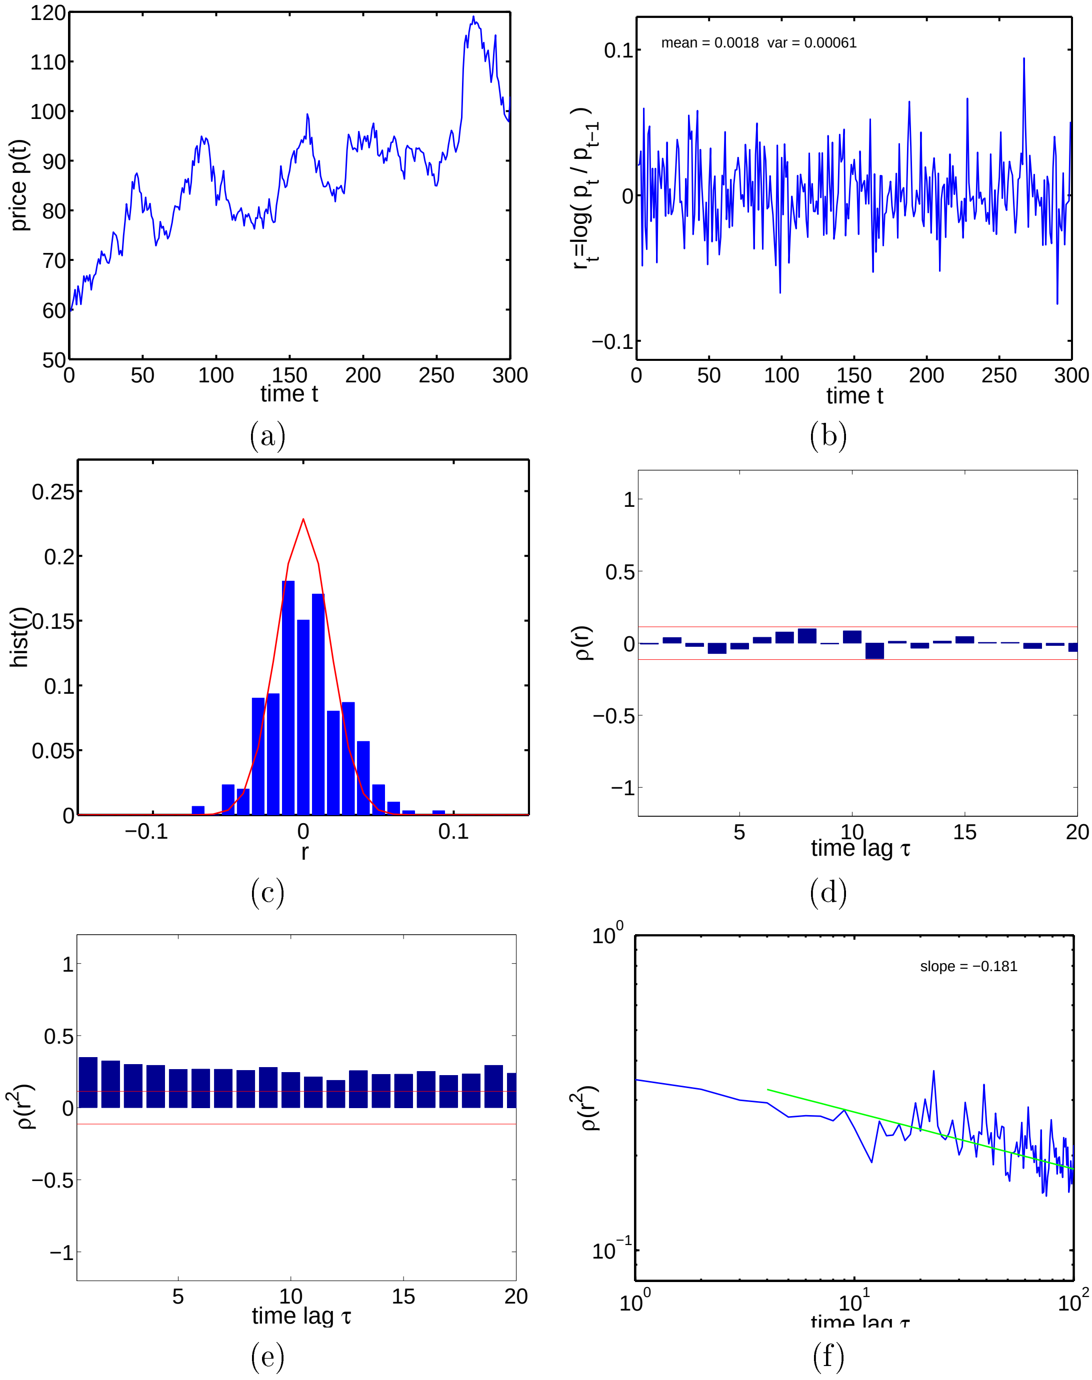

Figure 3: (a) Prices series MicroSoft for 11/20/1998–01/31/2000, (b) returns, (c) distribution of returns, (d) autocorrelation of returns, (e), (f) autocorrelation of squared returns.

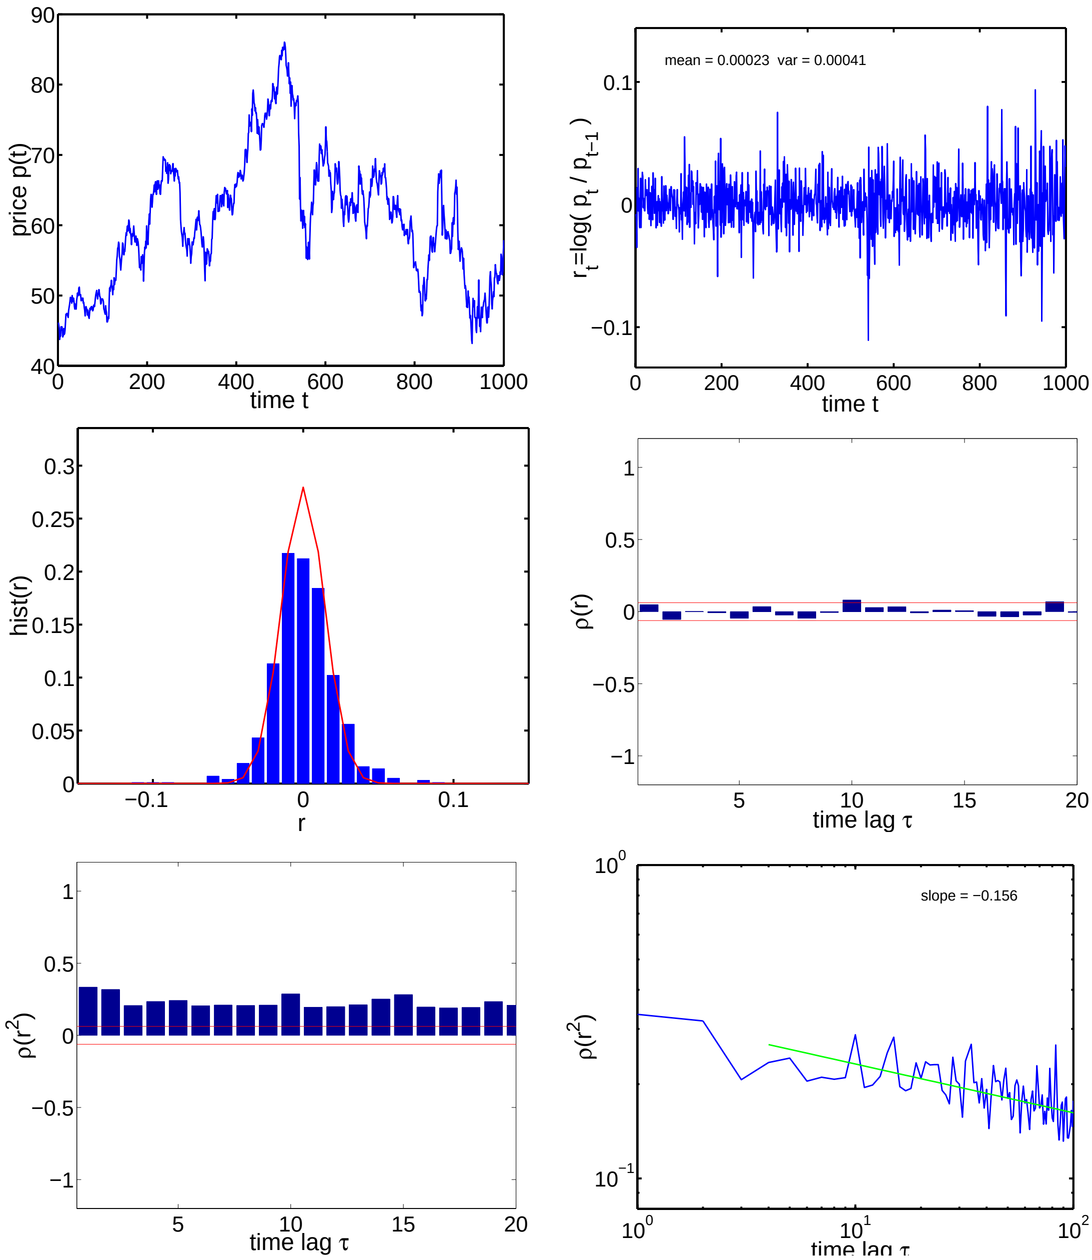

Figure 4: (a) Prices series Coca Cola for 07/09/1996–06/23/2000, (b) returns, (c) distribution of returns, (d) autocorrelation of returns, (e), (f) autocorrelation of squared returns.

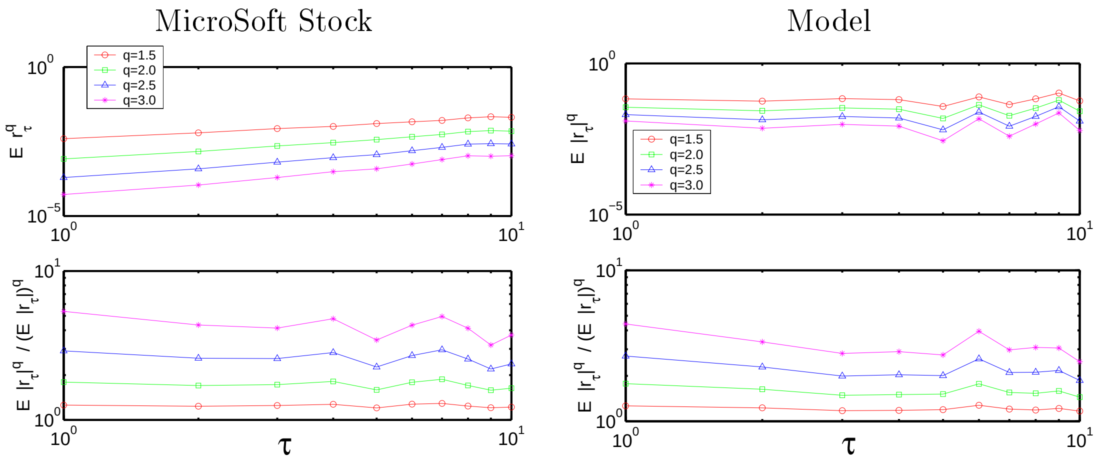

Figure 5: Comparison of higher moments of real (left) and model data from the example of Figure 1 (right). Top: Higher moments of the return process are scaling processes with respect to aggregation time τ. Bottom: $E|r_\tau|^q/(E|r_\tau|)^q$ is not a constant, indicating a multifractal underlying return process.