Financial Power Laws: Empirical Evidence, Models, and Mechanism

Thomas Lux[1]

Department of Economics
University of Kiel
lux@bwl.uni-kiel.de

Prepared for *Power Laws in the Social Sciences: Discovering Complexity and Non-Equilibrium Dynamics in the Social Universe*, Claudio Cioffi, editor.

Summary. Financial markets (share markets, foreign exchange markets and others) are all characterized by a number of universal power laws. The most prominent example is the ubiquitous finding of a robust, approximately cubic power law characterizing the distribution of large returns. A similarly robust feature is long-range dependence in volatility (i.e., hyperbolic decline of its autocorrelation function). The recent literature adds temporal scaling of trading volume and multi-scaling of higher moments of returns. Increasing awareness of these properties has recently spurred attempts at theoretical explanations of the emergence of these key characteristics form the market process. In principle, different types of dynamic processes could be responsible for these power-laws. Examples to be found in the economics literature include multiplicative stochastic processes as well as dynamic processes with multiple equilibria. Though both types of dynamics are characterized by intermittent behavior which occasionally generates large bursts of activity, they can be based on fundamentally different perceptions of the trading process. The present chapter reviews both the analytical background of the power laws emerging from the above data generating mechanism as well as pertinent models proposed in the economics literature.

I. Introduction

While research on power laws in income and wealth dates back to the nineteenth century (Pareto), the attention on power laws in financial data is relatively recent. The first ever manifestation of power laws in finance can probably be found in Mandelbrot's seminal "Variation of Certain Speculative Prices" published in the 1963 volume of the *Journal of Business* and followed by Eugene Fama's elaboration on "Mandelbrot and the Stable Paretian Hypothesis" published as the immediately succeeding paper in the same issue. This breakthrough very much dominated the discussion over the next thirty

(!) years or so with an immense number of papers dedicated to providing supporting or contradicting evidence for the Paretian or Levy stable hypothesis. While the dust has settled over the last decade and the power-law behavior of large price changes now counts as one of the most pervasive findings in financial economics, it had remained the only power law under discussion in this arena for quite some time.

Only recently was it joined by other candidates for Pareto-like behavior. By now well accepted within the scientific community is a second power law characterizing the temporal dependence structure of volatility. However one tries to proxy the unobservable quantity 'volatility' (most straightforwardly via the squares or absolute values of financial returns), the autocorrelations of these entities appear to decay hyperbolically, i.e. Pareto-like. Although this feature is linked to the long known clustering of volatility in financial markets, the fact that the dependency in the fluctuations is of a long-range type had only been realized in the nineties. Credit for this observation is probably due to Ding, Engle and Granger's paper "A Long Memory Property of Stock Market Returns and a New Model" which appeared in the *Journal of Empirical Finance* of 1993. Later on, several papers by physicists emphasized the power law nature of this finding and its potential root in complex market interactions (cf. Lux and Ausloos, 2002). The power law in returns and in volatility seem to be intimately related: none of them was ever observed without the other and it, therefore, seems warranted to interpret them as the joint essential characteristics of financial data.

Very recently additional power laws have entered the scene: transaction volume (which is strongly correlated to volatility) also appears to be characterized by long-range dependence (although it is not clear whether volatility and volume share the same degree of long memory). Availability of high-frequency tick-by-tick data has furthermore revealed other types of power-law behavior, such as a power law for the number of trades in the New York Stock Exchange Trades and Quotes Database (Gopikrishnan et al., 2001). Similar results are reported for the Japanese stock market.

The plan of the remainder of this chapter is the following: sec. II gives a more formal description of the main financial power laws characterizing returns and volatility together with a survey of pertinent literature. After having set the scene, we turn to explanatory models. Sec. III deals with the so-called rational bubble model which emerged as a potential explanation of financial power laws from the standard body of rational expectations models in economics. Interestingly, this approach points to multiplicative stochastic processes as a type of data generating process with generic power-laws. This interesting property of the underlying process notwithstanding, the rational bubble model makes grossly incorrect numerical predictions about the magnitude of the exponent. In sec. IV we, therefore, turn to more recently proposed models in the behavioral finance literature. From the diversity of available approaches and models, we try to single out the basic ingredients and mechanisms leading to true or at least apparent power laws in simulated data. Sec.

V attempts to draw some overall conclusions from the hitherto available body of literature on potential explanations of financial scaling laws.

II. Empirical Power Laws in Finance

The modern literature in this area starts with Mandelbrot (1963) and Fama (1963), who both proposed the so-called Paretian or Levy stable distributions as statistical models for financial returns[1] (cotton futures were analyzed in Mandelbrot's paper). The theoretical appeal of this family of distributions is its stability under aggregation. At the time of publication of these papers, it had already been known for some time that a Generalized Central Limit Law holds for distributions with non convergent (infinite) second moments: while existence of the second moment warrants convergence of sums of random variables (at least in the IID case and under weak dependence) towards the Gaussian, non-convergence of the variance implies convergence of the distribution of sums towards members of the family of Levy stable distributions. Under this perspective, the pronounced deviation of histograms for financial returns from the shape of the Normal distribution together with their apparent additivity (daily returns can be expressed as the sum of all intra-daily price changes) was interpreted as striking evidence in favor of the Levy hypothesis. The Levy distributions are characterized by an asymptotic power-law behavior of their tails with an index α (called the characteristic exponent) which implies a complementary cumulative density function of returns (denoted by *ret* in the following) which in the tails converges to:

$$Pr\left(|ret| > x\right) \approx x^{-\alpha} \tag{1}$$

The Levy hypothesis restricts the power-law for returns to the admissible range of $\alpha \in (0, 2]$ which indicates the mentioned non-convergence of the second moment (with $\alpha < 1$ not even the mean would converge). Empirical estimates based upon the Levy model typically found α hovering around 1.7.

While this result was confirmed again and again when the parameters of the Levy laws were estimated themselves, other studies raised doubts in the validity of the Levy hypothesis by questioning the stability-under-aggregation property of these estimates (Hall et al., 1989) or pointed to apparent convergence of sample second moments (Lau et al., 1990). From the early nineties, however, it became common practice to concentrate on the tail behavior of

[1] The quantity of interest in empirical research in financial economics is typically 'returns' defined as relative (or logarithmic) price changes over a certain time horizon. Research on the statistical properties of returns started with data at weekly or monthly frequencies but has moved on to high frequency data over time (daily and intra-daily data up to the highest frequencies at which all tick-by-tick changes are recorded).

the distribution itself and estimate its decay parameter via conditional maximum likelihood without assuming a particular distributional model (cf. Hill, 1975).[2] The pertinent literature gradually converged to the insight of an exponent significantly larger than 2 and mostly close to 3 (cf. Jansen and de Vries, 1991; Lux, 1996). These results nicely agree with estimates obtained by physicists via their typical log-log regression approach (Cont et al., 1997; Gopikrishnan et al., 1998). The approximate cubic form of the power-law of returns is by now accepted as a universal feature of practically all types of financial markets (from share markets and futures to foreign exchange and precious metal markets). Note that this finding implies rejection of the time-honored Levy hypothesis as $\alpha \approx 3$ means that the decay of the outer part of the distribution is faster than allowed by this family of distributions.

The second type of power laws relates to the opalescent concept of volatility of price fluctuations. Focusing on absolute returns, $|ret|$, as one of its most frequently analyzed manifestations, the pertinent power law applies to their autocovariance function:

$$E\left[|ret_t| \cdot |ret_{t-\Delta t}|\right] \approx \Delta t^{-\gamma} \tag{2}$$

Although the precise value of γ has received less publicity than that of α (maybe because it is not estimated directly but rather via its relation to the so-called Hurst exponent or related measures, cf Lux and Ausloos, 2002), reported statistics are also remarkably uniform across time series with typical values around $\gamma = 0.2$ - 0.3 (Ding, Engle and Granger, 1993; Lobato and Savin, 1998; Vandewalle and Ausloos, 1997, 1998). Fig. 1 provides an example for the typical behavior of financial data, using a large series of daily observations for the New York Stock Exchange Corporate Index. The somewhat low decay parameter $\gamma = 0.14$ for volatility is a consequence of fitting the whole ensemble of autocorrelations up to lag 200 by one single power law. Concentrating on longer lags we would have obtained a larger decay parameter in line with other estimates.

It is worth mentioning that recently mounting evidence speaks in favor of *multi-scaling* in the temporal dependence structure of financial fluctuations: rather than the simple scaling (power) law (2) we might indeed face a continuum of scaling laws for various powers of absolute returns:

$$E\left[|ret_t|^q \cdot |ret_{t-\Delta t}|^q\right] \approx \Delta t^{-\gamma(q)} \tag{3}$$

Note that any power q of absolute returns can be interpreted as an alternative measure of volatility. With the multi-scaling in eq. (3) we, thus, get a much more detailed picture of the temporal development of financial fluctuations (cf. Mandelbrot, 1999; Calvet and Fisher, 2002; Lux, 2004). Most excitingly, a non-linear dependence of the scaling parameter γ on the power q

[2] See chapter 3 for details on maximum likelihood estimation of limiting Pareto distributions and the underlying concepts from statistical extreme value theory.

is also a key characteristic of turbulent fluids and has motivated the development of so-called multi-fractal models in statistical physics. Recent research shows that these models may have some potential for modeling financial prices.

It is worth emphasizing that the power-law behavior of large returns and their fluctuations seem to be truly universal and can be found without exception in all financial data. This is quite in contrast to many other 'stylized facts' in economics, for example concerning macroeconomic data (such as GDP, inflation rates etc.). It should also be pointed out that the above power-laws are not at all esoteric concepts. Quite to the contrary, they are of tantamount practical importance in financial engineering: the probabilistic law governing large returns (eq. 1) can be applied directly for an assessment of the inherent risk of extreme events (i.e., crashes). Similarly, models covering the temporal dependence of volatility depicted in eqs. (2) and (3) are of immediate practical use in predicting the future extend of price fluctuations.

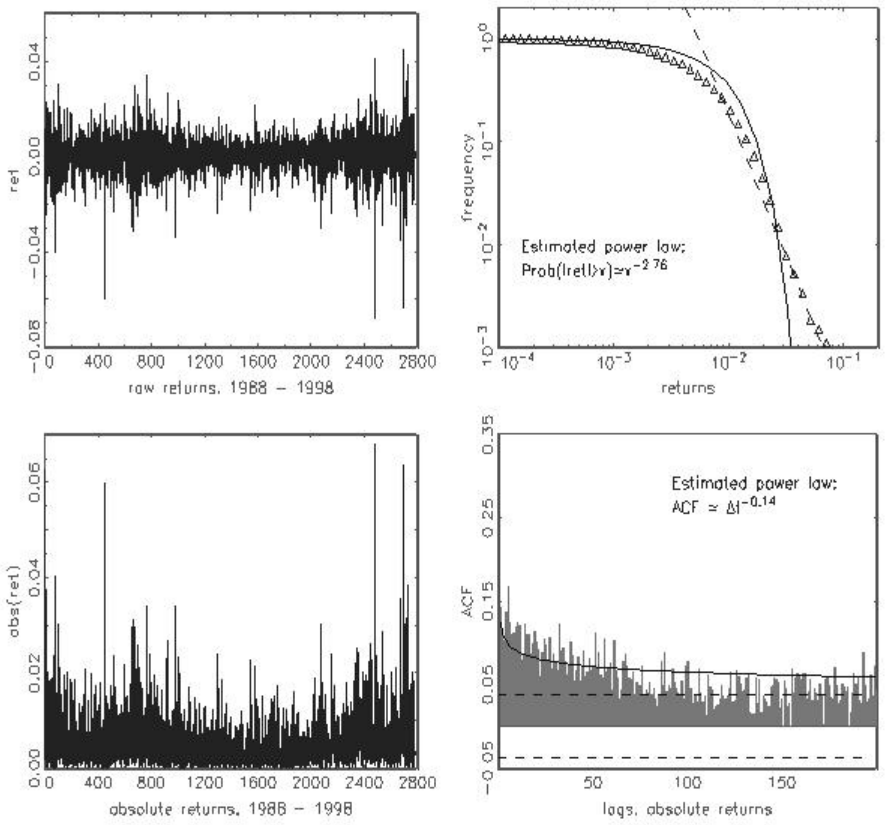

Fig. 1. Illustration of the main power law characteristics of financial data. Our data are daily recordings of the New York Stock Exchange Composite Index over the period 1988 to 1998 (a total of 2782 observations). Instead of the index itself, we show its returns defined as the log increments between adjacent trading days (upper left panel). These already show the two main 'stylized facts': a large number of extreme positive and negative realizations, and the clustering of volatility, i.e. characteristic switches between turbulent and more tranquil episodes. The power law nature of the first feature is underlined by inspecting the log histogram of (absolute) returns (upper right hand). In contrast to a Normal distribution with the same variance (solid line), the triangles for the empirical distribution show a prolongation in the tail which in its outer part has an almost linear shape in agreement with eq. (1). The lower panels show absolute returns as a straightforward measure of volatility (left) and their autocorrelations (right). The ACF decays extremely slowly and even after 200 lags has significantly positive entries (95% values are demarcated by broken lines). This hyperbolic rather than exponential decay is in line with eq. (2).

Some recent research points to additional power laws in the financial arena: Gopikrishnan et al. (2001) find that high-frequency data from the U.S. stock market exhibit the following regularities: trading volume (V) behaves like:

$$Pr\,(V > x) \approx x^{-1.5}, \tag{4}$$

and the number of trades (N) follows a law:

$$Pr\,(N > x) \approx x^{-3.4}, \tag{5}$$

Additional results on the distribution of inter-transaction times and related quantities can be found in several studies (cf. Takayasu, 2002). So far, these additional regularities have mainly been discussed in the physics literature but have hardly been acknowledged by economists. Because of the limited number of available studies, it is also not clear at present whether these findings are of a similarly universal structure as (1) to (3). The only exception here is temporal dependence of transaction volume for which substantial statistical evidence has already been gathered and for which pretty much the same pattern has been found as for absolute returns (Lobato and Velasco, 2000).

III. Power Laws in Orthodox Financial Economics: Stochastic Approaches and Rational Bubbles

Standard textbooks on theoretical and empirical finance (see e.g. O'Hara, 1995, for a comprehensive treatment of behavioral models and Campbell, Lo and MacKinley, 1997, for a similarly comprehensive survey of empirical techniques) lack explicit entries on the power-law behavior of financial data. It is only via stochastic processes with asymptotic power law behavior that they implicitly take into account the existence of the universal scaling laws highlighted in sec. II. Until very recently, financial power laws have, therefore, only been taken into account under a purely statistical perspective. The hallmark of this literature is the (G)ARCH (Generalized Autoregressive Conditional Heteroskedasticity) class of processes introduced in Engle's seminal (1983) paper. GARCH essentially models returns as a random process with a time-varying variance which shows autoregressive dependence, i.e.

$$\begin{aligned} ret_t &= \sigma_t \varepsilon_t \\ \sigma_t^2 &= \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \beta_1 \sigma_{t-1}^2, \end{aligned} \tag{6}$$

with $\varepsilon_t \sim N(0,1)$. As this type of auto-correlation is readily apparent in any plot of a financial time series (cf. Fig. 1), it is not too surprising that GARCH captures the short-run dynamics of volatility quite well. Implying exponential rather than hyperbolic decay of the volatility autocorrelations it nevertheless

falls short of providing a stochastic process in accordance with the long memory property depicted in eq. (2). However, even with a Gaussian distribution of the increments, the compound unconditional distribution resulting from (6) produces fat tails of the unconditional distribution and hence is in accordance with eq. (1). This relatively simple statistical recipe, therefore, already allows to reproduce *one* of the universal laws of returns. Refinements of the GARCH approach have, in fact, also covered the scaling in the autocovariances via imposition of an infinite number of lags with hyperbolically decaying weights in the difference equation governing the volatility dynamics (cf. the fractionally integrated GARCH model of Baillie et al., 1997).

Although these statistical models and the myriads of variations on this topic which have come out in the literature are important tools in financial engineering, they do not provide an avenue towards an *explanation* of the empirical regularities. In fact, until very recently, standard models in the theoretical literature were unable to explain even the phenomenological aspects covered by the GARCH models, let alone the asymptotic laws in eqs. (2) and (3) with their perplexingly precise numerical manifestation in the data. To be honest, these regularities were neither well-known among financial theorists nor did their standard models provide an easy avenue towards explanations of such power laws.

There is one important exception though in that one ingredient of the received body of models, in fact, produces power law statistics as an immediate consequence of its underlying model structure, which, however, has also been realized only very recently. This class of models is known as models of *speculative bubbles with rational expectations* (RE bubbles). This theory attempts to explain the rather obvious frequent deviation of market prices from its underlying fundamental value (also known as the 'intrinsic' value of an asset which is determined by current expectations about future earning prospects) without sacrificing the 'rationality' assumption of traditional asset pricing models.

Before proceeding to RE bubbles, let us introduce the standard asset pricing model of the textbook Efficient Market Paradigm. The starting point of this approach is the formula for *fair* or *arbitrage-free* valuation of an asset:

$$p_t = \delta E[p_{t+1} + d_{t+1}|I_t] \tag{7}$$

with p_t the price at time t, $\delta < 1$ the discount factor reflecting the time preference of agents, and d_t denoting dividends at time t. Eq. (6) says that a fair and arbitrage-free price should be identical to the discounted expected value (conditional on the current information set I_t) of next period's price plus the dividend paid out in that period. With identical expectations of all agents, the equilibrium price should converge to this benchmark. Otherwise, agents would sell/buy as long as an inequality prevails between the right-hand side and left-hand side of eq. (7). Imposing a so-called 'transversality condition',

$lim_{i\to\infty}\delta^i E[p_{t+i}|I_t] = 0$ one can replace p_{t+1} by the pertinent arbitrage-free pricing equation at period $t+1$ (of course, eq. (7) has to hold in all periods). Continuing along these lines by further substitution of the pertinent equations at $t+2$, $t+3$, $\cdots$, one obtains:

$$p_t = p_{f,t} = \sum_{i=1}^{\infty} \delta^i E[d_{t+i}|I_t] \tag{8}$$

Eq. (8) postulates that the price should always equal the discounted expected future stream of dividends which is what is also often called the fundamental value of the asset, $p_{f,t}$. This pricing formula can be seen as a manifestation of the *Efficient Market Paradigm* postulating the prices reflect all available information about the fundamental factors of the underlying asset in an unbiased manner and immediately react to forthcoming new information about these fundamentals. It is worthwhile to note in passing that the traditional view of the Efficient Market Hypothesis is not necessarily inconsistent with observations of power laws in financial data since it remains agnostic about the structure of the 'news process' driving returns. However, it would have to attribute their origin entirely to exogenous factors: all power laws would have to be explained by fundamental valuation factors, $E[d_{t+i}|I_t]$, exhibiting these same characteristics which are then reproduced by price changes reflecting changes in fundamentals. Unfortunately, 'fundamentals' are essentially unobservable and cover a bundle of diverse factors such as firm-specific events, political influences etc.

As pointed out by Blanchard and Watson (1982) it requires only a minor modification of the above assumption, to allow for deviations of prices from fundamental values. Namely, though mathematically convenient, the 'transversality condition' needed to proceed from (7) to (8) is by no means necessary from a theoretical perspective. Dropping it, however, opens a Pandora's box of possible asset price paths deviating from fundamental valuation. Blanchard and Watson (1982) first remarked that without the transversality condition, eq. (7) does *not* necessarily exclude per se such weird phenomena like speculative overvaluation and subsequent crashes. Quite the opposite, it allows the price to contain a *bubble* component B_t by which it deviates from the fundamental value $p_{f,t}$:

$$p_t = p_{f,t} + B_t \tag{9}$$

Earlier literature had avoided this possibility by assuming $B_t = 0$. Nevertheless, this framework is still far away from an agnostic view of "anything goes" in financial markets. Rather, the rational expectations assumption in eq. (7) still poses heavy restrictions on admissible bubble dynamics. In particular, only those bubbles are allowed which satisfy:

$$B_t = \delta E[B_{t+1}|I_t] \tag{10}$$

A fairly general class of processes obeying (10) can be written as:

$$B_t = a_t B_{t-1} + \varepsilon_t, \tag{11}$$

with $a_t \in A = \{ a_1, a_2, \cdots, a_n \}$ occurring with probabilities $\pi_1, \pi_2, ..., \pi_n$ and ε_t IID with mean zero. The only additional restriction which the a_i have to meet is $E[a_t] = \sum_{i=1}^{n} \pi_i a_i = \frac{1}{\delta}$ in order to be in harmony with (10). With both a_t and ε_t stochastic variables, the bubble process (11) is a so-called *multiplicative stochastic process*. This kind of random difference equation had already been studied comprehensively in Kesten (1975) who showed that *multiplicative processes are generic power-law generators.* [3] He also showed that the power-law exponent for the unconditional distribution of the underlying dynamic variable can be precisely determined from the distribution of the multiplicative component. In our notation, the power-law applies to deviations from the fundamental value $(B_t) : Prob(|B_t| > x) \sim x^{-\mu}$ with μ given by:

$$E[|a_t|^{\mu}] = 1 \tag{12}$$

Figs. 2a and b give an illustration of the resulting time series and the distribution of B_t which, of course, agrees with the theoretical result. Lux and Sornette (2002) show that the power-law in the bubble component carries over directly to price changes and also dominates the distribution of returns.

What generates the power-law tails of the bubble model? Since the mean value of a_t is $1/\delta > 1$, the set A has to include at least one element larger than unity (typically, the smallest a_i would be set to zero to allow for the possibility of a total crash of the bubble). Hence, the realisations of a_i will switch between values smaller and larger than one. In a deterministic setting, the former would guarantee convergence towards $B_\infty = 0$ while the later would yield an explosive dynamics (which would ensue also if we only allow for one single $a_t = 1/\delta > 1$). The time-variation of a_t, then, generates an intermittent amplification of fluctuations which, however, only continues as long as it is not interrupted by realizations of $a_t < 1$ (cf. Fig. 2 for an example). It is remarkable that, in general, the additive noise component, ε_t, has no influence at all on the emerging power-law exponent μ which is fully determined by the distribution of a_t. This means that the resulting exponent μ, in fact, only reflects the *intrinsic* dynamics of speculative bubbles (the way B_t depends on B_{t-1}) while in a sense external factors affecting ε_t in B_t as well as the fundamental factors affecting the component $p_{f,t}$ in the asset price given by eq. (9) have only a very subordinate role.

[3] This is true under very mild conditions on the structure of the random difference equation. A glance at the additional conditions stated in Kesten's theorem, in fact, shows that they will only be violated by very particular processes, cf. also Lux and Sornette (2002).

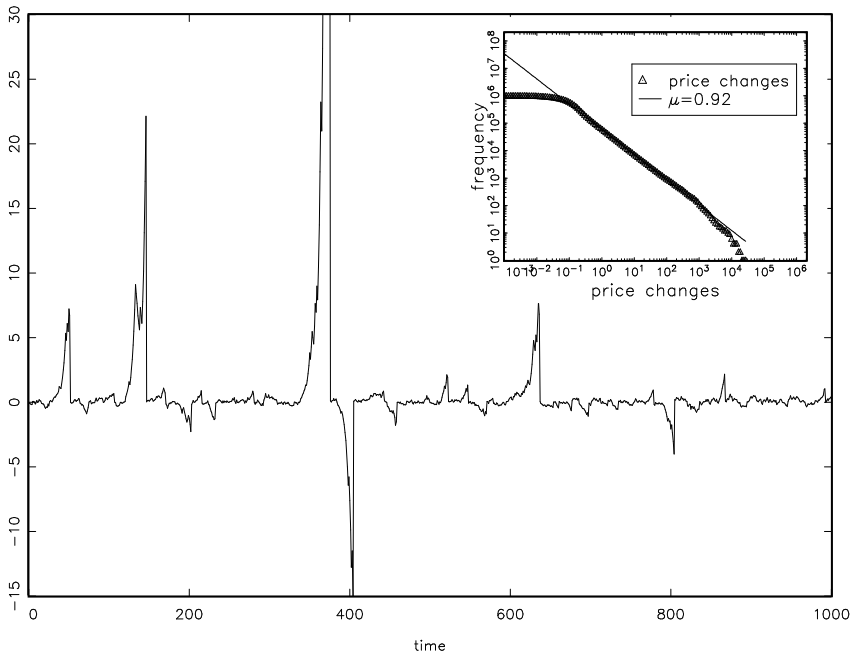

Fig. 2. Example of a RE bubble process with $A = \{0, 0.9, 1.3\}$ occurring with probabilities $\pi(0) = 0.1, \pi(0.9) = 0.4$ and $\pi(1.3) = 0.5$, and additive stochastic component ε_t following a Normal distribution with mean zero and standard deviation 0.1. The simulation illustrates how this framework generates repeated outbreaks of bubbles (periods of deviations form the fundamental values) which all end in crashes leading to a reversal of the price to its fundamental value. The inlet shows a loglog plot of the inverse of the cumulative distribution function of price changes from the same RE bubble process shown in Fig. 1. The theoretical (exact!) power-law index can be calculated from eq. (12). With a simulation over 10^6 time steps, the resulting $\mu = 0.92$ (for a discount rate $\delta = 0.99$) can be nicely observed over at least four orders of magnitude. The slight deviation for very large values is due to the scarcity of observations in the very far tails.

Eq. (11) with arbitrary distributions of its multiplicative and additive components, a_t and ε_t, is a multiplicative stochastic difference equation in its most general form. The only restriction on multiplicative processes that could be used as rational bubble processes in a finance setting is, therefore, the restriction stemming for the assumption of rationality or consistency of expectations (i.e., all agents have correct expectations about both the future development of the bubble as well as the fundemental conponent of the asset price): $E[a_t] = 1/\delta > 1$. Unfortunately, as pointed out by Lux and Sornette (2002), it is exactly this requirement which restricts the admissible range of

outcomes to $\mu < 1$ - far off from the empirical findings of an exponent around three. The important consequence is that this simple and appealing theory of rational bubbles rather than being in harmony with empirical power-laws produces results that are at odds with empirical observations. Since this result applies to the entire class of RE bubble models, it seems that economists have to accept deviations from the ideal case of perfect rationality in order to explain fat tails and clustered volatility. Furthermore, the assumption of 'rational expectations' in economic models has provoked mounting criticism as it requires a capacity of computation and information processing of agents far beyond imagination and immense efforts have been devoted recently to introduce boundedly rational behavior into economic models (cf Brenner, 1999).

IV. Multi-Agent Models in Behavioral Finance

IV.1 Overview

Although the rational bubble theory can be viewed as a generic power-law generator, it might appear uncomforting in more than one way. First, as we have seen, it produces power-laws which are definitely at odds with the empirical ones - and hence the perspective of analyzing scaling behavior, in fact, shows the limitations of this approach. Second, it shares the conceptual problems of economic models with 'fully' rational agents. As Mark Buchanan puts it: "To anyone who is not an economist, the orthodox perspective that sees people as rational agents who always work out their rational self-interests and act on them, seems more than a little peculiar" (Buchanan, 2000). Clearly, the 'full rationality' assumption appears much too strong a statement in the eyes of the layman and seems to be in so striking contrast to observed real-life behavior to let it appear almost ridiculous for those who have not been brought up in this tradition. Although there might be good reasons to use the rationality postulate when tackling certain questions[4] and although the question of rational vs. non-rational expectations might be of minor concern for various economic problems, modeling economic phenomena governed by not-fully-rational behavior requires a different approach. This is surely the case with financial markets for which it has been observed that their immense trading volume already raises doubts about common rationality and knowledge of such rationality on the part of other agents (Leroy, 1989). Survey studies confirm that market participants themselves contribute a large portion of price fluctuations to bandwagon effects, overreaction and speculative dynamics (cf. Cheung and Wong, 2000). The recently mounting literature on experimental asset markets adds creditability to this view by showing that

[4]Think of analyses of future effects of fiscal or monetary policy measures: a rational expectations framework allows to single out their purely intrinsic effects without mingling them with problems of misperceptions of agents of the economic environment.

human subjects typically produce price bubbles and crashes even in simple laboratory markets (cf. Noussair et al., 2001).

From the power-law perspective, the universal cubic law of price returns and the long-range correlations of volatility suggest to view financial data as results of a social process of interacting agents. Models founded on rational expectations including the RE bubbles theory contain nothing of that sort: agents themselves are typically invisible and enter via some equilibrium condition like eg. (7). Of course, one could add a description of the typical behavior of a rational agent in this framework (the so-called representative agent) but this approach excludes any interaction of traders and by its focus on steady state solutions does not provide an avenue towards replicating distributional properties of financial data. More recently, therefore, a rapidly growing research area has purposely allowed for deviations from 'rational' behavior formalizing financial markets as evolving ecologies of multi-agent populations of traders with diverse strategies and expectations.

Due to the obvious diversity of trading motives and strategies in real markets, such an approach had always been virtually existant: Keynes' famous beauty contest parabola of the stock market (Keynes, 1936) undoubtedly presupposes heterogeneous (and hence, non-rational) expectations. Early formal models of interacting groups of traders can be found in Baumol (1957) and Zeeman (1974). As has become standard in most of the subsequent literature, these authors distinguish between two types of traders: the first, so-called fundamentalists, views asset prices as being determined by fundamental factors alone. These traders would buy/sell if they consider the current market price to be below/above the rationally computed fundamental value (i.e. eq. 8). The second group, mostly called chartists or noise traders, would rather be convinced that asset markets are driven by systematic trends and that patterns exist that could be extracted by means of regressions, moving averages or more complex procedures. Typically a short-cut representative rule (like trend-following behaviour) is used to introduce this second component into the model. The market price, then, results from the interplay of the two groups and their respective demand and supply.

Hibernating over the highdays of the rational expectations paradigm, the recent literature has seen an enormous surge of work along the lines of these early pioneers.[5] Another limitation had, however, to be overcome to tackle the stylized facts of sec. 2. Namely, the interest in economic models has typically been to trace out the effect of changes of one (exogenous) economic variable on other, endogenous variables characterizing some sort of market equilibrium. Even with a diverse ensemble of traders, such a comparative static

[5]In fact, the boundaries between rational and non-rational behavioural approaches have become extremely fuzzy. For example, an interesting recent strain of literature analyses the implications of overconfidence of some traders and finds that this kind of misperception might even provide an evolutionary advantage, cf. Hirshleifer and Guo, (2001).

approach is ill suited for explaining distributional characteristics of the data. One, therefore, has to go beyond static models and beyond linear dynamics to account for power-law phenomena as important overall features characterizing financial time series. First attempts at analyses of complex data generating processes with behavioral foundations appeared in the early nineties: leading examples are Kirman (1991, 1993) and De Grauwe et al. (1993) who propose complex models of interacting speculators and study their overall statistical characteristics. While they did not focus on power-laws proper (then still not broadly acknowledged), they both showed that their models could mimic the random walk nature of financial prices although their data generating processes were both clearly different from a random walk. Notably, both papers study what we might describe as secondary stylized facts, typical results of certain - unsuccessful - tests of hypothesis of exchange rate formation, and find results similar to those obtained with empirical records.

Evidence for volatility clustering as an emergent property of a multi-agent model appeared first in Grannan and Swindle (1994). Ramsey (1996) offers a rather general perspective of how a statistical description of agents' behaviour could give rise to time-varying moments as emergent macroscopic characteristics of a market. Simultaneously, first attempts appeared at designing market models with heterogeneous autonomous agents who use artificial intelligence techniques (genetic algorithms, classifier systems, neural networks) as expectation formation mechanisms. Much of the early literature in this vein is preoccupied with analysis of convergence or not of the learning process of agents to homogeneous rational expectations equilibria (Arthur et al., 1997; Chen and Yeh, 1997; Arifovic, 1996). Only recently has explanation of empirical characteristics become a topic in this strand of literature (LeBaron et al., 1999; Chen and Yeh, 2002; Lux and Schornstein, 2004). Typically, some tendency towards fat tails and volatility clustering is observed, although numerical results are often far from empirical power laws.[6]

Another current emerging in the early nineties is microscopic models of financial markets constructed along the lines of multi-particle systems in statistical physics. The first example here is Takayasu et al. (1992). Broadly similar approaches are Levy, Levy and Solomon (1994, 1995), Bak, Paczuski and Shubik (1997) and Cont and Bouchaud (2000). Most of these contributions do either not focus on power-laws or generate scaling laws different from the empirical ones. Similarly, Sato and Takayasu (1998) derive a multiplicative stochastic difference equation as an approximation to the dynamics of a microscopic stock market model proposed by Takayasu et al. (1992). Their model consists of a fixed number of dealers who set limit prices at which they

[6]For example, in a comprehensive analysis of the seminal Santa Fe model of Arthur et. al., Wilpert (2003) observes that the decay of the auto-correlation of both volatility and volume is exponential rather than hyperbolic. Interestingly, results become 'better' if one includes typical elements of chartist/fundamentalist interaction.

are prepared to buy or sell stocks. These thresholds are then modified dynamically in reaction to past price trends. Trades occur when the difference between the maximum buy price and the minimum sell price is positive. The market price is recorded as the average between the marginal bid and ask. Since traders are assumed to follow the tendency of the market with their reservation prices, a self-amplifying mechanism comes into play. Its approximation via a stochastic difference equation leads to a structure similar to eq. (11). Without any further restriction, realistic exponents can be generated if the parameters of the model are appropriately chosen. Similar feedback in similar models brought about the same result in the framework of Bak et al. (1997). In the following, I review in more detail three dominant types of behavioral models which cover a large part of the hitherto available literature.

IV.2 Models Inspired By Statistical Physics

Cont and Bouchaud's (2000) model adopts a lattice-based percolation framework in which clusters of agents with synchronized behavior are formed. Allowing the clusters to enter on the demand or supply side of the market (with identical probability a) or to stay inactive (with probability *1-2a)* and adding a typical price adjustment equation for the reaction on imbalances between demand and supply one gets the following dependence of returns on the current configuration of clusters:

$$ret_t = \frac{p_t - p_{t-1}}{p_{t-1}} \sim \sum_{buy} n_s \cdot s - \sum_{sell} n_s \cdot s, \tag{13}$$

with s the cluster sizes and n_s the average number of clusters containing n sites on both the demand and supply side of the market. With this type of framework, it follows for small a that returns will inherit the well-known scaling law of the cluster size distribution which has been studied extensively in statistical physics. As Sornette, Stauffer and Takayasu (2002) put it in their sympathetic review of this approach: "This percolation model thus applies physics knowledge collected over decades, instead of inventing new models for financial fluctuations."

Despite being a known candidate for generating power-laws, the percolation model has its shortcomings: first, in the basic variant described here it is without memory and, hence, unable to produce volatility correlations. Second, quite the opposite of multiplicative random processes, the percolation model is *not* a generic power-law mechanism, it rather needs fine-tuning so that the probability for connection of lattice sites, say q, is close to the so-called *percolation threshold*, a critical value above which an infinite cluster (i.e. a cluster spanning the entire lattice) appears. What is more, the power-law emerging in Cont and Bouchaud's case of infinite connections between sites is unrealistic: it leads to a power-law index 1.5, some way apart from the "universal cubic law". Finite-size effects and variations of parameters could

generate alternative power-laws anywhere between the Levy and the Gaussian regime, but finding a cubic law would necessitate a particular choice of model design (e.g. q slightly above q_c, cf. Stauffer and Penna, 1998).

The well-known structure of percolation models in physics and wide-spread availability of pertinent simulation software has motivated a large number of variations on the Cont-Bouchaud topic. It has been shown that autocorrelation of volatility (and possibly of volume, too) can be obtained by sluggish evolution of cluster configurations over time (Stauffer et al., 1999). Similarly, if clusters dissolve or amalgamate after transactions, more realistic return distributions can be obtained (Eguiluz and Zimmermann, 2000). Additionally, fundamental values and feedback from market prices have been introduced to make the model more realistic (Chang and Stauffer, 1999). Focardi et al. (2002) consider latent connections which only become active in certain times, while Iori (2002) investigates an Ising-type lattice model with interaction between nearest neighbors, but without the group dynamics of the percolation approach. Still, one of the major drawbacks of this whole class of models is the extreme dependence of the resulting power laws on carefully adjusted parameters near criticality. Whether and how the market might self-organize towards these critical points remains unclear. Though sweeping a system back and forth through these critical states might yield an interesting perspective (Stauffer and Sornette, 1999), a behavioral underpinning of such a mechanism is missing.

Another recent model rooted in the statistical physics literature offers a very different avenue towards an explanation of financial power laws: Gabaix et al. (2003) provide a theoretical framework in which eqs. (1), (2), (4) and (5) are combined with the additional empirical observation of a Zipf law for the size distribution of mutual funds and a square root relationship between transaction volume and price changes ($\Delta p \approx V^{0.5}$). In this theory, scaling of price changes according to the cubic law (1) results from scaling of transactions which in turn is a consequence of the Zipfian size distribution of big investors. This explanation via multiple transmissions of power laws between various economic quantities is actually quite different form all other behavioural models. Since the ultimate source of financial power laws in Gabaix et al. is, therefore, the (exogeneous) Zipf distribution of large investors, this view is somewhat similar to the efficient market hypothesis (which attributes the power laws to the exogenous news arrival process), while all other models rather view these laws as an emergent manifestation of the intrinsic dynamics of speculative markets.

IV.3 Interaction Models of Financial Markets

One of the sources from which interacting agent models have developed is incorporation of herding and contagion phenomena into economic models. A first approach in this direction is Kirman (1991, 1993) who adapts a simple

stochastic model for information transmission in ant colonies. While the original set-up has ants exchanging information about the direction of food sources, the adaptation to a financial setting replaces them by foreign exchange dealers exchanging information about the accurateness of chartist and fundamentalist predictions of exchange rates. Agents who meet other traders adopt their strategy with a certain probability, but may also undergo autonomous changes of opinion without interaction. Similar models with interpersonal influences have been proposed by Lux (1995, 1997, 1998), Lux and Marchesi (1999, 2000), Chen et al. (2003), Aoki (2002) and Wagner (2003).

Alfarano and Lux (2004) have a stripped down version of an extremely parsimonious herding model which still appears to do the job of generating appropriate power laws for returns and volatility. They again assume that two different groups interact in the market: the well-known fundamentalists and a second group denoted by noise traders who are assumed to follow the current mood of the market. While the first group simply trades on the base of observed mispricing (i.e. differences between price p and fundamental value p_f), noise traders are assumed to be influenced by contagion dynamics. They are optimists (buyers) or pessimists (sellers) and switch between both sides of the market with simple probabilities reflecting the influence of the majority opinion:

$$prob(O \to P) = \nu \frac{N_P}{N}, \qquad prob(P \to O) = \nu \frac{N_O}{N} \tag{14}$$

with $N_P(N_O)$ the number of pessimistic (optimistic) agents, ν a time-scaling parameter and $N = N_O + N_P$.

Adding up excess demand by both fundamentalists and noise traders, the overall difference between demand and supply can be written as:[7]

$$ED = T_F(p_f - p) + T_c x, \qquad x = \frac{N_O - N_p}{N} \tag{15}$$

with T_F, T_C constants determining the trading volume of fundamentalists and noise traders (aka chartists). Assuming that market equilibrium is attained instantaneously, one can solve for the equilibrium price:

$$p = p_f + \frac{T_c}{T_f} x. \tag{16}$$

As can be seen from (16), price changes are generated by both (i) exogenous inflow of new information about fundamentals (p_f) and (ii) endogenous changes in demand and supply brought about via the herding mechanism (i.e. via changes of x). While traditional finance models only allow for the first component (for x = 0, we would end up with eq. (8)) and, therefore, have to

[7] Since optimistic (pessimistic) noise trader would buy (sell), their net contributions to excess demand simply depends on the difference $N_O - N_p$.

trace back all features of returns to similar features of fundamentals, behavioral finance models give a role to the intrinsic dynamics of financial markets. While this structure is similar to the RE bubble model (eq. 9), in contrast to rational bubbles the second component in eq. (16) does not have to obey a restriction for the rationality or consistency of expectations.

As illustrated in a typical simulation in Fig. 3, these few ingredients detailed above are sufficient to produce relatively realistic time series for returns whose distributional and temporal characteristics are quite close to empirical findings. What is the reason for this outcome? The herding mechanism of eq. (14) produces a bi-modal limiting distribution for the fraction of noise traders in the two groups of optimistic and pessimistic traders. Most of the time, one, therefore encounters a majority of the noise traders on either the supply or demand side of the market (which goes along with undervaluation or overvaluation of the asset price). However, the stochastic nature of the process also leads to recurrent switches from one majority to another. During these periods, large swings in average opinion lead to an increase of volatility which will last for some time until a lock-in occurs again to a stable optimistic or pessimistic majority.

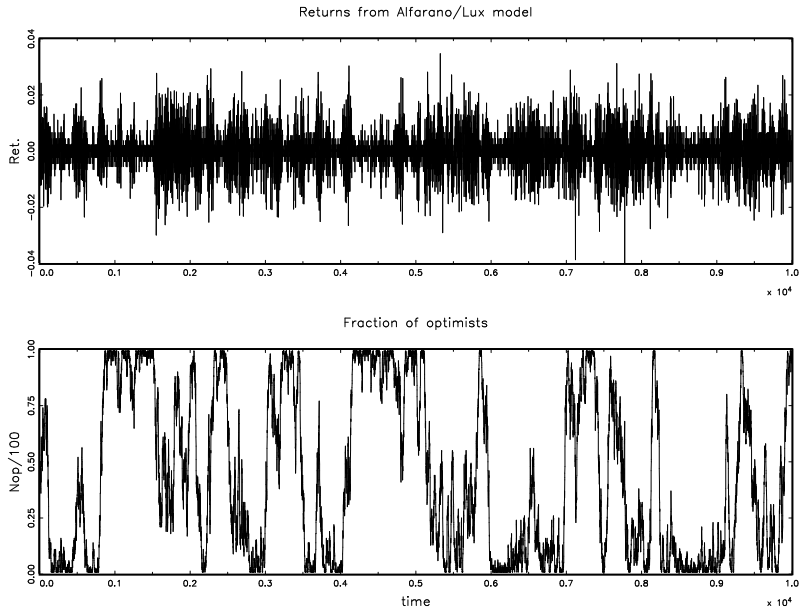

Fig. 3. Typical simulation of the model by Alfarano and Lux (2004). The driving force of this model are stochastic transitions of agents between trading strategies which lead to switches between high and low volatility episodes. The models can reproduce the phenomenological appearance of empirical series (see Fig. 1) to some extend and could also give the impression of power laws under standard statistical tests. However, it does not generate 'true' power laws in returns and volatility, but only apparent scaling behavior.

However, as outlined in Alfarano and Lux (2003), at least the temporal scaling of volatility in this model does not follow a true power-law. Because of the Markovian structure of the model, no 'true' long-term dependence can exist, although standard tests are typically unable to distinguish between the apparent power-law behavior of this model and true power-law behavior. Similar pre-asymptotic, spurious scaling might occur in the very similar models of Kirman (1991) and Kirman and Teyssière (2002). Wagner (2002) shows that a more detailed statistical analysis reveals differences between volatility scaling in real-life markets and the behavior of these simple models. For instance, the volatility clusters within a time window are all of the same size in these toy models and small size clustering while observed in empirical time series is absent in the simulated data.

Certain refinements of this approach provide simulated time series which come closer to empirical observations. Aoki (2002) considers a stochastic process of interaction of chartists and fundamentalists with entry and exit of traders and also achieves to derive a stochastic difference equation for the

ensuing returns dynamics with a tail index $\mu > 1$. Finally, Carvalho (2003) demonstrates that one can even do without chartist traders or similar self-reinforcing feedback in price formation. Even when only fundamentalists are present in the market, a multiplicative noise structure can result. Letting fundamentalists react on the difference between market returns and changes of fundamentals, and attributing to them identical market power but allowing for activation of each trader with a certain probability only, suffices to generate a multiplicative stochastic process via the probability distribution of the number of active traders and, hence, produces power law behavior of returns.

Another extension of the simple herding dynamics presented above is the earlier model of Lux and Marchesi (1999, 2000), who add transitions between noise traders and fundamentalists depending on the profitability of both strategies. Lux and Marchesi show that this relatively complex set-up could reproduce empirically realistic scaling laws for both returns and volatility in rather good numerical agreement with empirical data. Various sensitivity analyses also indicate that the numerical results are not very sensitive with respect to parameter variations. Furthermore, an analytical approximation via the Master equation framework adopted from statistical physics suggests a general robustness of the qualitative appearance of the dynamics. The characteristic switches between tranquil and volatile phases are triggered by recurring temporal deviations from an otherwise stable equilibrium in which the price is close to its fundamental value. The mechanism is this: in the neighborhood of the equilibrium, neither mispricing nor any detectable patterns in the price trajectories exist so that neither the chartist nor the fundamentalist strategy has an advantage. Agents, then, switch between these alternatives in a rather unsystematic manner which makes the population composition (in terms of strategies) follow a random walk. However, stability of the fundamental equilibrium depends crucially on the ratio of chartists and fundamentalists. This is an insight from practically all models with chartist-fundamentalist interaction and it also features prominently in the Lux/Marchesi framework[8]. The random walk in strategy space now makes the population sweep over this threshold every once in a while creating an onset of volatility.[9]

In fact, the deterministic approximation of the dynamic system shows that one can interpret the number of chartists and fundamentalists as a critical parameter of a simpler system in which changes of this variable lead to movements back and forth through a Hopf bifurcation scenario. What happens fits the definition of so-called *on-off intermittency* in physics which describes a dynamic system undergoing phases of temporal (intermittent) bursts of activity when the variation of one dynamical variable moves it through the bifurcation point of another dynamic variable. Typical examples are chaotic attractors coupled to stochastic motion or coupled oscillators (cf. Ott and Sommerer,

[8]For each combination of parameters, a threshold value can be computed for the fraction of chartists beyond which a loss of stability occurs.

[9]An illustration of the resulting dynamics is provided in Fig. 4.

1996, Venkatamarani *et al.*, 1990). Time series produced in this literature are often perplexingly similar to the phenomenological appearance of volatility clustering in financial data and several fractal properties of the time series have been analyzed in the above papers (albeit the ones of particular interest in finance are not discussed in this literature). Although no theoretical results are available so far for the model of Lux and Marchesi (1999), recent results on some simpler versions (Alfarano, Lux and Wagner, 2004) suggest that under some circumstances, interacting agent models of this type can generate true power laws for returns and volatility.

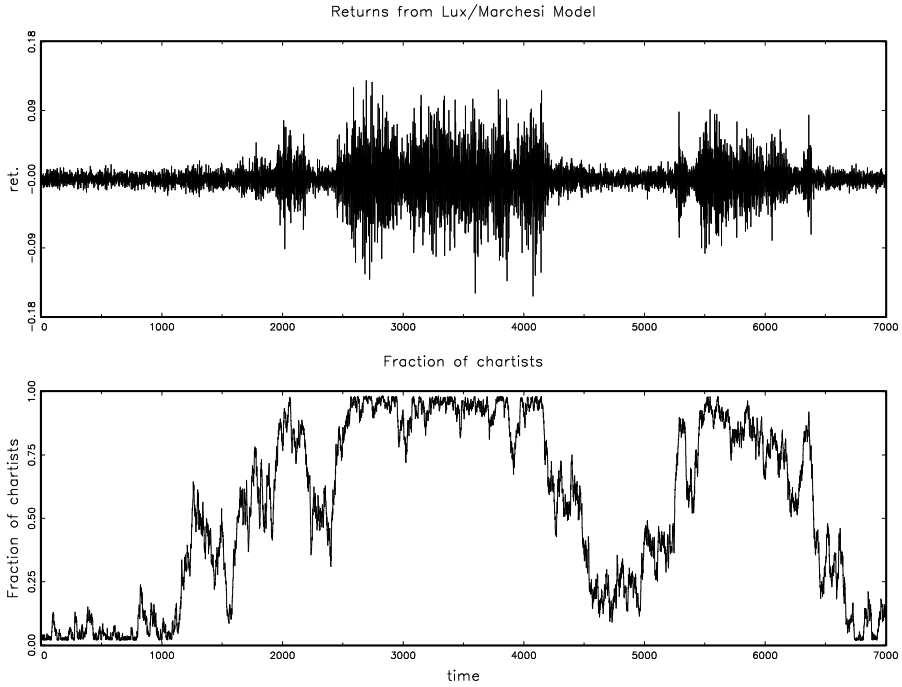

Fig. 4. Typical simulation of the model by Lux and Marchesi (1999). Here switching between the groups of chartists and fundamentalists is allowed in addition to the herding within the later group featuring in Alfarano and Lux (2004). As is clearly visible, large fluctuations occur if the majority of traders adopts the chartist strategy. Again, the phenomenology of the model comes close the empirical power-laws of sec. 2. Lux and Marchesi (1999) report estimates of power-law stochastic in agreement with the empirical stylised faxts (eq. (1) and (2)). However, it is currently still unclear whether it produces 'true' or only apparent power laws.

Lux and Marchesi (1999, 2000) argue that irrespective of the concrete details of the model, the indeterminateness of the population composition in

a market equilibrium might be a relatively general phenomenon (because of the absence of profitability of any trading strategy in such a steady state situation) and together with dependence of stability on the population composition, potential on-off intermittency should exist in a broad class of behavioral finance models. Some support for this argument is provided by Lux and Schornstein (2004) who consider a quite different multi-agent design of a foreign exchange market with agents endowed with artificial intelligence (genetic algorithm learning). They point out that a similar scenario like in Lux and Marchesi prevails in this framework which also generates behavior close to realistic power-laws in simulated data. Similarly, Giardina *et al.* (2002) allow more general strategies than Lux/Marchesi, but also find a random walk in strategy space to account for the emergence of realistic dynamics.

There is one important shortcoming of this whole class of models: their outcome usually depends in a sensitive way on the system size (i.e., the number of agents operating in the market). With the increase of the size of the population, the nice dynamic features and power-law statistics get lost (cf. Egenter *et al.*, 1999). The reason is that with an increasing number of autonomous agents, a law of large numbers comes into play and the stochastic dynamics effectively becomes equivalent to draws from a Normal distribution.

IV.4 Another Mechanism: Switching between Predictors and Attractors

A closely related approach has emerged from an adaptation of the seminal random utility framework for empirical analysis of discrete choice problems[10] (Manski and McFadden, 1981) for formalizing - again - the interplay of the notorious chartists and fundamentalists in speculative markets. This approach dates back to Brock and Hommes (1997) but its potential to generate realistic time series has only recently been revealed. Models of this type are very close in terms of economic intuition to those reviewed in the previous section, but have a somewhat different flavor from a dynamic systems perspective.

Let us illustrate the basic ingredients of these models via an example along the lines of Gaunersdorfer and Hommes (2003). Our starting point is the distinction between different groups of traders (mostly two) whose excess demand for a risky asset depends on their group-specific expectation of future price investments and dividends:

$$ED_{i,t} = \frac{E_{i,t}[p_{t+1} + d_{t+1} - (1+r)p_t]}{\mu\sigma^2} \tag{17}$$

with $ED_{i,t}$: demand of group i at time t, p_t: the price, d_t: the dividend, r: the risk-free rate available for a riskless asset (i.e. government bonds), μ: a parameter of risk aversion, and σ^2 the variance of the expectation term in the numerator. Eq. (17), therefore, formalizes demand depending on the

[10]For example, choosing the preferred one from two or more brands of whiskey.

risk-adjusted excess returns (the sum of dividends and capital gains) of a risky over a riskless investment. Economists know that this kind of demand function can be derived from simple mean-variance utility functions as well as from negative exponential utility functions under a few additional conditions. With some effort, the more phenomenological excess demand function of the previous chapter could also be cast into such a framework.

If there is no additional inflow of assets (no new issues for the time being), market equilibrium simply requires that:

$$\sum_i n_{i,t} ED_{i,t} = 0, \tag{18}$$

with $n_{i,t}$ the fraction of market participants subscribing to strategy i at time t. Allowing only for agents who compare the expectation of future price and distributions according to eq.(7), we would once more fall back on the fundamental price of eq. (8). However, introducing 'non-rational' speculative strategies deviations from this fundamental valuation become possible. As an example, Gaunersdorfer and Hommes (2003) analyze a model with two groups, $i = 1, 2$, using different expectation formation rules:

$$\begin{aligned} E_{1,t}[p_{t+1}] &= p_f + v(p_{t-1} - p_f), \\ E_{2,t}[p_{t+1}] &= p_{t-1} + \gamma(p_{t-1} - p_{t-2}), \end{aligned} \tag{19}$$

together with stationary expectations of dividends $E_{1,t}[d_{t+1}] = E_{2,t}[d_{t+1}] = \bar{d}$. Clearly, group 1 can be identified as a fundamentalist crowd while group 2 has a chartist prediction technique. The discrete choice framework comes into play when modeling the choice of one of these two predictors by the agents in the market. Like in Lux/Marchesi, choice of strategies is governed by objective economic factors. Gaunersdorfer and Hommes use accumulated profits:

$$\pi_{i,t} = R_t \cdot ED_{i,t-1} + \rho \pi_{i,t-1} \tag{20}$$

where R_t is the actual return achieved with investment in the period before and ρ is a parameter for the memory of past profitability. Other papers use risk-adjusted profits or utility measures instead of monetary profits as fitness criteria. However, it does not appear from the available literature that these choices would lead to grossly different outcomes. The last step is predictor choice based on the fitness criterion:

$$n_{i,t+1} = \frac{exp(\beta \pi_{i,t})}{\sum_i exp(\beta \pi_{i,t})} \quad , \tag{21}$$

i.e. through the standard discrete choice formalization. β is called the intensity of choice with the extreme cases $\beta = 0$ ($\beta \to \infty$) leading to constant fractions equal to 0.5 or total polarization (one $n_{i,t} = 1$, all others equal to zero if the slightest differences in fitness exist). Often, as in Gaunersdorfer

and Hommes (2003), additional factors enter the determination of fractions $n_{i,t}$.

Formally, the combined dynamics of predictor choice and price development make up a discrete deterministic system (in the above example it is a difference equation system of order three). Typically one would be able to derive analytical results concerning the number of possible equilibria, local stability conditions and bifurcations of the system dynamics. Because of the highly non-linear nature of these systems, many kinds of complex dynamics can arise: limit cycles, chaotic dynamics, homoclinic orbits as well as co-existence of different types of attractors may all happen for different sets of parameters. The later scenario is of particular interest and can be characterized as another type of 'intermittend dynamic behaviour': when two or more attractors exist for the deterministic system, adding noise (random shocks) superimposed on the deterministic dynamics may trigger switches between different basins of attraction of different limiting sets. For example, with coexistence of a fixed point-attractor and a limit cycle, the stochastic movements between both types of dynamics also goes along with switches between turbulent and tranquil market phases. An illustration is provided in Fig. 5 which exhibits a simulation of the example given by Gaunersdorfer and Hommes (2003). Quite intuitively plausible, the market exhibits small fluctuations as long as it hovers within the basin of attraction of the locally stable fixed point (where the price is equal to the fundamental value on average) and the majority of traders chooses a fundamentalist strategy, but onset of more violent fluctuations occurs once it traverses to the basin of the limit cycle (with its cyclic dynamics being governed by a prevailing chartist attitude). As can be seen from the resulting dynamics of returns, at least for certain choices of parameter values, the deterministic origin of the process is almost entirely concealed by its stochastic components. As Gaunersdorfer and Hommes show, a number of statistics give quite satisfactory agreement between their simulated data and empirical records. Although they do not estimate power law indices, it seems natural from the simulated time series that their process should be able to mimic the hyperbolic decay of the returns distribution and the long-term dependence of volatility.

Further examples are given by Gaunersdorfer (2001) who has a slightly different model set-up which could even lead to coexisting fixed points and chaotic attractors. Other recent contributions along very similar lines are Chiarella and He (2002, 2003), Fernandez-Rodriguez et al. (2002), Westerhoff (2003), and De Grauwe and Grimaldi (2003). The latter paper has an interesting variation of the attractors-switching avenue to intermittent fluctuations: introducing transaction costs for the acquisition of fundamental information they obtain a band of inactivity of fundamentalists around the fundamental equilibrium. The price process, then, follows different patterns inside and outside of this transaction cost band which apparently also generates intermittent volatility clustering. Both De Grauwe and Grimaldi and Westerhoff estimate

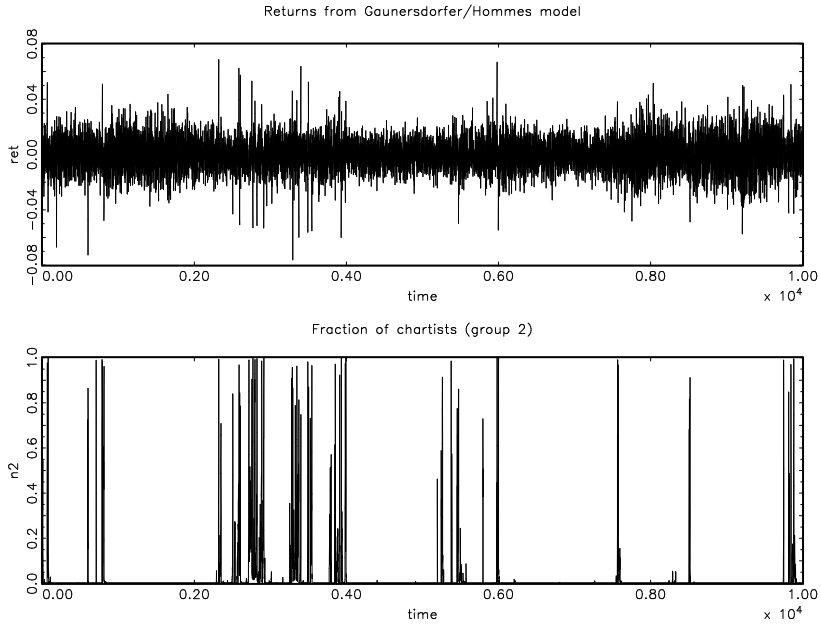

Fig. 5. Typical simulation from the model of Gaunersdorfer and Hommes. Similarly to the Lux/Marchesi model, large fluctuation are triggered by a chartist majority. However, here transitions between periods with fundamentalist and chartist dominance are sharper corresponding to the rather sudden change between attractors pointed out in the text. The model also reproduces the clustered volatility of empirical records but due its origin from a deterministic dynamics with additive noise will probably not generate 'true' power laws.

the tail indices for large returns and obtain numbers in agreement with the stylized facts, i.e. numbers in the vicinity of the cubic power law.

One of the distinguishing features of the above contributions is the assumption of an *infinite population of speculators* which allows to study the resulting dynamics via systems of deterministic difference equations derived for the infinite population limit. While this is an approach very much in line with traditional economic theorizing, it gives these models a somewhat different flavor compared to the finite-size multi-agent models reviewed before. In particular, the interplay of noise and deterministic factors is quite different in both approaches: while noise appears on the level of each individual agent in, for instance, percolation models or the Lux/Marchesi model throught the transition probabilities (eg. 14), the inherent fluctuations of the discrete choice approach are averaged out by the assumption of an infinite population in Gaunersdorfer and Hommes (2003). Thence, despite the randomness at the level of the agents and the use of random choice probabilities, eq. (21), for in-

stance, becomes a deterministic equation under the assumption of an infinite population. The noise component responsible for the switches between attractors has, therefore, to be superimposed on the market dynamics and enters on the level of the macroscopic system by, for example adding a stochastic term in the excess demand equation (18). However, this implies that the noise level has to be relatively large to obtain the results exhibited in Fig. 5. In fact, inspection of the simulations of this example shows that the added stochastic component in market excess demand is of almost the same size like fundamentalists' average excess demand (with chartists' average excess demand being equal to the sum of the two other components). In simulations with realistic time series properties, the 'signal-to-noise ratio' in this model is, therefore, practically equal to one. With hindsight, this may be a necessary requirement for obtaining 'realistic' time series as a high noise factor will be required to mask the inherent deterministic forces prevailing around both the fix point equilibrium and the limit cycle regime. Since, on the other hand, the noise level should not be too high so to not totally dominate the deterministic roots of the dynamics, the variance of the added stochastic term is probably a central ingredient of the model.

V. Conclusions

In contrast to one recently articulated opinion (Durlauf, 2003) power laws in finance have never been regarded as curiosities. On the contrary, both the scaling of the tails of the distribution of returns and the long-range dependence of volatility are all-important to practitioners and have motivated a vast statistical and econometric literature. Much of the development of the toolbox in empirical finance is due to these apparently universal properties of financial markets (with Engle's family of ARCH models the most prominent example). It is, however, also true that economic theory has been altogether silent on behavioral roots of these regularities until very recently. Besides having been regarded as merely statistical findings outside the realm of economic theory, typically they have mostly been described in a more phenomenological way via the shape of histograms or the apparent clustering of tranquil and turbulent episodes. While a power-law perspective has already appeared in Mandelbrot's seminal contribution (1963), whether it is really a power law has remained unclear until the nineties. By now, availability of high-frequency data and new analytical tools have led to a consensus of a relatively uniform exponent around three which seems to hold across markets, countries and time periods. Similarly homogeneous evidence is emerging for the long-range dependence of both volatility and volume.

Theoretical work trying to explain these features has followed diverse ways. As depicted in Table 1 we can, in principle, distinguish between exogenous and endogenous approaches for explaining financial power-laws. The first class has

only been dealt with in passing in the main text: its major member is the whole body of traditional efficient market models in finance which would have to attribute all time series properties of financial returns to the structure of news about fundamental factors. In this view, scaling in returns and volatility would only mirror scaling in increments and fluctuations of fundamental information. Needless to say, many proponents of behavioral models take the unsatisfactory nature of this explanation as their starting point. An interesting alternative exogenous explanation can be found in the recent paper by Gabaix et al. (2003) who derive the cubic law for returns from Zipf's law for the capital of large investors plus some other ingredients. However, the empirical basis of this theory has already been questioned by careful analysis of some of its ingredients. Even if it would go through it would leave one with an unexplained Zipf's law for the wealth of investors from which the other power laws are derived.

Most papers on the topic have, however, pursued a different approach focusing on the intrinsic dynamics of speculative interaction in financial markets. The first example in this second class is the RE bubble theory which, in fact, emerges from relaxation of a minor technical condition of traditional present value models. Although it has the fascinating property of defining a whole class of data-generating processes with generic power laws, numerically these laws are disappointingly far off from the empirical ones. The conclusion to be drawn from this result is that we can exclude this whole class of fully rational models of speculative activity. This leaves one essentially with either the choice of subscribing to the Efficient Market Hypothesis of price formation being exclusively governed by fundamentals, or resorting to one or the other brand of models of speculative activity with bounded rationality. Within the recent econophysics literature, percolation models adopted from statistical physics have attracted the interest of a sizable number of researchers. Their disadvantage is the lack of robustness: the model parameters have to be fine-tuned to arrive at the required power-laws. As another drawback, most parameters in these models are not easy to interpret, so that an economic assessment of the explanatory power of the resulting dynamics at the critical percolation threshold is difficult.

In the behavioral finance literature, several types of models with interaction of different trader groups have been proposed. Typically, interesting time series are obtained from some kind of 'intermittent' dynamics. A general conclusion from this body of literature is that some kind of self-amplification of fluctuations via herd behavior or technical trading is necessary (and often sufficient) to generate time series which are phenomenologically close to empirical records. One of the more important problems of these models is the relationship between system size, deterministic forces and stochastic elements. On the one hand, typical simulation models suffer from a critical dependence of their 'nice' results on the number of agents operating in the market. Models starting with an infinite population, on the other hand, have to adjust the

noise level in a way to counterbalance the deterministic core of their market dynamics.

	Class of Models	Source of Power Law	Problems
Exogenous explanations	Traditional finance	fundamentals	fundamentals are unobservable
	Gabaix et. al (2003)	Zipf's law for investment capital	neglect of behavioral roots? empirical validity
Endogenous explanations	RE bubbles	multiplicative dynamics	unrealistic exponents
	Percolation models, e.g. Cont/Bouchaud (2000)	cluster formation	power laws not robust
	Multi-agent models, e.g. Lux/Marchesi (1999)	intermittent dynamics	sensivity with respect to number of agents
	Discrete choice models, e.g. Gaunersdorfer/Hommes (2003)	switching between attractors	sensivity with respect to noise amplitude

Table 1. Sources of Power Laws in Finance

The recent introduction of concepts and models from statistical physics has also evoked interesting methodological discussions: while physicists forcefully argue in favor of building simple models of interacting economic agents and neglecting as far as possible details which are not in the center of interest (cf. Stauffer, 2003), some economists have criticized this approach for producing models that are not economically insightful (Durlauf, 2003). Of course, parsimony is also a concern in economic modeling. However, physicists and economists would often differ in their assessment of the essential model ingredients a truly parsimonious model should contain: while physicists would favor interactions, economists would traditionally prefer to emphasize the microeconomic foundations of agents' behavior. However, having to begin with a full-fletched microeconomic optimization approach makes modeling of interactions superimposed on the traditional microstructure an even more demanding task. It could also restrict the outcome by adhering to particular forms of utility functions, ways of information acquisition and information processing, trading strategies etc. The relative success of several simple models in explaining a good degree of the hitherto unexplained empirical characteristics of financial data casts doubts on the paradigm of micro-foundations in the sense of 'representative' individual optimization of agents in economic models (which is a classical example of a reductionist approach). Although analysis of individual optimization is in no way unimportant, exclusively focusing on this aspect of economic life comes with the danger of neglecting the equally important consequences of both market - mediated and social interactions of market participants.

Acknowledgements

The author gratefully acknowledges many helpful comments by the audience of the Lustrum Symposium at the Center for Nonlinear Dynamics in Economics and Finance, University of Amsterdam, 2003 and the 9^{th} Workshop on Economics with Heterogeneous Interacting Agents at the University of Kyoto, 2004. Particular thanks are also due to the participants of the workshop on "Power Laws in Social Sciences" at James Madison University and to Claudio Cioffi, Dietrich Stauffer and Frank Westerhoff for their detailed comments on earlier versions of the manuscript.

References

Alfarano, S. and T. Lux, A Simple Model of Herd Behavior with Realistic Time Series Properties, in: A. Kirman and G. Teyssière, eds., Long Memory in Economics and Econometrics, Berlin, Springer, 2004, in press.

Alfarano, S., T. Lux and F. Wagner, Time Variation of Higher Moments in Financial Markets with Heterogeneous Agents: An Analytical Approach, manuscript, University of Kiel, 2004.

Aoki, M. and H. Yoshikawa, Demand Saturation-Creation and Economic Growth, Journal of Economic Behavior and Organization 48, No.2, 2002, 127-154.

Arifovic, J., The Behavior of the Exchange Rate in the Genetic Algorithm and Experimental Economies, Journal of Political Economy 104, 1996, 510-541.

Arthur, W. B., J. Holland, B. LeBaron, R. Palmer and P. Tayler, Asset Pricing under Endogenous Expectations in an Artificial Stock Market, in: Arthur, W. B., S. Durlauf and D. Lane, eds., The Economy as an Evolving Complex System II, Addison-Wesley, Reading, 1997.

Baillie, R. T., T. Bollerslev and H. O. Mikkelsen, Fractionally Integrated Generalized Autoregressive Conditional Heteroskedasticity, Journal of Econometrics 74, 1996, 3-30.

Bak, P., M. Paczuski and M. Shubik, Price variations in a stock market with many agents, Physica A 246, 1997, 430-453.

Baumol, W. J., Speculation, Profitability and Stability, Review of Economics and Statistics 39, 1957, 263-271.

Blanchard, O. J. and M. W. Watson, Bubbles, Rational Expectations and Speculative Markets, in: Wachtel, P., ed., Crisis in Economic and Financial Structure: Bubbles, Bursts, and Shocks, Lexington Books, Lexington, 1982.

Brenner, T., Modelling Learning in Economics, Edward Elgar, Cheltenham, 1999.

Brock, W. A. and C. Hommes, A Rational Route to Randomness, Econometrica 65, 1997, 1059-1095.

Brock, W. A. and C. H. Hommes, Models of Complexity in Economics and Finance, in: Hey, C., J. M. Schumacher, B. Hanzon, and C. Praagman, eds., System Dynamics in Economic and Financial Models, pp. 3-44, Wiley, New York, 1997.

Buchanan, M., Ubiquity, Weidenfeld & Nicholson, 1999.

Calvet, L. and A. Fisher, Multi-Fractality in Asset Returns: Theory and Ev-

idence, Review of Economics & Statistics 84, 2002, 381-406.

Campbell, J., A. Lo and A. MacKinlay, The Econometrics of Financial Markets, Princeton University Press, Princeton, 1997.

Carvalho, R., The Dynamics of the Linear Random Farmer Model, manuscript, University of Lisbon, 2002.

Chang, I. and D. Stauffer, Fundamental Judgement in Cont-Bouchaud Herding Model of Financial Fluctuations, Physica A 264, 1999, 294-298.

Chen, S. H., T. Lux and M. Marchesi, Testing for Nonlinear Structure in an Artificial Financial Market, Journal of Economic Behavior and Organization, 46, 2001, 327-342

Chen, S. H. and C. H. Yeh, On the Emergent Properties of Artificial Stock Markets: The Efficient Market Hypothesis and the Rational Expectations Hypothesis, Journal of Economic Behavior and Organization 49, 2002, 217-239.

Cheung, Y. und C. Wong, A Survey of Market Practitioners' Views on Exchange Rate Dynamics, Journal of International Economics 51, 2000, 401-419.

Chiarella, C. and X. He, Heterogeneous Beliefs, Risk and Learning in a Simple Asset Pricing Model, Computational Economics 19, 2002, 95-132.

Chiarella, C. and X. He, Dynamics of Beliefs and Learning under al-Processes - the Heterogeneous Case, Journal of Economic Dynamics and Control 27, 2003, 503-531.

Cont, R., M. Potters and J.-P. Bouchaud, Scaling in Stock Market Data: Stable Laws and Beyond, in: Dubrulle, B., F. Graner and D. Sornette, eds., Scale Invariance and Beyond, Springer, Berlin, 1997.

Cont, R. and J.-P. Bouchaud, Herd behaviour and aggregate fluctuations in financial markets, Macroeconomic Dynamics 4, 170-196, 2000.

DeGrauwe, P., H. Dewachter, and M. J. Embrechts, Exchange Rate Theory: Chaotic Models of Foreign Exchange Market, Blackwell, Oxford, 1993.

DeGrauwe, P. and M. Grimaldi, Exchange Rate Puzzles: A Tale of Switching Attractor, manuscript, University of Leuven, 2003.

Ding, Z., R. Engle and C. Granger, A Long Memory Property of Stock Market Returns and a new Model, Journal of Empirical Finance 1, 1993, 83-106.

Durlauf, S., Complexity and Empirical Economics, Economic Journal, 2004 (in press).

Egenter, E., T. Lux and D. Stauffer, Finite-Size Effects in Monte Carlo Simulations of Two Stock Market Models, Physica A 268, 1999, 250-256.

Eguiluz V. and M. Zimmermann, Transmission of Information and Herd Behavior: An Application to Financial Markets, Physical Review Letters 85, 2000, 5659-5662.

Engle, R. F., Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of U.K. Inflation, Econometrica 50, 1982, 987-1008.

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

Farmer, J.D. and Lillo, F., On the Origin of Power Law Tails in Price Fluctuations, Quantitative Finance 4, 2004, C7 - C11.

Fernandez-Rodriguez, F., M.-D. Garca-Artiles and J. M. Martín-González, A Model of Speculative Behaviour with a Strange Attractor, Applied Mathematical Finance 9, 2002, 143-161.

Focardi, M., S. Cincotti and M. Marchesi, Self-Organization and Market Crashes, Journal of Economic Behavior & Organization 49, 2002, 241-267.

Gabaix, X. et al., A Theory of Power-Law Distributions in Financial Market Fluctuations, Nature 423, 2003, 267-270.

Gaunersdorfer, A., Endogenous Fluctuations in a Simple Asset Pricing Model with Heterogeneous Beliefs, Journal of Economic Dynamics & Control 24, 2000, 799-831.

Gaunersdorfer, A. and C. Hommes, A Nonlinear Structural Model for Volatility Clustering, manuscript, University of Vienna, 2001.

Giardina, I. and J.-P. Bouchaud, Bubbles, Crashes and Intermittency in Agent Based Market Models, European Physical Journal B 31, 2003, 421-437.

Gopikrishnan, P., M. Meyer, L. A. N. Amaral and H. E. Stanley, Inverse Cubic Law for the Distribution of Stock Price Variations, European Physical Journal B 3, 1998, 139-140.

Gopikrishnan, P., V. Plerou, X. Gabaix, L. A. N. Amaral and H. E. Stanley, Price fluctuations, Market Activity and Trading Volume, Quantitative Finance 1, 2001, 262-270.

Grannan, E. R. and G. H. Swindle, Contrarians and Volatility Clustering, Working Paper: Santa Fe Institute, 1994.

Hall, J. A., B. W. Brorsen and S. H. Irwin, The Distribution of Futures Prices: A Test of the Stable Paretian and Mixture of Normals Hypothesis, Journal of Finance and Quantitative Analysis 24, 1989, 105.

Hirschleifer, D and G. Luo, On the Survival of Overconfident Traders in a Competitive Securities Market, Journal of Financial Markets 4, 2001, 73-84.

Hill, B. M., A Simple General Approach to Inference about the Tail of a Distribution, Annals of Statistics 3, 1975, 1163.

Iori, G., A Microsimulation of Traders Activity in the Stock Market: The Role of Heterogeneity, Agents' Interactions and Trade Frictions, Journal of Economic Behavior and Organization 49, 2002, 269-285.

Jansen, D. W. and C. G. de Vries, On the Frequency of Extreme Stock Returns: Putting Booms and Busts into Perspectives, Review of Economics and Statistics 73, 1991, 18.

Kesten, H., Random Difference Equations and Renewal Theory for Products of Random Matrices, Acta Mathematica 131, 1973, 207-248.

Keynes, J. M., The General Theory of Employment, Interest and Money, London: Macmillan, 1936.

Kirman, A., Epidemics of Opinion and Speculative Bubbles in Financial Markets. In M. P. Taylor, Money and Financial Markets, 1991, 354-368, Blackwell, Cambridge.

Kirman, A., Ants, Rationality, and Recruitment, Quarterly Journal of Economics 108, 1993, 137-156.

Kirman, A. and G. Teysièrre, Microeconomic Models for Long Memory in the Volatility of Financial Time Series, Studies in Nonlinear Dynamics & Econometrics 5, 2002, 281-302.

Lau, A. H.-L., H.-S. Lau and J. R. Wingender, The Distribution of Stock Returns: New Evidence against the Stable Model, Journal of Business Economics and Statistics 8, 1990, 217.

LeBaron, B., W. B. Arthur, and R. Palmer, The Time Series Properties of an Artificial Stock Market, Journal of Economic Dynamics and Control 23, 1999, 1487-1516.

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

Levy, M., H. Levy and S. Solomon, A Microscopic Model of the Stock Market: Cycles, Booms, and Crashes, Economics Letters 45, 1994, 103-111.

Levy, M., H. Levy, and S. Solomon, Microscopic Simulation of the Stock Market: The Effect of Microscopic Diversity, Journal de Physique I 5, 1995, 1087-1107.

Lobato, I. N. and N. E. Savin, Real and Spurious Long-Memory Properties of Stock Market Data, Journal of Business and Economics Statistics 16, 1998, 261-283.

Lobato, I. N. and C. Velasco, Long Memory in Stock-Market Trading Volume, Journal of Business and Economic Statistics, 18, 410-427, 2000.

Lux, T., Herd Behaviour, Bubbles and Crashes, Economic Journal 105, 1995, 881-896.

Lux, T., The Stable Paretian Hypothesis and the Frequency of Large Returns: An Examination of Major German Stocks, Applied Financial Economics 6, 1996, 463-475.

Lux, T., Time Variation of Second Moments from a Noise Trader/Infection Model, Journal of Economic Dynamics and Control 22, 1997, 1-38.

Lux, T., The Socio-Economic Dynamics of Speculative Markets: Interacting Agents, Chaos, and the Fat Tails of Return Distributions, Journal of Economic Behavior and Organization 33, 1998, 143-165.

Lux, T., The Multi-Fractal Model of Asset Returns: Its Estimation via GMM and Its Use for Volatility Forecasting, manuscript, University of Kiel, 2004.

Lux, T. and M. Ausloos, Market Fluctuations I: Scaling, Multi-Scaling and Their Possible Origins *chap.*13 in: Bunde, A., J. Kropp and H. J. Schellnhuber, eds., The Science of Disasters: Climate Disruptions, Heart Attacks, and Market Crashes, 372 - 409, Springer 2002.

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

Lux, T. and M. Marchesi, Volatility Clustering in Financial Markets: A Micro-Simulation of Interacting Agents, International Journal of Theoretical and Applied Finance 3, 2000, 67-702.

Lux, T. and D. Sornette, On Rational Speculative Bubbles and Fat Tails, Journal of Money, Credit and Banking 34, No. 3, 2002, 589-610.

Lux, T. and S. Schornstein, Genetic Learning as an Explanation of Stylized Facts of Foreign Exchange Markets, 2004, forthcoming in: Journal of Mathematical Economics.

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

Mandelbrot, B., A Multifractal Walk down Wall Street, Scientific American, 1999, 50-53.

Manski E. and D. McFadden, Structural Analysis of Discrete Data with Economic Applications, MIT-Press, Cambridge, 1981.

Noussair, C. , S. Robin and B. Ruffieux, Price Bubbles in Laboratory Asset Markets with Constant Fundamental Values, Experimental Economics 4,

2001, 87-105.

O'Hara, M., Market Microstructure Theory, Blackwell, Cambridge, 1995.

Ott, E. and J. C. Sommerer, Blowout Bifurcations: The Occurrence of Riddled Basins and On-Off Intermittency, Physics Letters A 188, 1994, 39-47.

Ramsey, J. B., On the Existence of Macro Variables and of Macro Relationships, Journal of Economic Behavior and Organisation 30, 1996, 275-299.

Sato, A.-H. and H. Taskayasu, Dynamic Numerical Models of Stock Market Prices: From Microscopic Determinism to Macroscopic Randomness, Physica A 250, 1998, 231-252.

Sornette D., D. Stauffer and H. Takayasu, Market Fluctuations II: Multiplicative and Percolation Models, Size Effects, and Predictions, *chap.*14 in: Bunde, A., J. Kropp and H. J. Schellnhuber, eds., The Science of Disasters: Climate Disruptions, Heart Attacks, and Market Crashes, 411-436, Springer 2002.

Stauffer, D., Introduction to Statistical Physics outside Physics, Physica A 336, 2004, 15.

Stauffer, D., de Oliveira, P. and A. Bernardes, Monte Carlo Simulation of Volatility Clustering in Market Model with Herding, International Journal of Theoretical and Applied Finance 2, 1999, 83-94.

Stauffer, D. and T. Penna, Crossover in the Cont-Bouchaud Percolation Model for Market Fluctuations, Physica A 256, 1998, 284.

Stauffer, D. and D. Sornette, Self-Organized Percolation Model for Stock Market Fluctuation, Physica A 271, 1999, 496-506.

Takayasu, H., ed., Empirical Science of Financial Fluctuations: The Advert of Econophysics, Springer, Berlin, 2002.

Takayasu, H., H. Miura, T. Hirabayashi and K. Hamada, Statistical Properties of Deterministic Threshold Elements - the Case of Market Price, Physica A 184, 1992, 127-134.

Vandewalle, N. and M. Ausloos, Coherent and Random Sequences in Financial Fluctuations, Physica A 246, 1997, 454-459.

Vandewalle, N. and M. Ausloos, Sparseness and Roughness of Foreign Exchange Rates, International Journal of Modern Physics C 9, 1998, 711-719.

Venkatamarani, S. C., T. M. Antonsen Jr., E. Ott and J. C. Sommerer, On-Off Intermittency: Power Spectrum and Fractal Properties of Time Series, Physica D 96, 1996, 66-99.

Wagner, F., Volatility Cluster and Herding, Physica A 322, 2003, 607-619.

Westerhoff, F., Speculative Markets and the Effectiveness of Price Limits, Journal of Economic Dynamics & Control 28, 2003, 493-508.

Wilpert, M., Künstliche Aktienmarktmodelle auf der Basis von Classifier Systems, Ph.D. thesis, University of Kiel, 2003.

Youssefmir, M. and A. Huberman, Clustered Volatility in Multiagent Dynamics, Journal of Economic Behavior and Organization 32, 1997, 101-118.

Zeeman, E. C., On the Unstable Behaviour of Stock Exchanges, Journal of Mathematical Economics 1, 1974, 39-49.