Value at Risk and Self–Similarity

Olaf Menkens

School of Mathematical Sciences, Dublin City University,

Glasnevin, Dublin 9, Ireland

olaf.menkens@dcu.ie

January 10, 2007

Abstract

The concept of Value at Risk measures the "risk" of a portfolio and is a statement of the following form: With probability q the potential loss will not exceed the Value at Risk figure. It is in widespread use within the banking industry.

It is common to derive the Value at Risk figure of d days from the one of one–day by multiplying with $\sqrt{d}$. Obviously, this formula is right, if the changes in the value of the portfolio are normally distributed with stationary and independent increments. However, this formula is no longer valid, if arbitrary distributions are assumed. For example, if the distributions of the changes in the value of the portfolio are self–similar with Hurst coefficient H, the Value at Risk figure of one–day has to be multiplied by d^H in order to get the Value at Risk figure for d days.

This paper investigates to which extent this formula (of multiplying by $\sqrt{d}$) can be applied for all financial time series. Moreover, it will be studied how much the risk can be over– or underestimated, if the above formula is used. The scaling law coefficient and the Hurst exponent are calculated for various financial time series for several quantiles.

JEL classification: C13, C14, G10, G21.
Keywords: Square–root–of–time rule, time–scaling of risk, scaling law, Value at Risk, self–similarity, order statistics, Hurst exponent estimation in the quantiles.

1 Introduction

There are several methods of estimating the risk of an investment in capital markets. A method in widespread use is the Value at Risk approach. The concept of **Value at Risk** (VaR) measures the "risk" of a portfolio. More precisely, it is a statement of the following form: With probability q the potential loss will not exceed the Value at Risk figure.

Although this concept has several disadvantages (e.g. it is not subadditive and thus not a so–called coherent risk measure (see Artzner et al. [1]), however see also Daníelsson et al. [4]), it is in widespread use within the banking industry. It is common to derive the Value at Risk figure of d days from the one of one–day by multiplying the Value at Risk figure of one–day with $\sqrt{d}$. Even banking supervisors recommend this procedure (see the Basel Committee on Banking Supervision [2]).

Obviously, this formula is right, if the changes in the value of the portfolio are normally distributed with stationary and independent increments (namely a Brownian Motion). However, this formula is no longer valid, if arbitrary distributions are assumed. For example, if the distributions of the changes in the value of the portfolio are self–similar with Hurst coefficient H, the Value at Risk figure of one–day has to be multiplied by d^H in order to get the Value at Risk figure for d days.

In the following, it will be investigated to what extent this scaling law (of multiplying with $\sqrt{d}$) can be applied for financial time series. Moreover, it will be studied how much the risk can be over– or underestimated, if the above formula is used. The relationship between the scaling law of the Value at Risk and the self–similarity of the underlying process will be scrutinized.

The outline of the paper is the following; the considered problem will be set up in a mathematical framework in the second section. In the third section, it will be investigated how much the risk can be over- or underestimated, if the formula (2) (see below) is used. The fourth section deals with the estimation of the Hurst coefficient via quantiles, while the fifth section describes the used techniques. The sixth section considers the scaling law for some DAX–stocks and for the DJI and it's 30 stocks. In the seventh section the Hurst exponents are estimated for the above financial time series. Possible interpretations in finance of the Hurst exponent are given in section eight. The ninth section concludes the paper and gives an outlook.

2 The Set Up

Speaking in mathematical terms, the Value at Risk is simply the q–quantile of the distribution of the change of value for a given portfolio P. More specifically,

$$\mathrm{VaR}_{1-q}(P^d) = -F^{-1}_{P^d}(q), \tag{1}$$

where P^d is the change of value for a given portfolio over d days (the d–day return) and F_{P^d} is the distribution function of P^d. With this definition, this paper considers the commercial return

$$P^d_c(t) := \frac{P(t) - P(t-d)}{P(t-d)}$$

as well as the logarithmic return

$$P^d_l(t) := \ln(P(t)) - \ln(P(t-d)),$$

where $P(t)$ is the value of the portfolio at time t. Moreover, the quantile function F^{-1} is a "generalized inverse" function

$$F^{-1}(q) = \inf\{x : F(x) \geq q\}, \quad \text{for } 0 < q < 1.$$

Notice also that it is common in the financial sector to speak of the q-quantile as the $1-q$ Value at Risk figure. Furthermore, it is common in practice to calculate the overnight Value at Risk figure and derive from this the d–th day Value at Risk figure with the following formula

$$\mathrm{VaR}_{1-q}(P^d) = \sqrt{d} \cdot \mathrm{VaR}_{1-q}(P^1). \tag{2}$$

This is true, if the changes of value of the considered portfolio for d days P^d are normally distributed with stationary and independent increments and with standard deviation $\sqrt{d}$ (i.e., $P^d \sim \mathcal{N}(0, d)$). In order to simplify the notation the variance $\sigma^2 \cdot d$ has been set to d, meaning $\sigma^2 = 1$. However, the following calculation is also valid for $P^d \sim \mathcal{N}(0, \sigma^2 \cdot d)$.

$$F_{P^d}(x) \quad = \quad \int_{-\infty}^{x} \frac{1}{\sqrt{2\pi d}} \exp\left(-\frac{z^2}{2d}\right) dz$$

$$
\begin{aligned}
&= \int_{-\infty}^{\frac{x}{\sqrt{d}}} \frac{1}{\sqrt{2\pi d}} \exp\left(-\frac{w^2}{2}\right) \sqrt{d}\, dw \\
&= \int_{-\infty}^{\frac{x}{\sqrt{d}}} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{w^2}{2}\right) dw \\
&= F_{P^1}\left(d^{-\frac{1}{2}} x\right),
\end{aligned}
$$

where the substitution $z = \sqrt{d} \cdot w$ was used. Applying this to $F_{P^d}^{-1}$ yields

$$
\begin{aligned}
F_{P^d}^{-1}(q) &= \inf\{x : F_{P^d}(x) \geq q\} \\
&= \inf\{x : F_{P^1}(d^{-\frac{1}{2}} x) \geq q\} \\
&= \inf\{\sqrt{d} \cdot w : F_{P^1}(w) \geq q\} \\
&= \sqrt{d} \cdot F_{P^1}^{-1}(q).
\end{aligned}
$$

On the other hand, if the changes of the value of the portfolio P are self–similar with Hurst coefficient H, equation (2) has to be modified in the following way:

$$
\mathrm{VaR}_{1-q}(P^d) = d^H \cdot \mathrm{VaR}_{1-q}(P^1). \tag{3}
$$

To verify this equation, let us first recall the definition of self–similarity (see for example Samorodnitsky and Taqqu [16], p. 311, compare also with Embrechts and Maejima [10]).

Definition 2.1
A real–valued process $(X(t))_{t\in\mathbb{R}}$ is **self–similar with index** *$H > 0$ (H–ss) if for all $a > 0$, the finite–dimensional distributions of $(X(at))_{t\in\mathbb{R}}$ are identical to the finite–dimensional distributions of $\left(a^H X(t)\right)_{t\in\mathbb{R}}$, i.e., if for any $a > 0$*

$$
(X(at))_{t\in\mathbb{R}} \stackrel{d}{=} \left(a^H X(t)\right)_{t\in\mathbb{R}}.
$$

This implies

$$
\begin{aligned}
F_{X(at)}(x) &= F_{a^H X(t)}(x) \quad \text{for all } a > 0 \text{ and } t \in \mathbb{R} \\
&= P\left(a^H X(t) < x\right) \\
&= P\left(X(t) < a^{-H} x\right) \\
&= F_{X(t)}(a^{-H} x).
\end{aligned}
$$

Thus, the assertion (3) has been verified. So far, there are just three papers known to the author, which also deal with the scaling behavior of Value at Risk (see Diebold et al. [7] or [6], Dowd et al. [8], and Daníelsson and Zigrand [5]).

For calculating the Value at Risk figure, there exist several possibilities, such as the historical simulation, the variance–covariance approach and the Monte Carlo simulation. Most recently, the extreme value theory is also taking into consideration for estimating the Value at Risk figure. In the variance–covariance approach the assumption is made that the time series (P^d) of an underlying financial asset is normally distributed with independent increments and with drift μ and variance σ^2, which are estimated from the time series. Since this case assumes a normal distribution with stationary and independent increments, equation (2) obviously holds. Therefore, this case will not be considered in this paper. Furthermore, the Monte Carlo simulation will not be considered either, since a particular stochastic model is chosen for the simulation. Thus the self–similarity holds for the Monte Carlo simulation, if the chosen underlying stochastic model is self–similar. The extreme value theory approach is a semi–parametric model, where the tail thickness will be estimated by empirical methods (see for example Daníelsson and de Vries [3] or Embrechts et al. [9]). However, this tail index estimator already determines the scaling law coefficient.

There exists a great deal of literature on Value at Risk, which covers the variance–covariance approach and the Monte Carlo simulation. Just to name the most popular, see for example Jorion [14] or Wilmott [20]. For further references see also the references therein.

However, in practice banks often estimate the Value at Risk via order statistics, which is the focus of this paper. Let $G_{j:n}(x)$ be the distribution function of the j–th order statistics. Since the probability, that exactly j observations (of a total of n observations) are less or equal x, is given by (see for example Reiss [15] or Stuart and Ord [19])

$$\frac{n!}{j! \cdot (n-j)!} F(x)^j \, (1 - F(x))^{n-j} \, ,$$

it can be verified, that

$$G_{j:n}(x) = \sum_{k=j}^{n} \frac{n!}{k! \cdot (n-k)!} F(x)^k \, (1 - F(x))^{n-k} \, . \tag{4}$$

This is the probability, that at least j observations are less or equal x given a total of n observations.

Equation (4) implies, that the self–similarity holds also for the distribution function of the j–th order statistics of a self–similar random variable. In this case, one has

$$G_{j:n,P^d}(x) = G_{j:n,P^1}(d^{-H} \cdot x) \, .$$

It is important, that one has for (P^1) as well as for (P^d) n observations, otherwise the equation does not hold. This shows that the j–th order statistics preserves – and therefore shows – the self–similarity of a self–similar process. Thus the j–th order statistics can be used to estimate the Hurst exponent as it will be done in this paper.

3 Risk Estimation for Different Hurst Coefficients

This section investigates how much the risk is over– or underestimated if equation (2) is used although equation (3) is actually the right equation for $H \neq \frac{1}{2}$. In this case, the difference $d^H - \sqrt{d}$ determines how much the risk will be underestimated (respectively overestimated, if the difference is negative). For example, for $d = 10$ days and $H = 0.6$ the underestimation will be of the size 0.82 or 25.89% (see Table 1). This underestimation will even extent to 73.7% if the one year Value at Risk is considered (which is the case $d = 250$). Here, the relative difference has been taken with respect to that value (namely $\sqrt{d}$), which is used by the banking industry.

Most important is the case $d = 10$ days, since banks are required to calculate not only the one–day Value at Risk but also the ten–day Value at Risk. However, the banks are allowed to derive the ten–day Value at Risk by multiplying the one–day Value at Risk with $\sqrt{10}$ (see the Basel Committee on Banking Supervision [2]). The following table shows (see Table 2), how much the ten–day Value at Risk is underestimated (or overestimated), if the considered time series are self–similar with Hurst coefficient H.

Table 1: Value at Risk and Self–Similarity I

Days	$H = 0.55$			$H = 0.6$		
d	$d^{0.55}$	$d^{0.55} - d^{\frac{1}{2}}$	Relative Difference in Percent	$d^{0.6}$	$d^{0.6} - d^{\frac{1}{2}}$	Relative Difference in Percent
5	2.42	0.19	8.38	2.63	0.39	17.46
10	3.55	0.39	12.2	3.98	0.82	25.89
30	6.49	1.02	18.54	7.7	2.22	40.51
250	20.84	5.03	31.79	27.46	11.65	73.7

This table shows d^H, the difference between d^H and $\sqrt{d}$, and the relative difference $\frac{d^H - \sqrt{d}}{\sqrt{d}}$ for various days d and for $H = 0.55$ and $H = 0.6$.

Table 2: Value at Risk and Self–Similarity II

H	10^H	$10^H - 10^{\frac{1}{2}}$	Relative Difference in Percent	H	10^H	$10^H - 10^{\frac{1}{2}}$	Relative Difference in Percent
0.35	2.24	- 0.92	- 29.21	0.4	2.51	- 0.65	- 20.57
0.45	2.82	- 0.34	- 10.87	0.46	2.88	- 0.28	- 8.8
0.47	2.95	- 0.21	- 6.67	0.48	3.02	- 0.14	- 4.5
0.49	3.09	- 0.07	- 2.28	0.5	3.16	0	0
0.51	3.24	0.07	2.33	0.52	3.31	0.15	4.71
0.53	3.39	0.23	7.15	0.54	3.47	0.31	9.65
0.55	3.55	0.39	12.2	0.56	3.63	0.47	14.82
0.57	3.72	0.55	17.49	0.58	3.8	0.64	20.23
0.59	3.89	0.73	23.03	0.6	3.98	0.82	25.89
0.61	4.07	0.91	28.82	0.62	4.17	1.01	31.83
0.63	4.27	1.1	34.9	0.64	4.37	1.2	38.04
0.65	4.47	1.3	41.25	0.66	4.57	1.41	44.54

This table shows 10^H, the difference between 10^H and $\sqrt{10}$, and the relative difference $\frac{10^H - \sqrt{10}}{\sqrt{10}}$ for various Hurst exponents H.

4 Estimation of the Hurst Exponent via Quantiles

The Hurst exponent is often estimated via the p–th moment with $p \in \mathbb{N}$. This can be justified with the following

Proposition 4.1
Suppose $Y(k) = m(k) + X(k)$ with a deterministic function $m(k)$ and $X(k)$ is a stochastic process with all moments $\mathbb{E}\left[|X(k)|^p\right]$ existing for $k \in \mathbb{N}$ and distributions $F_k(x) := Prob(\{\omega \in \Omega \,:\, X(k,\omega) \leq x\})$ symmetric to the origin. Then the following are equivalent:

1. *For each $p \in \mathbb{N}$ holds:*
$$\mathbb{E}\left[|Y(k) - \mathbb{E}\left[Y(k)\right]|^p\right] = c(p) \cdot \sigma^p |k|^{pH} \tag{5}$$
2. *For each k the following functional scaling law holds on $SymC_0^0(\mathbb{R})$:*
$$F_k(x) = F_1(k^{-H}x)\,, \tag{6}$$
where $SymC_0^0(\mathbb{R})$ is the set of symmetric (with respect to the y–axis) continuous functions with compact support.

This has basically been shown by Singer et al. [18].

Example 4.2
Let Y be a normal distributed random variable with variance σ. It is well–known that

$$\mathbb{E}\left[(Y-\mathbb{E}[Y])^p\right]=\left\{\begin{array}{ll}0 & \textit{if } p \textit{ is odd.}\\ \sigma^p\,(p-1)\,(p-3)\cdot\ldots\cdot 3\cdot 1 & \textit{else}\end{array}\right\}.$$

Hence, in this case Proposition 4.1 holds with $H=\frac{1}{2}$ and

$$c(p)=\left\{\begin{array}{ll}0 & \textit{if } p \textit{ is odd.}\\ (p-1)\,(p-3)\cdot\ldots\cdot 3\cdot 1 & \textit{else}\end{array}\right\}.$$

Proposition 4.1 states that the p–th moment obeys a scaling law for each p given by equation (5) if a process is self–similar with Hurst coefficient H and the p–th moment exists for each $p\in\mathbb{N}$. In order to check whether a process is actually self–similar with Hurst exponent H, it is most important, that H is independent of p. However, often the Hurst exponent will be estimated just from one moment (mostly $p=1$ or 2), see Evertsz et al. [11]. For more references on the Hurst coefficient see also the references therein. This is, because the higher moments might not exist (see for example Samorodnitsky and Taqqu [16], p. 18 and p. 316). Anyway, it is not sufficient to estimate the Hurst exponent just for one moment, because the important point is, that the Hurst exponent $H=H(p)$ is equal for all moments, since the statement is for each $p\in\mathbb{N}$ in the proposition. Thus equation (5) is a necessary condition, but not a sufficient one, if it is verified only for some $p\in\mathbb{N}$, but not for all $p\in\mathbb{N}$.

However, even if one has shown, that equation (5) hold for each p, one has just proved, that the one–dimensional marginal distribution obeys a functional scaling law. Even worse is the fact that this proves only that this functional scaling law holds just for symmetric functions. In order to be a self–similar process, a functional scaling law must hold for the finite–dimensional distribution of the process (see Definition 2.1, p. 3).

The following approach for estimating the Hurst coefficient is more promising, since it is possible to estimate the Hurst coefficient for various quantiles. Therefore, it is possible to observe the evolution of the estimation of the Hurst coefficient along the various quantiles. In order to derive an estimation of the Hurst exponent, let us recall, that

$$\mathrm{VaR}_{1-q}\left(P^d\right)=d^H\cdot\mathrm{VaR}_{1-q}\left(P^1\right),$$

if (P^d) is H–ss. Given this, it is easy to derive that

$$\log\left(\mathrm{VaR}_{1-q}(P^d)\right)=H\cdot\log(d)+\log\left(\mathrm{VaR}_{1-q}\left(P^1\right)\right). \tag{7}$$

Thus the Hurst exponent can be derived from the gradient of a linear regression in a log–log–plot.

4.1 Error of the Quantile Estimation

Obviously, (7) can only be applied, if $\mathrm{VaR}_{1-q}(P^d)\neq 0$. Moreover, close to zero, a possible error in the quantile estimation will lead to an error in (7), which is much larger than the original error from the quantile estimation.

Let l be the number, which represents the q–th quantile of the order statistics with n observations. With this, x_l is the q–th quantile of an ordered time series X, which consists of n observations with $q=\frac{l}{n}$. Let X be a stochastic process with a differentiable density function $f > 0$. Then Stuart and Ord [19] showed, that the variance of x_l is

$$\sigma_{x_l}^2 = \frac{q \cdot (1-q)}{n \cdot (f(x_l))^2}\,,$$

where f is the density function of X and f must be strict greater than zero.

The propagation of errors are calculated by the total differential. Thus, the propagation of this error in (7) is given by

$$\begin{aligned}\sigma_{\log(x_l)} &= \frac{1}{x_l} \cdot \left(\frac{q \cdot (1-q)}{n \cdot (f(x_l))^2}\right)^{\frac{1}{2}} \\ &= \sqrt{\frac{q \cdot (1-q)}{n}} \cdot \frac{1}{x_l \cdot f(x_l)}\,.\end{aligned}$$

For example, if $X \sim \mathcal{N}(0, \sigma^2)$ the propagation of the error can be written as

$$\begin{aligned}\sigma_{\log(x)} &= \sqrt{\frac{q \cdot (1-q)}{n}} \cdot \frac{\sqrt{2\pi} \cdot \sigma}{x \cdot \exp\left(-\frac{x^2}{2\sigma^2}\right)} \\ &= \sqrt{\frac{q \cdot (1-q)}{n}} \cdot \frac{\sqrt{2\pi}}{y \cdot \exp\left(-\frac{y^2}{2}\right)}\,,\end{aligned}$$

where the substitution $\sigma \cdot y = x$ has been used. This shows, that the error is independent of the variance of the underlying process, if this underlying process is normally distributed (see also Figure 1).

Similarly, if X is Cauchy with mean zero, the propagation of the error can be shown to be

$$\begin{aligned}\sigma_{\log(x)} &= \sqrt{\frac{q \cdot (1-q)}{n}} \cdot \frac{\pi \cdot (x^2 + \sigma^2)}{\sigma \cdot x} \\ &= \sqrt{\frac{q \cdot (1-q)}{n}} \cdot \frac{\pi \cdot (y^2 + 1)}{y}\,,\end{aligned}$$

where the substitution $\sigma \cdot y = x$ has also been used. Once again the error is independent of the scaling coefficient σ of the underlying Cauchy–process. Since for Levy–processes with Hurst coefficient $\frac{1}{2} < H < 1$ closed forms for the density functions do not exit, the error can not be calculated explicitly as in the normal and in the Cauchy case.

Figure 1 shows, that the error is minimal around the five percent quantile in the case of a normal distribution, while for a Cauchy distribution, the error is minimal around the 20 percent quantile (see Figure 2). Furthermore, the minimal error is in the normal case even less than half as large as in the Cauchy case.

This error analysis shows already the major drawback of estimating the Hurst exponent via quantiles. Because of the size of the error, it is not possible to estimate the Hurst exponent around the 50 percent quantile. However, it is still possible to estimate the Hurst coefficient in the (semi–)tails. Moreover, it is possible to check, whether the Hurst exponent remains constant for various quantiles.

Hartung et al. [13] state that the $1 - \alpha$ confidence interval for the q–th quantile of an order statistics, which is based on n points, is given approximately by $[x_r; x_s]$. Here r and s are the next higher natural numbers of

$$\begin{aligned}r^* &= n \cdot q - u_{1-\alpha/2}\sqrt{n \cdot q(1-q)} \qquad \text{and} \\ s^* &= n \cdot q + u_{1-\alpha/2}\sqrt{n \cdot q(1-q)}\,, \text{ respectively.}\end{aligned}$$

Figure 1: Error Function for the Normal Distribution, Left Quantile

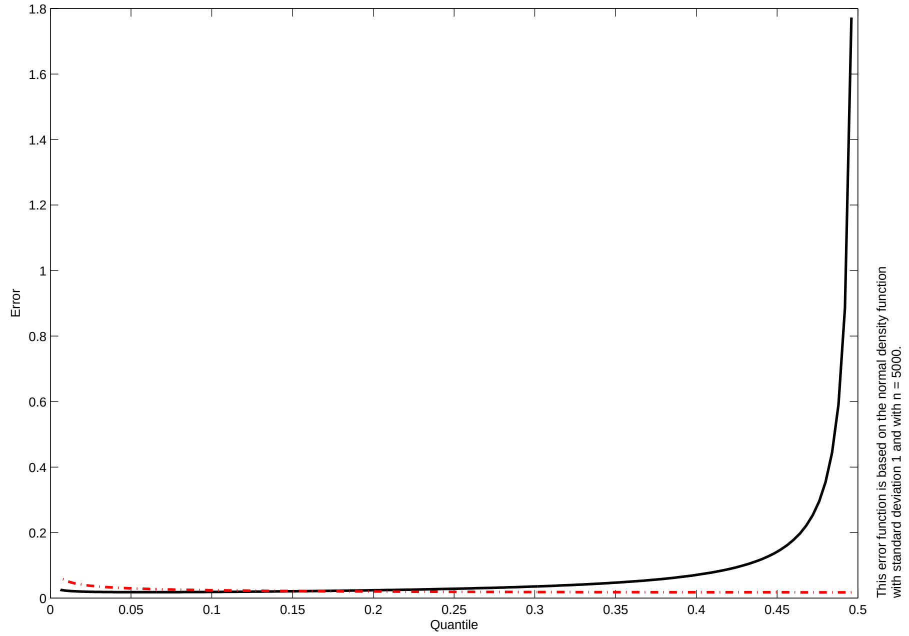

This figure shows the error curve of the logarithm of the quantile estimation (black solid line). Moreover, the red dash–dotted line depicts the error curve of the quantile estimation.

The notation u_α has been used for the α–quantile of the $N(0,1)$–distribution. Moreover, Hartung et al. [13] say, that this approximation can be used, if $q \cdot (1-q) \cdot n > 9$. Therefore this approximation can be used up to $q = 0.01$ (which denotes the 1 percent quantile and will be the lowest quantile to be considered in this paper), if $n > 910$, which will be the case in this paper.

Obviously, these confidence intervals are not symmetric, meaning that the distribution of the error of the quantile estimation is not symmetric and therefore not normally distributed. However, the error of the quantile estimation is asymptotically normally distributed (see e.g. Stuard and Ord [19]). Thus for large n the error is approximately normally distributed. Bearing this in mind, an error σ_{x_l} for the q–th quantile estimation will be estimated by setting $u_\alpha = 1$ and making the approximation

$$\sigma_{x_l} \approx \frac{x_s - x_r}{2} \qquad \text{with } l = n \cdot q\,.$$

5 Used Techniques

Since the given financial time series do not have enough sample points to consider independent 2^j–day returns for $j = 1 \ldots 4$, this paper uses overlapping data in order to get more sample points.

Figure 2: Error Function for the Normal and the Cauchy Distribution, a Zoom In

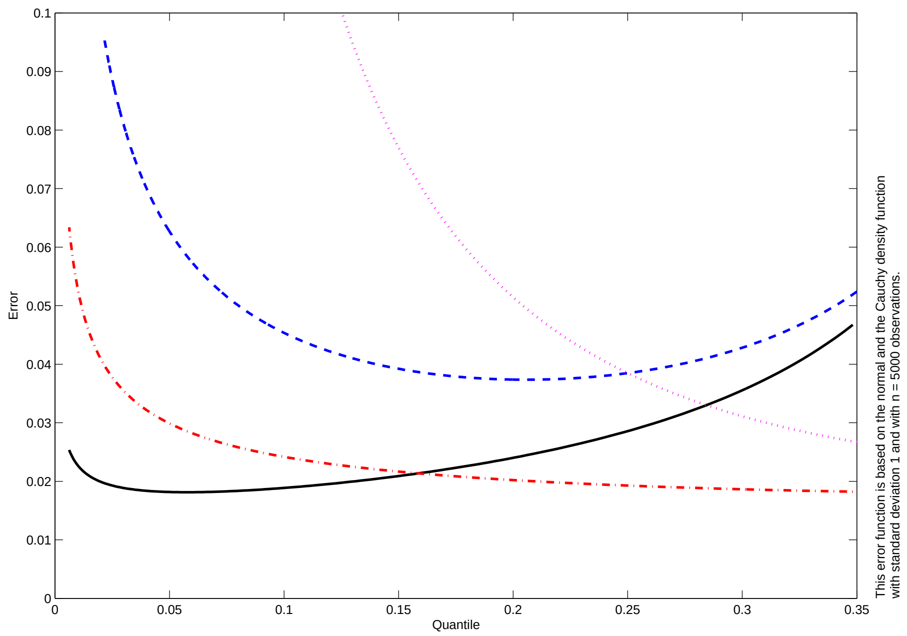

This figure shows the error curves of the logarithm of the quantile estimation for the normal distribution (black solid line) and for the Cauchy distribution (blue dashed line). Moreover, the red dash–dotted line depicts the error curve of the quantile estimation in the case of the normal distribution and the magenta dotted line is the error curve of the quantile estimation in the case of the Cauchy distribution.

5.1 Detrending

Generally, it is assumed, that financial time series have an exponential trend. This drift has been removed from a given financial time series X in the following way.

$$^{d}X_t = \exp\left(\log(X_t) - \frac{t}{T}\left(\log(X_T) - \log(X_0)\right)\right), \tag{8}$$

where ^{d}X will be called the *detrended financial time series associated to the original financial time series* X. This expression for ^{d}X will be abbreviated by the phrase **detrended financial time series**. In order to understand the meaning of this detrending method let us consider $Y_t := \log(X_t)$, which is the cumulative logarithmic return of the financial time series (X_t). Assume that Y_t has a drift this means, that the drift in Y has been removed in ^{dl}Y and thus the exponential drift in X has been removed in $^{d}X_t = \exp\left(^{dl}Y_t\right)$. Observe, that $^{d}X_t$ is a bridge from X_0 to X_0. Correspondingly, $^{dl}Y_t$ is a bridge from Y_0 to Y_0.

Observe, that this method of detrending is not simply subtracting the exponential drift from the given financial time series, rather it is dividing the given financial time series by the exponential of the drift of the underlying logarithmic returns as it can be seen from the formula. By subtracting the drift one could get negative stock prices, which is avoided with the above described method of detrending.

It is easy to derive, that building a bridge in this way is the same as subtracting its mean from the one–day logarithmic return. In order to verify this, let us denote with $P_l^1(t)$ the one–day logarithmic return of the given time series X (as it has been defined in section 1). Therefore, one has $P_l^1(t) = Y_t - Y_{t-1}$. Keeping this in mind, one gets

$$\begin{aligned}
{}^{dl}Y_t - {}^{dl}Y_{t-1} &= Y_t - Y_{t-1} - \frac{t}{T}(Y_T - Y_0) + \frac{t-1}{T}(Y_T - Y_0) \\
&= P_l^1(t) - \frac{1}{T}(Y_T - Y_0) \\
&= P_l^1(t) - \frac{1}{T}\sum_{n=1}^{T}(Y_n - Y_{n-1}) \\
&= P_l^1(t) - \frac{1}{T}\sum_{n=1}^{T}P_l^1(n)\,.
\end{aligned}$$

5.2 Considering Autocorrelation

In order to calculate the autocorrelation accurately no overlapping data has been used. The major result is that the autocorrelation function of returns of the considered financial time series is around zero. The hypothesis that the one–day returns are white noise can be rejected for most of the time series considered in this paper to both the 0.95 confidence interval and the 0.99 confidence interval. Considering the ten–day returns however, this is no longer true. This indicates that the distributions of the one–day returns are likely to be different from the distributions of the ten–day returns. Hence, it is not likely to find a scaling coefficient for the above distributions. However, it is still possible to calculate the scaling coefficients for certain quantiles as it will be done in the following.

5.3 Test of Self–Similarity

In order to be self–similar, the Hurst exponent has to be constant for the different quantiles. Two different tests are introduced in the sequel.

5.3.1 A First Simple Test

This first simple test tries to fit a constant for the given estimation for the Hurst coefficient on the different quantiles. The test will reject the hypothesis (that the Hurst coefficient is constant, and thus the time series is self–similar), if the goodness–of–fit is rejected. Fitting a constant to a given sample is a special case of the linear regression by setting $b = 0$. Apply the goodness–of–fit test in order to decide if the linear regression is believable and thus if the time series might be self–similar.

However, for the goodness–of–fit test it is of utmost importance, that the estimation of the Hurst coefficient for the different quantiles are independent of each other. Obviously, this is not the case for the quantile estimation, where the estimation of the Hurst coefficient is based on.

5.3.2 A Second Test

The second test tries to make a second linear regression for the given estimation for the Hurst coefficient on the different quantiles. y_i is the estimation of the Hurst coefficient for the given quantile, which will be x_i. Moreover, σ_i is the error of the Hurst coefficient estimation, while N is the number of considered quantiles.

The null hypothesis is than, that $b = 0$. The alternative hypothesis is $b \neq 0$. Thus the hypothesis will be rejected to the error level γ, if

$$\left|\frac{b}{\sigma_b}\right| > t_{N-2,1-\frac{\gamma}{2}}\,,$$

where $t_{\nu,\gamma}$ is the γ–quantile of the t_ν–distribution (see Hartung [13]).

If the hypothesis is not rejected, a is the estimation of the Hurst exponent and σ_a is the error of this estimation. Again, this test is based on the assumption, that the estimation of the Hurst coefficient for the different quantiles are independent of each other.

Both tests lead to the same phenomena which has been described by Granger and Newbold [12]. That is both tests mostly reject the hypothesis of self–similarity. And this not only for the underlying processes of the financial time series, but also for generated self–similar processes such as the Browian Motion or Levy processes.

It remains for future research to develop some test on self–similarity on the quantiles which overcome these obstacles.

6 Estimating the Scaling Law for Some Stocks

A self–similar process with Hurst exponent H can not have a drift. Since it is recognized that financial time series do have a drift, they can not be self–similar. Because of this the wording **scaling law** instead of *Hurst exponent* will be used when talking about financial time series, which have *not* been detrended.

Since the scaling law is more relevant in practice than in theory, only those figures are depicted which are based on commercial returns.

6.1 Results for Some DAX–Stocks

The underlying price processes of the DAX–stocks are the daily closing prices from January 2nd 1979 to January 13th 2000. Each time series consists of 5434 points.

Figure 3 shows the estimated scaling laws of 24 DAX–stocks in the lower quantiles. Since this figure is not that easy to analyze Figure 4 combines Figure 3 by showing the mean, the mean plus/minus the standard deviation, the minimal and the maximal estimated scaling law of the 24 DAX–stocks over the various quantiles. Moreover, the following figures will show only the quantitative characteristics of the estimation of the scaling laws for 24 DAX–stocks (see Figure 4 to 7) for the various quantiles, since it is considered more meaningful.

By doing so one has to be well aware of the fact, that the 24 financial time series are not several realisations of one stochastic process. Therefore, one has to be very careful with the interpretation of the graphics in the case of financial time series. The interpretation will be that the graphics show the overall tendencies of the financial time series. Furthermore, the mean of the estimation of the scaling law is relevant for a well diversified portfolio of these 24 stocks. The maximum of the estimation of the scaling law is the worst case possible for the considered stocks.

The estimation of the scaling law on the lower (left) quantile for 24 DAX–stocks, which is based on commercial returns, shows that the shape of the mean is curved and below 0.5 (see Figure 4). The interpretation of this is that a portfolio of these 24 DAX–stocks which is well diversified has a scaling law below 0.5. However, a poorly diversified portfolio of these 24 DAX–stocks can obey a scaling law as high as 0.55. This would imply an underestimation of 12.2% for the ten–day Value at Risk.

On the upper (right) quantile, the mean curve is sloped, ranging from a scaling law of 0.66 for the 70 percent quantile to 0.48 for the 99 percent quantile. The shape of the mean curve of the right quantile is totally different from the one of the left quantile, which might be due to a drift and/or asymmetric distribution of the underlying process (compare Figure 4 with Figure 5). In particular, the mean

Figure 3: Estimation of the Scaling Law for 24 DAX–Stocks, Left Quantile

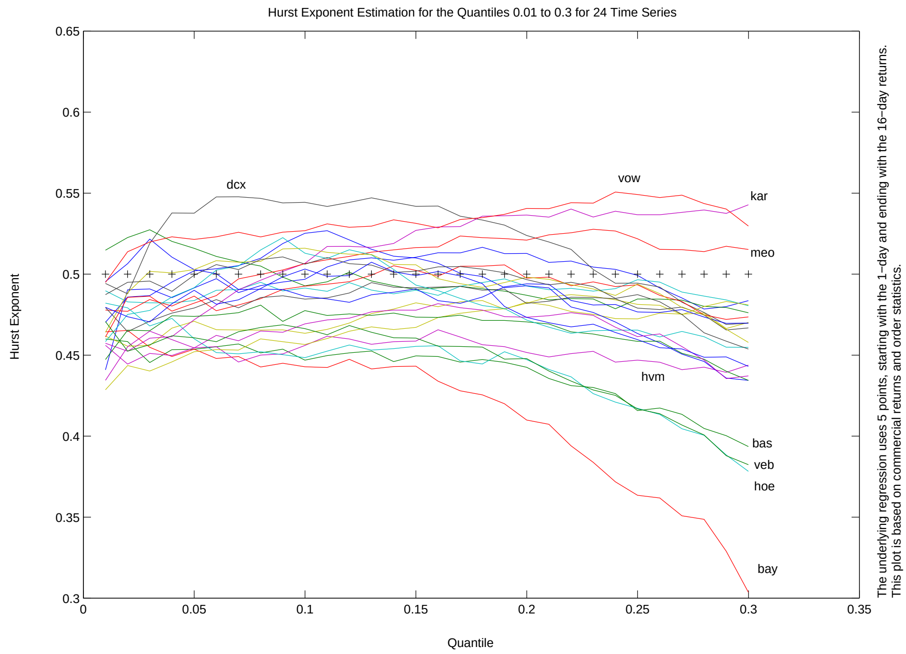

Shown are the estimation for the scaling law for 24 DAX–stocks. The underlying time series is a commercial return. Shown are the lower (left) quantiles. The following stocks are denoted explicitly: DaimlerChysler (dcx), Karstadt (kar), Volkswagen (vow), Metro (meo), Hypovereinsbank (hvm), BASF (bas), Veba (veb), Hoechst (hoe), and Bayer (bay).

is only in the 0.99–quantile slightly below 0.5. For all other upper quantiles, the mean is above 0.5.

Obviously, the right quantile is only relevant for short positions. For example, the Value at Risk to the 0.95–quantile will be underestimated by approximately 9.6% for a well diversified portfolio of short positions in these 24 DAX–stocks and can be underestimated up to 23% for some specific stocks.

The curves of the error of the estimation are also interesting (see Figure 6 and 7). First notice that the minimum of the mean of the error curves is in both cases below 0.1, which is substantially below the minimal error in the case of a normal distribution (compare with Figure 1) and in the case of a Cauchy distribution (see Figure 2).

However, the shape of the mean of the error curves of the left quantile is like the shape of the error curve of a normal distribution. Only the minimum of the mean is in the 0.3–quantile. Thus even more to the left than in the case of a normal distribution. The shape of the mean of the error curves of the right quantile looks like some combination of the error curves of the normal distribution and the Cauchy distribution.

Assuming that the shape of the error curve is closely related to the Hurst exponent of the underlying process, this would imply that the scaling law for the left quantile is less or equal 0.5 and for the right quantile between 0.5 and 1 – as it has been observed. However, the relationship between the shape of the error curve and the Hurst exponent of the underlying process has still to be verified.

The situation does not change much, if logarithmic returns are considered. The mean on the lower quantile is not as much curved as in the case of commercial

Figure 4: Estimation of the Scaling Law for 24 DAX–Stocks, Left Quantile

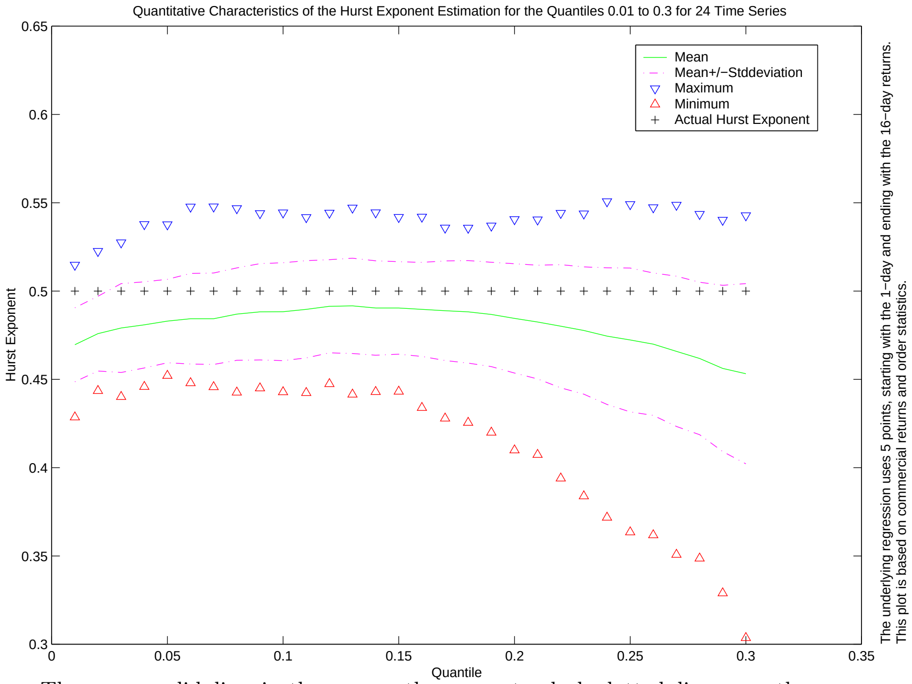

The green solid line is the mean, the magenta dash–dotted lines are the mean plus/minus the standard deviation, the red triangles are the minimum, and the blue upside down triangles are the maximum of the estimation for the scaling law, which are based on 24 DAX–stocks. The underlying time series is a commercial return. Shown are the lower (left) quantiles.

return. However, it is still curved. Moreover, in both cases (of the logarithmic and the commercial return) the mean is below 0.5 on the lower quantile. On the upper quantile, the mean curve in the case of logarithmic return is somewhat below the one of commercial return, but has otherwise the same shape. Therefore, the concluding results are the same as in the case of commercial return.

Figure 5: Estimation of the Scaling Law for 24 DAX–Stocks, Right Quantile

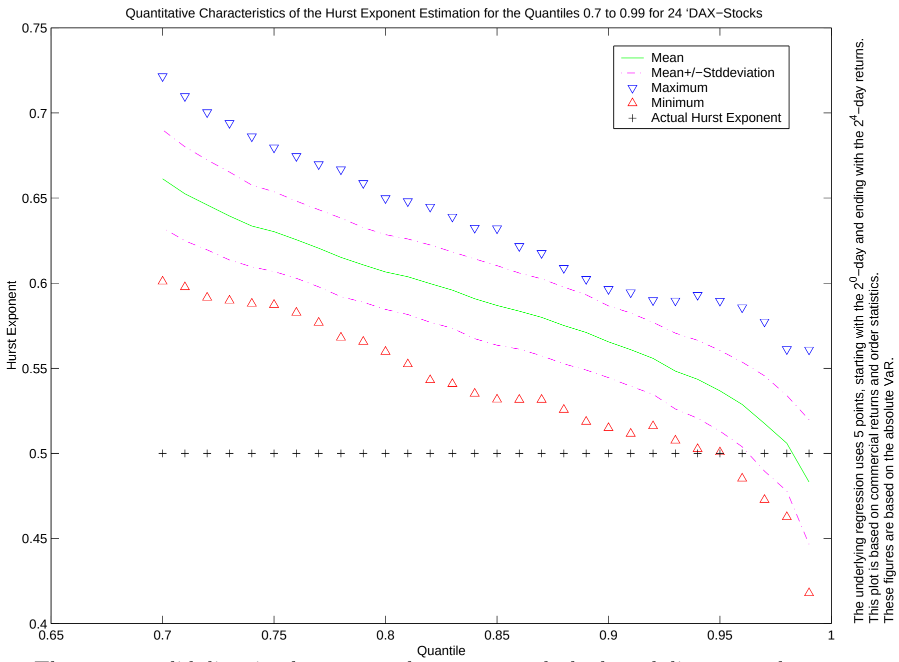

The green solid line is the mean, the magenta dash–dotted lines are the mean plus/minus the standard deviation, the red triangles are the minimum, and the blue upside down triangles are the maximum of the estimation for the scaling law, which are based on 24 DAX–stocks. The underlying time series is a commercial return. Shown are the upper (right) quantiles.

Figure 6: Error of the Estimation of the Scaling Law for 24 DAX–Stocks

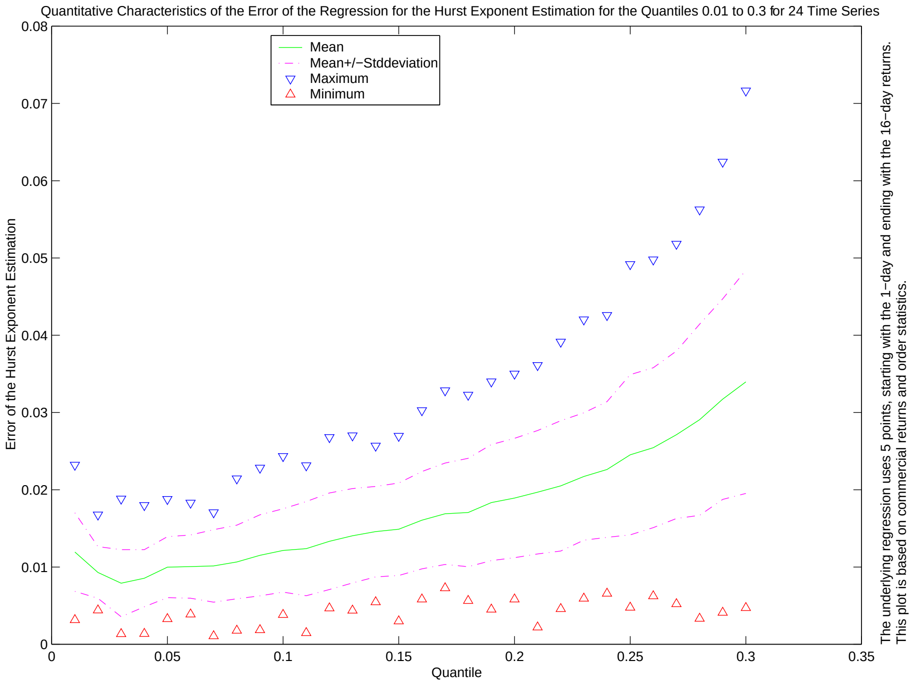

Figure 7: Error of the Estimation of the Scaling Law for 24 DAX–Stocks

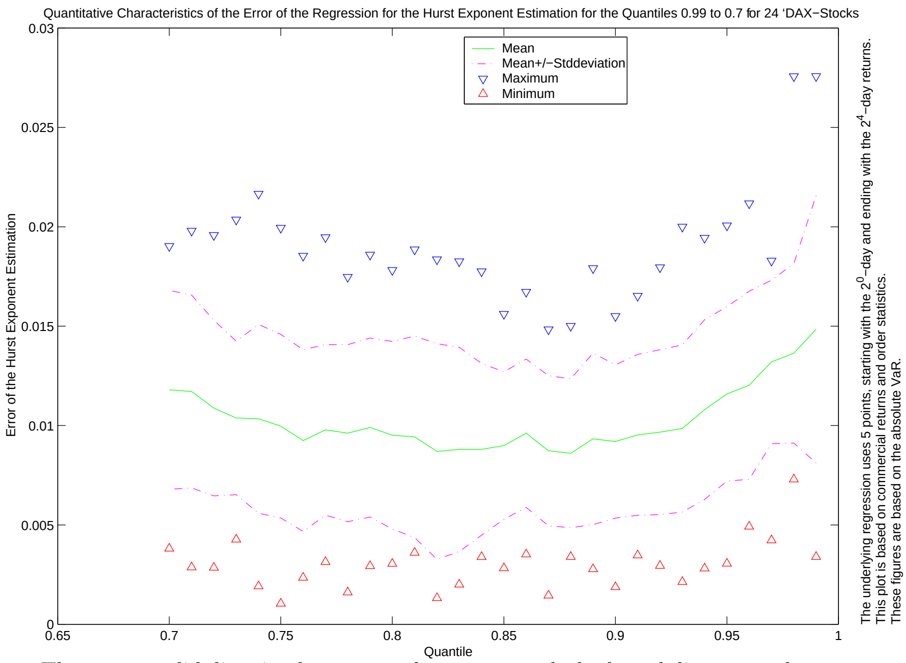

The green solid line is the mean, the magenta dash–dotted lines are the mean plus/minus the standard deviation, the red triangles are the minimum, and the blue upside down triangles are the maximum of the error curves of the estimation for the scaling law, which are based on 24 DAX–stocks. The underlying time series is a commercial return. Shown are the lower (left) quantiles (see Figure 6) and the upper (right) quantiles (see Figure 7).

6.2 Results for the Dow Jones Industrial Average Index and its Stocks

The estimation of the scaling law for the Dow Jones Industrial Average Index (DJI) and its 30 stocks is based on 2241 points of the underlying price process, which dates from March, 1st 1991 to January, 12th 2000. As in the case of the 24 DAX–stocks the underlying price processes are the closing prices.

Figure 8: Estimation of the Scaling Law for the DJI and its Stocks

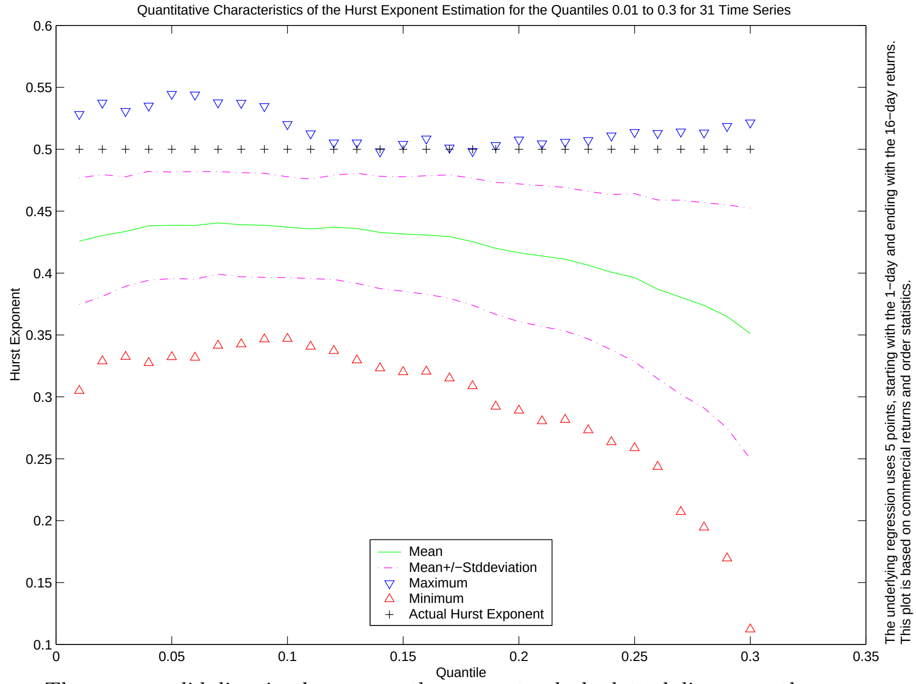

The green solid line is the mean, the magenta dash–dotted lines are the mean plus/minus the standard deviation, the red triangles are the minimum, and the blue upside down triangles are the maximum of the estimation for the scaling law, which are based on the DJI and its 30 stocks. The underlying time series is a commercial return. Shown are the lower (left) quantiles.

The results for the DJI and its 30 stocks are surprising, since the mean of the estimated scaling law is substantially lower than in the case of the 24 DAX–stocks. Moreover, for both considered cases, that is for the logarithmic return as well as for the commercial return, the mean of the estimation of the scaling law is below 0.5 and has a curvature on the lower quantile (see Figure 8).

The mean of the estimation of the scaling law for the upper quantile is sloped (as in the case of the 24 DAX–stocks) and is below 0.5 in the quantiles which are greater or equal 0.95.

On the left quantile, the shape of the mean of the error curves in the case of commercial and logarithmic returns is comparable with the shape of the mean of the error curve for the DAX–stocks and therefore comparable with the shape of the mean error curve of the normal distribution. The shape of the mean of the error curves on the right quantile is again like a combination of the error curve of the normal distribution and the error curve of the Cauchy distribution. However, the level of the mean of the error curves are of the same height as the level of the error curve of the normal distribution and thus substantially higher than the mean of the error curves of the corresponding 24 DAX–stocks.

Figure 9: Estimation of the Scaling Law for the DJI and its Stocks

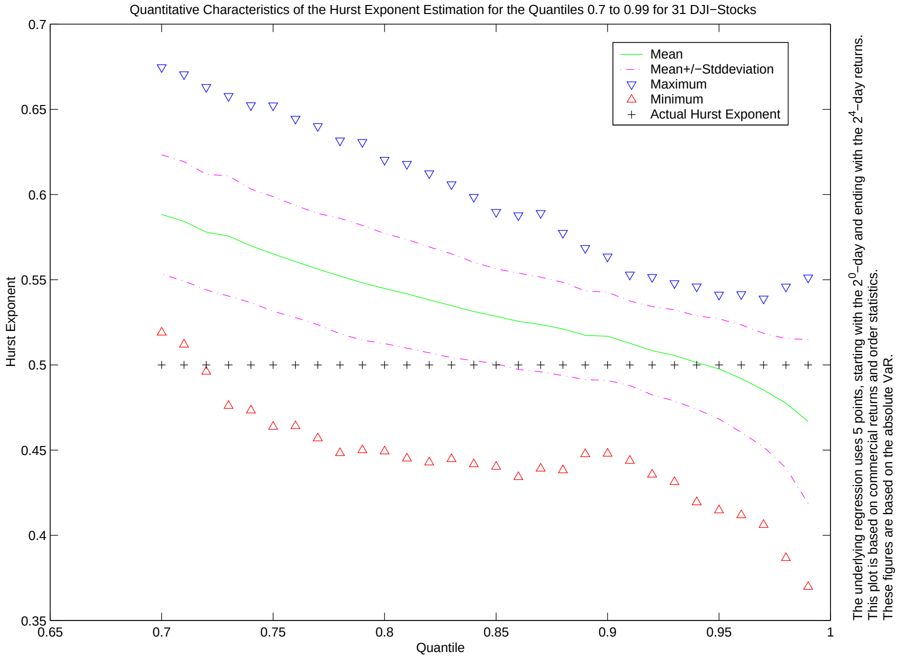

The green solid line is the mean, the magenta dash–dotted lines are the mean plus/minus the standard deviation, the red triangles are the minimum, and the blue upside down triangles are the maximum of the estimation for the scaling law, which are based on the DJI and its 30 stocks. The underlying time series is a commercial return. Shown are the upper (right) quantiles.

7 Determining the Hurst Exponent for Some Stocks

It has been already stated, that the financial time series can not be self–similar. However, it is possible that the detrended financial time series are self–similar with Hurst exponent H. This will be scrutinized in the following where the financial time series have been detrended according to the method described in section 5.1.

Since the Hurst exponent is more relevant in theory than in practice, only those figures are shown which are based on logarithmic returns.

7.1 Results for Some DAX–Stocks

The results are shown in Figure 10 and 11. The mean of the estimation of the Hurst exponent on the left quantile is considerably higher for the detrended time series than for the time series which have not been detrended. However, the mean shows for both, the logarithmic as well as the commercial return, a slope (see Figure 10). On the right quantile, the mean curve has for the detrended time series the same shape as in the case of the non–detrended time series for both, the commercial return and the logarithmic return. The slope is for the detrended time series is not as big as for the non–detrended time series and the mean curve of the detrended time series lies below the mean curve of the corresponding non–detrended time series.

The shape of the mean curve of the upper quantiles is comparable to the one of the lower quantiles. This is valid for both the commercial and the logarithmic return. However, for example in the case of commercial return, the slope is much stronger (the mean Hurst exponent starts at about 0.6 for the 70 percent quan-

Figure 10: Hurst Exponent Estimation for 24 DAX–Stocks, Left Quantile

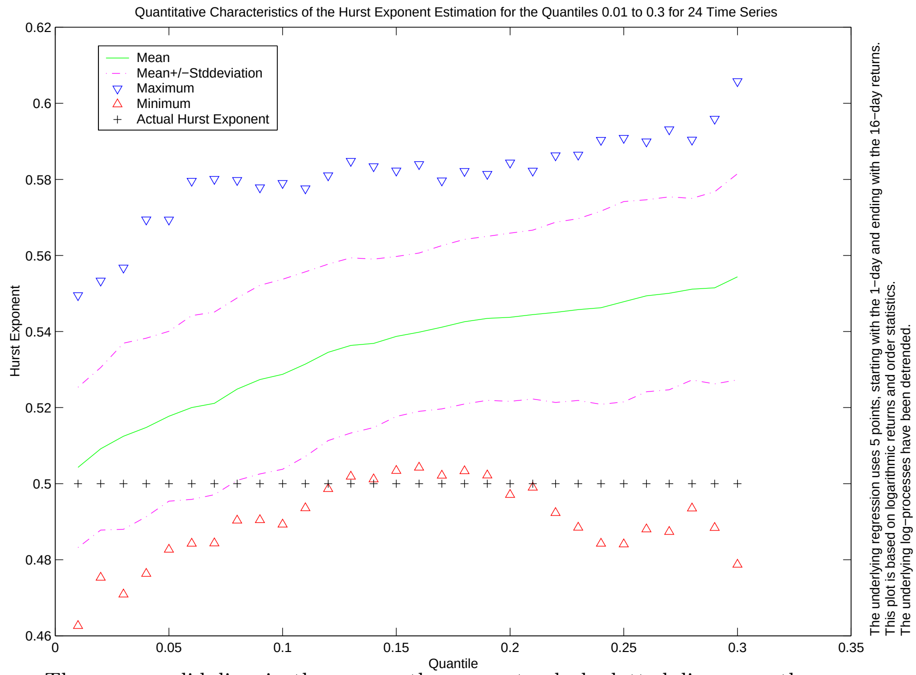

The green solid line is the mean, the magenta dash–dotted lines are the mean plus/minus the standard deviation, the red triangles are the minimum, and the blue upside down triangles are the maximum of the estimation for the Hurst exponent, which are based on 24 DAX–stocks. The underlying time series are logarithmic returns, which have been detrended. Shown are the lower (left) quantiles.

tile and ends at about 0.47 for the 99 percent quantile compared to 0.55 for the 30 percent quantile and 0.48 for the 1 percent quantile). This indicates that the distribution of the underlying process might not be symmetric. Moreover, the wide spread of the mean Hurst exponent over the quantiles indicates, that the detrended time series are not self–similar as well.

The mean of the error curves is not much affected by the detrending. The shape of the mean of the error curves on the left quantiles are in both cases similar to the shape of the theoretical error of the normal distribution. The shape of the error curves on the upper quantiles is totally different from the ones on the lower quantiles and looks like a combination of the error curve of a normal distribution and a Cauchy distribution.

Altogether, these results indicate, that the considered financial time series are not self–similar. However, the tests on self–similarity introduced in section 5.3 are not sensitive enough to verify these findings. In order to verify these results, it is necessary to develop a test for self–similarity which is sensitive enough.

7.2 Results for the Dow Jones Industrial Average Index and Its Stocks

The mean of the estimation of the Hurst exponent on the left quantile is considerably higher for the detrended time series than for the time series which have not been detrended. However, the mean shows for both, the logarithmic as well as the commercial return, a curvature on the lowest quantiles (see Figure 12). The shape of the mean curve of the detrended time series is similar to the mean curve of the

Figure 11: Hurst Exponent Estimation for 24 DAX–Stocks, Right Quantile

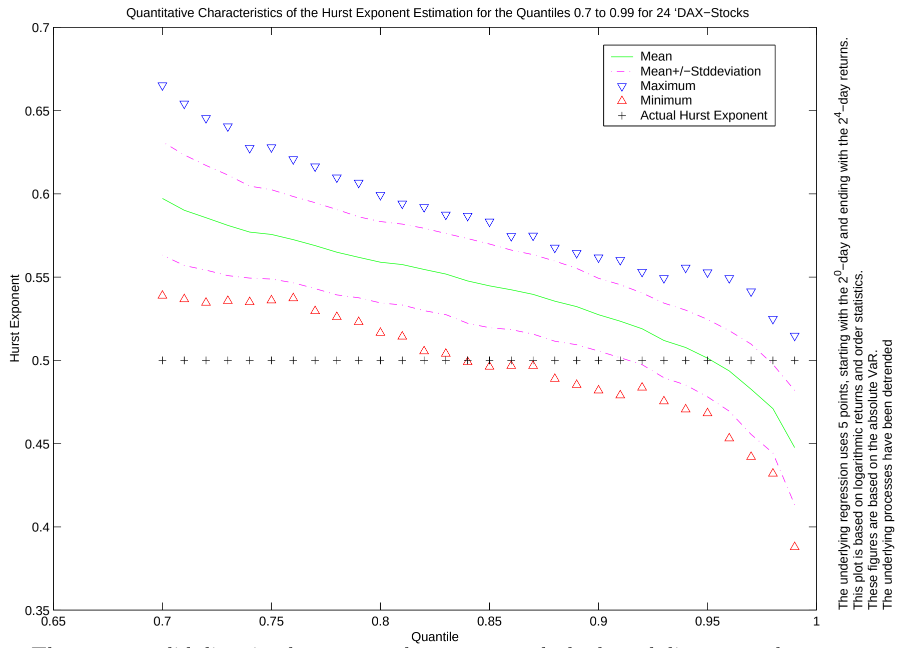

The green solid line is the mean, the magenta dash–dotted lines are the mean plus/minus the standard deviation, the red triangles are the minimum, and the blue upside down triangles are the maximum of the estimation for the Hurst exponent, which are based on 24 DAX–stocks. The underlying time series are logarithmic returns, which have been detrended. Shown are the upper (right) quantiles.

corresponding non–detrended time series.

On the right quantile, the mean curve has for the detrended time series the same shape as in the case of the non–detrended time series for both, the commercial return and the logarithmic return. The slope is for the detrended time series is not as big as for the non–detrended time series and the mean curve of the detrended time series lies below the mean curve of the corresponding non–detrended time series.

The shape of the mean curve of the upper quantiles is comparable to the one of the lower quantiles only on the outer quantiles. This is valid for both the commercial and the logarithmic return. Moreover, for example in the case of commercial return, the mean Hurst exponent ranges from about 0.52 for the 70 percent quantile to about 0.45 for the 99 percent quantile compared to the range of 0.48 for the 30 percent quantile and 0.44 for the 1 percent quantile. This indicates that the distribution of the underlying process might not be symmetric. Moreover, the spread of the mean Hurst exponent over the quantiles might indicate, that the detrended time series are not self–similar as well.

The mean of the error curves is not much affected by the detrending for the right quantiles. However, the mean of the error curves on the left quantiles are in both cases about constant up to the lowest quantiles where the curves go up. The shape of the error curves on the upper quantiles are not as constant as the ones on the lower quantiles and look like some combination of the error curve of a normal distribution and a Cauchy distribution.

Altogether, these results indicate, that the considered financial time series are not self–similar. However, the tests on self–similarity introduced in section 5.3 are not sensitive enough to verify these findings as it has already been stated.

Figure 12: Hurst Exponent Estimation for the DJI and its Stocks

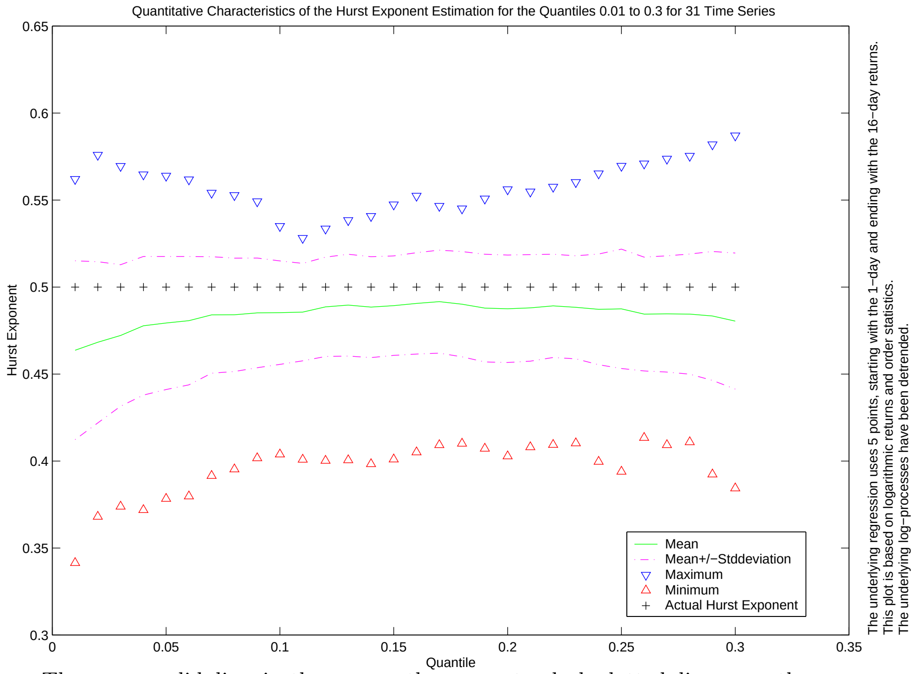

The green solid line is the mean, the magenta dash–dotted lines are the mean plus/minus the standard deviation, the red triangles are the minimum, and the blue upside down triangles are the maximum of the estimation for the Hurst exponent, which are based on the DJI and its 30 stocks. The underlying time series are logarithmic returns, which have been detrended. Shown are the lower (left) quantiles.

8 Interpretation of the Hurst Exponent for Financial Time Series

First, let us recall the meaning of the Hurst exponent for different stochastic processes. For example, for a fractional Brownian Motion with Hurst coefficient H, the Hurst exponent describes the persistence or anti–persistence of the process (see for example Shiryaev [17]). For $1 > H > \frac{1}{2}$ the fractional Brownian Motion is persistent. This means that the increments are positively correlated. For example, if an increment is positive, it is more likely that the succeeding increment is also positive than that it is negative. The higher H is, the more likely is that the successor has the same sign as the preceding increment. For $\frac{1}{2} > H > 0$ the fractional Brownian Motion is anti–persistent, meaning that it is more likely that the successor has a different sign than the preceding increment. The case $H = \frac{1}{2}$ is the Brownian Motion, which is neither persistent nor anti–persistent (see Shiryaev [17]).

This is, however, not true for Levy processes with $H > \frac{1}{2}$, where the increments are independent of each other. Therefore, the Levy processes are as the Brownian Motion neither persistent nor anti–persistent. In the case of Levy processes, the Hurst exponent H tells, how much the process is heavy tailed.

Considering financial time series, the situation is not at all that clear. On the one hand, the financial time series are neither fractional Brownian Motions nor Levy processes. On the other hand, the financial time series show signs of persistence and heavy tails.

Assuming the financial time series are fractional Brownian Motions, then a Hurst

Figure 13: Hurst Exponent Estimation for the DJI and its Stocks

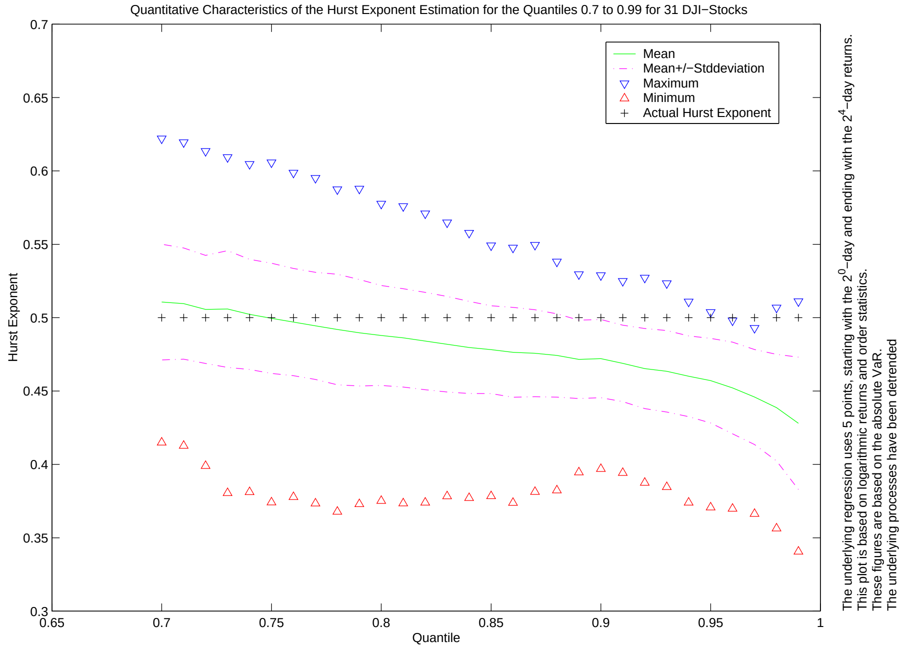

The green solid line is the mean, the magenta dash–dotted lines are the mean plus/minus the standard deviation, the red triangles are the minimum, and the blue upside down triangles are the maximum of the estimation for the Hurst exponent, which are based on the DJI and its 30 stocks. The underlying time series are logarithmic returns, which have been detrended. Shown are the upper (right) quantiles.

exponent $H > \frac{1}{2}$ would mean that the time series are persistent. The interpretation of the persistence could be, that the financial markets are either rather slow to incorporate the actual given information or this could indicate, that insider trading is going on in the market. The first case would be a contradiction of the efficient market hypothesis, while the second case would be interesting for the controlling institutions as the SEC and the BAFin (the German analog of the SEC). A Hurst exponent of $H < \frac{1}{2}$ would mean that the financial market is constantly overreacting.

For this interpretation compare also the findings for the DJI–stocks with the results for the 24 DAX–stocks. The average estimation of the Hurst coefficient of the 24 detrended DAX–stocks is substantially higher than the one of the detrended DJI–stocks. Thus this interpretation would support the general believe, that the US–american financial market is one of the most efficient market of the world, while the german market is not that efficient, which is often cited as the "Deutschland AG"–phenomena.

Given that this interpretation is right, one could check whether a market (or an asset) has become more efficient. If its corresponding Hurst exponent gets closer to 0.5 over the time, then the market (or the asset) is getting more efficient.

Assuming, that the Hurst coefficient of financial time series reflects persistence, the results of the detrended financial time series can be interpreted in the following way. While the financial market is in normal market situations rather slow in incorporating the actual news, it tends to overreact in extreme market situation.

Not much can be said, if one assumes that the Hurst coefficient of financial time series reflects a heavy tail property. It can not be verified, that large market movements occur more often in the german financial market than in the US–american financial market. However, in both financial markets do big market movement occur much more often than in the case of a Brownian Motion. Therefore the financial time series are heavy tailed.

9 Conclusion and Outlook

The main results are that

- the scaling coefficient 0.5 has to be used very carefully for financial time series and
- there are substantial doubts about the self–similarity of the underlying processes of financial time series.

Concerning the scaling law, it is better to use a scaling law of 0.55 for the left quantile and a scaling law of 0.6 for the right quantile (the short positions) in order to be on the safe side. It is important to keep in mind that these figures are only based on the (highly traded) DAX– and Dow Jones Index–stocks. Considering low traded stocks might yield even higher maximal scaling laws. These numbers should be set by the market supervision institutions like the SEC.

However, it is possible for banks to reduce their Value at Risk figures if they use the correct scaling law numbers. For instance, the Value at Risk figure of a well diversified portfolio of Dow Jones Index–stocks would be reduced in this way about 12%, since it would have a scaling law of approximately 0.44.

Regarding the self–similarity, estimating the Hurst exponent via the quantiles might be a good alternative to modified R/S–statistics, Q–Q–plots and calculating the Hurst exponents via the moments. However, it remains to future research to develop a test on self–similarity on the quantiles which overcomes the phenomena already described by Granger and Newbold [12].

Finally, Daníelsson and Zigrand [5] mentioned that the square–root–of–time–rule is also used for calculating volatilities. The presented results indicate that the appropriate scaling law exponents for volatilities is most likely higher than the estimated scaling law exponents for the quantiles. However, giving specific estimates for this situation is left for future research.

10 Acknowledgment

I like to thank Jean–Pierre Stockis for some very valuable discussions and hints. The discussions with Peter Singer and Ralf Hendrych have always been very inspiring. Finally, I appreciated the comments and advices of Prof. Philippe Jorion, Prof. Klaus Schürger, and an anonymous referee.

References

[1] Philippe Artzner, Freddy Delbaen, Jean-Marc Eber, and David Heath. Coherent measure of risk. *Mathematical Finance*, 9(3):203 – 228, 1999.

[2] Basel Committee on Banking Supervision. Overview of the amendment to the capital accord to incorporate market risk. Basel Committee Publications No. 23, Bank for International Settlements, CH–4002 Basel, Switzerland, January 1996. See also `http://www.bis.org/bcbs/publ.htm`.

[3] Jón Daníelsson and Casper G. de Vries. Value–at–risk and extreme returns. *Annales d'Economie et de Statistique*, 60(Special Issue):236 – 269, 2000. See also `http://www.riskresearch.org/`.

[4] Jón Daníelsson, Bjørn N. Jorgensen, Gennady Samorodnitsky, Mandira Sarma, and Casper G. de Vries. Subadditivity re–examined: the case for value–at–risk. Working Paper, see `http://www.riskresearch.org/`, November 2005.

[5] Jón Daníelsson and Jean-Pierre Zigrand. On time–scaling of risk and the square–root–of–time rule. FMG Discussion Paper – Financial Markets Group dp439, London School of Economics and Political Science, London, United Kingdom, March 2003. See also `http://www.riskresearch.org/`.

[6] Francis Diebold, Andrew Hickman, Atsushi Inoue, and Til Schuermann. Converting 1–day volatility to h–day volatility: Scaling by $\sqrt{h}$ is worse than you think. Working Paper – Financial Institutions Center 97–34, The Wharton School – University of Pennsylvania, Philadelphia, USA, July 1997. See `http://fic.wharton.upenn.edu/fic/papers/97/9734.pdf`.

[7] Francis Diebold, Atsushi Inoue, Andrew Hickman, and Til Schuermann. Scale models. *Risk Magazine*, pages 104–107, January 1998.

[8] Kevin Dowd, David Blake, and Andrew Cairns. Long–term value at risk. Discussion paper pi–0006, The Pensions Institute.

[9] Paul Embrechts, Claudia Klüppelberg, and Thomas Mikosch. *Modelling Extremal Events*, volume 33 of *Applications of Mathematics*. Springer, Berlin, third edition, 1997. ISBN 3-540-60931-8.

[10] Paul Embrechts and Makoto Maejima. An introduction to the theory of self–similar stochastic processes. *International Journal of Modern Physics B*, 14(12 & 13):1399 – 1420, 2000.

[11] Carl Evertsz, Ralf Hendrych, Peter Singer, and Heinz-Otto Peitgen. *Komplexe Systeme und Nichtlineare Dynamik*, chapter Zur fraktalen Geometrie von Börsenzeitreihen, pages 400–419. Springer, Berlin, 1999.

[12] Clive W. J. Granger and Paul Newbold. Spurious regressions in econometrics. *Journal of Econometrics*, 2:111 – 120, 1974.

[13] Joachim Hartung, Bärbel Elpelt, and Karl-Heinz Klösener. *Statistik*. Oldenbourg, München, 10 edition, 1995.

[14] Philippe Jorion. *Value at Risk: The New Benchmark for Controlling Market Risk*. McGraw–Hill, New York, 1997.

[15] Rolf-Dieter Reiss. *Approximate Distributions of Order Statistics*. Springer, New York, 1989.

[16] Gennady Samorodnitsky and Murad S. Taqqu. *Stable Non–Gaussian Random Processes*. Chapman & Hall, New York, 1994.

[17] Albert N. Shiryaev. *Essentials of Stochastic Finance*. World Scientific, Singapore, 1999.

[18] Peter Singer, Ralf Hendrych, Carl Evertsz, and Heinz-Otto Peitgen. Semi–stable processes in finance and option pricing. Preprint, 2000.

[19] Alan Stuart and J. Keith Ord. Distribution theory. In *Kendall's Advanced Theory of Statistics*, volume 1. Arnold, London, sixth edition, 1994.

[20] Paul Wilmott. *Derivatives: The Theory and Practise of Financial Engineering*. John Wiley & Sons, Inc., Chichester, 1998.