

Empirical properties of asset returns: stylized facts and statistical issues

Rama Cont¹

Centre de Mathématiques Appliquées, Ecole Polytechnique, F-91128
Palaiseau, France

E-mail: Rama.Cont@polytechnique.fr

Received 28 October 2000

Abstract

We present a set of stylized empirical facts emerging from the statistical analysis of price variations in various types of financial markets. We first discuss some general issues common to all statistical studies of financial time series. Various statistical properties of asset returns are then described: distributional properties, tail properties and extreme fluctuations, pathwise regularity, linear and nonlinear dependence of returns in time and across stocks. Our description emphasizes properties common to a wide variety of markets and instruments. We then show how these statistical properties invalidate many of the common statistical approaches used to study financial data sets and examine some of the statistical problems encountered in each case.

Although statistical properties of prices of stocks and commodities and market indexes have been studied using data from various markets and instruments for more than half a century, the availability of large data sets of high-frequency price series and the application of computer-intensive methods for analysing their properties have opened new horizons to researchers in empirical finance in the last decade and have contributed to the consolidation of a data-based approach in financial modelling.

The study of these new data sets has led to the settlement of some old disputes regarding the nature of the data but has also generated new challenges. Not the least of them is to be able to capture in a synthetic and meaningful fashion the information and properties contained in this huge amount of data. A set of properties, common across many instruments, markets and time periods, has been observed by independent studies and classified as ‘stylized facts’. We present here a pedagogical overview of these stylized facts. With respect to previous reviews [10, 14, 16, 50, 95, 102, 109] on the same subject, the aim of the present paper is to focus more on the properties of empirical data than on those of statistical models and introduce the reader to some new insights provided by methods based on statistical techniques recently applied in empirical finance.

Our goal is to ‘let the data speak for themselves’ as much as possible. In terms of statistical methods, this is achieved by using so-called *non-parametric* methods which make only qualitative assumptions about the properties of the stochastic process generating the data: they do not assume that they belong to any prespecified parametric family.

Although non-parametric methods have the great theoretical advantage of being model free, they can only provide qualitative information about financial time series and in order to obtain a more precise description we will sometimes resort to semi-parametric methods which, without completely specifying the form of the price process, imply the existence of a parameter which describes a property of the process (for example the tail behaviour of the marginal distribution).

Before proceeding further, let us fix some notations. In the following, $S(t)$ will denote the price of a financial asset—a stock, an exchange rate or a market index—and $X(t) = \ln S(t)$ its logarithm. Given a *time scale* Δt , which can range from a few seconds to a month, the log return at scale Δt is defined as:

$$r(t, \Delta t) = X(t + \Delta t) - X(t). \quad (1)$$

In many econometric studies, Δt is set implicitly equal to one in appropriate units, but we will conserve all along the variable Δt to stress the fact the properties of the returns depend

¹ Web address: <http://www.cmap.polytechnique.fr/~rama>

(strongly) on Δt . Time lags will be denoted by the greek letter τ ; typically, τ will be a multiple of Δt in estimations. For example, if $\Delta t = 1$ day, $\text{corr}[r(s + \tau, \Delta t), r(s, \Delta t)]$ denotes the correlation between the daily return at period s and the daily return τ periods later. When Δt is small—for example of the order of minutes—one speaks of ‘fine’ scales whereas if Δt is large we will speak of ‘coarse-grained’ returns.

1. What is a stylized fact?

As revealed by a casual examination of most financial newspapers and journals, the view point of many market analysts has been and remains an *event-based* approach in which one attempts to ‘explain’ or rationalize a given market movement by relating it to an economic or political event or announcement [27]. From this point of view, one could easily imagine that, since different assets are not necessarily influenced by the same events or information sets, price series obtained from different assets and—*a fortiori*—from different markets will exhibit different properties. After all, why should properties of corn futures be similar to those of IBM shares or the Dollar/Yen exchange rate? Nevertheless, the result of more than half a century of empirical studies on financial time series indicates that this is the case if one examines their properties from a *statistical* point of view: the seemingly random variations of asset prices *do* share some quite non-trivial statistical properties. Such properties, common across a wide range of instruments, markets and time periods are called *stylized empirical facts*.

Stylized facts are thus obtained by taking a common denominator among the properties observed in studies of different markets and instruments. Obviously by doing so one gains in generality but tends to lose in precision of the statements one can make about asset returns. Indeed, stylized facts are usually formulated in terms of *qualitative properties* of asset returns and may not be precise enough to distinguish among different parametric models. Nevertheless, we will see that, albeit qualitative, these stylized facts are so constraining that it is not easy to exhibit even an (*ad hoc*) stochastic process which possesses the same set of properties and one has to go to great lengths to reproduce them with a model.

2. Stylized statistical properties of asset returns

Let us start by stating a set of stylized statistical facts which are common to a wide set of financial assets.

- 1. Absence of autocorrelations:** (linear) autocorrelations of asset returns are often insignificant, except for very small intraday time scales ($\simeq 20$ minutes) for which microstructure effects come into play.
- 2. Heavy tails:** the (unconditional) distribution of returns seems to display a power-law or Pareto-like tail, with a tail index which is finite, higher than two and less than five for most data sets studied. In particular this excludes stable laws with infinite variance and the normal distribution. However the precise form of the tails is difficult to determine.

- 3. Gain/loss asymmetry:** one observes large drawdowns in stock prices and stock index values but not equally large upward movements².
- 4. Aggregational Gaussianity:** as one increases the time scale Δt over which returns are calculated, their distribution looks more and more like a normal distribution. In particular, the shape of the distribution is not the same at different time scales.
- 5. Intermittency:** returns display, at any time scale, a high degree of variability. This is quantified by the presence of irregular bursts in time series of a wide variety of volatility estimators.
- 6. Volatility clustering:** different measures of volatility display a positive autocorrelation over several days, which quantifies the fact that high-volatility events tend to cluster in time.
- 7. Conditional heavy tails:** even after correcting returns for volatility clustering (e.g. via GARCH-type models), the residual time series still exhibit heavy tails. However, the tails are less heavy than in the unconditional distribution of returns.
- 8. Slow decay of autocorrelation in absolute returns:** the autocorrelation function of absolute returns decays slowly as a function of the time lag, roughly as a power law with an exponent $\beta \in [0.2, 0.4]$. This is sometimes interpreted as a sign of long-range dependence.
- 9. Leverage effect:** most measures of volatility of an asset are negatively correlated with the returns of that asset.
- 10. Volume/volatility correlation:** trading volume is correlated with all measures of volatility.
- 11. Asymmetry in time scales:** coarse-grained measures of volatility predict fine-scale volatility better than the other way round.

3. Some issues about statistical estimation

Before proceeding to present empirical results let us recall some general issues which are implicit in almost any statistical analysis of asset returns. These issues have to be kept in mind when interpreting statistical results, especially for scientists with a background in the physical sciences where orders of magnitude may be very different.

3.1. Stationarity

‘Past returns do not necessarily reflect future performance’. This warning figures everywhere on brochures describing various funds and investments. However the most basic requirement of any statistical analysis of market data is the existence of *some* statistical properties of the data under study which remain stable over time, otherwise it is pointless to try to identify them.

The invariance of statistical properties of the return process in time corresponds to the *stationarity* hypothesis:

² This property is not true for exchange rates where there is a higher symmetry in up/down moves.

which amounts to saying that for any set of time instants $t_1, \dots, t_k$ and any time interval τ the joint distribution of the returns $r(t_1, T), \dots, r(t_k, T)$ is the same as the joint distribution of returns $r(t_1 + \tau, T), \dots, r(t_k + \tau, T)$. It is not obvious whether returns verify this property in calendar time: seasonality effects such as intraday variability, weekend effects, January effects. . . . In fact this property may be taken as a definition of the time index t , defined as the proper way to ‘deform’ calendar time in order to obtain stationarity. This time deformation is chosen to correct for seasonalities observed in calendar time and is therefore usually a cumulative measure of market activity: the number of transactions (tick time) [3], the volume of transactions [19] or a sample-based measure of market activity (see work by Dacorogna and coworkers [97, 105] and also [2]).

3.2. Ergodicity

While stationarity is necessary to ensure that one can mix data from different periods in order to estimate moments of the returns, it is far from being sufficient: one also needs to ensure that empirical averages do indeed converge to the quantities they are supposed to estimate! For example one typically wants to identify the sample moment defined by

$$\langle f(r(t, T)) \rangle = \frac{1}{N} \sum_{t=1}^N f(r(t, T)) \quad (2)$$

with the theoretical expectation $Ef(r(t, T))$ where E is the expectation (ensemble average) with respect to the distribution F_T of $r(t, T)$. Stationarity is necessary to ensure that F_T does not depend on t , enabling the use of observations at different times to compute the sample moment. But it is not sufficient to ensure that the sum indeed converges to the desired expectation. One needs an *ergodic* property which ensures that the time average of a quantity converges to its expectation. Ergodicity is typically satisfied by IID observations but it is not obvious—in fact it may be very hard to prove or disprove—for processes with complicated dependence properties such as the ones observed in asset returns (see below). In fact failure of ergodicity is not uncommon in physical systems exhibiting long-range dependence [12]. This may also be the case for some multifractal processes recently introduced to model high-frequency asset returns [90, 100], in which case the relation between sample averages and model expectations remains an open question.

3.3. Finite sample properties of estimators

Something which seems obvious to any statistician but which is often forgotten by unsophisticated users of statistics is that a statistical estimator, which is defined as a sample average, has no reason to be equal to the quantity it estimates, which is defined as a moment of the theoretical (unknown) distribution of the observations (ensemble average).

This confusion is frequent in fields where ‘sample sizes’ are large: for example, in statistical mechanics where one frequently identifies the sample average of a microscopic quantity and its expected value or *ensemble average*. Since




Figure 1. Daily returns of BMW shares on the Frankfurt Stock Exchange, 1992–1998.

in a typical macroscopic system the number of particles is around the Avogadro number $N = 10^{23}$, the relative error is of order $N^{-1/2} \simeq 10^{-12}$! The issue is completely different when examining a data set of five years of daily returns of a stock index where $N \simeq 10^3$. In this case, a statistic without a ‘confidence interval’ becomes meaningless and one needs to know something about the distribution of the estimator itself.

There is a long tradition of ‘hypothesis testing’ in financial econometrics in which one computes the likelihood of a model to hold given the value of a statistic and rejects/accepts the model by comparing the test statistics to a threshold value. With a few exceptions (see [8, 87, 88]), the large majority of statistical tests are based on a central limit theorem for the estimator from which the asymptotic normality is obtained. This central limit theorem can be obtained by assuming that the noise terms (innovations) in the return process are ‘weakly’ dependent [30]. In order to obtain confidence intervals for finite samples, one often requires the residuals to be IID and some of their higher-order (typically fourth order) moments to be well defined (finite). As we shall see below, the properties of empirical data—especially the heavy tails and nonlinear dependence—do not seem to, in general, validate such hypotheses, which raises the question of the meaning and relevance of such confidence intervals. As we will discuss below, this can have quite an impact on the significance and interpretation of commonly used estimators (see also discussions in [1, 31, 107]).

4. The distribution of returns: a tale of heavy tails

Empirical research in financial econometrics in the 1970s mainly concentrated on modelling the unconditional distribution of returns, defined as:

$$F_T(u) = P(r(t, T) \leq u). \quad (3)$$




Figure 2. Kernel estimator of the density of 30 minute price increments. S&P 500 index futures.

The probability density function (PDF) is then defined as its derivative $f_T = F'_T$. As early as the 1960s, Mandelbrot [80] pointed out the insufficiency of the normal distribution for modelling the marginal distribution of asset returns and their heavy-tailed character. Since then, the non-Gaussian character of the distribution of price changes has been repeatedly observed in various market data. One way to quantify the deviation from the normal distribution is by using the kurtosis of the distribution F_T defined as

$$\kappa = \frac{\langle (r(t, T) - \langle r(t, T) \rangle)^4 \rangle}{\sigma(T)^4} - 3, \quad (4)$$

where $\sigma(T)^2$ is the variance of the log returns $r(t, T) = x(t + T) - x(t)$. The kurtosis is defined such that $\kappa = 0$ for a Gaussian distribution, a positive value of κ indicating a ‘fat tail’, that is, a slow asymptotic decay of the PDF. The kurtosis of the increments of asset prices is far from its Gaussian value: typical values for $T = 5$ minutes are (see table 1): $\kappa \simeq 74$ (US\$/DM exchange rate futures), $\kappa \simeq 60$ (US\$/Swiss Franc exchange rate futures), $\kappa \simeq 16$ (S&P500 index futures) [16, 21, 22, 102].

One can summarize the empirical results by saying that the distribution $f_{\Delta t}$ tends to be non-Gaussian, sharp peaked and heavy tailed, these properties being more pronounced for




Figure 3. Second empirical moment of five minute price changes as a function of sample size. S&P index futures.

intraday values of T ($T < 1$ day). Figure 2 gives an example of the density of returns for $T = 30$ minutes.

These features are not sufficient for identifying the distribution of returns and leave a considerable margin for the choice of the distribution. Fitting various functional forms to the distribution of stock returns and stock price changes has become a popular pastime: there are dozens of parametric models proposed in the literature, starting with the normal distribution, stable distributions [80], the Student distribution [9, 72], hyperbolic distributions [37, 104], normal inverse Gaussian distributions [7], exponentially truncated stable distributions [11, 21] are some of the parametric models which have been proposed. From the empirical features described above, one can conclude that, in order for a parametric model to successfully reproduce all the above properties of the marginal distributions it must have at least four parameters: a location parameter, a scale (volatility) parameter, a parameter describing the decay of the tails and eventually an asymmetry parameter allowing the left and right tails to have different behaviours. For example, normal inverse Gaussian distributions [7], generalized hyperbolic distributions [104] and exponentially truncated stable distributions [11, 21] meet these requirements. The choice among these classes is then a matter of analytical and numerical tractability.

4.1. How heavy are the tails of the distribution?

The non-Gaussian character of the distribution makes it necessary to use other measures of dispersion than the standard deviation in order to capture the variability of returns. One can consider for example higher-order moments or cumulants as measures of dispersion and variability. However, given the heavy-tailed nature of the distribution, one has to know beforehand whether such moments are well defined. The tail index k of a distribution may be defined as the order of the highest absolute moment which is finite. The higher the

Table 1. Descriptive statistics for five minute price increments.

Data	μ/σ	Skewness	Kurtosis
S&P 500 futures	0.003	−0.4	15.95
Dollar/ DM futures	0.002	−0.11	74
Dollar/ Swiss Franc futures	0.002	−0.1	60
IID 95% confidence interval	—	0.018	0.036




Figure 4. Fourth empirical moment of a Student distribution with four degrees of freedom as calculated from a data set obtained from a random number generator.

tail index, the thinner the tail; for a Gaussian or exponential tail, $k = +\infty$ (all moments are finite), while for a power-law distribution with exponent α , the tail index is equal to α . But, as we shall see below, a distribution may have a finite tail index α without being a power-law distribution. Measuring the tail index of a distribution gives a measure of how heavy the tail is.

A simple method, suggested by Mandelbrot [80,89], is to represent the sample moments (or cumulants) as a function of the sample size n . If the theoretical moment is finite then the sample moment will eventually settle down to a region defined around its theoretical limit and fluctuate around that value. In the case where the true value is infinite the sample moment will either diverge as a function of sample size or exhibit erratic behaviour and large fluctuations. Applying this method to time series of cotton prices in [80], Mandelbrot conjectured that the theoretical variance may be infinite since the sample variance did not converge to a particular value as the sample size increased and continued to fluctuate incessantly.

Figure 3 indicates an example of the behaviour of the sample variance as a function of sample size. The behaviour of sample variance suggests that the variance of the distribution is indeed finite: the sample variance settles down to a limit value after a transitory phase of wild oscillations. A systematic analysis on a wide range of US and French stocks yields similar results [22].

The behaviour of the fourth moment is usually more erratic. The standard deviation of the sample kurtosis involves a polynomial containing the theoretical moments up to order eight! The eighth moment of the distribution having a very large numerical value, it is not surprising to see the fourth moment fluctuate wildly. As an illustration of this phenomenon we have estimated (figure 4) the fourth moment for a numerically generated series of IID random variables with a Student distribution with four degrees of freedom which displays a tail behaviour similar to many asset returns, with a




Figure 5. The 1987 crash: evolution of the S&P500 index.

power-law decay of exponent four. It can be seen from figure 4 that the statistical fluctuations are very strong even for such a familiar distribution, suggesting that fourth or higher-order moments are not numerically stable as quantitative measures of risk. This is linked to the fact that the behaviour of sample moments is controlled by higher-order theoretical moments, which may be infinite. Going beyond the graphical analysis described above, the next section describes how *extreme value theory* may be used to estimate the tail index of returns.

4.2. Extreme values

One of the important characteristics of financial time series is their high variability, as revealed by the heavy-tailed distributions of their increments and the non-negligible probability of occurrence of violent market movements. These large market movements, far from being discardable as simple outliers, focus the attention of market participants since their magnitude may be such that they compose an important fraction of the return aggregated over a long period: figure 5 illustrates such an example. These observations have motivated numerous theoretical efforts to understand the intermittent nature of financial time series and to model adequately the tails of the distribution of returns. Not only are such studies of direct relevance for risk management purposes but they are rendered necessary for the calculation of the *Value-at-Risk*, which is required to determine regulatory capital requirements. Value-at-Risk (VaR) is defined as a high quantile of the loss distribution of a portfolio over a certain time horizon

Δt :

$$P(W_0(r(t, \Delta t) - 1) \leq \text{VaR}(p, t, \Delta t)) = p \quad (5)$$

where W_0 is the present market value of the portfolio, $r(t, \Delta t)$ its (random) return between t and $t + \Delta t$. Δt is typically taken to be one day or ten days and $p = 1\%$ or 5% . Calculating VaR implies a knowledge of the tail behaviour of the distribution of returns. In recent years there has been an upsurge of interest in modelling the tails of the distributions of stock returns using the tools of *extreme value theory*, a branch of probability theory dealing precisely with the probabilities of extreme events. To our knowledge, the first application of extreme value theory to financial time series was given by Jansen and de Vries [70], followed by Longin [76], Dacorogna *et al* [28], Lux [77] and others.

Given a series of n non-overlapping returns $r(t, \Delta t)$, $t = 0, \Delta t, 2\Delta t, \dots, n\Delta t$, the extremal (minimal and maximal) returns are defined as:

$$m_n(\Delta t) = \min\{r(t + k\Delta t, \Delta t), k \in [1, n]\}, \quad (6)$$

$$M_n(\Delta t) = \max\{r(t + k\Delta t, \Delta t), k \in [1, n]\}. \quad (7)$$

In economic terms, $m_n(\Delta t)$ represents the worst relative loss over a time horizon Δt of an investor holding the portfolio $P(t)$. A relevant question is to know the properties of these extremal returns, for example the distribution of $m_n(\Delta t)$ and $M_n(\Delta t)$. More generally one is interested in the properties of large price fluctuations (not only minima and maxima). Obviously if one knew the stochastic process generating the returns, one could also evaluate the distribution of the extremes, but this is unfortunately not the case as attested by the zoology of parametric models used to fit the marginal distribution of returns! This is where extreme value theory comes into play: in this approach, one looks for a distributional limit of $m_n(\Delta t)$ and $M_n(\Delta t)$ as the sample size n increases. If such a limit exists, then it is described by the Fisher–Tippett theorem in the case where the returns are IID.

Extreme value theorem for IID sequence [38].

Assume the log returns $(r(t, \Delta t))_{t \geq 0}$ form an IID sequence with distribution $F_{\Delta t}$. If there exist normalizing constants (λ_n, σ_n) and a non-degenerate limit distribution H for the normalized maximum return:

$$P\left(\frac{M_n - \lambda_n}{\sigma_n} \leq x\right) \xrightarrow{x \rightarrow \infty} H(x) \quad (8)$$

then the limit distribution H is either a Gumbel, Weibull or Fréchet distribution (see table 2).

The three distributional forms can be parametrized in the following unified form, called the Cramer–von Mises parametrization:

$$H_\xi(x) = \exp[-(1 + \xi x)^{-1/\xi}] \quad (9)$$

where the sign of the shape parameter ξ determines the extremal type: $\xi > 0$ for Fréchet, $\xi < 0$ for Weibull and $\xi = 0$ for Gumbel. This result implies that one need not know the exact parametric form of the marginal distribution

of returns F to evaluate the distribution of extremal returns. The value of ξ *only* depends on the tail behaviour of the distribution $F_{\Delta t}$ of the returns: a distribution $F_{\Delta t}$ with finite support gives $\xi < 0$ (Weibull) while a distribution $F_{\Delta t}$ with a power-law tail with exponent α falls in the Fréchet class with $\xi = 1/\alpha > 0$. The Fréchet class therefore contains most ‘heavy-tailed’ distributions. All other distributions fall in the Gumbel class $\xi = 0$ which plays a role for extreme values analogous to that of the normal distribution for sums of random variables: it is the typical limit for the distribution of IID extremes. For example, the normal, log-normal and exponential distribution fall in the Gumbel class, as well as most distributions with an infinite tail index (see section 4.1 for a definition).

This theorem also provides a theoretical justification for using a simple parametric family of distributions for estimating the extremal behaviour of asset returns. The estimation may be done as follows: one interprets the asymptotic result above as

$$P(M_n \leq u) = H_\xi\left(\frac{u - \lambda_n}{\sigma_n}\right) = H_{\xi, \lambda_n, \sigma_n}(x). \quad (10)$$

The estimation of the distribution of maximal returns then is reduced to a parameter estimation problem for the three-parameter family $H_{\xi, \lambda, \sigma}$. One can estimate these parameters by the so-called *block* method [38, 76]: one divides the data into N subperiods of length n and takes the extremal returns in each subperiod, obtaining a series of N extremal returns $(x_i)_{i=1, \dots, N}$, which is then assumed to be an IID sequence with distribution $H_{\xi, \lambda, \sigma}$. A maximum likelihood estimator of (ξ, λ, σ) can be obtained by maximizing the log-likelihood function:

$$L(\lambda, \sigma, \xi) = \sum_{i=1}^N l(\lambda, \sigma, \xi, x_i) \quad (11)$$

where l is the log density obtained by differentiating equation (9) and taking logarithms:

$$l(\lambda, \sigma, \xi, x_i) = -\ln \sigma - \left(1 + \frac{1}{\xi}\right) \ln \left[1 + \xi \left(\frac{x_i - \lambda}{\sigma}\right)\right] - \left[1 + \xi \left(\frac{x_i - \lambda}{\sigma}\right)\right]^{1/\xi}. \quad (12)$$

If $\xi > -1$ (which covers the Gumbel and Fréchet cases), the maximum likelihood estimator is asymptotically normal and well behaved [38].

These methods, when applied to daily returns of stocks, market indices and exchange rates, yield a positive value of ξ between 0.2 and 0.4, which means a tail index $2 < \alpha(T) \leq 5$ [64, 70, 76, 77]. In all cases, ξ is bounded away from zero, indicating heavy tails belonging to the Fréchet domain of

Table 2. Limit distributions for extreme values. Here $1_{x>0}$ and $1_{x \leq 0}$ are indicator functions.

Gumbel	$H(x) = \exp(-e^{-x})$
Fréchet	$H(x) = \exp(-x^{-\alpha}) 1_{x>0}$
Weibull	$H(x) = \exp(-(-x)^{-\alpha}) 1_{x \leq 0} + 1_{x>0}$




Figure 6. Autocorrelation function of USD/Yen exchange rate returns. Time scale: ticks.

attraction but the tail index is found to be larger than two—which means that the variance is finite and the tails lighter than those of stable Lévy distributions [41], but compatible with a power-law (Pareto) tail with (the same) exponent $\alpha(T) = 1/\xi$. These studies seem to validate the power-law nature of the distribution of returns, with an exponent around three, using a direct log–log regression on the histogram of returns [52]. Note however that these studies do not allow us to pinpoint the exponent with more than a single significant digit. Also, a positive value of ξ does not *imply* power-law tails [12] but is compatible with any regularly varying tail with exponent $\alpha = 1/\xi$ [38]:

$$F_{\Delta t}(x) \underset{x \rightarrow \infty}{\sim} \frac{L(x)}{x^\alpha} \quad (13)$$

where $L(\cdot)$ verifies $\forall y > 0, L(xy)/L(x) \rightarrow 1$ when $x \rightarrow \infty$. L is then called a slowly-varying function, the logarithm and the constant function being particular examples. Any choice of L will give a different distribution $F_{\Delta t}$ of returns but the *same* extremal type $\xi = 1/\alpha$, meaning that, in the Fréchet class, the extremal behaviour only identifies the tail behaviour up to a (unknown!) slowly-varying function which may considerably influence in turn the results of the log–log fit on the histogram!

A more detailed study on high-frequency data using different methods [28] indicates that the tail index varies only slightly when the time resolution moves from an intraday (30 minutes) to a daily scale [64], indicating a relative stability of the tails. However, the IID hypothesis underlying these estimation procedure has to be treated with caution given the dependence present in asset returns (see section 5 and [25]).

5. Dependence properties of returns

5.1. Absence of linear autocorrelation

It is a well-known fact that price movements in liquid markets do not exhibit any significant autocorrelation: the




Figure 7. Autocorrelation function of tick by tick returns on KLM shares traded on the NYSE. Time scale: ticks.

autocorrelation function of the price changes

$$C(\tau) = \text{corr}(r(t, \Delta t), r(t + \tau, \Delta t)) \quad (14)$$

(where corr denotes the sample correlation) rapidly decays to zero in a few minutes (see figures 6 and 7): for $\tau \geq 15$ minutes it can be safely assumed to be zero for all practical purposes [21]. The absence of significant linear correlations in price increments and asset returns has been widely documented [43, 102] and is often cited as support for the ‘efficient market hypothesis’ [44]. The absence of correlation is intuitively easy to understand: if price changes exhibit significant correlation, this correlation may be used to conceive a simple strategy with positive expected earnings; such strategies, termed *statistical arbitrage*, will therefore tend to reduce correlations except for very short time scales, which represent the time the market takes to react to new information. This correlation time is typically several minutes for organized futures markets and even shorter for foreign exchange markets. Mandelbrot [85] expressed this property by stating that ‘arbitrage tends to whiten the spectrum of price changes’. This property implies that traditional tools of signal processing which are based on second-order properties, in the time domain—autocovariance analysis, ARMA modelling—or in the spectral domain—Fourier analysis, linear filtering—cannot distinguish between asset returns and white noise. This points out the need for nonlinear measures of dependence in order to characterize the dependence properties of asset returns.

In high-frequency return series of *transaction* prices, one actually observes a negative autocorrelation at very short lags (typically, one or a few trades). This is traditionally attributed to the bid–ask bounce [16]: transaction prices may take place either close to the ask or closer to the bid price and tend to bounce between these two limits. However, one also observes negative autocorrelations at the first lag in bid or ask prices themselves, suggesting a fast mean reversion of the price at the tick level. This feature may be attributed to the action of a market maker [47].

The absence of autocorrelation does not seem to hold systematically when the time scale Δt is increased: weekly and monthly returns *do* exhibit some autocorrelation. However given that the sizes of the data sets are inversely proportional to Δt the statistical evidence is less conclusive and more variable from sample to sample.

5.2. Volatility clustering and nonlinear dependence

The absence of autocorrelations in return gave some empirical support for ‘random walk’ models of prices in which the returns are considered to be independent random variables. However it is well known that the absence of serial correlation does not imply the independence of the increments: independence implies that any nonlinear function of returns will also have no autocorrelation. This property does not hold however: simple nonlinear functions of returns, such as absolute or squared returns, exhibit significant positive autocorrelation or *persistence*. This is a quantitative signature of the well-known phenomenon of *volatility clustering*: large price variations are more likely to be followed by large price variations. Figure 1 illustrates this phenomenon on daily returns of BMW shares. Log prices are therefore not random walks [17, 21].

A quantity commonly used to measure volatility clustering is the autocorrelation function of the squared returns:

$$C_2(\tau) = \text{corr}(|r(t + \tau, \Delta t)|^2, |r(t, \Delta t)|^2). \quad (15)$$

Empirical studies using returns from various indices and stocks indicate that this autocorrelation function remains positive and decays slowly, remaining significantly positive over several days, sometimes weeks [10, 20–22, 34, 35, 39]. This is sometimes called the ‘ARCH effect’ in the econometric literature since it is a feature of (G)ARCH models [10, 39] but it is important to keep in mind that it is a model-free property of returns which does not rely on the GARCH hypothesis. This persistence implies some degree of predictability for the amplitude of the returns as measured by their squares. In the same way one can study autocorrelation functions of various powers of the returns:

$$C_\alpha(\tau) = \text{corr}(|r(t + \tau, \Delta t)|^\alpha, |r(t, \Delta t)|^\alpha). \quad (16)$$

Comparing the decay of C_α for various values of α , Ding and Granger [34, 35] remarked that, for a given lag τ , this correlation is highest for $\alpha = 1$, which means that absolute returns are more predictable than other powers of returns. Several authors [11, 20–22, 54, 55, 59, 105] have remarked that the decay of $C_\alpha(\tau)$ as τ increases is well reproduced by a power law:

$$C_\alpha(\tau) \sim \frac{A}{\tau^\beta} \quad (17)$$

with a coefficient $\beta \in [0.2, 0.4]$ for absolute or squared returns [21, 22, 74]. This slow decay is sometimes interpreted as a sign of long-range dependence in volatility and motivated the development of models integrating this feature (see below). More generally, one can ask what is the nonlinear function of the returns which yields the highest predictability, i.e., which maximizes the one-lag correlation? This question, which is




Figure 8. Behaviour of some nonlinear correlation functions of price changes.

the object of ‘canonical correlation analysis’, can yield more insight into the dependence properties of returns [29]. Some examples of autocorrelations of different nonlinear transforms of returns are compared in figure 8. These autocorrelations are actually weighted sums of covariances of various integer powers of returns, weighted by the coefficients of the Taylor expansion of the nonlinear transform considered [22]. Recent work [90, 100] on multifractal stochastic volatility models has motivated yet another measure of nonlinear dependence based on correlations of the logarithm of absolute returns:

$$C_0(\tau) = \text{corr}(\ln |r(t + \tau, \Delta t)|, \ln |r(t, \Delta t)|). \quad (18)$$

Muzy *et al* [100] show that this function also exhibits a slow decay, which they represent by a logarithmic form over a certain range of values:

$$C_0(\tau) = a \ln \frac{b}{\Delta t + \tau}. \quad (19)$$

Another measure of nonlinear dependence in returns is the so-called ‘leverage effect’: the correlation of returns with subsequent squared returns defined by

$$L(\tau) = \text{corr}(|r(t + \tau, \Delta t)|^2, r(t, \Delta t)) \quad (20)$$

starts from a *negative* value and decays to zero [13, 102], suggesting that negative returns lead to a rise in volatility. However this effect is asymmetric $L(\tau) \neq L(-\tau)$ and in general $L(\tau)$ is negligible for $\tau < 0$.

The existence of such nonlinear dependence, as opposed to absence of autocorrelation in returns themselves, is usually interpreted by stating that there is correlation in ‘volatility’ of returns but not the returns themselves. These observations motivate a decomposition of the return as a product

$$r(t, \Delta t) = \sigma(t, \Delta t)\epsilon(t) \quad (21)$$

where $\epsilon(t)$ is a white noise, uncorrelated in time, and $\sigma(t, \Delta t) > 0$ a conditional volatility factor whose dynamics should be specified to match the empirically observed dependences. Examples of models in this direction are GARCH models [10,39] and long-memory stochastic volatility models [20,59,100]. Note however that in this decomposition the volatility variable $\sigma(t, \Delta t)$ is not directly observable, only the returns $r(t, \Delta t)$ are. Therefore, the definition of ‘volatility’ is model dependent and ‘volatility correlations’ are not observable as such, whereas the correlations of absolute returns *are* computable.

5.3. How reliable are autocorrelation functions?

Autocorrelation functions (ACF) were originally developed as a tool for analysing dependence for Gaussian time series and linear models, for which they adequately capture the dependence structure of the series under study. This is less obvious when one considers nonlinear, non-Gaussian time series such as the ones we are dealing with. In particular, the heavy-tailed feature of these time series can make the interpretation of a sample ACF problematic.

As shown by Davis and Mikosch [31], the sample ACF of heavy-tailed nonlinear time series can have non-standard statistical properties which invalidate many econometric testing procedures used for detecting or measuring dependence. In particular, if the marginal distribution of the returns has an infinite fourth moment—a property which is suggested by studies using extreme value techniques (see section 4.2)—then, although the sample ACF remains a consistent estimator of the theoretical ACF, the rate of convergence is slower than $\sqrt{n}$ and, more importantly, asymptotic confidence bands for sample ACFs are wider than classical ones. The situation is even worse for sample ACFs of the *squares* of the returns, which are classically used to measure volatility clustering. For example, the autocorrelation coefficient of the squares of the returns is often used as a moment condition for fitting GARCH models to financial time series [48]. First, in order for autocorrelations of squared returns to be well defined, one needs finiteness of fourth moments of returns. On the other hand, the statistical analysis of large returns (see section 4.2) indicates that the tail index obtained for most assets is typically close to four (sometimes less) which means that the fourth moment is not a well-defined numerically stable quantity. This means that there exists a great deal of variability in sample autocorrelations of squared returns, which raises some doubts about the statistical significance of quantitative estimates derived from them.

These criticisms can be quantified if one considers analogous quantities for some time series models with fat-tailed marginals, such as GARCH. In a critical study of GARCH models, Mikosch and Starica [96] show that the ACF of the squared returns in GARCH(1,1) models can have non-standard sample properties and generate large confidence bands, which raises serious questions about the methods used to fit these models to empirical data.

To summarize, for such heavy-tailed time series, estimators of the autocorrelation function of returns and

their squares can be highly unreliable and even in cases where they are consistent they may have large confidence intervals associated with them. Therefore, one should be very careful when drawing quantitative conclusions from the autocorrelation function of powers of the returns.

6. Cross-asset correlations

While the methods described above are essentially univariate—they deal with one asset at a time—most practical problems in risk management deal with the management of portfolios containing a large number of assets (typically more than a hundred). The statistical analysis of the risk of such positions requires information on the joint distribution of the returns of different assets.

6.1. Covariances and correlations of returns

The main tool for analysing the interdependence of asset returns is the *covariance matrix* $\mathcal{C}$ of returns:

$$\mathcal{C}_{ij} = \text{cov}(r_i(t, T), r_j(t, T)). \quad (22)$$

The covariance between two assets may be seen as a product of three terms: the two assets’ volatilities and their *correlation* ρ_{ij} :

$$\mathcal{C}_{ij} = \sigma_i \sigma_j \rho_{ij} \quad \rho_{ij} \in [-1, 1]. \quad (23)$$

Obviously, the heteroskedastic nature of individual asset returns results in the instability in time of covariances: the covariance $\mathcal{C}_{ij}$ may vary, not because the correlation between the two assets change, but simply because their individual volatilities change. This effect can be corrected by considering the correlation matrix $C = [\rho_{ij}]$ instead of the covariance:

$$\mathcal{C} = \Sigma(t) C \Sigma(t) \quad (24)$$

where $\Sigma(t) = \text{diag}(\sigma_1(t), \dots, \sigma_n(t))$ is the diagonal matrix of conditional standard deviations. The matrix C may be estimated from time series of asset returns: $\mathcal{C}_{ij}$ is estimated by the sample correlation between assets i and j .

The most interesting features of the matrix C are its eigenvalues λ_i and its eigenvectors e_i , which have been usually interpreted in economic terms as factors of randomness underlying market movements.

In a recent empirical study of the covariance matrix of 406 NYSE assets, Laloux *et al* [78] (see also [103]) showed that among the 406 available eigenvalues and principal components, apart from the highest eigenvalue (whose eigenvector roughly corresponds to the market index) and the next few (ten) highest eigenvalues, the other eigenvectors and eigenvalues do not seem to contain any information: in fact, their marginal distribution closely resembles the spectral distribution of a positive symmetric matrix with random entries [94] whose distribution is the ‘most random possible’—i.e., entropy maximizing. These results strongly question the validity of the use of the sample covariance matrix as an input for portfolio optimization, as suggested by classical methods such as mean-variance optimization, and support the rationale behind factor models such as the CAPM and APT, where the

correlations between a large number of assets are represented through a small number of factors. To examine the residual correlations once the common factors have been accounted for, one can define *conditional* correlations by conditioning on an aggregate variable such as the market return before computing correlations [18].

6.2. Correlations of extreme returns

Independently of the significance of its information content, the covariance matrix has been criticized as a tool for measuring dependence because it is based on an averaging procedure which emphasizes the centre of the return distribution whereas correlations, which are used for portfolio diversification, are mainly useful in circumstances when stock prices undergo large fluctuations. In these circumstances, a more relevant quantity is the conditional probability of a large (negative) return in one stock given a large negative movement in another stock:

$$F_{ij}(x, y) = P(r_i < -x | r_j < -y). \quad (25)$$

For example, one can consider high (95% or 99%) level quantiles q_i for each asset i

$$e_{ij}(q) = P(r_i < -q_i | r_j < -q_j). \quad (26)$$

It is important to remark that two assets may have extremal correlations while their covariance is zero: covariance does not measure the correlation of extremes. Some recent theoretical work has been done in this direction using copulas [108] and multivariate extreme value theory [64, 112, 113], but a lot remains to be done on empirical grounds. For a recent review with applications to foreign exchange rate data see Hauksson *et al* [64].

7. Pathwise properties

The risky character of a financial asset is associated with the irregularity of the variations of its market price: risk is therefore directly related to the (un)smoothness of the trajectory and this is one crucial aspect of empirical data that one would like a mathematical model to reproduce.

Each class of stochastic models generates sample paths with certain local regularity properties. In order for a model to represent adequately the intermittent character of price variations, the local regularity of the sample paths should try to reproduce those of empirically observed price trajectories.

7.1. Hölder regularity

In mathematical terms, the regularity of a function may be characterized by its *local Hölder exponents*. A function f is h -Hölder continuous at point t_0 iff there exists a polynomial of degree $< h$ such that

$$|f(t) - P(t - t_0)| \leq K_{t_0} |t - t_0|^h \quad (27)$$

in a neighborhood of t_0 , where K_{t_0} is a constant. Let $C^h(t_0)$ be the space of (real-valued) functions which verify the above

property at t_0 . A function f is said to have local Hölder exponent α if for $h < \alpha$, $f \in C^h(t_0)$ and for $h > \alpha$, $f \notin C^h(t_0)$. Let $h_f(t)$ denote the local Hölder exponent of f at point t . If $h_f(t_0) \geq 1$ then f is differentiable at point t_0 , whereas a discontinuity of f at t_0 implies $h_f(t_0) = 0$. More generally, the higher the value of $h_f(t_0)$, the greater the local regularity of f at t_0 .

In the case of a sample path $X_t(\omega)$ of a stochastic process X_t , $h_{X(\omega)}(t) = h_\omega(t)$ depends on the particular sample path considered, i.e., on ω . There are however some famous exceptions: for example for fractional Brownian motion with self-similarity parameter H , $h_B(t) = 1/H$ almost everywhere with probability one, i.e., for almost all sample paths. Note however that no such results hold for sample paths of Lévy processes or even stable Lévy motion.

7.2. Singularity spectrum

Given that the local Hölder exponent may vary from sample path to sample path in the case of a stochastic process, it is not a robust statistical tool for characterizing signal roughness: the notion of a *singularity spectrum* of a signal was introduced to give a less detailed but more stable characterization of the local smoothness structure of a function in a ‘statistical’ sense.

Definition. Let $f : R \rightarrow R$ be a real-valued function and for each $\alpha > 0$ define the set of points at which f has local Hölder exponent h :

$$\Omega(\alpha) = \{t, h_f(t) = \alpha\}. \quad (28)$$

The *singularity spectrum* of f is the function $D : R^+ \rightarrow R$ which associates to each $\alpha > 0$ the Hausdorff–Besicovich dimension³ of $\Omega(\alpha)$:

$$D(\alpha) = \dim_{HB} \Omega(\alpha). \quad (29)$$

Using the above definition, one may associate to each sample path $X_t(\omega)$ of a stochastic process X_t its singularity spectrum $d_\omega(\alpha)$. If d_ω is ‘strongly dependent’ on ω then the empirical estimation of the singularity spectrum is not likely to give much information about the properties of the process X_t .

Fortunately, this turns out not to be the case: it has been shown that, for large classes of stochastic processes, the singularity spectrum is the same for almost all sample paths. Results due to Jaffard [68] show that a large class of Lévy processes verifies this property.

As defined above, the singularity spectrum of a function does not appear to be of any practical use since its definition involves first the continuous time ($\Delta t \rightarrow 0$) limit for determining the local Hölder exponents and second the determination of the Hausdorff dimension of the sets $\Omega(\alpha)$ which, as remarked already by Halsey *et al* [58], may be intertwined fractal sets with complex structures and impossible to separate on a point by point basis. The interest of physicists

³ The Hausdorff–Besicovich dimension is one of the numerous mathematical notions corresponding to the general concept of ‘fractal’ dimension. For details see [40].

and empirical researchers in singularity spectra was ignited by the work of Parisi and Frisch [101] and subsequently of Halsey *et al* [58] who, in different contexts, proposed a formalism for empirically computing the singularity spectrum from sample paths of the process. This formalism, called the multifractal formalism [58, 66, 67, 101], enables the singularity spectrum to be computed from sample moments ('structure functions') of the increments. More precisely, if the sample moments of the returns verify a scaling property

$$\langle |r(t, T)|^q \rangle \sim T^{\zeta(q)} \quad (30)$$

then the singularity spectrum $D(\alpha)$ is given by the Legendre transform of the scaling exponent $\zeta(q)$

$$\zeta(q) = 1 + \inf(q\alpha - D(\alpha)). \quad (31)$$

$\zeta(q)$ may be obtained by regressing $\log \langle |r(t, T)|^q \rangle$ against $\log T$. In the case of multifractal processes for which the scaling in equation (30) holds exactly, the Legendre transform (31) may be inverted to obtain $D(\alpha)$ from $\zeta(q)$. The technique was subsequently refined [62, 99] using the wavelet transform [92], who proposed an algorithm for determining the singularity spectrum from its wavelet transform [66, 67].

7.3. Singularity spectra of asset price series

These methods provide a framework to investigate pathwise regularity of price trajectories [4, 26, 45, 100]. A first surprising result is that the shape of the singularity spectrum does not depend on the asset considered: all series exhibit the same, 'inverted parabola' shape also observed by Fisher *et al* [45] on the USD/DEM high-frequency exchange rate data using the structure function method [101]. The spectra have a support ranging from 0.3 to 0.9 (with some variations depending on the data set chosen) with a maximum centred around 0.55–0.6. Note that 0.55–0.6 is the range of values of the 'Hurst exponent' reported in many studies of financial time series using the R/S or similar techniques, which is not surprising since the maximum of $D(\alpha)$ represents the 'almost-everywhere' Hölder exponent which is the one detected by 'global' estimators such as the R/S statistic. It should be noted that this non-trivial spectrum is very different from what one would expect from diffusion processes, Lévy processes or jump-diffusion processes used in continuous-time finance, for which the singularity spectrum is theoretically known. In fact, it indicates no discontinuous ('jump') component in the signal since the Hölder exponent does not extend down to zero. The rare examples of stochastic processes for which the singularity spectrum resembles the one observed in empirical data are stochastic cascades [90] or their causal versions, the multifractal random walks [6, 100].

One drawback of the singularity spectrum is that its finite sample properties are not well known. Veneziano *et al* [114] have investigated, in the context of study of width functions of river basins, the results obtained when a multifractal formalism is applied to a non-fractal data set such as the graph of a simple deterministic function (a parabola in [114]). They obtain a concave nonlinear shape for $\zeta(q)$ which is supposedly

the signature of multifractality and a non-trivial multifractal spectrum $D(\alpha)$ with an inverse parabolic shape, while the real spectrum in their case is reduced to two points! This leads one to believe that one will generically obtain non-trivial multifractal spectra even if the data set studied is not multifractal, in which case the results indicated above should be interpreted with extreme caution. Moreover, *even* in the case of a genuine multifractal process the convergence of the empirical spectrum to its true value can be rather slow. In any case the subject merits further study in order to obtain a suitable characterization of finite sample behaviour of the estimators proposed for the various quantities of interest.

As in [45], one can supplement such studies by applying the same techniques to Monte Carlo simulations of various stochastic models used in finance in order to check whether the peculiar shapes of the spectra obtained are not artefacts due either to small sample size or discretization. Our preliminary results [26] seem to rule out such a possibility. In addition, the three different multifractal formalisms yield similar estimators for the singularity spectrum.

8. Conclusion

In the preceding sections, we have tried to present, in some detail, a set of statistical facts which emerge from the empirical study of asset returns and which are common to a large set of assets and markets. The properties mentioned here are model free in the sense that they do not result from a parametric hypothesis on the return process but from rather general hypotheses of qualitative nature. As such, they should be viewed as *constraints* that a stochastic process has to verify in order to reproduce the statistical properties of returns accurately. Unfortunately, most currently existing models fail to reproduce all these statistical features at once, showing that they are indeed very constraining.

Finally, we should point out several issues we have *not* discussed here. One important question is whether a stylized empirical fact is relevant from an economic point of view. In other words can these empirical facts be used to confirm or rule out certain modelling approaches used in economic theory? Another question is whether these empirical facts are useful from a practitioner's standpoint. For example, does the presence of volatility clustering imply anything interesting from a practical standpoint for volatility forecasting? If it does, can this be put to use to implement a more effective risk measurement/management approach? Can one exploit such correlations to implement a volatility trading strategy? Or, how can one incorporate a measure of irregularity such as the singularity spectrum or the extremal index of returns in a measure of portfolio market risk? We leave these questions for future research.

Acknowledgments

The author thanks Rokhsaneh Ghandi for constant encouragement and Doyne Farmer for many helpful comments on this manuscript. This work has been supported by the CNRS Research Program on Quantitative Methods in Financial Modeling (FIQUAM).

References

- [1] Adler R, Feldman R and Taqqu M (eds) 1996 *A Practical Guide to Heavy Tails: Statistical Techniques for Analyzing Heavy Tailed Distributions* (Boston: Birkhauser)
- [2] Andersen T G and Bollerslev T 1997 Intraday periodicity and volatility persistence in financial markets *J. Empirical Finance* **4** 115–58
- [3] Ané T and Geman H 1999 Stochastic volatility and transaction time: an activity based volatility estimator *J. Risk* 57–69
- [4] Arnéodo A, Muzy J F and Sornette D 1998 Causal cascade in the stock market from the infrared to the ultraviolet *Euro. Phys. J. B* **2** 277–82
- [5] Ausloos M and Vandewalle N 1997 Coherent and random sequences in financial fluctuations *Physica A* **246** 454–9
- [6] Bacry E, Muzy J F and Delour J 2001 Multifractal random walks *Phys. Rev. E* at press
- [7] Barndorff-Nielsen O E 1997 Normal inverse Gaussian distributions and the modelling of stock returns *Scandinavian J. Statistics* **24** 1–13
- [8] Beran J 1992 Statistical methods for data with long range dependence *Stat. Sci.* **7** 404–27
- [9] Blattberg R and Gonnede N 1974 A comparison of stable and Student distribution as statistical models for stock prices *J. Business* **47** 244–80
- [10] Bollerslev T, Chou R C and Kroner K F 1992 ARCH modelling in finance *J. Econometrics* **52** 5–59
- [11] Bouchaud J-P and Potters M 1997 *Théorie des Risques Financiers* (Aléa-Saclay: Eyrolles)
- [12] Bouchaud J-P 2001 Power laws in economics and finance *Quantitative Finance* **1** 105–12
- [13] Bouchaud J-P, Matacz A and Potters M 2001 The leverage effect in financial markets: retarded volatility and market panic *Preprint* <http://xxx.lptthe.jussieu.fr/abs/cond-mat/0101120>
- [14] Brock W A and de Lima P J F 1995 Nonlinear time series, complexity theory and finance *Handbook of Statistics Volume 14: Statistical Methods in Finance* ed G Maddala and C Rao (New York: North-Holland)
- [15] Campbell J, Grossmann S and Wang J 1993 Trading volume and serial correlation in stock returns *Q. J. Economics* **108** 905–39
- [16] Campbell J, Lo A H and McKinlay C 1997 *The Econometrics of Financial Markets* (Princeton, NJ: Princeton University Press)
- [17] Campbell J, Lo A H and McKinlay C 1999 *A Non-Random Walk Down Wall Street* (Princeton, NJ: Princeton University Press)
- [18] Cizeau P, Potters M and Bouchaud J-P 2000 Correlations of extreme stock returns within a non-Gaussian one-factor model *Science & Finance Working Paper*
- [19] Clark P K 1973 A subordinated stochastic process model with finite variance for speculative prices *Econometrica* **41** 135–55
- [20] Comte F and Renault E 1996 Long memory continuous time models *J. Econometrics* **73** 101–50
- [21] Cont R, Potters M and J-P Bouchaud 1997 Scaling in stock market data: stable laws and beyond *Scale Invariance and Beyond (Proc. CNRS Workshop on Scale Invariance, Les Houches, 1997)* ed Dubrulle, Graner and Sornette (Berlin: Springer)
- [22] Cont R 1998 Statistical Finance: empirical and theoretical approaches to the statistical modelling of price variations in speculative markets *Doctoral Thesis, Université de Paris XI*
- [23] Cont R 1999 Statistical properties of financial time series *Mathematical Finance: Theory and Practice. Lecture Series in Applied Mathematics* ed R Cont and J Yong vol 1 (Beijing: Higher Education Press)
- [24] Cont R 1999 Modelling economic randomness: statistical mechanics of market phenomena *Statistical Physics on the Eve of the 21st Century* ed M Batchelor and L T Wille (Singapore: World Scientific)
- [25] Cont R 1998 Clustering and dependence in extreme market returns *Extremes: Risk and Safety. Statistical Extreme Value Theory and Applications, Gothenburg, Sweden, 18–22 August 1998*
- [26] Cont R 2000 Multiresolution analysis of financial time series *Working Paper*
- [27] Cutler D M, Poterba J M and Summers L 1989 What moves stock prices? *J. Portfolio Management* 4–12
- [28] Dacorogna M M, Müller U A, Pictet O V and de Vries C G 1992 The distribution of extremal foreign exchange rate returns in large data sets *Olsen and Associates Internal document UAM. 19921022*
- [29] Darolles S, Florens J P and Gouriéroux Ch 1999 Kernel based nonlinear canonical analysis *CREST Working Paper*
- [30] Davidson J 1997 *Stochastic Limit Theory* (Oxford: Oxford University Press)
- [31] Davis R A and Mikosch T 1998 Limit theory for the sample ACF of stationary process with heavy tails with applications to ARCH *Ann. Statistics* **26** 2049–80
- [32] Davis R A and Mikosch T 1999 The sample autocorrelations of financial time series models *EURANDOM Working Paper* 99–039
- [33] Davison A C and Cox D R 1989 Some simple properties of sums of random variables having long range dependence *Proc. R. Soc. A* **424** 255–62
- [34] Ding Z, Granger C W J and Engle R F 1983 A long memory property of stock market returns and a new model *J. Empirical Finance* **1** 83
- [35] Ding Z X and Granger C W J 1994 Stylized facts on the temporal distributional properties of daily data from speculative markets *University of California, San Diego, Working Paper*
- [36] Ding Z and Granger C W J 1996 Modelling volatility persistence of speculative returns: a new approach *J. Econometrics* **73** 185–216
- [37] Eberlein E, Keller U and Prause K 1998 New insights into smile, mispricing and value at risk: the hyperbolic model *J. Business* **71** 371–405
- [38] Embrechts P, Kluppelberg C and Mikosch T 1997 *Modelling Extremal Events*, (New York: Springer)
- [39] Engle R F 1995 *ARCH: Selected Readings* (Oxford: Oxford University Press)
- [40] Falconer K 1990 *Fractal Geometry: Mathematical Foundations and Applications* (New York: Wiley)
- [41] Fama E F 1963 Mandelbrot and the stable paretian hypothesis *J. Business* **36** 420–29
- [42] Fama E F 1965 The behavior of stock market prices *J. Business* **38** 34–105
- [43] Fama E F 1971 Efficient capital markets: a review of theory and empirical work *J. Finance* **25** 383–417
- [44] Fama E F 1991 Efficient capital markets: II *J. Finance* **46** 1575–613
- [45] Fisher A, Calvet L and Mandelbrot B 1998 Multifractal analysis of USD/DM exchange rates *Yale University Working Paper*
- [46] Ghashghaie S *et al* 1996 Turbulent cascades in foreign exchange markets *Nature* **381** 767
- [47] Goodhart C E and O'Hara M 1997 High frequency data in financial markets: issues and applications *J. Empirical Finance* **4** 73–114
- [48] Gouriéroux C 1997 *ARCH Models and Financial Applications* (Berlin: Springer)
- [49] Gouriéroux C, Jasiak J and Lefol G 1999 Intra-day market activity *J. Financial Markets* **2** 193–226

- [50] Gouriéroux Ch and Jasiak J 2001 *Econometrics of finance Manuscript*
- [51] Ghashghaie S *et al* 1996 Turbulent cascades in foreign exchange markets *Nature* **381** 767
- [52] Gopikrishnan P, Meyer M, Amaral L A N and Stanley H E 1998 Inverse cubic law for the distribution of stock price variations *Euro. Phys. J. B* **3** 139–40
- [53] Granger C W J and Morgenstern O 1970 *Predictability of Stock Market Prices* (Lexington, KY: Heath)
- [54] Granger C W J 1977 Long term dependence in common stock returns *J. Financial Economics* **4** 339–49
- [55] Granger C W J and Ding Z 1996 Varieties of long memory models *J. Econometrics* **73** 61–77
- [56] Guillaume D M, Dacorogna M M, Davé R R, Müller U A, Olsen R B and Pictet O V 1997 From the birds eye to the microscope: a survey of new stylized facts of the intra-day foreign exchange markets *Finance Stochastics* **1** 95–130
- [57] Hall J, Brorsen B and Irwin S 1989 The distribution of future prices: a test of the stable paretian and mixture of normals hypothesis *J. Financial Quantitative Anal.* **24** 105–16
- [58] Halsey T C, Jensen M H, Kadanoff L P, Procaccia I and Shraiman B L 1986 Fractal measures and their singularities: the characterization of strange sets *Phys. Rev. A* **33** 1141–51
- [59] Harvey A C 1998 Long memory in stochastic volatility *Forecasting volatility in financial markets* ed Knight and Satchell (Oxford: Butterworth-Heinemann) pp 307–20
- [60] Hsieh D A 1988 The statistical properties of daily exchange rates *J. Int. Economics* **13** 171–86
- [61] Hsieh D A 1989 Testing for nonlinear dependence in foreign exchange rates *J. Business* **62** 339–668
- [62] Hwang W and Mallat S 1994 Characterization of self-similar multifractals with wavelet maxima *Appl. Computational Harmonic Anal.* **1** 316–28
- [63] Härdle W 1990 *Applied Non Parametric Regression* (Cambridge: Cambridge University Press)
- [64] Hauksson H A, Dacorogna M M, Domenig T, Müller U and Samorodnitsky G 2001 Multivariate extremes, aggregation and risk estimation *Quantitative Finance* **1** 79–95
- [65] Hosking J R M 1996 Asymptotic distributions of the sample mean, autocovariances and autocorrelations of long memory time series *J. Econometrics* **73** 261–85
- [66] Jaffard S 1997 Multifractal formalism for functions I: results valid for all functions *SIAM J. M. Anal.* **28** 944–70
- [67] Jaffard S 1997 Multifractal formalism for functions II: self-similar functions *SIAM J. Math. Anal.* **28** 971–98
- [68] Jaffard S 1997 The multifractal nature of Lévy processes *Preprint*
- [69] Jain P C and Joh G-H 1988 The dependence between hourly prices and trading volume *J. Financial Quantitative Anal.* **23** 269–84
- [70] Jansen D W and de Vries C G 1991 On the frequency of large stock returns: putting booms and busts into perspective *Rev. Economics Statistics* **73** 18–24
- [71] Johansen A and Sornette D 1998 Stock market crashes are outliers *Euro. Phys. J. B* **1** 141–3
- [72] Kon S 1984 Models of stock returns: a comparison *J. Finance* **XXXIX** 147–65
- [73] Lamoureux C G and Lastrapes W D 1990 Heteroskedasticity in stock return data: volume versus GARCH effects *J. Finance* **XLV** 221–9
- [74] Liu Y, Cizeau P, Meyer M, Peng C K and Stanley H E 1997 Correlations in economic time series *Physica A* **245** 437
- [75] Lo A W 1991 Long term memory in stock market prices *Econometrica* **59** 1279–313
- [76] Longin F 1996 The asymptotic distribution of extreme stock market returns *J. Business* **63** 383–408
- [77] Lux T 1997 The limiting extremal behavior of speculative returns: an analysis of intraday data from the Frankfurt Stock exchange *University of Bonn Working Paper*
- [78] Laloux L, Cizeau P, Potters M and Bouchaud J-P 1999 Random matrix theory *Risk* **12** 69
- [79] Laloux L, Cizeau P, Bouchaud J-P and Potters M 1999 Noise dressing of financial correlation matrices *Phys. Rev. Lett.* **83** 1467
- [80] Mandelbrot B 1963 The variation of certain speculative prices *J. Business* **XXXVI** 392–417
- [81] Mandelbrot B 1966 Forecasts of future prices, unbiased markets and martingale models *J. Business* **39** 242–55
- [82] Mandelbrot B and Wallis J 1968 Noah, Joseph and operational hydrology *Water Resources Res.* **4** 909–18
- [83] Mandelbrot B and Van Ness J W 1968 Fractional Brownian motion, fractional Brownian noises and applications *SIAM Rev.* **10** 422–37
- [84] Mandelbrot B 1968 Forecasts of future prices, unbiased markets and martingale models *J. Business* **39** 242–55
- [85] Mandelbrot B 1971 When can prices be arbitrated efficiently? A limit to the validity of random walk and martingale models *Rev. Economics Statistics* **53** 225–36
- [86] Mandelbrot B 1974 Intermittent turbulence in self-similar cascades *J. Fluid Mechanics* **62** 331
- [87] Mandelbrot B 1975 Limit theorems on the self normalized range for weakly and strongly dependent processes *Zeitschrift für Wahrscheinlichkeitstheorie und verwandete Gebiete* **31** 271–85
- [88] Mandelbrot B and Taqqu M 1979 Robust R/S analysis of long serial correlation *Bull. Int. Stat. Institute* **II** pp 59–104
- [89] Mandelbrot B 1997 *Fractals and Scaling in Finance: Discontinuity, Concentration and Risk* (New York: Springer)
- [90] Mandelbrot B, Fisher A and Calvet L 1997 The multifractal model of asset returns *Cowles Foundation for Economic Research Working Paper*
- [91] Mandelbrot B 2001 Scaling in financial prices: tails and dependence *Quantitative Finance* **1** 113–23
- [92] Mallat S 1996 *A Wavelet Tour of Signal Processing* (New York: Academic)
- [93] Mantegna R N and Stanley H E 1995 Scaling behavior of an economic index *Nature* **376** 46–9
- [94] Mehta M L 1990 *Random Matrices* (New York: Academic)
- [95] Maddala and Rao (ed) 1997 *Handbook of Statistics: Statistical Methods in Finance* vol 14 (Amsterdam: North-Holland)
- [96] Mikosch T and Starica C 2000 Limit theory for the sample autocorrelations and extremes of a GARCH(1,1) process *Ann. Statistics* **24** no 4
- [97] Müller U A, Dacorogna M, Davé R D, Olsen R B, Pictet O V and von Weizsacker J E 1997 Volatilities of different time resolutions—analyzing the dynamics of market components *J. Empirical Finance* **4** 213–39
- [98] Müller U A, Dacorogna M A and Pictet O V 1996 Heavy tails in high-frequency financial data *Olsen and Associates, Zürich, Preprint*
- [99] Muzy J F, Bacry E and Arneodo A 1994 The multifractal formalism revisited with wavelets *Int. J. Bifurcation Chaos* **4** 245
- [100] Muzy J F, Delour J and Bacry E 2000 Modelling fluctuations of financial time series: from cascade process to stochastic volatility model *Euro. Phys. J. B* at press
- [101] Parisi G and Frisch U 1985 On the singularity structure of fully developed turbulence *Proc. Int. School of Physics Enrico Fermi 1983* ed Ghil *et al* (Amsterdam: North-Holland) pp 84–7
- [102] Pagan A 1996 The econometrics of financial markets *J. Empirical Finance* **3** 15–102
- [103] Plerou V, Gopikrishnan P, Rosenow B, Amaral L A N and Stanley H E 1999 Universal and nonuniversal properties of cross correlations in financial time series *Phys. Rev. Lett.*

- 83** 1471
- [104] Prause K 1998 The generalized hyperbolic model *PhD Dissertation, University of Freiburg*
- [105] Pictet O V, Dacorogna M, Müller U A, Olsen R B and Ward J R 1997 Statistical study of foreign exchange rates, empirical evidence of a price change scaling law and intraday analysis *J. Banking Finance* **14** 1189–208
- [106] Rebonato R 1999 *Volatility and Correlation* (New York: Wiley)
- [107] Resnick S 1998 Why nonlinearities can ruin a heavy tail modelers day *A Practical Guide to Heavy Tails: Statistical Techniques and Applications* ed R J Adler, R E Feldman and M S Taqqu (Boston, MA: Birkhauser)
- [108] Bouye E, Durrleman V, Nikeghbali A, Riboulet G and Roncalli T 2000 Copulas for finance: a reading guide and some applications *Groupe de Recherche Operationnelle, Credit Lyonnais*
- [109] Shephard N 1996 Statistical aspects of ARCH and stochastic volatility *Time Series Models* ed D R Cox, D V Hinkley and O E Barndorff-Nielsen (London: Chapman & Hall)
- [110] Shiller R 1989 *Market Volatility* (Cambridge, MA: MIT Press)
- [111] Smith R and Weissman I 1994 Estimating the extremal index *J. R. Stat. Soc. B* **56** 515–28
- [112] Starica C 1998 Multivariate extremes for models with constant conditional correlations *J. Empirical Finance* **6** 515–53
- [113] Straetmans S 1998 Extreme financial returns and their comovements *PhD Thesis, Erasmus Universiteit, Rotterdam*
- [114] Veneziano D, Moglen G E and Bras R L 1995 Multifractal analysis: pitfalls of standard procedures and alternatives *Phys. Rev. E* **52** 1387–98