ELSEVIER

Physica A 283 (2000) 486–528

PHYSICA A

www.elsevier.com/locate/physa

Scaling in the Norwegian stock market

Johannes A. Skjeltorp[1]

Research Department, Norges Bank (Central Bank of Norway), Bankplassen 2, Oslo, Norway

Received 3 January 2000

Abstract

The main objective of this paper is to investigate the validity of the much-used assumptions that stock market returns follow a random walk and are normally distributed. For this purpose the concepts of *chaos theory* and *fractals* are applied. Two independent models are used to examine price variations in the Norwegian and US stock markets. The first model used is the *range over standard deviation* or *R/S* statistic which tests for *persistence* or *antipersistence* in the time series. Both the Norwegian and US stock markets show significant *persistence* caused by long-run "memory" components in the series. In addition, an average non-periodic cycle of four years is found for the US stock market. These results are not consistent with the random walk assumption. The second model investigates the *distributional scaling behaviour* of the high-frequency price variations in the Norwegian stock market. The results show a remarkable *constant scaling behaviour* between different time intervals. This means that there is no intrinsic time scale for the dynamics of stock price variations. The relationship can be expressed through a scaling exponent, describing the development of the distributions as the time scale changes. This description may be important when constructing or improving pricing models such that they coincide more closely with the observed market behaviour. The empirical distributions of high-frequency price variations for the Norwegian stock market is then compared to the *Lévy stable distribution* with the relevant scaling exponent found by using the *R/S*- and *distributional scaling* analysis. Good agreement is found between the Lévy profile and the empirical distribution for price variations less than ±6 standard deviations, covering almost three orders of magnitude in the data. For probabilities larger than ±6 standard deviations, there seem to be an exponential fall-off from the Lévy profile in the tails which indicates that the second-moment may be finite.

Keywords: Econophysics; Scaling; *R/S*-analysis; Lévy distributions; Financial economics

E-mail address: johannes-a.skjeltorp@norges-bank.no (J.A. Skjeltorp).

[1] This paper is based on the author's Master of Science thesis "Fractal Scaling Behaviour in the Norwegian Stock Market" at the Norwegian School of Management, N-1301 Sandvika, Norway in 1996. Views and conclusions expressed are the responsibility of the author alone and cannot be attributed to Norges Bank.

0378-4371/00/$ - see front matter
PII: S0378-4371(00)00212-0

1. Introduction and problem statement

1.1. Motivation

Economies and the capital markets are regarded as complex systems. Large efforts have been made over the years in order to understand how and how fast information flows from one investor to another and then is incorporated into prices. The models that have been developed to explain the capital markets are, of necessity, simplifications of reality. From this, an entire analytic framework has been created, but the models do not work too well. They explain some of the structures, but leave many questions unanswered. Economists often find that their forecasts have limited empirical validity.

In the traditional scientific thinking, one has tried to understand complicated conditions and processes by breaking up the systems into smaller parts. By studying these parts individually, one has attempted to gain insight into the total, accumulated behaviour by putting the pieces together and work out averaged developments. Chaos theory is an emerging science where one attempts to study complexity as an interplay and self-organization of many interacting parts. One may say that "more is different" in the sense that the co-operative behaviour of many parts may not be predicted from the behaviour of the single parts. This is in deep contrast to the *rational behaviour* or *representative agent* concepts in economics today.

A distinct feature with complexity is that in many situations one has what is called *sensitive dependence upon initial conditions*. What would, for example, have happened if Bill Gates and Paul Allen, the founders of Microsoft, had not met each other about 30 years ago? We would probably not have had Microsoft Windows or Word, but perhaps something even better? We will of course never know what would have come instead. These types of complex chain reactions limit our ability to predict the future, but it also incorporates a kind of long-term memory effect in different processes which presently is ignored in the basic theoretical framework of finance and economics.

As we shall see, the messiness of the stock market may be studied using chaos theory and fractal concepts which embody their own kind of simple laws. Although we still only get approximate results, they will prove to reproduce the data much more accurately than model results obtained using conventional statistical averages and Gaussian normal distributions. Through chaos theory and fractal geometry we have access to tools and a new way of thinking that has been widely used within the physical sciences to describe complex systems and processes. If we believe that the financial markets are complex non-linear dynamic systems, then these tools will be very useful to obtain a better description of the financial markets. This again may be used to improve those models already at hand and even make new models that conform more closely to observed market behaviour. Econophysics [1] is an interdisciplinary field of research applying methods from physics to analyse economic systems. The field has gained increased practical and theoretical interest the recent years. From a theoretical viewpoint, it offers a fresh look at existing theories in finance. Due to the increased availability of high-frequency data, the need for new tools to analyse the data has become crucial.

The main objective of this paper is to apply the concepts of chaos, fractals and "new" analytical techniques to explore the possibility that scaling phenomena occur in stock market returns. The analysis will mainly be concentrated on the Norwegian stock market. In particular, the concept of efficient markets, which is the bedrock of quantitative capital market theory will be examined, as well as the much used assumption of stock prices following a random walk (Brownian motion) or related independent processes like the martingale, sub-martingale, etc. The paper contains six main parts. In the remainder of this first part, there will be a review of the existing model framework and a discussion why there may be a need for a new viewpoint in economics. Then, in Section 2, there will be a survey of the theoretical background of chaos theory and fractals. In Section 3 an explanation and review of the methods which will be used to examine the Norwegian and US stock markets will be given. In Sections 4 and 5 the findings and their implications will be discussed followed by a more general conclusion.

1.2. A review of the traditional model framework

1.2.1. Perfect capital markets

There is one main purpose of capital markets. This is to allocate funds efficiently between lenders and borrowers, where lenders often are savers and borrowers may be viewed as the producers in the economy. The characteristic feature of an *allocationally efficient market* [2] is that the risk-adjusted marginal rates of return for all producers and savers are equated for all states of nature. Thus, in such a market all savings are optimally allocated to productive investments in a way that it benefits everyone. If we have perfect markets, then there are certain conditions that must be met. These are:

- *Frictionless markets*: no transaction costs or taxes, perfect divisibility and marketability of all assets, no regulations.
- *Perfect competition* in product and securities markets.
- *Markets are informationally efficient*: information is costless and received simultaneously by all individuals.
- *All individuals are rational*: they maximize their expected utility.

When all these conditions are met, we have an *allocationally and operationally* efficient market. An operationally efficient market means that the cost of transferring funds is assumed to be zero. The notion of perfect capital markets is a very restrictive assumption, and more restrictive than the concept of *capital market efficiency*. One may say that we have efficient markets when some of the perfect capital market assumptions are relaxed.

1.2.2. Capital market efficiency

According to Fama [3], in an efficient capital market, all the information in some information set $\boldsymbol{\Phi}$ is fully, and more or less instantaneously, reflected in the securities prices. Thus, prices are accurate signals for capital allocation. However, in contrast to perfect capital markets, we may still have market efficiency even though we have allocational inefficiencies in the product markets and/or costly information. In the 1960s, Fama operationalized the notion of market efficiency. He proposed to distinguish between three versions of the efficient market hypothesis depending on the specification of the information set $\boldsymbol{\Phi}$. These were:

- *Weak-form efficiency*: it is not possible to develop any trading rules based on historical price information to achieve excess returns in the future. $\boldsymbol{\Phi}$ includes only historical prices. In other words, it is not possible to earn excess returns through technical analysis or other active trading strategies.
- *Semistrong-form efficiency*: it is not possible for any investor to earn abnormal returns from trading rules based on any *publicly available information*. $\boldsymbol{\Phi}$ is broadened to include also all information that is publicly available.
- *Strong-form efficiency*: no investor can earn excess return by using any information available, whether it is *insider*- or *public information*. $\boldsymbol{\Phi}$ is thus broadened even further to also include all insider information as well.

This means that given an information set at time t, then today's price P_t is the best estimate of tomorrow's discounted price plus dividends

$$P_t = E\left[\left.\frac{\tilde{P}^*_{t+1}}{1+\rho}\right|\boldsymbol{\Phi}_t\right] . \tag{1}$$

Here $\tilde{P}^*_{t+1}$ is the $\tilde{P}_{t+1} + \tilde{d}_{t+1}$, $\tilde{d}_{t+1}$ the expected dividends, $\boldsymbol{\Phi}_t$ the current information set at t, ρ the discount rate.

A crucial point when discussing market efficiency is therefore how much and how fast information is captured and incorporated in prices. In this regard it is often assumed that when information is costly, the *net gain* obtained from collecting information must be zero in equilibrium. The capital market is therefore efficient relative to a given information set only after consideration of the costs of getting to a particular information structure is taken into consideration. So, what does this mean for the *behaviour* of stock market prices? This will be discussed in the following section.

1.2.3. The Gaussian hypothesis of stock price behaviour

The assumption of market efficiency implies that since all fundamental- and price history information ($\boldsymbol{\Phi}_t$) known up to and including time t, is reflected and discounted in the current price, prices only move when new information arrives in the market. Because the stock market is a large system that has a large number of degrees of freedom or investors, it is assumed that current prices must reflect the information everyone already has. In other words, today's change in price is caused only by today's unexpected new information. This means that today's price, given all relevant information available today, is the best forecast of tomorrow's discounted price plus dividends (1). Therefore, yesterday's news is no longer important, and today's returns are unrelated to yesterday's returns; i.e., the returns are *independent*. If we collect enough historical data of stock prices and returns, the distribution should, according to theory, in the

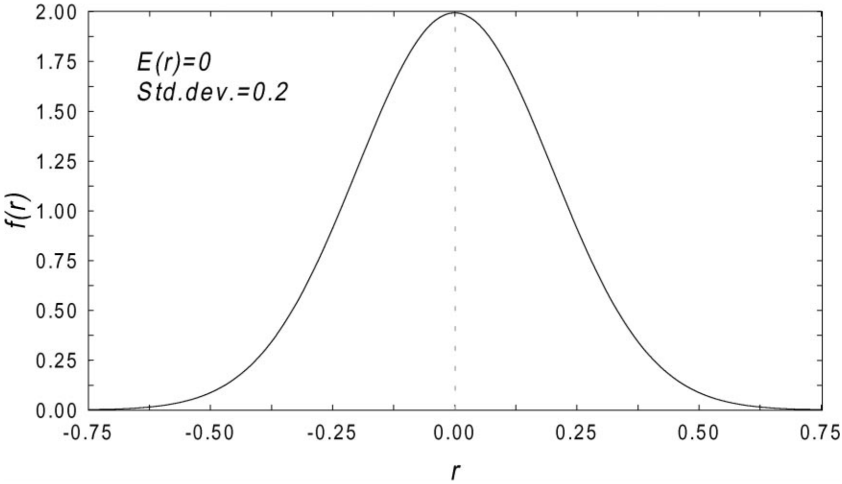

Fig. 1. The familiar normal (Gaussian) distribution for the frequency of returns when $E(r) = 0$ and the standard deviation is 0.2.

limit of large numbers of data approach the normal-distribution with a stable mean and a finite variance (Fig. 1). The equation for the frequency of returns, r, which are symmetrically and normally distributed with no skewness or kurtosis, is usually expressed as

$$f(r) = \frac{1}{\sigma_i\sqrt{2\pi t}} \exp\left[-\frac{1}{2t}\left(\frac{r - E(r)}{\sigma_i}\right)^2\right], \tag{2}$$

where $E(r)$ is the mean return, σ_i the standard deviation for security i, t the time index.

It is important to note that the efficient market hypothesis (EMH) does not necessarily imply a random walk, but conversely a random walk does imply market efficiency. However, the assumption of independence is a very crucial and deeply rooted aspect of the EMH. In the traditional literature treating the efficiency of capital markets, there are three main theories of the time series behaviour of prices: *the random-walk model* (*Brownian motion*), *the martingale- and the fair-game model*. The main differences between these models is whether they allow for dependencies in the higher-order moments of the distribution of returns.

1.3. Why there may be a need for a revisited model framework

One of the great scientists of this century, Nobel Laureate and Norwegian-born Lars Onsager once commented on doing important science which will matter in the long run: “Look at the problem and then choose your ‘weapons’ to solve it”, he said. It is “poor” science to attempt the opposite by constructing nice theories and then “look around” for the phenomena fitting the theory. As we shall see below, it may be claimed that the random-walk model applied to economics to a certain extent falls into this category. In capital market theory, the assumption of normality and finite variance, as well as models based on those assumptions, were developed even as empirical evidence continued to contradict theory [4]. In other words, it seems like the theory was constructed to justify the methods instead of the empirical facts themselves.

The current paradigm, with the EMH as the core of it, basically says that investors react to the arrival of information in a *linear* fashion. This implies that investors react to information instantaneously as the information is received, and not in a cumulative fashion to a series of events. This linear paradigm, which is built into the *rational investor* concept, is based on the assumption of past information already being reflected in security prices. Thus, the linear paradigm implies that returns should be approximately identically normally distributed and independent. However, the new paradigm that is discussed and applied in this paper, is the possibility of *non-linear* reaction to information by investors and traders. It also looks at the trading horizon as an important dimension when analysing the distribution of stock returns, and the interaction between the investors and traders with different investment horizons. Today, it is generally recognized that returns are not normally distributed [1,5–10]. The distribution of security returns are *leptokurtic*, meaning that the empirical distribution is higher around the mean and has fatter tails than predicted by the normal distribution. In Fig. 2, a plot of the empirical distribution of daily Dow–Jones Industrial Index of one day changes in the logarithm of prices for the period 1962 through 1993 is compared to the normal distribution with the same variance. The figure clearly reveals the leptokurtosis in the frequency distribution. The non-normal shape of the distribution is not peculiar to only the Dow–Jones series, but is, in fact, the case for most financial time series [5,7,11–13]. From Fig. 2, we see an indication that markets do not follow a random walk. Such fat-tailed distributions are normally attributed to long "memory" effects generated by a non-linear stochastic process which both exhibit periods of stability as well as sudden large movements on all time horizons. If this is the case, we may be over- or understating our risk and return potential on all trading horizons by using the current model framework.

Now, if we look more closely at the empirical versus the theoretical distribution, we see that theory predicts that the probability of a larger than three-sigma event is 0.3%. However, if we look at the empirical distribution based on actual data, the probability of such an event occurring (positive or negative) is about 1.5%. Thus, standard theory fails by a factor of five! This discrepancy increases even more if we look at the probability of a larger than four-sigma event occurring. Theory predicts that this will occur on average with a probability of 0.0064%, while empirically it seems to be a probability of 0.66% of such an event to occur, or more than 100 times the probability estimated theoretically. This difference in the theoretical and empirical probability of large events occurring can be seen from Fig. 3 where the tails of the normal and empirical distributions are compared.

It therefore seems that events that theoretically should not occur, actually occurs relatively quite "often". From this we can see that stock market returns are not

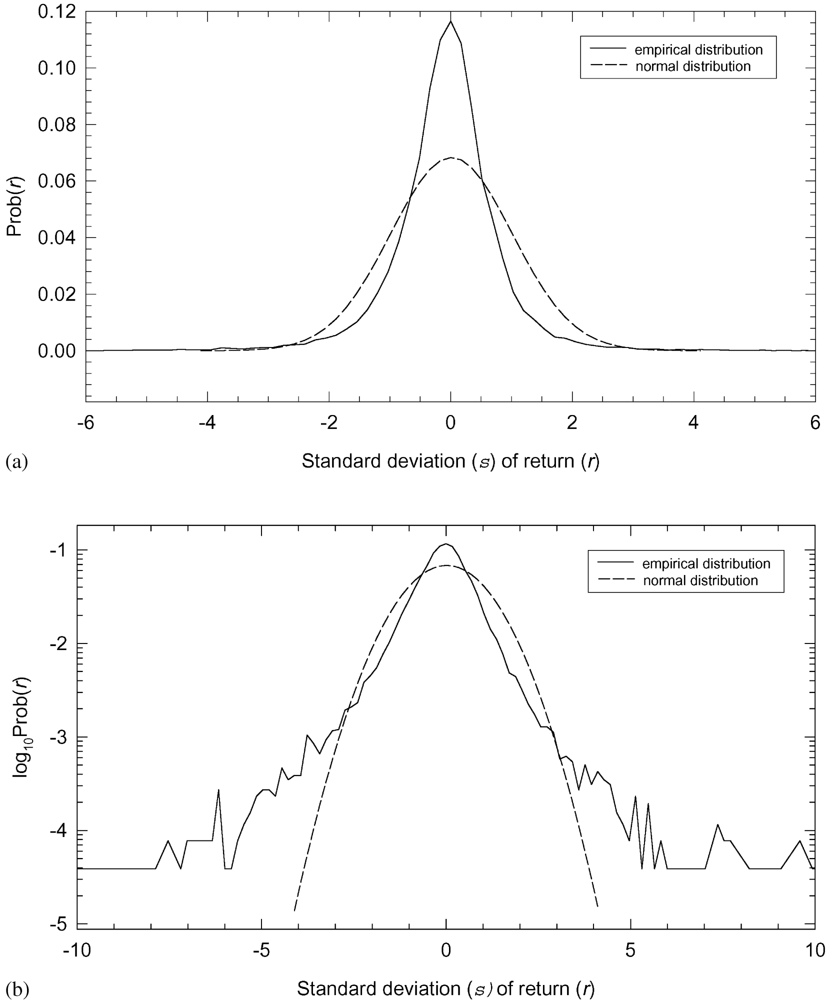

Fig. 2. (a) The empirical distribution for the daily returns for the Dow–Jones Industrial Average (DJIA) during the period 1962–1993 compared to the Gaussian normal distribution. (b) The same distribution as in (a), but in logarithmic probabilities to amplify the wings of the distributions.

normally distributed. In fact, they are not even approximately normally distributed as Osbourne assumed in his famous 1964 article [4]. The consequences of this are then that much statistical analysis, particularly diagnostics such as correlation coefficients and t-statistics, is seriously weakened and may give misleading results.

The underlying reasons for this leptokurtosis have been widely discussed as well as whether the random-walk theory is applicable or not. The most common explanation for the fat tails and high peak at the mean is that information shows up in infrequent clusters rather than in a smooth and continuous fashion. When information arrives in the market in such a fashion, there are periods of low and high volatility, which gives

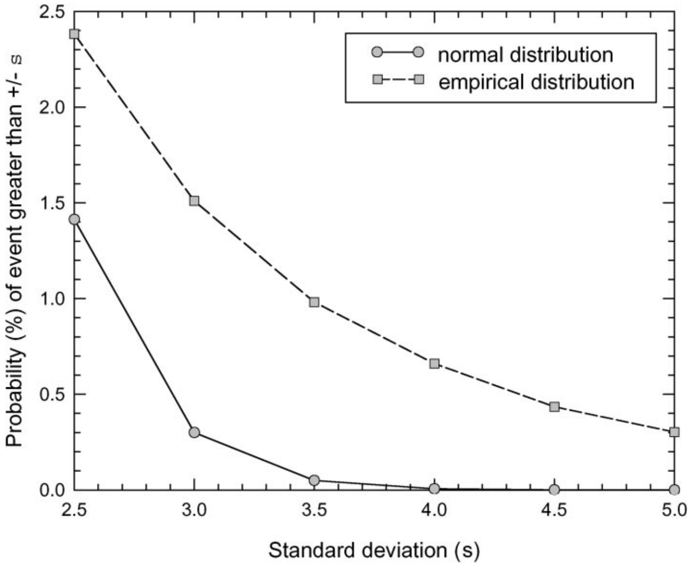

Fig. 3. The difference between the normal distribution and the empirical frequency distribution of returns for the Dow–Jones Ind. Average when we examine the accumulated probability of returns being larger than different standard deviations (both positive and negative).

rise to relatively high values of the probability densities both in the centre and the tails. Thus, because the distribution of information is leptokurtic, the distribution of price changes is leptokurtic as well. The most direct argument against the random-walk model is based on plain observations of actual data. Stock prices need not to be continuous and they are in fact conspicuously discontinuous due to periods of low fluctuations which are interrupted by periods of turbulent burst; such a behaviour is often referred to as *intermittency* [12,14]. In contrast, Brownian motion, which is the basis for the random walk and martingale models, is a continuous process. There is no reason whatsoever to expect that a natural phenomenon like Brownian motion, governed by Newtonian physics should apply to economics. The effects of the actions of economists, investors, speculators and politicians on stock prices are simply not governed by the laws of physics. People have a *free will* to influence outcomes. This is incompatible with the "passive" evolution of natural phenomena described by, e.g. the Brownian motion. The random-walk model may at best be characterized as being able to catch some features of actual economic data, but fails per definition miserably in perhaps the most important aspect in economics; for having some predictive possibility.

The fact that prices, returns, etc. are discontinuous, hardly seems to contain any predictive ability in itself. But, as we shall see, it has the effect of forcing us to look at these phenomena with completely new eyes. In so doing, the new "weapons" we have to choose, in fact reveals new features and even new predictive powers. The new models may still not be "perfect", but will be much better than the random-walk model.

Mandelbrot [6,12] suggested in 1964 that the price movements follow a family of distributions called Stable Paretian. These types of distributions have high peaks at the mean and fat tails, much like the observed behaviour of stock price movements (ref. Fig. 5). However, for these distributions, variance is infinite or undefined, something which mostly is found intuitively and practically unacceptable by economists. On the other hand, phenomena governed by non-linear behaviour were notoriously difficult to handle until the chaos theory and fractals came into the picture. This is a completely new paradigm in economics, which was started in the physical sciences several decades ago, and is now beginning to be applied to finance and economics. The theoretical background for this emerging science and applications to finance and economics will be explained in the next part.

2. Chaos theory and fractals in finance

2.1. Fractal time-series

The fractal dimension of an object, such as the Sierpinski triangle and the Cantor set [15], says something about the extent to which the object fills space. On the other hand, the fractal dimension of a time series measures how jagged the time series is and how it scales statistically in time. A characteristic feature of financial price records like that shown in Fig. 4, is that it is virtually impossible to distinguish a daily price record from a monthly price record when the axes are not labeled.

This apparent statistical self-similarity seen in Fig. 4 is qualitatively similar to that found in ordinary Brownian motion. Different time series can be classified quantitatively by looking at their fractal dimension. If a time series has a fractal dimension $D_f = 1$ it is a smooth line. A time series with $D_f = 1.5$ is a random walk since it scales with the square-root of time, see Eq. (3). There is a 50/50 chance of prices rising or falling. However, for a time series with $1 < D_f < 1.5$, the curve is smoother than

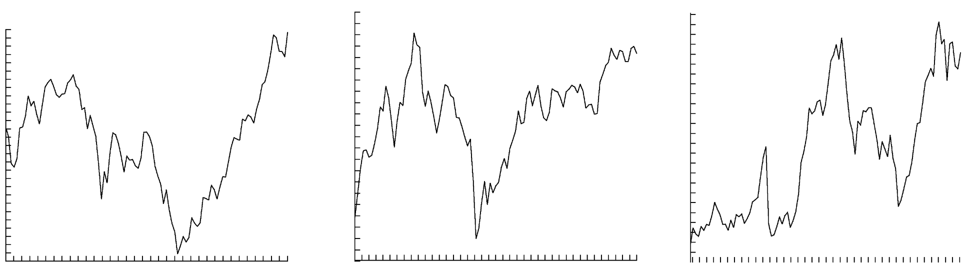

Fig. 4. A 100-day daily price record, a 100-week weekly price record and a 100-month monthly price record of the Oslo-stock exchange general index. Due to the self-similar structure in the records, it is impossible to say which is which without any labels on the axes. Test yourself before you check for the correct labeling. (a) is weekly records, (b) is daily records and (c) is monthly records.

that for a random walk, but more jagged than a straight line. This means that the underlying process is somewhere between deterministic (a smooth line with $D_f = 1$) or totally unpredictable (a random walk with $D_f = 1.5$). Furthermore, if $1.5 < D_f < 2$, the process has more reversals (is more jagged) than a random walk would imply. The important point to note is that a time series with a fractal dimension different from 1.5 would also possess non-Gaussian statistics. Thus, by using fractal geometry we have a way of determining quantitatively the extent to which a time series deviates from a pure Brownian motion. This in turn is a measure for a long "memory" effect different from memory processes that have been proposed for financial time series such as, e.g. ARCH and GARCH models [16,17], which do not take into account the scaling property of the process. This memory effect thus reveals itself with a power-law behaviour different from that of a random walk.

The most widely studied function in economics is probably the Brownian motion function discussed earlier. Let us now return to this process and take a closer look at the particle moving in the plane shown in Fig. 2. If we zoom in on the trajectory in this figure and increase the time resolution by reducing the time steps Δt, we will see a statistically, self-similar random walk. In Fig. 4a, we have plotted the particle trajectory in Fig. 2 in the x-direction versus time. The time axis is therefore included as an extra dimension. This time record here is said to be *self-affine* rather than *self-similar*. The distinction between these two notations will be discussed in the following. The displacement r_t in Fig. 4b is defined as a stochastic variable with zero mean, $\langle r_t \rangle = 0$, where the angle brackets $\langle\ \rangle$ denote the mean value of the enclosed quantity, and a variance of,

$$\langle r_t^2 \rangle \propto t \,. \tag{3}$$

In order to have the time record in Fig. 4b look "the same" under a change of resolution we must have what is called *scale invariance*. It may be shown [15] that the Brownian random process is invariant in distribution under a transformation that changes the time scale by a factor b and the length scale by a factor $b^{0.5}$. This means that we observe the particle position only at intervals $b\tau$ where τ is the time step and b is some arbitrary number. To get a time record that "looks the same" as that of Fig. 4b, the length scale has to be multiplied by $b^{0.5}$ i.e., $r_t \to r_t \cdot b^{0.5}$. A "side-by-side comparison" of two self-affine curves with $b = 4$ is shown in Fig. 5.

The *fractional Brownian motion* (FBM) [5,7,12,15] is an extension of the concept of the usual Brownian motion discussed so far. The return, r_t, for such a non-normal process in one dimension is, by definition, *a stochastic variable with zero average and with variance*,

$$\langle r_t^2 \rangle \propto t^{2D_f} \,. \tag{4}$$

Here, D_f is the fractal dimension of the time trace. We see that a special case of the *fractional Brownian motion* is the normal Brownian motion when $D_f = 1.5$, because

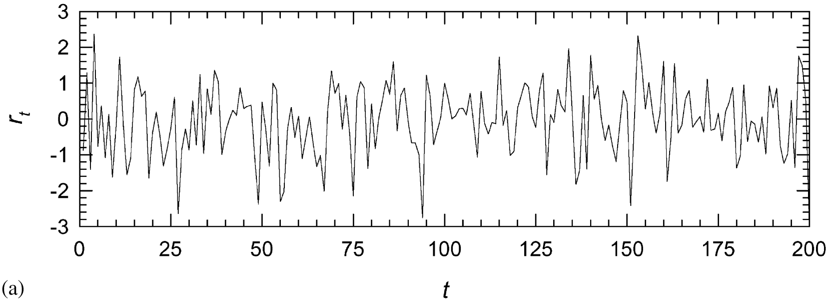

(a)

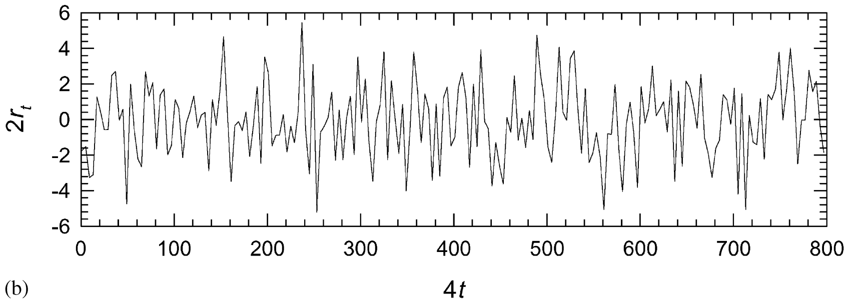

(b)

Fig. 5. Two self-affine curves of a Brownian motion process. We cannot expect r_t and t to scale by the same ratio. The concept of self-affinity is therefore important when analysing time series. When scaling up r_t by a factor of two we need to scale the time index, t, by a factor of 4 to preserve the relationship of r_t scaling by $t^{0.5}$ for a Brownian motion.

then Eq. (4) is equal to Eq. (3) and we have normal scaling behaviour. When we have a fractal dimension between 1.5 and 2 the time series is said to be *antipersistent*. If the fractal dimension is between 1 and 1.5 the time series is said to be *persistent*. Examples of simulations of fractional Brownian motion series are shown in Fig. 6. We see that for the persistent series in Fig. 6 (1a), positive changes are more likely to be followed by positive changes and vice versa. The time series is smoother and less jagged than in 1b and 1c which has a higher fractal dimension. There is less noise in the system and the “trends” are more pronounced as D_f gets smaller. By increasing the fractal dimension in the simulation, we see that the time series becomes more jagged. Thus, the fractal dimension of a time series measures how smooth or deterministic the time series is. A perfectly deterministic system would produce a smooth curve with $D_f = 1$, and a purely random system has a $D_f = 1.5$. Thus, if we have a persistent time series, the system is somewhere between a purely deterministic and completely random system.

One method used to estimate the fractal dimension of a time series is called the *Rescaled Range Analysis or R/S analysis* [11]. This is *one* of several methods that will be used in this paper to examine the fractal structure of the Norwegian stock market and the US market. The model framework and a detailed explanation and discussion of the R/S analysis will be presented in Section 3.

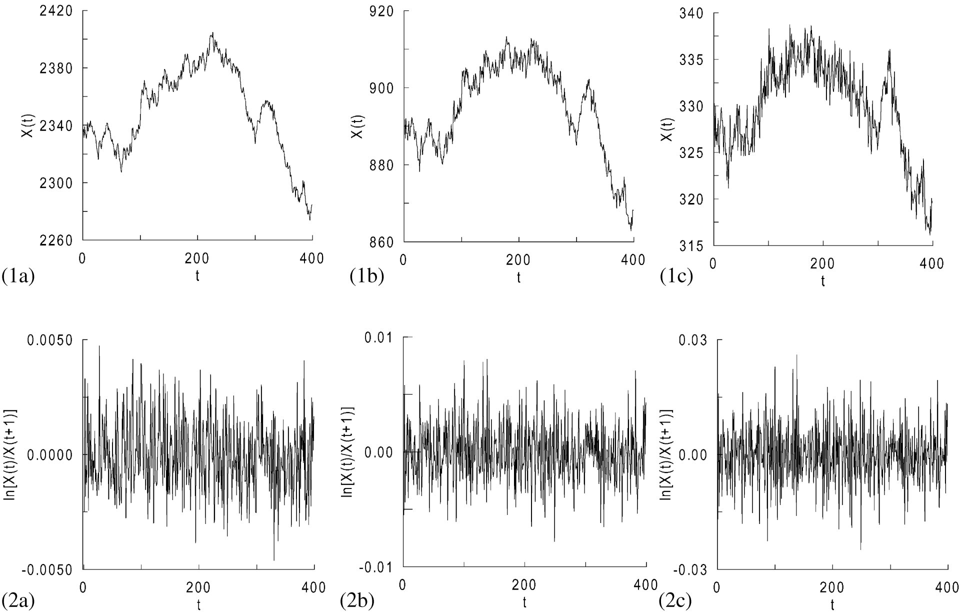

Fig. 6. Examples of simulated fractional Brownian motion time series for different fractal dimensions D_f as discussed in the text. In 1a we have a persistent time series with a fractal dimension of $D_f = 1.3$. In 1b we have a random walk with $D_f = 1.5$, and in 1c we have an antipersistent time series with $D_f = 1.7$. The corresponding returns are shown in graphs 2a, 2b and 2c.

2.2. Chaos and financial economics

In economics and capital market theory, researchers have long used the "Newtonian" assumption that a system when left alone tends to equilibrium. Thus, when modelling the capital market and currency market it has been modelled as being naturally at equilibrium unless perturbed by an exogenous shock. In other words, there is a natural balance between supply and demand of stocks or currencies, unless an exogenous shock changes the supply and/or demand, which will cause the system to seek a new equilibrium and stabilize itself if no further information is added to the system. The emerging, and new capital market paradigm treats the markets as complex, interdependent systems where the state of the system is continually fluctuating, with no natural equilibrium state. The market has long-term correlations and trends due to feedback effects, and can be more or less erratic under certain conditions, conditioned upon earlier events. The efficient market hypothesis [3,18] has one primary goal; to justify the use of probability calculus in analyzing capital markets. If markets are non-linear dynamic systems, then the use of standard statistical analysis can give misleading results. This is particularly true if a random-walk or martingale model is used, which is often the case. Some characteristic features one would expect for non-random, non-linear behaviour are: long-term correlations and trends (feedback effects/memory effects), erratic markets under certain conditions and certain times (volatility clustering), the series may have self-similar and self-affine characteristics, a fractal structure and less reliable forecasts due to the sensitivity of such a system on initial conditions.

3. Model framework

In this part we will review the *two* independent models used in this paper to investigate the possible fractal structure of the Norwegian and American stock markets. The first model uses the rescaled range analysis to estimate the fractal dimension and the degree of persistence or anti-persistence of a time series. This can be used to see how the time-series under study differs from the normal-Gaussian alternative. The second model looks more at the scaling behaviour of the distribution of returns for different time horizons and is here called the distributional scaling approach, which was first proposed by Mandelbrot [6,12].

3.1. Rescaled Range Analysis (R/S analysis)

3.1.1. Background

The rescaled range analysis or R/S analysis, was developed by Hurst [11,15], a hydrologist who worked on the problem of reservoir control on the Nile River dam project around 1907. His problem was to determine the ideal design of a reservoir based upon the given record of observed river discharges. An ideal reservoir never empties or overflows. In constructing the model, it was common to assume that the uncontrollable part of the system, which in this case was the influx due to rainfall, followed a random walk due to the many degrees of freedom in the weather. When Hurst examined this assumption he gave us a new statistical measure, *the Hurst exponent* (H). As we will see, it is also closely connected to the fractal dimension discussed above. His statistical method is very robust and has few underlying assumptions. The Hurst statistics can be used to classify time series into random or non-random series. The analysis is very robust with respect to the underlying distribution of the process. As noted by Mandelbrot and Wallis [19], even extremely non-Gaussian *independent* processes which have a log-normal, hyperbolic or gamma distribution will give an $H = 0.5$. Using R/S analysis one also finds the average non-periodic cycle, if there is any, and the degree of persistence in trends due to long "memory" effects.

Standard Gaussian statistics works best under very restrictive assumptions. The events measured must be *independent and identically distributed* (IID). Thus, events must not influence one another, and they must all be equally likely to occur. The normality or near-normality assumption is often made when analysing and describing large complex systems with very many degrees of freedom so that standard statistical analysis can be applied. However, if the system is not IID, we need a non-parametric method. Such a method is the R/S analysis.

3.1.2. The R/S methodology

The main idea behind using the R/S analysis for our purpose is that one looks at the scaling behaviour of the rescaled cumulative deviations from the mean, or the distance the system travels as a function of time. This is compared to the null-hypothesis of a random walk. As mentioned before, for an independent system, the distance covered increases, on average, by the square-root of time. If the system covers a larger distance than this, it cannot be independent by definition, and the changes must be influencing each other; they have to be correlated. Although there may be autoregressive (AR) processes present that can cause short-term correlations, we will see that when adjusting for such short-term correlations (serial correlations), there may be other forms of memory effects present which need to be examined. Next, the steps needed to do the analysis is reviewed.

We first start with a time series in prices of length M. This time series is then converted into a time series of logarithmic ratios or returns of length $N = M - 1$ such that

$$N_i = \log\left(\frac{M_{i+1}}{M_i}\right), \quad i = 1, 2, \ldots, (M-1)\,. \tag{5}$$

Divide this time period into A contiguous sub-periods of length n, such that $A \cdot n = N$. Each sub-period is labelled I_a, with $a = 1, 2, \ldots, A$. Then, each element in I_a is labelled $N_{k,a}$ such that $k = 1, 2, \ldots, n$. For each sub-period I_a of length n the average is calculated as

$$e_a = \frac{1}{n}\sum_{k=1}^{n} N_{k,a}\,. \tag{6}$$

Thus, e_a is the average value of the N_i contained in sub-period I_a of length n. Then, we calculate the time series of *accumulated departures from the mean* ($X_{k,a}$) for each sub-period I_a, defined as

$$X_{k,a} = \sum_{i=1}^{k} (N_{i,a} - e_a), \quad k = 1, 2, \ldots, n\,. \tag{7}$$

As can be seen from Eq. (7), the series of accumulated departures from the mean always will end up with zero. Now, the range that the time series covers relative to the mean within each sub-period is defined as

$$R_{I_a} = \max(X_{k,a}) - \min(X_{k,a}), \quad 1 < k < n\,. \tag{8}$$

The next step is to calculate the standard deviation for each sub-period I_a,

$$S_{I_a} = \sqrt{\frac{1}{n}\sum_{k=1}^{n} (N_{k,a} - e_a^2)}\,. \tag{9}$$

Then, the range for each sub-period (R_{I_a}) is rescaled/normalized by the corresponding standard deviation (S_{I_a}). Recall that we had A contiguous sub-periods of length n. Thus, the average R/S value for length or "box" n is

$$(R/S)_n = \frac{1}{A}\sum_{a=1}^{A}\left(\frac{R_{I_a}}{S_{I_a}}\right) . \quad (10)$$

Now, the calculations from Eq. (5) to Eq. (10) must be repeated for different time horizons. This is achieved by successively increasing n and repeating the calculations until we have covered all integer n's. One can say that R/S analysis is a special form of "box-counting" for time series. However, the method was developed long before the concepts of fractals.

After having calculated R/S values for a large range of different time-horizons n, we plot $\log(R/S)_n$ against $\log(n)$. By performing a least-squares regression with $\log(R/S)_n$ as the dependent variable and $\log(n)$ as the independent one, we find the slope of the regression which is the estimate of the *Hurst exponent* (H). The relationship between the fractal dimension D_f discussed earlier and the Hurst exponent (H) can be expressed as [7]

$$D_f = 2 - H . \quad (11)$$

3.1.3. Interpreting the Hurst exponent

According to theory, $H = 0.5$ means that the time series is independent, but as mentioned above the process need not be Gaussian. If $H = 0.5$, the process may in fact be a non-Gaussian process as e.g. the Student-t or gamma. If $H \in \langle 0.5, 1.0]$ it implies that the time series is persistent which is characterized by long "memory" effects on all time scales. I.e., all daily price changes are correlated with all future daily price changes; all weekly price changes are correlated with all future weekly price changes and so on. This is one of the key characteristics of fractal time series as discussed earlier. It is also a main characteristic of non-linear dynamic systems that there is a *sensitivity to initial conditions* which implies that such a system in theory would have an infinite memory. The persistence implies that if the series has been up or down in the last period then the chances are that it will continue to be up or down, respectively, in the next period. This behaviour is also independent of the time scale we are looking at. The strength of the trend-reinforcing behaviour, or persistence, increases as H approaches 1.0. This impact of the present on the future can be expressed as a correlation function (C),

$$C = 2^{(2H-1)} - 1 . \quad (12)$$

In the case of $H = 0.5$ the correlation C equals zero, and the time series is uncorrelated. However, if $H = 1.0$ we see that $C = 1$, indicating perfect positive correlation. On the other hand, when $H \in [0, 0.5\rangle$ we have *antipersistence*. This means that whenever the time series have been up in the last period, it is more likely that it will be down in the next period. Thus, an antipersistent time series will be more choppier than a pure random walk with $H = 0.5$.

As mentioned above, the R/S analysis can also uncover average non-periodic cycles in the system under study. If there is a long "memory" process at work, for a natural system this memory is often finite, even though long "memory" processes theoretically are supposed to last forever, as was the case for mathematical fractals and the logistic map. When the long-term memory is lost, or the memory of the *initial conditions* has vanished, the system begins to follow a random walk; this is also called the *crossover point*. Thus, a crucial point in the estimation of the Hurst exponent is to use the proper range for which there is non-normal scaling behaviour. This is the range for which the scaling behaviour is "linear" in the $\log(R/S)_n$ versus $\log(n)$ plot. If there is a crossover-point, this can be seen as a "break" in the plot where the slope changes for a certain value, $\log(n_{\max})$. If this is the case, it is an indication of a non-periodic cycle with average cycle length equal to $n_{\max}$.

3.1.4. Weaknesses with the method

The Scaling Range and small sample sizes. Persistent time series, with $0.5 < H \leqslant 1.0$, are fractal in the sense that they can be represented as *fractional Brownian motions* (FBM). In a fractional Brownian motion there is a correlation between events across time scales, as described by Eq. (12). However, as n gets very large it is expected that the series will converge to the value $H = 0.5$, because the memory effect will diminish to a point where it becomes unmeasurable. The regression to estimate the Hurst exponent should therefore be performed on the data prior to the convergence of H to 0.5, which we called the crossover point, in the discussion above. There is also a lower cutoff limit $n_{\min}$ for n below which the data are not useful in the regression analysis. As a rule of thumb this is the case for $n_{\min} = 10$. The reason for this is that small values of n produce unstable estimates when sample sizes are small. This is often the case for financial data sets. In order to get a well-defined Hurst exponent, it is therefore important to establish the proper scaling range $n_{\min} < n < n_{\max}$ before running the regression.

Short-term dependencies in the data. Another weakness with the R/S-analysis method is that it is sensitive to short-term dependencies, which can bias our estimate of H. Financial time series of high frequency (daily or more frequent observations) generally exhibit significant autoregressive tendencies. This is due to the fact that the high-frequency data are primarily trading data, and traders do influence each other. Thus, an empirical investigation of long-term "memory" effects in stock returns must first take into account the presence of high-frequency autocorrelations [7,16,20]. When doing the R/S analysis it is therefore important to try to eliminate or, at least minimize, such linear dependencies, since it can bias the Hurst exponent and classify a process as having a long-term "memory" when it is, in reality, a short-term "memory" effect. By taking the AR(1) residuals we minimize the bias, a method which is called *detrending* or *prewhithening*, used by Peters [7] and Lo [20]. This will eliminate, or at least reduce, serial correlation as well as inflationary growth. The former is a problem with very high-frequency data while the latter is a problem when we are dealing with low-frequency data. However, for R/S analysis the short-memory process is more of a problem than the inflationary growth problem [7]. To do a detrending we begin with a series of logarithmic returns,

$$S_t = \log\left(\frac{P_t}{P_{t-1}}\right) , \tag{13}$$

where S_t is the logarithmic return at time t, P_t the price at time t.

Now, S_t is regressed as the dependent variable against S_{t-1} as the independent variable. We then obtain the intercept, a, and slope, b. Then the AR(1) residual of S_t subtracts out the dependence of S_t on S_{t-1},

$$X_t = S_t - (a + b \cdot S_{t-1}) . \tag{14}$$

Here; $X_t =$ the AR(1) residual of S at time t.

Now, the R/S analysis is done according to the procedure outlined above in Eqs. (5)–(10), except that we use X_t which has been adjusted for AR(1) residuals as returns and start with Eq. (6) instead of Eq. (5). This procedure is taken from Peters [7].

3.1.5. Examples of the method

Below we shall show some examples of how to estimate the Hurst exponent on a few simulated fractional Brownian motion processes as well as a pure random walk. By constructing a series of Gaussian random numbers it is possible to see if the rescaled range analysis gives a satisfactory result with $H = 0.5$. In the example below, a series of 8000 pseudorandom numbers was constructed in Mathcad, and then an R/S analysis was performed on the series. In Fig. 7 a log/log plot of the results is shown. As may be seen, the simulated results are well described by the linear regression form $n > 20$ with a Hurst exponent $H = 0.515$ which is slightly higher than the theoretical value

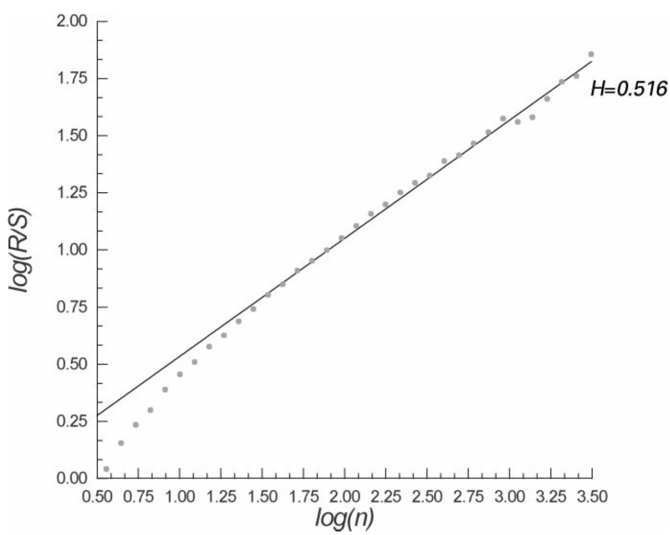

Fig. 7. R/S as a function of time lag n for a series consisting of 8000 random Gaussian numbers. The solid line is the fitted curve $R/S_n = an^H$ for $n > 20$ ($\log(n) > 1.3$) with $H = 0.516$ and intercept $a = 0.0178$.

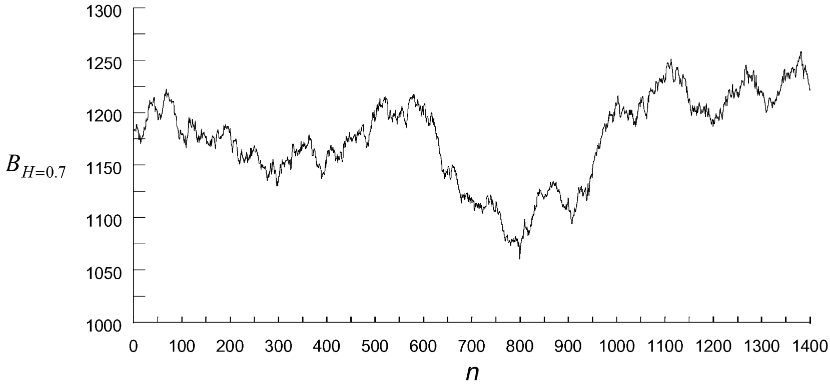

Fig. 8. A simulated fractal Brownian function B_H with $H = 0.7$ and a long memory of $M = 200$ observations and $n = 5$.

$H = 0.5$. For $n < 20$, the simulated results fall significantly below the linear fit. That such deviations could be expected was already pointed out by Mandelbrot and Wallis [19]. We shall return to this point later in Section 3.1.6 where the "power" of the R/S analysis is reviewed.

We shall now turn to the Mandelbrot and Van Ness [21] method, it is possible to simulate a fractional Brownian motion and thus approximate a long memory process. In this example we take 8000 Gaussian random numbers and convert them into 1400 fractional Brownian numbers shown in Fig. 8. Each biased increment is made up of $n = 5$ random numbers and has a memory of the last 200 biased random numbers. A relatively short memory of $M = 200$ is selected because of the huge increase in computation time by choosing a longer memory.

In Fig. 9 we have performed an R/S analysis of the data in Fig. 8 and we obtain $H = 0.65$ for $4 < n < 200$. This is somewhat lower than the theoretical $H = 0.7$, and may be due to the limited set of only 1400 observations. The estimate would probably improve if a larger amount of data had been used, but due to the huge increase in computation time required for simulating a larger data set, this has not been attempted. The present example serves its purpose as an illustration. From Fig. 9 we see that there is a "break" in the plot at about $\log(n) = 2.3$, or 200. This is exactly the long memory effect of $M = 200$ which was used in the simulation. When $\log(n) > 2.3$, the slope of the log/log plot drops to $H = 0.51$, which means that there is no measurable memory effect left and the series starts to follow a random walk.

3.1.6. *The "power" of the R/S analysis*

To evaluate the significance of the R/S analysis results, a confidence test is needed. Such a confidence test has been developed by Peters [7] on the basis of Monte Carlo simulations and previously developed asymptotic theory. For small values of n the R/S will scale at a faster rate than $H = 0.5$ which is the expected scaling rate for an independent series, see Fig. 7. This is due to the fact that for small n the standard

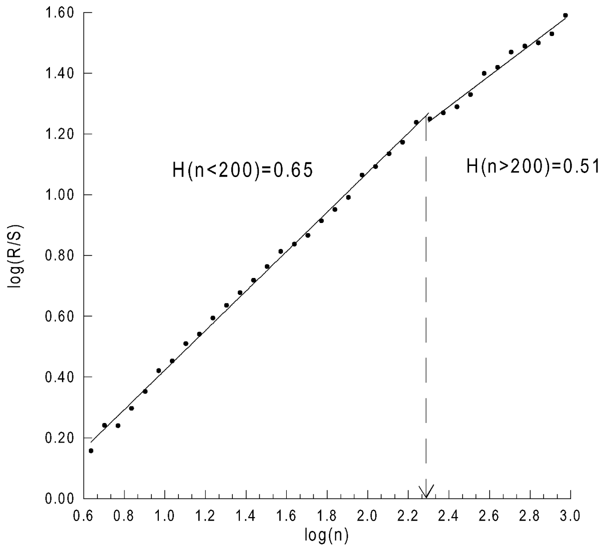

Fig. 9. $\log(R/S)$ as a function of $\log(n)$ for the fractal Brownian function $B_H(n)$ with $H=0.7$. The solid line for $n<200$ is a linear fit to $R/S=an^H$ with $H=0.65$ and intercept $a=-0.228$. The asymptotic fit for $n>200$ produces $H=0.51$ and $a=0.061$. The break in the plot at $n=200$ is clearly seen. Thus, the long time memory effect of $M=200$ has been captured by the R/S-analysis.

deviation will scale at a relatively slower rate than the range [19] as defined in Eqs. (8) and (10). Mandelbrot and Wallis referred to the range of small n as "transient" because n was not large enough for the proper behaviour to be seen. They pointed out that the R/S analysis tends to *overestimate* H for $H<0.72$ and *underestimate* H for $H>0.72$.

Thus, when assessing the significance of our findings it is important to have a correct measure of the random null hypothesis for all n which takes into account this bias. In 1976 Anis and Lloyd [22] developed an equation for calculating the expected R/S values which took this bias into account. By simulating a huge amount of scrambled pseudo random numbers 300 times and calculating the average R/S values for all n, it is possible to approximate the true behaviour of the R/S-analysis for a Gaussian random walk. For small n, the expected R/S values will scale at a faster rate than $H=0.5$ as was seen in Fig. 7, but converge asymptotically to $H=0.5$ as n goes to infinity. The expected R/S values can be expressed as [7]

$$E(R/S)_n=\left[\left(\frac{n-0.5}{n}\right)\left(n\frac{\pi}{2}\right)\right]^{-0.5}\sum_{r=1}^{n-1}\sqrt{\frac{n-r}{r}}\,. \tag{15}$$

By using Eq. (15) we can now generate expected values of the Hurst exponent by taking a regression of $\mathrm{E}(R/S)$ for the same range of n as for our empirical R/S values. Because the R/S values are normally distributed random variables we would also expect that the values of H are normally distributed with expected variance $1/T$, where T is the total number of observations in the sample. This has been thoroughly tested by

Peters [7] and has been shown to be a good approximation. Thus, by using Eq. (15) as the null hypothesis of a random walk, we can determine the significance of the results we get. This is done by calculating how many standard deviations away from the null hypothesis $E(H)$ the empirical Hurst exponent H is through the expression

$$\sigma_H = \frac{H - E(H)}{\sqrt{1/T}} \,. \tag{16}$$

Next, there will be a review of the second model which extends the scaling to the entire distribution of returns.

3.2. Lévy stable distributions (Fractal distributions)

The R/S analysis discussed in the previous section is the simplest quantitative measure of the scaling behaviour and self-similarity of a time series. A stronger form of self-similarity is obtained when we look at not only the scaling of the *rescaled range* (R/S), but also the scaling property of the full distributions of returns for different time intervals or lags Δt. And then probe whether these distributions scale at a rate similar to those found by using the R/S-method. This additional analysis will be important to complement the R/S-method. By looking at fractal distributions and fractal statistics we are not only able to examine the characteristics of the process, but also the statistics which is a very important aspect in financial economics. If we find a Hurst exponent significantly larger than 0.5 this can be due to either long-tailed distributions of returns, or due to long-range positive correlations, or a combination of both. In other words, since $H > 0.5$ indicates that the mean absolute deviations increases at a faster rate than what is expected for the Gaussian case, we must have that there are both large jumps in the series and persistence such that the series is "forced" to travel a larger distance than a random walk would imply.

As was seen in Figs. 5 and 6 there was a clear leptokurtosis in the daily returns of the Dow–Jones industrial average index (DJIA), with longer tails and higher mean than the normal distribution. Standard financial models assume that investors are risk averse, and the amount of risk that the investors take on is reflected by the distribution of returns. Thus, if an investor assumes that the returns are drawn from a normal distribution, the risk taken on by holding a portfolio for which the returns are drawn from a leptokurtic distribution is likely to be much higher than expected. Option pricing models, like the Black–Scholes model [23,24], are heavily based on the Gaussian assumption. In particular, the shape of the Gaussian normal distribution and the scaling exponent $H = 0.5$ are used. The distributional self-similarity approach that will be discussed here could be used for e.g. in more realistic option pricing models.

3.2.1. Stable distributions/ fractal distributions

The basic question of fractals is: *when does the whole look like its parts*? The French mathematician Paul Lévy asked himself: *when does the probability* $PN(X)$ *for the sum of* N *steps* $X = X_1 + X_2 + \cdots + X_N$ *have the same distribution* $p(x)$ *up to a* *scale factor as the individual steps*, or when do we have distributional self-similarity. Usually, the standard answer is that the $p(x)$ should be a Gaussian since the sum of N Gaussian numbers is again Gaussian, but with N times the variance of the original. However, Lévy proposed a more general approach with the Gaussian as only a special case. The representation of the Lévy stable distribution which will be used here is [10]

$$L_\alpha(Z,\Delta t) \equiv \frac{1}{\pi}\int_0^\infty \exp(-\gamma\Delta t q^\alpha)\cos(qZ)\,\mathrm{d}q\,, \tag{17}$$

where α is the characteristic exponent (fractal dimension of the probability space) $0<\alpha\leqslant 2$, Z the return, γ the scale factor, and Δt the time interval.

The Lévy stable distribution rescales under the following transformations:

$$Z_S \equiv \frac{Z_{\Delta t}}{(\Delta t)^{1/\alpha}} \tag{18}$$

and

$$L_\alpha(Z_S,1) \equiv \frac{L_\alpha(Z,\Delta t)}{(\Delta t)^{-1/\alpha}}\,. \tag{19}$$

Thus, it is expected that the rescaled empirical distributions for different Δt will collapse on the $\Delta t = 1$ distribution due to the normalization when we use Eqs. (18) and (19) and α, which may be estimated indirectly through R/S analysis. Our objective will be to compare the empirical distributions with the theoretical Lévy distribution in Eq. (17), and see if the Lévy distribution is a good description of the data.

The exponent α characterizes how "peaked" and "fat-tailed" the distribution is. If $\alpha = 2.0$ the distribution is Gaussian, and we have a finite second moment. On the other extreme with $\alpha = 1$ we have the Cauchy distribution with both undefined first and second moments. For both these cases the parameter values can be defined in terms of closed-form mathematical expressions and they are special cases of the stable Lévy distribution. In the region for which $1.0 < \alpha < 2.0$, the second moment becomes infinite or undefined, but with a stable mean. There are no closed-form mathematical solutions for this case, only numerical ones found using computers [6,7]. We expect to find an α within this region for the Norwegian stock market, but with a new technique described below it is possible to determine the parameters required in a rather simple and straightforward fashion.

The concept of infinite or undefined variance is something which normally is an unappealing property for financial economics since this creates problems when one attempts to determine risk. The concept of infinite variance has thus been the main critique for applying stable distributions to financial models. Despite this, the stable Lévy distributions have a number of nice features that reflect important aspects of the observed market behaviour, as we will see later in this paper.

The fractal dimension of the probability space, α, used in Eqs. (17)–(19) is related to the Hurst exponent of the time series in the following way:

$$\alpha = \frac{1}{H}\,. \tag{20}$$

The Hurst exponent H is thus a measure of the fractal dimension of the time trace through Eq. (11), while α characterizes the statistical self-similarity of the probability space.

3.2.2. Determination of α

Fama [25] and Mandelbrot [6,12] describe a number of different ways to measure α. These methods mainly investigate the tails of the distributions, which is difficult especially because larger values of Δt imply a reduced number of data and thus makes the estimates unreliable. There are also some more robust ways of estimating α, where the most obvious one is by using R/S analysis due to the relationship between α and the Hurst exponent shown in Eq. (20). It now appears that R/S analysis offers the most reliable method for calculating α. Another method recently proposed by Mantegna and Stanley [10] proves also to be a reliable method. In this paper both of these methods will be applied to determine whether there are scaling relationships that can be established in the stock market, and if there is consistency between the fractal dimension for the time trace and the distributions of returns as expressed through Eq. (20). Since the R/S method has already been described in Section 3.1, we will only discuss the Mantegna/Stanley [10] method here. The method is straightforward, but is very data intensive. First, one selects from a data set, the complete set of non-overlapping records separated by a time interval Δt. The values of the prices are denoted as $P(t)$ while the successive variations $Z(\Delta t)$ are denoted as

$$Z(\Delta t) \equiv P(t) - P(t - \Delta t) . \tag{21}$$

We can now determine the probability distributions $\Pr(Z)$ of price variations for different values of Δt. The empirical probability distributions are calculated by using relative frequencies. This is done by counting the number of returns falling within different ranges of $Z(\Delta t)$ and dividing by the total number of data in the sample. This ensures that both axioms of probability are satisfied. Since Eq. (21) implies that large Δt imply a reduced number of data, especially when investigating the wings of the distributions which is the traditional method, Mantegna and Stanley use another approach. They study the scaling behaviour of the "probability of zero return", $\Pr[Z(\Delta t)=0]$, as a function of Δt. This is chosen because we then estimate the point of each probability distribution that is least affected by the finiteness of the data set. If $\Pr[Z(\Delta t) = 0]$ is plotted against Δt in a log/log plot, the slope will show if we have a normal or a non-normal scaling behaviour depending on whether the slope (λ) is equal to $\lambda = -0.5$ or not. This slope is also related statistically to the Hurst exponent through

$$H = -\lambda . \tag{22}$$

Thus, by combining Eqs. (20) and (22), α can be determined as follows:

$$\alpha = \frac{1}{-\lambda} . \tag{23}$$

To summarize, if we have a normal scaling behaviour with $H = 0.5$ or $\lambda = -0.5$, the fractal dimension of the probability space is $\alpha = 2$, and we have the special case of

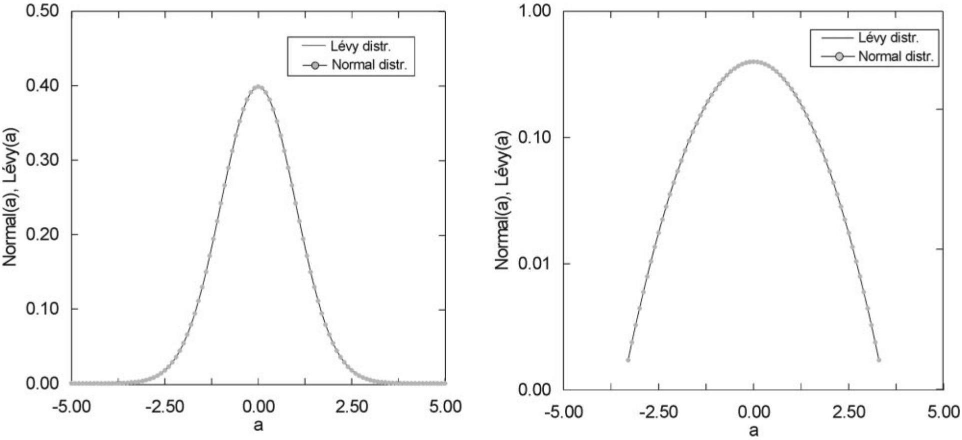

Fig. 10. Left: The normal distribution with mean 0 and standard deviation equal to 1, compared to the Lévy stable distribution with $\alpha = 2.0$ and standard deviation of 1. Right: The same distributions in a logarithmic plot.

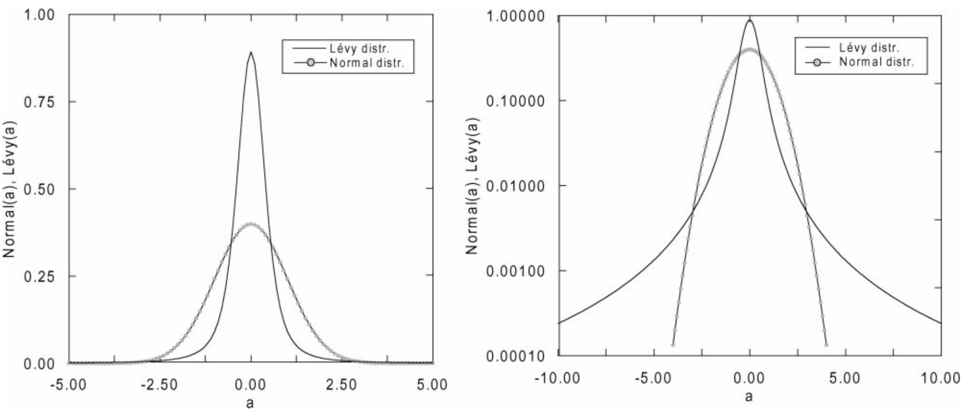

Fig. 11. Left: The normal distribution with zero mean and standard deviation equal to 1, compared to the Lévy distribution with $\alpha = 1.4286$ (corresponding to $H = 0.7$). Right: The same comparison, but in a logarithmic plot to enhance the tails.

the Gaussian normal distribution with finite variance. Fig. 10 shows that the normal distribution and the Lévy distribution indeed are the same for this case. On the other hand, if we have non-normal scaling behaviour with $H \neq 0.5$ and thus $\lambda \neq -0.5$, the second moment of the distribution becomes infinite and the leptokurtosis in the distribution increases. Fig. 11 shows a Lévy distribution with $\alpha = 1.4286$ ($H = 0.7$) compared to the normal distribution. The fatter tails and higher probability around the mean (leptokurtosis) are clearly seen.

3.2.3. Lévy distributions and observed market behaviour

Self-Similarity. One important property of stable Lévy distributions is that they are self-similar. This means that the probabilities of return are the same for all time frames once we adjust for the time scale. This self-similar property is why stable Lévy distributions often are referred to as fractal distributions. This means that a trader with a 1 min time frame (noise trader) faces the same risk as a 100 min trader or a 1000 min trader in his time frame when adjusted for scale. For example, if a 100 min trader faces a 4 standard deviation event in his time frame it may be a disaster for him. However, this is of far less importance in absolute returns for a 1000 min trader if this had happened to him in his time frame. The characteristic exponent, α, takes this scaling relationship into account.

Discontinuities. As mentioned before, in a "Gaussian market" there cannot be any large price jumps or discontinuities since a Gaussian process is everywhere continuous [2] and differentiable. This is not the case for a fractal process like the fractional Brownian motion or Lévy processes, for which such discontinuities are allowed features. When looking at how the market behaves on all time scales, it is not uncommon to have large price jumps which might be due to fears of capital or opportunity loss. In fact, such jumps seem to appear relatively often; and surely more often than "never". The market prices, are, in general, characterized by calm periods followed by sharp breaks and discontinuities. It is therefore not satisfactory to look at them as anomalies which should never occur and keep them out of the model. Instead, one should make them a part of the model framework for better or for worse.

The Truncated Lévy distribution. A very promising approach with respect to its application to mathematical finance [24], is the truncated Lévy distribution [1,26,27]. The truncated Lévy distribution (TLD) takes into account the slow convergence towards a Gaussian process as the time horizon (Δt) increases. Compared to the Lévy distribution which has an infinite second moment, the TLD has finite variance and it scales in a finite interval. For financial time series, this seems to be a good approximation, since there seems to be a convergence to a Gaussian regime for long horizons [28]. Due to its finite variance property, as well as scaling characteristics, the TLD has been an attractive candidate for option pricing models [29,30]. However, one limitation of the TLD is that it does not describe the time-dependent volatility observed in financial market data.

4. Results from the R/S analysis

4.1. The Data

4.1.1. The Oslo stock exchange general index (TOTX)

The first data set [3] that will be analysed here is the daily closing prices of the total-index at Oslo stock exchange (OSE) for 13 years of data ranging from 1983 to 1995, or a total of 3190 observations. This index is a weighted average of the stock

[2] The Gaussian assumption is used under the general heading of *continuous-time finance* where option pricing is one example.

[3] The data was obtained from associate professor Bernt Ødegaard at Norwegian School of Management (BI).

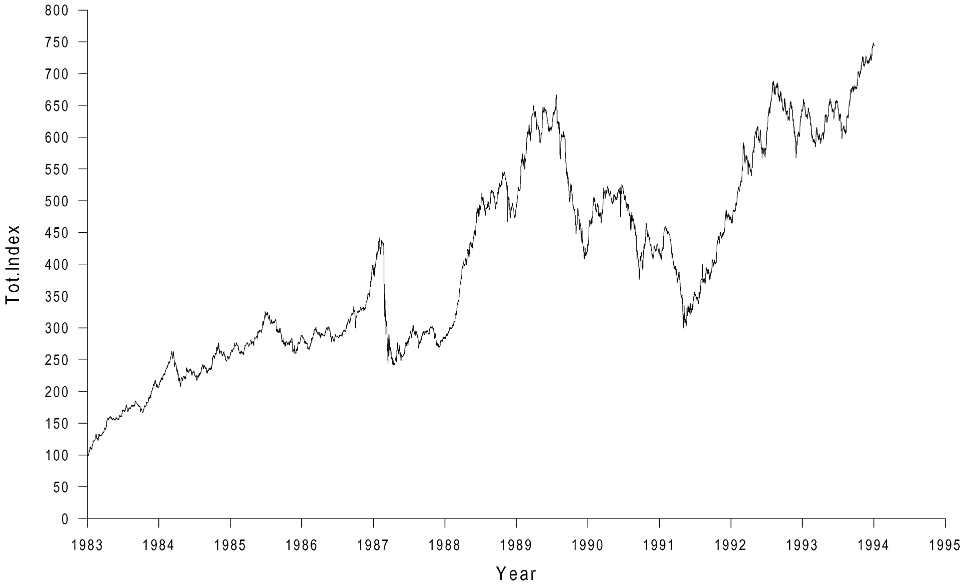

Fig. 12. Daily observations of the OSE general index for the period 1983 to 1995, or 3190 daily closing prices.

prices for all companies registered at OSE. Each stock's weight is calculated as the fraction of the total value of the stock to the total value of all the companies at the stock exchange. [4] The data set only contains trading days and not holidays and weekends. This means that five-day returns may not necessarily need to be for a Monday-to-Friday week, but five consecutive trading days. Therefore, when we talk about 30-day returns this is not "monthly" returns, but rather 30 trading days. A plot of the daily closing prices for the entire period is shown in Fig. 12.

4.1.2. The Dow–Jones Industrial Average (DJIA)

The second data set that will be examined is the daily closing prices of the Dow–Jones industrial average for the period 1962–1993, or 31 years of daily data. [5] The index itself consists of 30 large US stocks. The DJIA is not a "true" average as the name indicates, but rather the price one has to pay to get hold of one of each of the 30 stocks, adjusted by a factor of about 0.346. As was the case for the OSE series, weekends and holidays are also removed from this data set. A plot of the entire data set is shown in Fig. 13.

---

[4] State-owned companies only count 50% of their total value when their weights are calculated. These companies are, e.g. DNB, Kreditkassen, Sydvaranger, Raufoss, Hydro.

[5] The data was obtained from the Ohio State University database and consists of about 8000 observations.

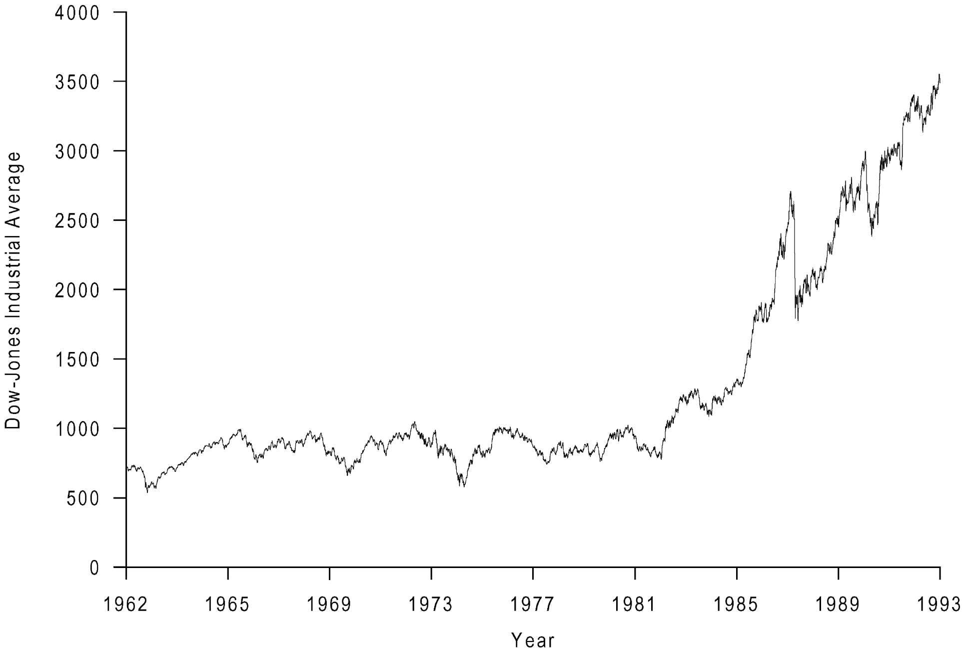

Fig. 13. Daily observations of the Dow–Jones Industrial Average for the period 1962 to 1993, about 8000 daily observations.

4.2. R/S -analysis of the OSE general index

4.2.1. Analysis results

By using the model framework described in Section 3, we are now able to examine whether the OSE general index behaves like a "pure" Brownian motion or not. Fig. 14 shows the log/log plot of R/S values as a function of time lag n for one-day returns of the OSE general index. Also plotted are the $E(R/S)$ values excepted for a Gaussian random-walk calculated by using Eq. (15) for the same range of n. As may be seen there is clearly a systematic deviation between the empirical R/S curve and the Gaussian random walk curve. However, there does not seem to be any "break" in the plot indicating the absence of any non-periodic cycles in the data.

By performing a linear regression on the empirical R/S values versus n we get a Hurst exponent equal to $H = 0.6064$. When doing a similar regression for the *expected* R/S values we find that the expected Hurst exponent is $E(H) = 0.535$. The regression results are shown in Table 1, and we see from the table that the regression on the empirical R/S data has an $R^2 = 0.99$, which is a remarkably good fit.

To determine the significance of the results, we first calculate the variance of the expected Hurst exponent for Gaussian random variables. As discussed in Section 3.1.5 the variance of $E(H)$ is $1/T$, where T is the total number of observations. Thus, the variance of $E(H)$ is 0.0003134(1/3190), or a standard deviation of 0.0177. By using

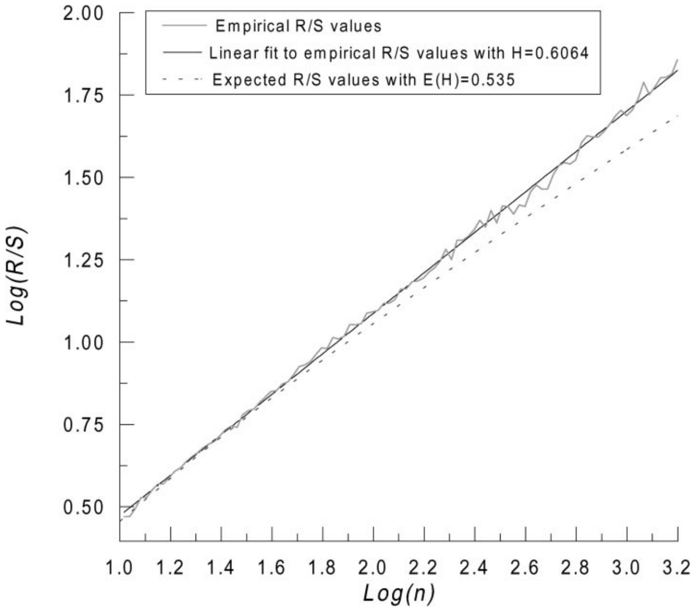

Fig. 14. R/S values as a function of time lag n for a series of daily returns of the OSE general index covering the period 1983 to 1995. The dotted line is the fitted curve to $R/S = an^H$ with $H = 0.6064$. The expected value of H is $E(H) = 0.535$ for the same range of n.

Table 1
Linear regression results from R/S analysis in Fig. 14, both for the empirical and expected R/S values for the range $10 < n < 1550$ or $1 < \log(n) < 3.2$

	OSE (daily)	E(R/S)
Constant	−0.1406	−0.0181
Slope (*Hurst exponent*)	**0.6140**	**0.5351**
Number of points	99	158
Fit range	$10 < n < 1550$	$10 < n < 1550$
Average X	2.1077	2.7740
Average Y	1.1539	1.4655
Regression sum of squares	15.1344	7.592
Residual sum of squares	0.0232	0.0069
Coeff. of determination, R^2	0.9984	0.9991
Residual mean square	0.00024	0.000045
(σ_H) std.dev. $H > E(H)$	**4.50847**	

Eq. (16), the Hurst exponent for the OSE data is calculated to be about 4.51 standard deviations above its expected value. This is a highly significant result at the 95% level.

From the above analysis we see that the one-day changes in the prices of the OSE general index are characterized by a *persistent* Hurst process with a Hurst exponent $H=0.614$. This is significantly different from a random walk. Because the series consists of AR(1) residuals achieved by using the *prewhithening* method described earlier, we know that a true long-memory process is at work. However, since there is no apparent

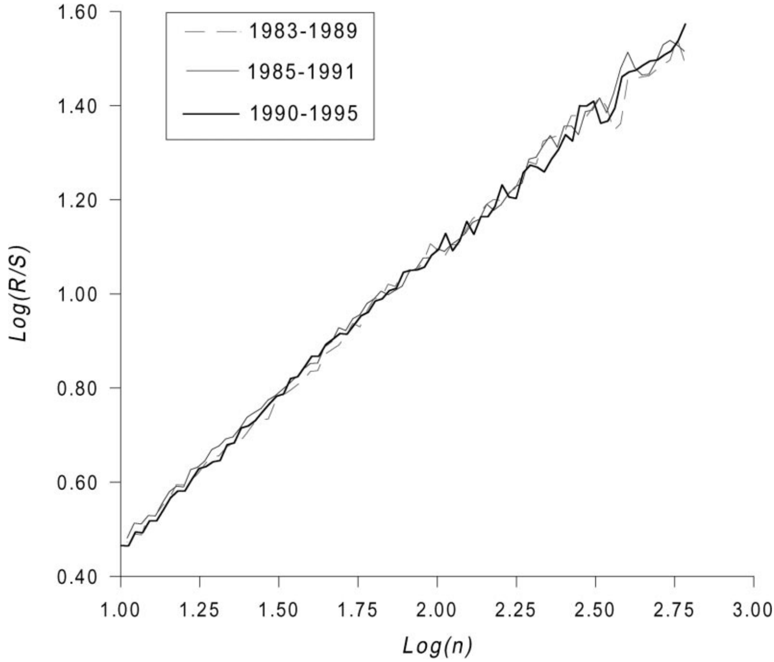

Fig. 15. Stability analysis of the Hurst exponent for the OSE general index data. The Hurst exponent is very stable despite the changing economic conditions for the three periods.

"break" in the plot, the persistent scaling does not seem to have a time limit. This means that there is no average cycle length in the data.

4.2.2. Stability analysis

An interesting point to examine, is the stability of the Hurst exponent for different time periods. The data set was therefore divided into three overlapping periods, and an R/S analysis was performed on each of the data sets for comparison. The periods that were chosen were 1983–1989 (1757 observations), 1985–1991 (1434 observations) and 1990–1995 (1752 observations). A log/log plot of the R/S values versus n for the three data sets is shown in Fig. 15. We see that the three data sets follow approximately the same scaling behaviour.

The linear regression results are shown in Table 2. We see that the Hurst exponent varies slightly between $H=0.604$ and $H=0.610$. All three Hurst exponents lies about 2 standard deviations above the expected Hurst exponent of $E(H)=0.557$, estimated for the same range of n. This means that they all are significant at the 95% level. Thus, the Hurst exponent is remarkably stable for the three periods analysed here. This despite the fact that the data sets cover periods of different underlying economic conditions. The data covering the period 1983–1989 include both the boom in the 1980s, due to the liberalization of the Norwegian credit market, and the stock market crash in 1987. The data set for the period 1990–1995 contains both the Gulf crisis and the

Table 2
Linear regression results from the stability analysis in Fig. 15 for three different periods of the Oslo Stock Exchange General Index

	OSE: 1983–1989	OSE: 1985–1991	OSE: 1990–1995	$E(R/S)$
Constant	−0.1304	−0.1174	−0.1140	−0.0659
Slope (*Hurst exponent*)	**0.6061**	**0.6044**	**0.6102**	**0.5571**
Number of points	80	80	81	29
Fit range (n)	$10 < n < 800$	$10 < n < 800$	$10 < n < 800$	$10 < n < 800$
Average X	1.9004	1.9004	1.8930	2.1811
Average Y	1.0214	1.0311	1.0187	1.1491
Regression sum of squares	7.7761	7.7319	8.1818	1.8139
Residual sum of squares	0.0351	0.0214	0.0249	0.0023
R-squared	0.9955	0.9972	0.9969	0.9987
Residual mean square	0.0004497	0.000274	0.0003154	0.0000846
(σ_H) std.dev. $H > E(H)$	**2.0547**	**1.9824**	**2.2279**	

steady positive developments of the Norwegian economy, while the data set covering the period 1985–1991 includes both the crash of 1987 and the Gulf crisis.

4.2.3. Summary of the results

Clear evidence of persistence and non-random behaviour in the Norwegian stock market is found. We obtain a Hurst exponent of $H = 0.614$ which is significantly higher than the expected Hurst exponent for the same range which is $E(H) = 0.534$. On the other hand, there does not seem to be any "break" in the R/S plot in Fig. 15, which indicates that there is no average cycle in the data set. The stability of the Hurst exponent is remarkable for the three periods examined, despite the fact that they are all taken from different periods of a 13-year time span with quite different underlying economic conditions.

4.3. R/S -analysis of the Dow–Jones industrial average

4.3.1. Analysis results

The Dow–Jones industrial average (DJIA) data set covers a much longer time span than the OSE series, and consists on average of much larger companies than the Norwegian general index. It will therefore be interesting to see if there are any differences between the R/S results for the two stock exchanges. There are two main features that we can expect to find different from the first data set. First, the average non-periodic cycle should show up, if there is any, due to the fact that this is a much longer data set. Second, the persistence could be less pronounced due to the large US market, perhaps physiologically less vulnerable than the much smaller Norwegian market.

The results from the R/S analysis are shown in Fig. 16, where we can see that the DJIA index has two scaling ranges divided by the vertical line segment at the "breakpoint". To the left of the vertical line the slope of the R/S values is obviously higher than for the R/S values to the right of the vertical line. This indicates that there is a non-periodic component in the time series with an average frequency of

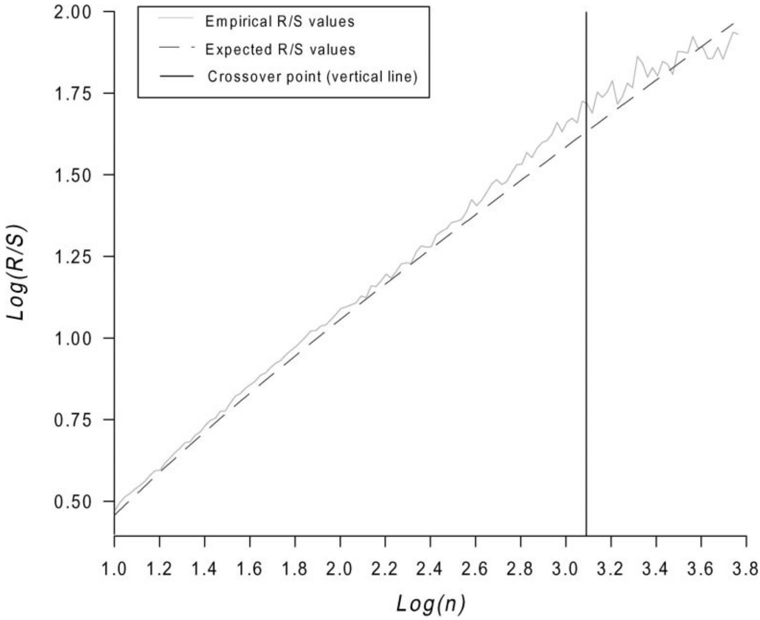

Fig. 16. R/S and $E(R/S)$ values plotted against the time lag n for daily observations on the Dow–Jones ind.average index for the period 1962–1993. The empirical R/S values follow two scaling regimes separated by the vertical line.

Table 3
Linear regression results from R/S analysis for the first scaling range $10 < n < 1260$, both for the empirical and expected R/S values

	DJIA	$E(R/S)$
Constant	−0.0828	−0.0525
Slope (*Hurst exponent*)	**0.5783**	**0.5493**
Number of points	95	40
Fit range	$10 < n < 1260$	$10 < n < 1260$
Average X	2.0477	2.34618
Average Y	1.1014	1.2363
Regression sum of squares	11.8562	3.5059
Residual sum of squares	0.01579	0.00411
Coeff. of determination, R^2	0.9987	0.9988
Residual mean square	0.0001699	0.0001081
(σ_H) std.dev. $H > E(H)$	**2.5916**	

approximately $10^{3.1} = 1260$ trading days, or about 4 years. To determine whether the Hurst exponent is the expected Hurst exponent, we do the analysis separately for the two regions. The results from the first regression for the range $10 < n < 1260$ are shown in Table 3.

From Table 3 we see that the Hurst exponent is $H = 0.578$ and the expected value of H for the same region is $E(H)$=0.5493. This does not seem significant, but the standard

Table 4
Linear regression results from R/S analysis for the second scaling range $1260 < n < 4000$ days, both for the empirical and expected R/S values

	DJIA	$E(R/S)$
Constant	0.7993	0.0589
Slope (*Hurst exponent*)	**0.2997**	**0.5090**
Number of points	30	12
Fit range	$1260 < n < 4000$	$1260 < n < 4000$
Average X	3.4399	3.4201
Average Y	1.8302	1.7999
Regression sum of squares	0.1002	0.1509
Residual sum of squares	0.0269002	4.66136E-007
Coeff. of determination, R^2	0.788322	0.999997
Residual mean square	0.0009607	4.66136E-008
(σ_H) std.dev. $H > E(H)$	**−18.7222**	

deviation of $E(H)$ is 0.0112 for 8000 observations. Thus, the Hurst exponent for the daily DJIA is 2.59 standard deviations above its expected value. This is therefore also a very significant result as was the case for the OSE general index. For the second range, covering $1260 < n < 4000$, a Hurst exponent of $H = 0.2997$ was found, which is 18.7 standard deviations below the expected Hurst exponent of $E(H) = 0.509$. This means that the process becomes antipersistent after the 4-year cycle. An interpretation of these findings will be done in Section 4.3.3. The regression results are shown in Table 4.

4.3.2. Stability analysis

A stability analysis is also performed on the DJIA series to see if the Hurst exponent shows the same degree of stability for the US stock market as the Norwegian market. The DJIA series covers a much longer time span (1962–1993). By dividing it into three non-overlapping sub-periods of equal length, we cover three decades of different underlying economic conditions. The first sub-period is 1962–1972, the second is 1972–1982 and the third is 1982–1993. The R/S analysis results for the three periods are plotted in Fig. 17. As can be seen from the plot, the slopes of the R/S values for the three time periods vary much more than for the OSE data. The regression results is shown in Table 5.

From Table 5 it can be seen that the Hurst exponent for the first period (1962–1971) is equal to $H = 0.6068$ which is 3.11 standard deviations above the expected Hurst exponent of $E(H) = 0.5445$. Thus, the result is very significant. However, for the other two periods (1972–1982, 1983–1993), the Hurst exponents are not significantly different from what is expected for a random process. These results are therefore quite different from those we got for the OSE series, which showed a very stable Hurst exponent. The most plausible explanation for this is probably the fact that the OSE series covered a much shorter time span with too small differences in the underlying economic conditions to affect the Hurst exponent. If we look at the last period covering 1983–1993, this is almost the same period as the OSE data set. We see from

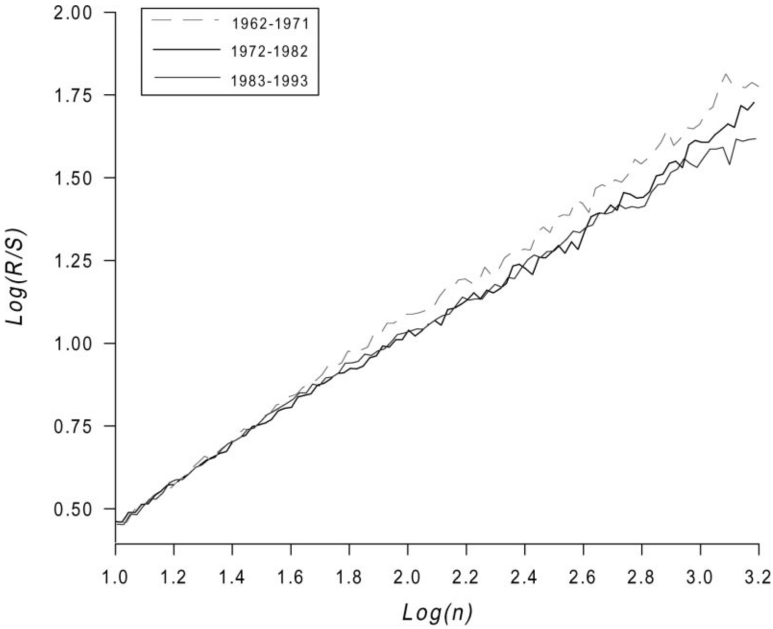

Fig. 17. R/S values versus time lag n for three different time periods for the Dow–Jones industrial average data.

Table 5
Regression results for the stability analysis

	DJIA 1962–1971	DJIA 1972–1982	DJIA 1983–1993	$E(R/S)$
Constant	−0.1433	−0.0871	−0.0551	−0.04143
Slope (*Hurst exponent*)	**0.6068**	**0.5549**	**0.5383**	**0.5445**
Number of points	94	95	93	42
Fit range (n)	$10 < n < 1300$	$10 < n < 1300$	$10 < n < 1300$	$10 < n < 1300$
Average X	2.0515	2.0473	2.0531	2.4311
Average Y	1.1015	1.0490	1.0502	1.2823
Regression sum of squares	12.6458	10.9160	9.6376	3.5176
Residual sum of squares	0.0372	0.0347	0.0274	0.0042
R-squared	0.99707	0.99683	0.99717	0.99882
Residual mean square	0.0004038	0.0003733	0.0003005	0.0001037
(σ_H) std.dev. $H > E(H)$	**3.1113**	**0.5488**	**−0.3187**	

Table 5 that the Hurst exponent for the DJIA is $H = 0.538$ which is not significantly different from a random walk. However, for the OSE series we have a Hurst exponent of $H = 0.614$ which is significantly higher than the expected Hurst exponent for a random process (see Table 1). This indicates that while the OSE General Index shows a persistent behaviour for this period, the DJIA does not.

4.3.3. Summary of the DJIA results

We find that the U.S. stock market represented by the Dow–Jones industrial average Index does not follow a random walk as also is the case for the OSE data. We find a Hurst exponent of $H = 0.578$ for the range $10 < n < 1260$ and an $H = 0.299$ for the range $1260 < n < 4000$, where n is the number of trading days. Thus, there seem to be a non-periodic average cycle of about 1260 trading days or approximately 4 years. It is tempting to couple this cycle to the political elections periods in the US. As the memory effect disappears for $n > 1260$, the series becomes antipersistent with more reversals than a random walk since $H < 0.5$. This conforms to Fama and French's findings [28] that returns are "mean reverting" in the long term.

4.4. Interpretation of the rescaled range analysis results

4.4.1. Main findings

There are three main differences between the results found for the Norwegian and US stock market data. First, the Hurst exponent was much higher for the OSE general index than for the DJIA. Second, there did not appear to be any average cycle in the OSE data, while a four-year average non-periodic cycle was found for the US data. Third, the Hurst exponent for the Norwegian stock market was much more stable than for the US stock market. This is most likely due to larger differences in the underlying economic conditions for the US data that cover a longer time span.

4.4.2. Discussion of the results

So, what may be the explanation for this observed persistent behaviour in the Norwegian and American stock markets? In a persistent market, with $H > 0.5$, capital market returns are influenced by the past, and this influence goes across time scales. One six-week period influences all subsequent six-week periods, one six-month period influences all subsequent six-month periods and so on. These long-term memory effects may be caused by investor bias and market sentiment, fads or fashions and create market trends and non-periodic cycles on all time scales. Thus, the Hurst exponent can be said to be a measure of the impact of market sentiment, generated by past events, upon future returns in the capital markets. The most obvious way to use this information is as the basis of momentum analysis and other forms of technical analysis. This persistence may be related to the "excess volatility" effects discussed by Shiller [31] and LeRoy [32]. They also correspond to the **U**-shaped patterns in first-order autocorrelations across increasing return horizons found by Fama and French [28].

From a *fractal*, *non-linear*, viewpoint, the main reason may lie in what fashion the market reacts to information. As discussed in Section 1.3, the efficient market hypothesis (EMH) postulates that the market reacts to information in a linear fashion. This means that all investors react in the same way to new information instantly as it is received, and they maximize their return on basis of it; they are rational. Thus, theory states that the aggregate market is the equivalent of the typical rational investor, so that the market can value information instantly and efficiently. However, what if investors react to information in a discontinuous, *non-linear* fashion? The new paradigm states that the main movement in markets comes from people observing what others around them with the same investment horizons are doing, and then reacting to it. Thus, there is a kind of dynamic interaction between investors, rather than a specific universal response by all investors identifiable purely to external news arrivals. Some people react to new information as soon as it is received, while others wait until they get confirming information and do not react until a trend is clearly established. Thus, information is accumulated and suddenly reacted on. This creates the long tails in the distribution of price variations on all time horizons. The amount of confirming information necessary to validate a trend varies, but the uneven assimilation of information may cause a biased random walk. This is of course not in agreement with the EMH, but may be a more realistic way of describing markets than in a linear fashion.

The next step in this paper will be to extend the analysis to include an examination of the scaling behaviour of the entire probability distribution of returns. There will also be an examination if stock price returns are well described by the family of Lévy stable distributions often referred to as fractal distributions.

5. Results from the Distributional Scaling Analysis

5.1. The data

The data set used for this exercise is ticker data on the OBX-index at OSE, and consists of almost 333.000 observations covering all trading days in the period 1990–1994. [6] The OBX-index consists of the 25 most frequently traded company stocks on OSE. Which stocks to include in the index is evaluated four times a year. A plot of the entire OBX data set is shown in Fig. 18. The time intervals between successive records are not fixed, but are on average close to 1 min.

5.2. Estimation of α

5.2.1. Estimation of α by using the Mantegna–Stanley method

From the data set, we select the complete set of non-overlapping records separated by a time interval Δt. The value of the OBX index is denoted as $P(t)$ while the successive variations $Z(\Delta t)$ are defined in Eq. (21). We now first determine the probability distributions $\Pr[Z(\Delta t)]$ of index variations for different values of Δt. Here, we have selected logarithmically equally spaced values of Δt ranging from $\Delta t = 1$ to 316 min. Fig. 19 shows a semilogarithmic plot of $\Pr[Z(\Delta t)]$ obtained for six different values of Δt. As expected, the distributions are roughly symmetrical and are spreading when Δt increases.

[6] The OBX data-set was obtained from A. Hagen at Oslo Stock Exchange.

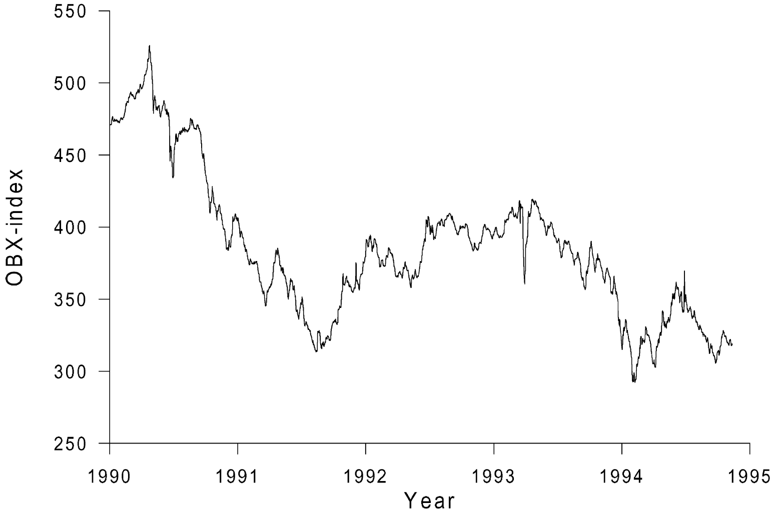

Fig. 18. Daily observations on the OBX-index for the period 1990–1994.

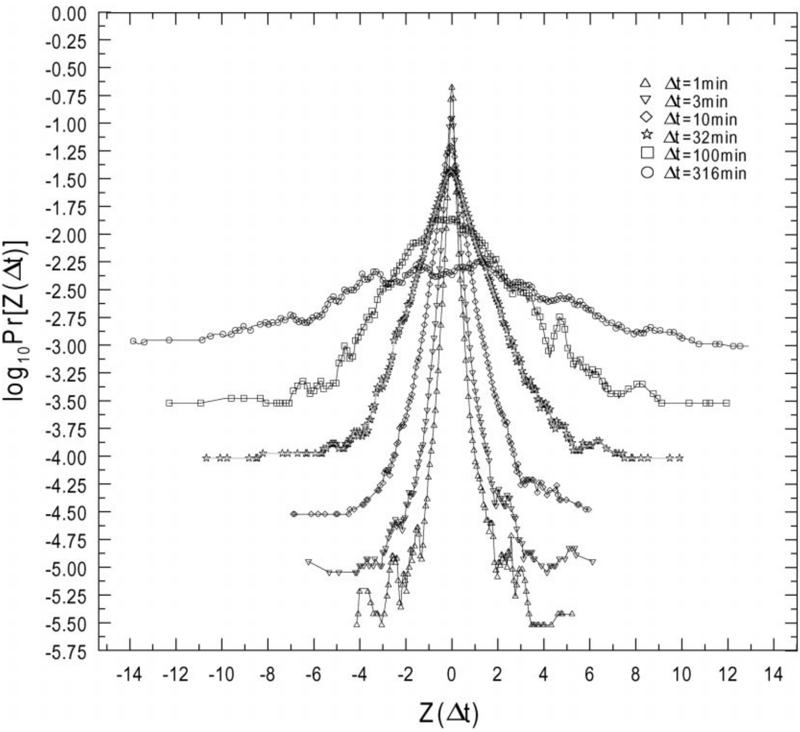

Fig. 19. Probability distributions $\Pr[Z(\Delta t)]$ of the OBX index variations $Z(\Delta t)$ observed at time intervals Δt, which range from 1 to 316 min. As Δt is increased a spreading of the probability distribution is observed.

Table 6
The probability of zero return $\Pr[Z(\Delta t) = 0]$ as a function of the time horizon Δt used in Fig. 20

Δt	$\Pr[Z(\Delta t) = 0]$	$\log(\Delta t)$	$\log Pr[Z(\Delta t) = 0]$
1	0.22320	0	−0.65130
3	0.10782	0.4771	−0.96730
10	0.06133	1	−1.21230
32	0.03373	1.5052	−1.47195
100	0.01356	2	−1.86765
312	0.00631	2.4942	−2.19978

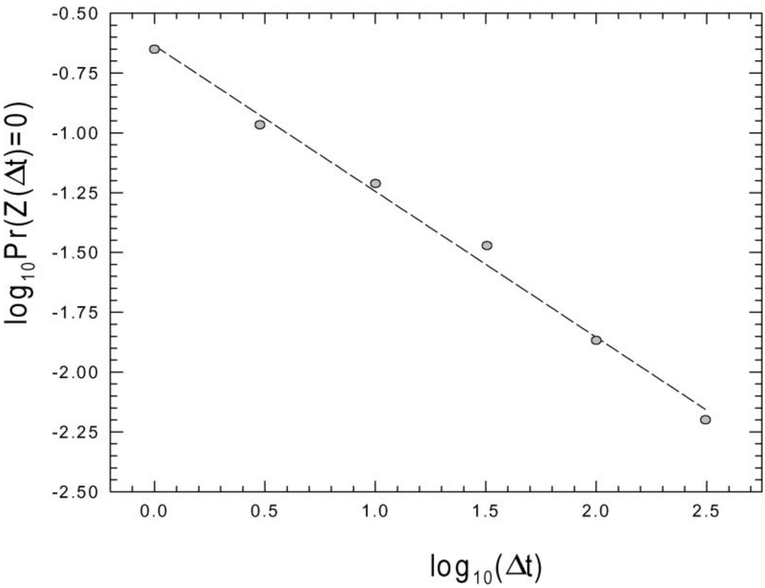

Fig. 20. Probability of zero return $\Pr[Z(\Delta t) = 0]$ as a function of the sampling interval Δt. A scaling law behaviour is observed for time intervals spanning almost three orders of magnitude. The slope of the best-fit straight line is $\lambda = -0.6097$. By using Eq. (23) we obtain the fractal dimension of the probability space of $\alpha = 1.6402$.

To estimate the fractal dimension of the probability space, α, we now use the Mantegna/Stanley method described in Section 3.2.2. Thus, by plotting the probability of zero return, $\Pr[Z(\Delta t) = 0]$, as a function of the time horizon, Δt, we are able to estimate α. The values of $\Pr[Z(\Delta t) = 0]$ for different Δt values are shown in Table 6. In Fig. 19 the same relationship is shown in a log/log plot. The slope of the best-fit straight line is $\lambda = -0.6097$. Thus, we observe a non-normal scaling behaviour ($\lambda \neq -0.5$) in an interval of trading time ranging from 1 to 316 min. From this scaling parameter we estimate α to be $\alpha = 1.64$ through Eq. (23) (Fig. 20).

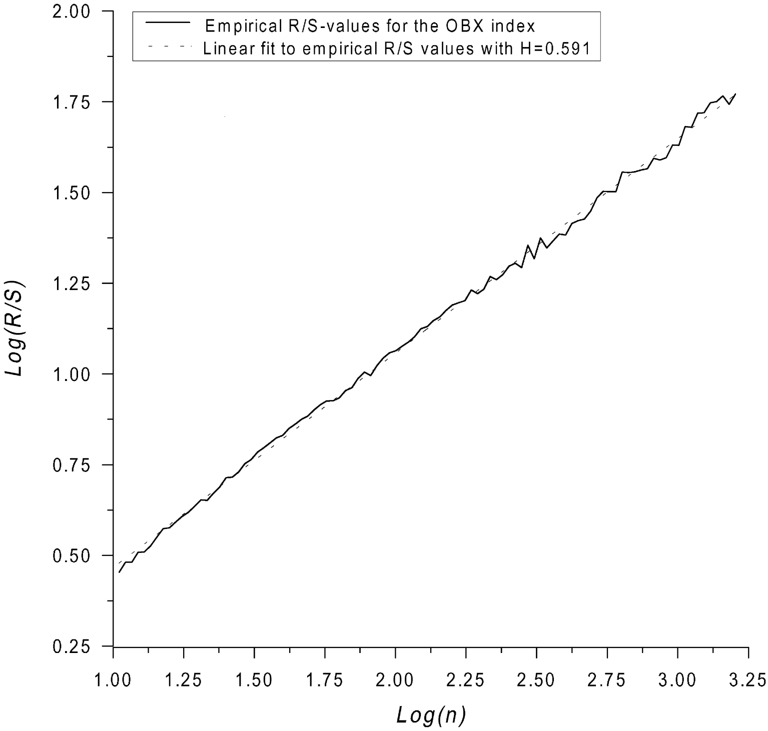

Fig. 21. R/S values as a function of time lag n for 100 min returns (about $\frac{1}{3}$ of a day) on the OBX-index for the period 1990 to 1994. The dotted line is a fitted curve to $R/S = an^H$ with $H = 0.5913$. The expected value of H for a random walk is $E(H) = 0.5417$ for the same range of n.

5.2.2. Estimation of α through R/S analysis

It is of interest to see if we obtain the same estimate of α by using R/S analysis as was obtained by using the Mantegna/Stanley method. It is expected that the Hurst exponent we get here will be approximately similar to the Hurst exponent we found for the total-index in Section 4.2 since these two data sets cover the Norwegian stock market. A description of the OBX data set was done in Section 5.1. We will here use 100 min returns for the R/S analysis due to the fact that a higher frequency implies both an increase in computation time and serial correlation which can bias the R/S analysis, while a lower frequency will mean few observations to perform the analysis. In Fig. 21 the log/log plot of the empirical R/S values as a function of the time lag n is shown. The dotted line is the best-fit straight line with a slope $H = 0.5913$, and an expected Hurst exponent of $E(H)=0.5437$. Since we have about $T=3330$ observations, the expected variance of H is $1/3330=0.0003$ or a standard deviation of 0.0173. Thus, the Hurst exponent lies about 2.75 standard deviations[7] above the expected value for a true Gaussian process. The linear regression results are shown in Table 7. Using

[7] $(0.5913-0.5437)/0.0173 = 2.7468$ standard deviations.

Table 7
Linear regression results from the R/S analysis in Fig. 21, both for the empirical and expected R/S values covering the range $10 < n < 1665$

	OBX-index	$E(R/S)$
Constant	−0.1252	−0.0399
Slope (*Hurst exponent*)	**0.5914**	**0.5438**
Number of points	99	43
Fit range	$10 < n < 1665$	$10 < n < 1665$
Average X	2.1127	2.4488
Average Y	1.1242	1.2917
Regression sum of squares	14.0303	3.67421
Residual sum of squares	0.020455	0.004301
Coeff. of determination, R^2	0.998544	0.998831
Residual mean square	0.0002197	0.0001049
(σ_H) # of std.dev. that Hurst exponent $> E(H)$	**2.7468**	

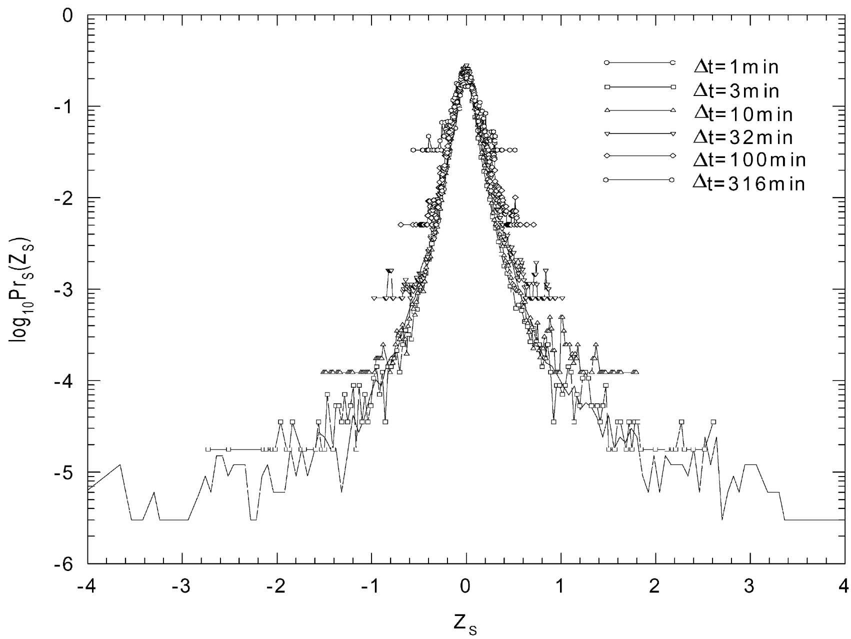

Fig. 22. Scaled plot of the probability distributions shown in Fig. 21. All the data collapse on the $\Delta t = 1$ min distribution by using the scaling transformations of Eqs. (27) and (28) with $\alpha = 1.64$. The flat line of points outside the average behaviour defines the noise level of that specific distribution.

Eq. 20, a value $H = 0.5913$ produces an $\alpha = 1.69$. This is in good agreement with $\alpha = 1.64$ found above using the Mantegna/Stanley method (Fig. 22).

5.3. Applying the estimated α

We have just shown that two completely different methods produced remarkably similar values for α, the fractal dimension of the probability space, for the OBX-index.

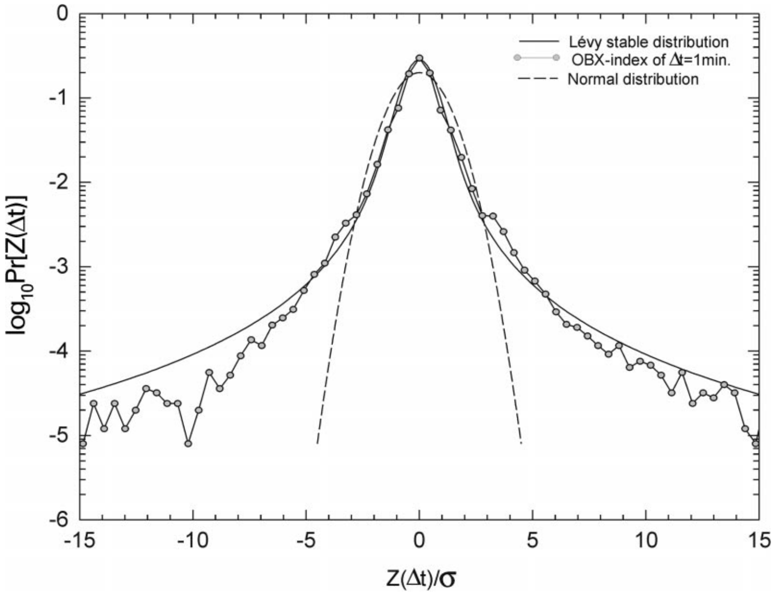

Fig. 23. Comparison of the $\Delta t = 1$ min empirical probability distribution with the symmetrical Lévy stable distribution with $\alpha = 1.64$ and scale factor $\gamma = 0.1167$. The normal distribution shows poor agreement with the data, especially for the wings.

By using the Mantegna/Stanley method examining the scaling of the "probability of zero return", we obtained an $\alpha = 1.64$. By using R/S analysis we got a somewhat higher alpha of $\alpha = 1.69$. However, both estimates indicate that the process is different from a pure random walk. As discussed in Section 3.2.1, $\alpha < 2.0$ is evidence of a non-normal process with larger wings and higher probability around the mean than for a Gaussian process. In fact, these results are evidence of a non-linear fractal system, which support and supplement the results under Section 4.2.

So far, we have tested if there is a scaling behaviour in the "probability of zero return". We now want to check if the scaling extends over the entire probability distribution as well as for $\Pr[Z(\Delta t) = 0]$. As mentioned under Section 3.2, the Lévy stable distributions rescale under the transformations in Eqs. (18) and (19). Fig. 23 shows the rescaled versions of the distributions in Fig. 21. As can be seen, all the data collapse on the $\Delta t = 1$ min distribution using Eqs. (18) and (19) with the estimated $\alpha = 1.64$. We may therefore conclude that a Lévy stable distribution describes well the dynamics of the probability distributions $\Pr[Z(\Delta t)]$ over time intervals spanning almost three orders of magnitude.

In Fig. 23, a comparison is made between the $\Delta t = 1$ min empirical probability distribution for the OBX-index and the Lévy stable distribution of $\alpha = 1.64$. As can be seen from Eq. (16) we need to solve for the scale factor γ. This is done by using the estimated α and $\Pr[Z(\Delta t)=0]$ in Eq. (17) for $\Delta t=1$ min such that

$$\Pr[Z(\Delta t)=0]=\frac{\Gamma(1/\alpha)}{\pi\alpha(\gamma\Delta t)^{1/\alpha}} \quad \text{or} \quad 0.2232=\frac{\Gamma(1/1.64)}{3.14\cdot 1.64(\gamma)^{1.64}}\,. \tag{24}$$

When solved with respect to γ this yields $\gamma=0.1167$. Thus, Fig. 22 shows that by using the estimated α with the corresponding γ we get good agreement between the Lévy distribution and the empirical distribution for the OBX-index, in contrast to the Gaussian normal distribution (dotted line) which poorly describes the data (Fig. 23).

5.4. *Discussion of the Lévy stable distribution analysis results*

5.4.1. *Main findings*

In the distributional scaling analysis performed above, there is evidence of the OBX-index following a scaling law similar to those for non-linear chaotic systems. We found an $\alpha=1.64$ and 1.69 when using the Mantegna/Stanley and R/S analysis methods, respectively. These estimates are remarkably close to each other, especially when taking into account the fact that we are using two completely different statistical methods only similar in the way they use the concept of scaling. This relationship has never, to our knowledge, been shown to hold for financial data. We then checked whether the scaling extended over the entire probability distribution for all Δt examined by using the transformation functions in Eqs. (17) and (18) with $\alpha=1.64$. From Fig. 22 it seems like the dynamics of the distributions for $\Delta t=1$ to 316 min are well described by a scaling parameter of $\alpha=1.64$ which also is the fractal dimension of the probability space. Finally, from Fig. 23 it also seems like the Lévy distribution as well as capturing the dynamics of the distribution, also is able to describe the empirical probabilities of return much better than the Gaussian normal distribution. The findings are also in good agreement with the findings by Stanley and Mantegna [10] for the S&P 500 index.

5.4.2. *Discussion of the results*

Proper modelling of the distribution of stock returns is important because of its significance for investors' management of risk. Standard financial models assume that investors are risk averse. The amount of risk that an investor takes on by holding, e.g. the OBX index is reflected by the distribution of the index returns. Hence, if an investor wrongly assumes that returns are drawn from a normal distribution, the risk taken on by holding the OBX index is likely to be much higher than expected. However, the fact that Lévy distributions with $\alpha<2$ have an infinite or undefined second moment due to the fat tails is a very unappealing property for financial economics. This is because the estimation of risk in the traditional sense is dependent on a finite variance. This was already pointed out by Mandelbrot [6] in 1964 when he observed that the second moment of probability distributions did not tend to any limit. In other words, the volatility of prices varied widely from one period to the next. However, the findings here show that the Lévy distribution is in good agreement with the empirical frequency distribution for the OBX index variations less than ±6 standard deviations (Fig. 23). For returns more than ±6 standard deviations away from the mean there seems to be an exponential "fall-off" of the tails compared to what is predicted by the Lévy distribution. This may indicate that the variance tends to a limit, but at a much slower rate than the Normal distribution. Thus, one approach for proper modelling of the distribution of the variations in the OBX index may be to simulate the tails of the distribution separately from the central part to conserve the finite variance assumption. This approach was first proposed by Mantegna and Stanley for the S&P 500 index [1,10].

Different types of models have been proposed to describe the statistical characteristics of price differences of financial indices as, for example, the behaviour of volatility. Prominent approaches are models using mixtures of distributions and ARCH/GARCH-type [16,17] models. Even though these models appear to produce leptokurtic distributions, they do not produce the observed scaling behaviour of frequency distributions of stock price variations. R/S analysis and distributional scaling analysis of ARCH and GARCH models show that they fail to reproduce the fractal scaling property observed for stock returns [7,10]. In fact, Mantegna/Stanley [10] report that the scaling property of stock returns in their sample is not well approximated by a GARCH(1,1) model. They find that the time evolution of the probability density functions (PDFs) for the GARCH(1,1) scale according to $\lambda = -0.53$, or has a Hurst exponent of $H = 0.53$. Their fractal dimension estimate for a GARCH(1,1) process is thus $\alpha = 1.89$ which is much higher than the $\alpha = 1.64$ that was found for the OBX-index earlier in this chapter.

Changes in institutions or in economic regime may also account, at least partially, for the observed leptokurtosis in the distribution of stock returns. This view allows for the possibility that large outliers in the tail of the distribution of stock returns are drawn from a different distribution than the observations in the centre. However, the Lévy distribution approach explains the events during very large market movements as a reasonable 'draw' from a distribution that also describes the price dynamics during more normal times. The distributional scaling approach used here should not be regarded as a substitute for existing models such as the ARCH. Instead, it is likely that the two approaches can be fruitfully combined to improve the models even further so that the scaling property, leptokurtic distribution and higher moment dependencies in stock returns are all preserved. Option pricing models based on truncated Lévy distributions [1,29,30] seem to be a promising path, though one needs to extend these models further, to also account for the observed time dependence in volatility.

6. General conclusions and future research

In the economic literature dealing with chaos theory, almost no focus has been put on the concept of fractals which is an essential tool for describing and characterizing the dynamics of non-linear processes without any intrinsic time scale. In this paper the concept of fractals has been applied to Norwegian and US stock market data. The analysis has revealed some interesting features of both the Norwegian and US Stock Index dynamics. Both markets show a fractal scaling behaviour significantly different from what a random walk would produce. These findings imply that there are patterns, or trends in returns that persist over time and different time scales. This provides a theoretical platform supporting the use of technical analysis and active trading rules to produce above average returns. The findings may also be used to improve the current models or to make new ones which use the concept of fractal scaling.

Chaos theory has been around for some time now, and critics point out that little of practical value has been achieved. One promise of chaos theory was to offer simple, deterministic models to predict the path of financial markets. So far, little empirical evidence has been found of the type of deterministic chaos that would enable such forecasting models to be found. This paper attacks the problem from a different angle which may prove to be a more powerful and correct approach in describing the stock market dynamics.

A natural continuation of this work would be to compare the results of this analysis with similar analyses performed on other stock exchanges. The use of *multifractals* [12,33,34], which is a kind of generalized fractal analysis, carries the analysis of dynamic scaling behaviour even one step further. This concept has been applied successfully in physical sciences to characterize, e.g. turbulence and oil well logs in the North Sea as a tool for predicting and finding oil deposits. This approach could provide a new framework within which to develop new and better economic and financial models, but this is beyond the scope of this paper. Some preliminary analysis has been performed in [35].

References

[1] R.N. Mantegna, H.E. Stanley, An Introduction to Econphysics: Correlations and Complexity in Finance, Cambridge University Press, Cambridge, 2000.
[2] C.F. Huang, R.L. Litzenberger, Allocative efficiency and the valuation of state contingent securities, Foundations for Financial Economics, North-Holland, Amsterdam, 1988, pp. 83–117 (Chapter 5).
[3] E.F. Fama, Efficient capital markets: a review of theory and empirical work, J. Finance (1970) 383–417.
[4] P. Cootner, The Random Character of Stock Market Prices, MIT Press, Cambridge, 1964.
[5] E.E. Peters, Chaos and Order in the Capital Markets, Wiley, New York, 1991.
[6] B. Mandelbrot, The variation of certain speculative prices, in: P. Cootner (Ed.), The Random Character of Stock Prices, MIT Press, Cambridge, 1964.
[7] E.E. Peters, Fractal Market Analysis, Wiley, New York, 1994.
[8] T.E. Copeland, J.F. Weston, Financial Theory and Corporate Policy, 3rd Edition, Addison-Wesley, Reading, MA, 1992, pp. 208–209.
[9] A.W. Lo, A.C. MacKinlay, The Econometrics of Financial Markets, Princeton University Press, Princeton, NJ, 1997.
[10] R.N. Mantegna, H.E. Stanley, Scaling behaviour in the dynamics of an economic index, Nature 376 (1995) 46–49.
[11] H.E. Hurst, Long-term storage of reservoirs, Trans. Amer. Soc. Civil Eng. 116 (1951) 770–808.
[12] B. Mandelbrot, Fractals and Scaling in Finance: Discontinuity, Concentration, Risk, Springer, Berlin, 1997.

[13] U.A. Müller, M.M. Dacorogna, R.B. Olsen, O.V. Pictet, M. Schwarz, C. Morgenegg, Statistical study of foreign exchange rates, empirical evidence of price change scaling law and intraday analysis, J. Banking and Finance 14 (1995) 1189–1208.
[14] J.L. McCauley, Chaos, Dynamics and Fractals: an algorithmic approach to deterministic chaos, Cambridge Nonlinear Science Series, Vol. 2, Cambridge University Press, Cambridge, 1993.
[15] J. Feder, Fractals, Plenum Press, New York, 1988.
[16] J.D. Hamilton, Time series analysis, Princeton University Press, Princeton, NJ, 1994.
[17] R. Engle, Autoregressive conditional heteroskedasticity with estimates of the variance of U.K. inflation, Econometrica 50 (1982) 987–1007.
[18] E.F. Fama, Efficient capital markets: II, J. Finance 46 (1991) 1575–1617.
[19] B. Mandelbrot, J.R. Wallis, Robustness of the rescaled range R/S in the measurement of non-cyclic long run statistical dependence, Water Resour. Res. 5 (1969) 909.
[20] A.W. Lo, Long-term memory in stock market prices, Econometrica 59 (5) (1991) 1279–1313.
[21] B. Mandelbrot, J.Van Ness, Fractional Brownian motion, fractional noises and applications, SIAM Rev. 10 (1968) 422.
[22] A.A. Anis, E.H. Lloyd, The expected value of the adjusted rescaled Hurst range of independent normal summands, Biometrika 63 (1976) 111.
[23] F. Black, M. Scholes, The pricing of options and corporate liabilities, J. Pol. Economy (1973).
[24] D. Duffie, Dynamic Asset Pricing Theory, Princeton University Press, Princeton, NJ, 1996.
[25] E.F. Fama, The behaviour of stock market prices, J. Bus. 38 (1965).
[26] R.N. Mantegna, H.E. Stanley, Stochastic process with ultraslow convergence to a gaussian: the truncated Lévy flight, Phys. Rev. Lett. 73 (1994) 2916–2919.
[27] R.N. Mantegna, H.E. Stanley, Modeling of financial data: Comparison of the truncated lévy flight and the ARCH(1) and GARCH(1,1) processes, Physica A 254 (1998) 77–84.
[28] E.F. Fama, K.R. French, Permanent and temporary components of stock prices, J. Pol. Economy 21 (1988) 246–267.
[29] A. Matacz, Financial modeling and option theory with the truncated Lévy process, Working paper at the Economic Working Paper Archive at WUSTL, 1997.
[30] S.I. Boyarchanko, S.Z. Levendorksi, Option pricing for truncated Lévy Processes under transaction costs, Appl. Math. Finance (1999) submitted for publication.
[31] R.J. Shiller, Do stock prices move too much to be justified by subsequent changes in dividends, Amer. Econom. Rev. 71 (1981) 421–436.
[32] S.F. LeRoy, Efficient capital markets and martingales, J. Econom. Lit. 27 (1989) 1583–1621.
[33] L. Calvet, A. Fisher, B. Mandelbrot, Large deviations and the distribution of price changes, Cowles Foundation Discussion Paper no. 1165 (1997).
[34] L. Calvet, A. Fisher, B. Mandelbrot, Multifractality of Deutschemark/US Dollar exchange rates, Cowles Foundation Discussion Paper No 1165 (1997).
[35] J.A. Skjeltorp, Fractal scaling behaviour in the Norwegian stock market, appendix 13, M.Sc. thesis, Norwegian School of Management (BI), 1996.