Long Memory vs. Structural Change in Financial Time Series[1]

by

Walter Krämer, Philipp Sibbertsen and Christian Kleiber
Fachbereich Statistik, Universität Dortmund, D-44221 Dortmund, Germany

Version July 2001

Abstract

The paper discusses structural change as possible mechanism that generates the appearance of long memory in economic time series. It shows that there are no long memory effects in German stock returns and that long memory in squares of German stock returns disappears once shifting means are properly accounted for.

1 Introduction and Summary

It is a well known stylized fact that many financial time series such as squares or absolute values of returns, even returns themselves behave as if they had long memory (Ding et al. 1993, Baillie et al. 1996, Lobato and Savin 1998 and many others). On the other hand, it is also well known that long memory is easily confused with structural change, in the sense that the slow decay of empirical autocorrelations which is typical for a time series with long memory is also produced when a short–memory time series exhibits structural breaks

[1]Research supported by Deutsche Forschungsgemeinschaft under SFB 475; we are grateful to Christoph Helwig, Tamara Könning and Klaus Nordhausen for expert computational assistance. Stock returns were obtained from Deutsche Finanzdatenbank (DFDB), Karlsruhe.

(Boes and Salas 1978, Diebold and Inoue 1999, Granger and Hyung 1999, Gourieroux and Jasiak 2001). Therefore it is of considerable theoretical and empirical interest to discriminate between these sources of slowly decaying empirical autocorrelations.

This is done below for German stock returns. Using daily data for various individual stocks and the German stock–price index DAX, we show that both the returns themselves and their squares are probably best modelled as short memory processes disturbed by breaking trends.

2 Long Memory as an Artefact of Structural Change

Let $\{x_t\}$ be the time series under investigation. Following Heyde and Yang (1997), we say that $\{x_t\}$ has "long memory" if

$$\frac{(x_1 + \ldots + x_n)^2}{x_1^2 + \ldots x_n^2} \xrightarrow{p} \infty \quad \text{as} \quad n \to \infty. \tag{1}$$

For a second order stationary process, this definition encompasses the more popular conditions

$$\sum_{h=0}^{\infty} \gamma(h) = \infty \quad \text{and} \tag{2}$$

$$\lim_{\lambda \to 0} f(\lambda) = \infty \tag{3}$$

as special cases, where $\gamma(h) = cov(x_t, x_{t+h})$ and $f(\lambda)$ is the spectral density of $\{x_t\}$. Definition (1) is slightly more general; in particular, it also applies to cases where second moments do not exist.

For statistical inference, it is helpful to further specify the rates of divergence in (2) and (3) such that

$$f(\lambda) \sim \frac{1}{\lambda^{2d}} \quad (d > 0, \lambda \to 0), \tag{4}$$

which is equivalent to

$$\gamma(h) \sim h^{2d-1} , \tag{5}$$

where d is the fractional differencing parameter in the popular long memory model (see e.g. Beran 1994, pp. 59 ff)

$$(1-B)^d\{x_t\} = \{y_t\}, \; y_t \sim ARMA(p,q). \tag{6}$$

From (2), it is obvious that an appearance of long memory in the data is easily produced by any mechanism that makes for a slow decay of empirical autocorrelations. As an illustration, figure 1 shows an $MA(2)$-process

$$x_t = \varepsilon_t + \frac{1}{2}\varepsilon_{t-1} + \frac{1}{3}\varepsilon_{t-2} \quad (t = 1, \ldots, 1000),$$

superimposed by a deterministic series

$$d_t = \begin{cases} 1 & t \leq 500 \\ -1 & t > 500 \end{cases} .$$

The MA process is stationary by construction, with $\gamma(h) = 0$ for $h \geq 3$. However, for anyone familiar with plots of long-memory time series the series very much looks like exhibiting long memory.

Figure 1

Figure 2 shows empirical autocorrelations computed from this series up to order $h = 50$. There is no tendency to tend to zero. In fact, it is easily seen that for second-order stationary processes superimposed by a deterministic series,

$$d_t = \begin{cases} \Delta & t \leq \frac{n}{2} \\ -\Delta & t > \frac{n}{2} \end{cases} . \tag{7}$$

we have

$$\hat{\rho}(h) := \frac{\sum_{i=1}^{n-h}(x_t - \bar{x})(x_{t+h} - \bar{x})}{\sum_{i=1}^{n}(x_t - \bar{x})^2} \xrightarrow{p} \frac{\Delta^2 + \gamma(h)}{\Delta^2 + \gamma(0)} \ . \tag{8}$$

In our example above, we have $\Delta = 1, \sigma_\varepsilon^2 = 4, \gamma(0) = 5.44$, so $\hat{\gamma}(0)/\hat{\gamma}(h) \xrightarrow{p} 0.155$ for all $h \geq 3$.

Figure 2

Figure 3 shows $\hat{\gamma}(h)/\hat{\gamma}(0)$ ($h = 1, \ldots, 50$) for sample sizes $n = 4000$, $n = 16000$ and $n = 32000$. It is seen that the autocorrelation function becomes flat very fast, as predicted by (8).

Figure 3

The appearance of long memory becomes less obvious when the importance of the time-varying deterministic component diminishes as sample size increases. Giraitis et al. (2000a) show that the Mandelbrot/Wallis Rescaled-Range test indicates long memory whenever the trend d_t decays with a rate slower than $\frac{1}{\sqrt{n}}$. For d_t decaying faster, the R/S statistic and related formal tests have the same limiting distributions as under short memory. In the general case the trend is given by

$$d_t^{(n)} = \begin{cases} k_1^{(n)} & \text{if} \quad 1 \leq t \leq t^* \\ k_2^{(n)} & \text{if} \quad t^* < t \leq n \end{cases} , \tag{9}$$

where $t^* = [\tau^* n]$ and $0 < \tau^* < 1$. If $\Delta_n := k_1^{(n)} - k_2^{(n)} \neq 0$ is of order $\Delta_n = \delta_n n^{-\frac{1}{2}}$, $\delta_n \to 0$, the limiting behaviour of the test is the same as under H_0. Otherwise $\Delta_n = \delta_n n^{-\frac{1}{2}}, \delta_n \to \infty$ leads to a rejection of the short-memory hypothesis.

If $\Delta_n = \Delta \neq 0$, does not depend on sample size, the R/S-statistic converges to the same constant independent of whether the disturbance has short or long-term memory.

Bhattacharya et al. (1983) show that a monotonic trend of the form

$$d_t = c(m+t)^{\beta}\ , -\frac{1}{2} < \beta < 0\ , \tag{10}$$

can also produce the Hurst effect. Giraitis et al. (2000a) generalize this result by specifying the rate of decay of the trend necessary for confusing the trend and long-range dependence.

Yet another type of structural change is considered by, among others, Boes and Salas (1978) or Gourieroux and Jasiak (2001), who investigate a series generated by

$$x_t = m_t + y_t, \tag{11}$$

where y_t is stationary short memory and where m_t remains constant over long stretches of time, i.e.

$$m_t = \begin{cases} k_1 & t \leq \tau_1 \\ k_2 & \tau_1 < t \leq \tau_1 + \tau_2 \\ k_3 & \tau_1 + \tau_2 < t \leq \tau_1 + \tau_2 + \tau_3 \\ \vdots & \vdots \end{cases}\ , \tag{12}$$

where $\{k_i\}$ and $\{\tau_i\}$ are i.i.d., τ_i integer-valued and positive. In models like these, whether or not $\{x_t\}$ exhibits long memory crucially depends on the tail probabilities of τ.

As an illustration, consider the case where k takes values $+m$ and $-m$, each with probability $\frac{1}{2}$, so that $E(k) = 0$ and $Var(k) = \sigma_k^2$. Assume further that $\tau_i \sim$ i.i.d. discrete $Pareto(\alpha)$, i.e. $P(\tau > c) \sim c^{-\alpha}$, for some $\alpha > 0$, and, for simplicity, that $y_t \sim$ i.i.d. $(0, \sigma^2)$ and that the sequences $\{\tau_i\}$, $\{k_i\}$, and $\{y_t\}$ are independent. Put $S_j = \sum_{i=1}^{j} \tau_i$, $S_0 = 0$, so that $\{N(t)\}$, $N(t) = \sup\{j : S_j \leq t\}$, is a renewal process in discrete time. Now m_t may be written as

$$m_t = \sum_{j=1}^{t} k_j \mathbb{1}_{(S_{j-1},S_j]}(t) \; . \tag{13}$$

By assumption $E(k_t) = 0 = E(y_t)$, so that

$$cov(x_t, x_{t+h}) = cov(m_t + y_t, m_{t+h} + y_{t+h}) = E(m_t m_{t+h}) \; . \tag{14}$$

Now

$$\begin{aligned} E(m_t m_{t+h}) &= E\left(\sum_{i=1}^{t} k_i \mathbb{1}_{(S_{i-1},S_i]}(t) \sum_{j=1}^{t+h} k_j \mathbb{1}_{(S_{j-1},S_j]}(t+h)\right) \\ &= \sum_{i=1}^{t}\sum_{j=1}^{t+h} E\left(k_i k_j \mathbb{1}_{(S_{i-1},S_i]}(t) \mathbb{1}_{(S_{j-1},S_j]}(t+h)\right) \\ &= \sum_{i=1}^{t}\sum_{j=1}^{t+h} E\left(E[k_i k_i \mathbb{1}_{(S_{j-1},S_j]}(t) \mathbb{1}_{(S_{j-1},S_j]}(t+h)] | S_1, S_2, \ldots\right) \\ &= \sigma_k^2 \cdot \sum_{i=1}^{t}\sum_{j=1}^{t+h} E\left(E[\mathbb{1}_{(S_{j-1},S_j]}(t) \mathbb{1}_{(S_{j-1},S_j]}(t+h)] | S_1, S_2, \ldots\right) \; . \end{aligned} \tag{15}$$

Observing that the sum equals the probability of no structural change in the interval $[t, t+h)$ one obtains

$$\begin{aligned} &\sigma_k^2 \cdot \sum_{j=0}^{t} P(S_j < t, S_{j+1} \geq t+h) \\ &= \sigma_k^2 \cdot \sum_{j=0}^{t} P(\tau_{j+1} > h) \\ &\sim \sigma_k^2 \cdot t \cdot h^{-\alpha} \; . \end{aligned} \tag{16}$$

In the case $t = 1$ this simplifies to $\sigma_k^2 \cdot h^{-\alpha}$, so finally

$$cov(x_1, x_{1+h}) \sim \sigma_k^2 h^{-\alpha} \; , \tag{17}$$

which should be compared to (5). Thus slowly decaying autocovariances are possible in the above framework if the waiting time distribution has sufficiently heavy tails, more specifically if $\alpha < 1$, i.e. $E(\tau) = \infty$.

Clearly, the derivation and results remain essentially unchanged if k does not have a two-point distribution and $\{y_t\}$ is a stationary short-memory process.

It is of interest to compare the above shifting-levels process to the Markov-switching or hidden Markov model introduced by Lindgren (1978). In a Markov-switching model with stationary transition probabilities sojourn times are geometrically distributed and long memory cannot occur (see Rydén et al., 1998). However, in a hidden Markov model with nonstationary transition probabilities one can likewise obtain the appearance of long memory (Diebold and Inoue 1999).

3 Long Memory in Stock Returns

In an efficient market, excess stock returns (i.e. actual returns minus required (= expected) returns) form a martingale difference sequence, so they are uncorrelated and cannot have any memory at all. However, Greene and Fielitz (1977), using the R/S statistic, find significant long memory in return series of 200 common stocks listed on the New York Stock Exchange. Peters (1992), Goetzmann (1993), Mills (1993) and Crato (1994) confirm this finding for other stocks and for indices like the S & P 500. Lo (1991) and Ambrose et al. (1993) dispute the significance of these results and propose an alternative test statistic based on an estimator of the return variance that is robust to short-term autocorrelation. The significance of the observed Hurst effect then tends to disappear. However, Willinger et al. (1999) reopen the debate by showing that Lo's approach is open to criticism as well, and find that long memory might still occur even after adjusting for possible short term autocorrelation.

Below we apply the methods used above to German data. We also use the recent Giraitis et al. V/S statistic, which has not been applied to stock returns before. Our approach is similar in spirit to Aydogan and Booth (1988) and Chow et al. (1995), who purge the raw returns from effects which might induce spurious long memory. Aydogan and Booth (1988) use the market model to decompose the total return in a systematic and a nonsystematic element and show that there is less evidence for the Hurst effect in the nonsystematic element. Chow et al. (1995) control for calendar effects, which are one of many possible types of "trend", and show that the significance of the observed long memory effects is thereby reduced.

Below we follow Mills (1993) by decomposing a stock return $r_t := \ell n(P_t/P_{t-1})$, where P_t is price (adjusted for dividends, stock splits and so on) into an expected and an unexpected part:

$$r_t = E_{t-1}(r_t) + r_t - E_{t-1}(r_t) = \bar{r}_t + r_t^* \ , \tag{18}$$

where we approximate $\bar{r}_t$ by the arithmetic mean.

It is only r_t^* which is by theory required to follow a martingale difference sequence.

Below, we approximate $\bar{r}_t$ by the six-month money rate, plus a risk premium of 3 percentage points (on an annual basis). Our data are daily returns of various individual German stocks and the German stock price index DAX from January 4, 1960 up to April 30, 1998.

For daily returns, $\bar{r}_t$ is close to zero, with little variation across time, which might explain why it has been neglected in previous research. The hypothesis we want to test is that it can make a slight difference if the time series is long enough.

The first test we use is the Mandelbrot/Wallis R/S test defined by

$$Q_n := \frac{R_n}{S_n} \ , \tag{19}$$

where

$$R_n = \max_{1\leq k\leq n}[\sum_{i=1}^{k}(x_i - \bar{x}_n)] - \min_{1\leq k\leq n}[\sum_{i=1}^{k}(x_i - \bar{x}_n)] \tag{20}$$

is the "range" and

$$S_n := \sqrt{\frac{1}{n}\sum_{i=1}^{n}(x_i - \bar{x}_n)^2} \tag{21}$$

is the sample standard deviation. A plot of $\ell n\ Q_k$ against $\ell n\ k$ scatters around a straight line with slope $\frac{1}{2}$ for independent or short-memory processes and it scatters around a straight line with slope $H \in (\frac{1}{2}, 1)$ in the long-memory case. As an illustration, figure 4 shows $\ell n(Q_k)$ and ℓnk for the DAX returns. It also plots the OLS regression line which has a slope slightly larger than 0.5.

Figure 4

Lo (1991) modifies the R/S statistic by allowing for short-term autocorrelation under H_0. The modified scale is defined by

$$S^2_{n,q} = \frac{1}{n}\sum_{i=1}^{n}(x_i - \bar{x}_n)^2 + 2\sum_{j=1}^{q}\omega_j(q)\hat{\gamma}_j \ , \tag{22}$$

where

$$\omega_j(q) = 1 - \frac{j}{q+1} \quad \text{and} \quad \hat{\gamma}_j = \frac{1}{n}\sum_{i=1}^{n-j}(x_i - \bar{x}_n)(x_{i+j} - \bar{x}_n),\ 0 \leq j \leq n. \tag{23}$$

The standard R/S statistic is obtained for $q = 0$.

To obtain better power properties Giraitis et al. (2000b) introduce the related V/S statistic. Defining

$$S^*_\ell := \sum_{j=1}^{\ell}(x_i - \bar{x}_n) \quad \text{and} \quad \hat{Var}(S^*_1, \ldots, S^*_n) = \frac{1}{n}\sum_{i=1}^{n}(S^*_i - S^*_n)^2 \ , \tag{24}$$

the V/S statistic has the form

$$M_n := n^{-1}\frac{\hat{Var}(S^*_1, \ldots, S^*_n)}{S^2_n} \ , \tag{25}$$

where S_n is defined in (21).

We also apply log-periodogram regression estimates introduced by Geweke/Porter-Hudak (1983). This estimator uses the special shape of the spectral density of a long-memory process at the origin. The spectral density of a long-memory process $\{x_t\}$ is of the form

$$f(\lambda) \sim c\lambda^{-2d} . \tag{26}$$

Taking logs and adding the log-periodogram calculated at the Fourier frequencies $\lambda_{k,n} = \frac{2\pi k}{n}, k = 1, \ldots, [n^{\frac{4}{5}}]$ on both sides gives

$$\log I(\lambda_{k,n}) \approx \log c - 2d \log \lambda_{k,n} + \log \xi_k \ , \tag{27}$$

where ξ_k denotes a sequence of identically standard exponentially distributed random variables.

Equation (27) is a linear regression model and the memory parameter d can be estimated by least squares.

Table 1 shows the Hurst coefficient H as estimated from the R/S statistic, the R/S statistic itself, both in standard form and its modified version R/S^*, computed with a bandwidth of $q = 90$ as suggested by Lo (1991), plus the more recent V/S variant suggested by Giraitis et al. (2000b), plus an estimate $\hat{d}$ for d computed by applying the GPH estimator, for various individual stocks and for the stock price index DAX. The GPH estimator is computed with the optimal bandwidth of $n^{\frac{4}{5}}$ (Hurvich et al., 1998). Hurvich et al. (1998) show that $n^{\frac{4}{5}}(\hat{d} - d)$ is asymptotically normal with variance $\frac{\pi^2}{24}$, which gives an approximate standard error of $\sigma_{GPH} = 0.016$ for our data. Returns are in raw form. Confirming the findings cited above, the Hurst coefficient is uniformly larger than 0.5, but not significantly so. Only for BMW and BASF the estimate $\hat{d}$ for d is significantly different from 0.

Table 1

Estimated Hurst coefficient and tests for long memory for raw returns

	H	R/S	R/S^*	V/S	$\hat{d}$
Daimler	0.59	0.77	1.08	0.02	-0.002
BMW	0.60	0.59	0.96	0.02	0.114
Hoechst	0.55	0.86	1.43	0.03	0.067
BASF	0.54	0.86	1.46	0.03	0.14
Deutsche Bank	0.57	0.69	0.90	0.14	-0.019
DAX	0.56	0.97	1.50	0.42	0.066

The 5% significance levels of the R/S, R/S^* and V/S-statistics respectively, are 1.747, 1.747 and 0.1869.

Table 2 shows the same statistics as computed from the excess returns r_t^*. It is seen that the Hurst coefficient is almost identical for most stocks (small changes after 2 digits) and slightly closer to 0.5 for BMW and that the significance of the test results is even somewhat increased. The data therefore reject our hypothesis that taking excess rather than raw returns has an impact on measures of long memory. This result is robust to alternative choices of the risk premium (2% or 4% per year) and the risk-free interest rate. No matter which type of returns we take, the evidence for long memory is statistically insignificant; given a sample size of more than 9000, this means that long memory in returns does not exist.

Table 2

Estimated Hurst coefficient and tests for long memory for adjusted returns

	H	R/S	R/S^*	V/S	$\hat{d}$
Daimler	0.59	0.80	1.11	0.04	0.03
BMW	0.57	0.61	0.98	0.05	0.13
Hoechst	0.55	0.94	1.55	0.10	0.08
BASF	0.54	0.94	1.58	0.12	0.15
Deutsche Bank	0.57	0.69	1.02	0.03	0.01
DAX	0.56	1.07	1.62	0.13	0.08

4 Long Memory in Squares of Stock Returns

Contrary to stock returns themselves, squares (or absolute values and various other functions) of stock returns have long been known to exhibit considerable autocorrelation. The only issue is whether or not the autocorrelation function dies out fast enough to preclude long memory.

In GARCH-models of the type

$$r_t = z_t\sqrt{h_t} \tag{28}$$
$$h_t = f(r_{t-1}^2, r_{t-2}^2, \ldots, h_{t-1}, h_{t-2}, \ldots)\ , \tag{29}$$

where z_t is i.i.d. (0,1) and f is linear, it is easily seen that the autocorrelations of the squared returns must die out exponentially if they exist (see e.g. Mikosch and Starica 1999). In the GARCH (1,1) model

$$h_t = w + \alpha r_{t-1}^2 + \beta h_{t-1}\ , \tag{30}$$

one often obtains estimates of α and β which sum to unity, in which case there are no finite fourth moments of r_t and the autocorrelations of r_t^2 do not exist. It is an open problem which we address in a separate paper whether in such models, and in the more general FIGARCH model of Baillie et al. (1996), condition (1) applies. In any case, empirical autocorrelations of squared returns, which can always be computed, whether population counterparts exist or not, often indicate the possibility of long memory.

This appearance of long memory can be explained either by employing a nonlinear function f in (29), or by entertaining the possibility of structural changes in the function f (i.e. changes in Er_t^2, see e.g. Mikosch and Starica 1999). If there are structural changes, the appearance of long memory should decrease when we consider subperiods of the sample:

This is what we do below. Table 3 gives the analogous statistics from tables 1 and 2, as applied to the squares of the returns. In addition, it also shows the ML estimates for α and β. The table shows that, more often than not, the data indicate long memory. The R/S, R/S^*, V/S and $\hat{d}$ statistics are all significant at 5%. Also, the estimated GARCH coefficients sum to almost unity.

Table 3
Estimated Hurst coefficient and tests for long memory for squared returns

	H	R/S	R/S^*	V/S	$\hat{d}$	α	β	$\alpha+\beta$
Daimler	0.66	4.04	2.34	0.48	0.42	0.14	0.84	0.98
BMW	0.69	2.83	2.40	0.43	0.28	0.14	0.85	0.99
Hoechst	0.67	3.76	2.76	0.55	0.27	0.12	0.84	0.96
BASF	0.66	3.31	2.55	0.38	0.28	0.13	0.82	0.95
Deutsche Bank	0.71	3.52	2.17	0.35	0.34	0.11	0.88	0.99
DAX	0.71	3.19	2.25	0.32	0.22	0.15	0.83	0.98

Table 4 shows the same statistics, as computed for the eight 5-year periods separately. The relevant approximate standard errors for $\hat{d}$ are now 0.036. To save space, only averages of the respective coefficients are indicated. It is obvious that the appearance of long memory is considerably reduced.

Table 4
Estimated Hurst coefficient and tests for long memory for squared returns
(averages over 5-year subperiods)

	H	R/S	R/S^*	V/S	$\hat{d}$	α	β	$\alpha+\beta$
Daimler	0.68	2.89	1.25	0.14	0.26	0.15	0.80	0.95
BMW	0.68	3.14	1.49	0.18	0.17	0.18	0.76	0.94
Hoechst	0.67	2.85	1.45	0.15	0.08	0.13	0.84	0.97
BASF	0.68	2.89	1.43	0.15	0.15	0.17	0.61	0.77
Deutsche Bank	0.71	3.25	1.31	0.15	0.21	0.17	0.61	0.90
DAX	0.73	2.73	1.20	0.13	0.20	0.15	0.97	0.95

The significance of the long-memory effects can be further reduced by taking still shorter periods (one year, two years) or by subdividing the sample at time points where structural changes are known to have occurred (oil crisis, 1987 crash), or by first estimating structural changes and subdividing the sample accordingly. This latter strategy suffers from unresolved problems concerning the null distribution of popular tests for structural changes (Krämer/Sibbertsen 2000) and is not followed here.

5 Conclusion

We show that long memory in German stock or squares of German stocks returns is basically a non-issue which results from either insignificant samples or structural change in the expectation of the series. An open question is whether non-existing second moments can likewise make a series appear as if it had long memory. We hope to add some insights here in later work.

References

Ambrose, B., Ancel, E. and Griffiths, M. (1993): "Fractal structure in the capital market revisited." *Financial Analysts Journal*, May/June, 73 – 77.

Aydogan, K. and Booth, G.G. (1988): "Are there long cycles in stock returns?" *Southern Economic Journal* 55, 141 – 149.

Baillie, R.T., Bollerslev, T. and Mikkelsen, H.O. (1996): "Fractionally integrated generalized autoregressive conditional heteroskedasticity." *Journal of Econometrics* 74, 3 – 30.

Beran, J. (1994): *Statistics for long-memory processes.* London: Chapman & Hall.

Bhattacharya, R.N., Gupta, V.K. and Waymire, E. (1983): "The Hurst effect under trends." *Journal of Applied Probability* 20, 649 – 662.

Boes, D.C. and Salas, J.D. (1978): "Nonstationarity of the mean and the Hurst phenomenon." *Water Resources Research* 14, 135 – 143.

Chow, K.V., Denning, K.C., Ferris, S. and Noronha, G. (1995): "Long and short-term price movements in the stock market." *Economics Letters* 49, 287 – 293.

Crato, N. (1994): " Some international evidence regarding the stochastic memory of stock returns." *Applied Financial Economics* 4, 33 – 39.

Diebold, F.X. and Inoue, A. (1999): "Long memory and structural change." Stern School of Business discussion paper.

Ding, Z., Engle, R.F. and Granger, C.W.J. (1993): "A long memory property of stock market returns and a new model." *Journal of Empirical Finance* 1, 83 – 106.

Geweke, J. and Porter-Hudak, S. (1983): "The estimation and application of long memory time series models." *Journal of Time Series Analysis* 4, 221 – 238.

Giraitis, L., Kokoszka, P. and Leipus, R. (2000a): "Testing for long memory in the presence of a general trend." Discussion paper, London School of Economics.

Giraitis, L., Kokoszka, P. and Leipus, R. (2000b): "Rescaled variance and related tests for long-memory in volatilities and levels." Discussion paper, London School of Economics.

Goetzmann, W. (1993): "Patterns in three centuries of stock market prices." *Journal of Business* 66, 249 – 269.

Gourieroux, C. and Jasiak, J. (2001): "Memory and infrequent breaks." *Economics Letters* 70, 29 – 41.

Granger, C.W.J. and Hyung, N. (1999): "Occasional structural breaks and long memory." UCSD discussion paper 99 – 14.

Greene, M.T. and Fielitz, B.D. (1977): "Long-term dependence in common stock returns." *Journal of Financial Economics* 4, 339 – 349.

Heyde, C.C. and Yang, Y. (1997): "On defining long–range dependence." *Journal of Applied Probability*, 34, 939 – 944.

Hiemstra, C. and Jones, J.D. (1997): "Another look at long memory in common stock returns." *Journal of Empirical Finance* 4, 373 – 401.

Hurvich, C., Deo, R. and Brodsky, J. (1998): "The mean square error of Geweke and Porter-Hudak's estimator of the memory parameter of a long-memory time series." *Journal of Time Series Analysis* 16, 17 – 41.

Krämer, W. and Sibbertsen, P. (2000): "Testing for structural change in the presence of long-memory." Technical Report 31/2000, SFB 475, Universität Dortmund.

Lindgren, G. (1978): "Markov regime models for mixed distributions and switching regressions." *Scandinavian Journal of Statistics* 5, 81 – 91.

Lo, A.W. (1991): "Long–term memory in stock market prices." *Econometrica* 59, 1279 – 1313.

Lobato, I.N. and Savin, N.E. (1998): "Real and spurious long-memory properties of stock market data (with discussion and reply)." *Journal of Business and Economic Statistics* 16, 261 – 283.

Mikosch, T. and Starica, C. (1999): "Change of structure in financial time series, long range dependence and the GARCH model." Preprint 99–5–06, University of Groningen.

Mills, T. (1993): "Is there long-term memory in UK stock returns?" *Applied Financial Economics* 3, 303 – 306.

Peters, E. (1992): "R/S analysis using logarithmic returns." *Financial Analysis Journal*, Nov./Dez., 81 – 82.

Rydén, T., Teräsvirta, T. and Åsbrink, S. (1998): "Stylized facts of daily return series and the hidden Markov model." *Journal of Applied Econometrics* 13, 217 – 244.

Willinger, W., Taqqu, M.S., and Tevorovsky, V. (1999): Stock market prices and long–range dependence, *Finance and Stochastics* 3, 1 – 13.

Figure 1

A short-memory time series superimposed by a broken trend

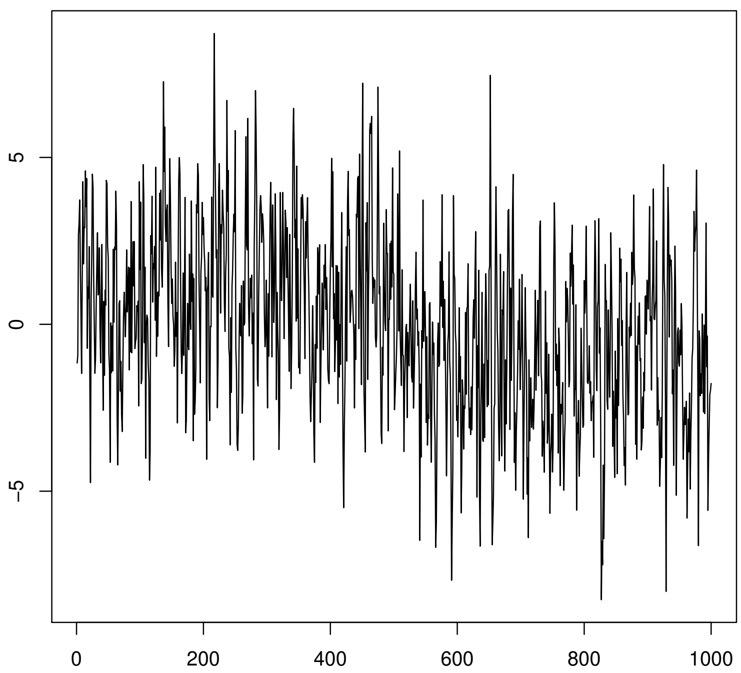

Figure 2

Empirical autocorrelations of the time series from Figure 1

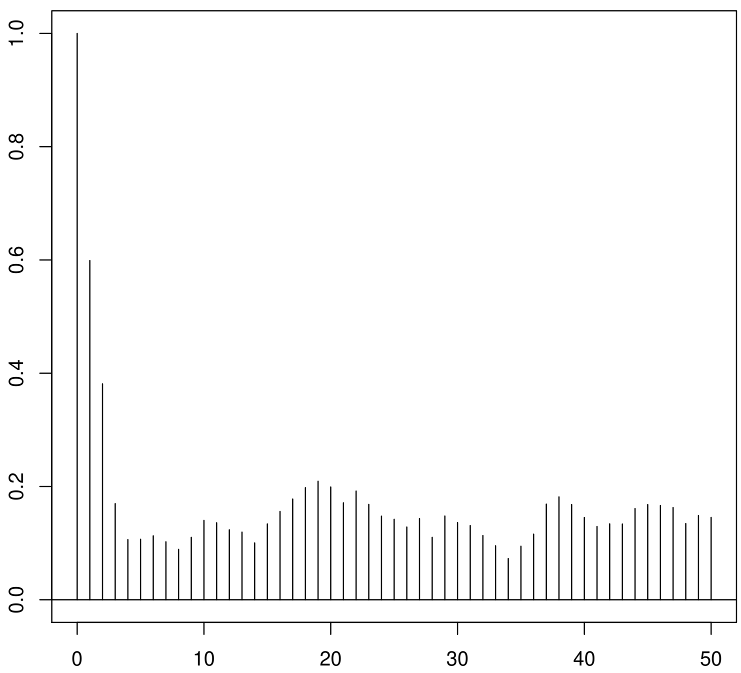

Figure 3

Empirical autocorrelations as sample size increases, $n = 4000, 16000, 32000$ (from top)

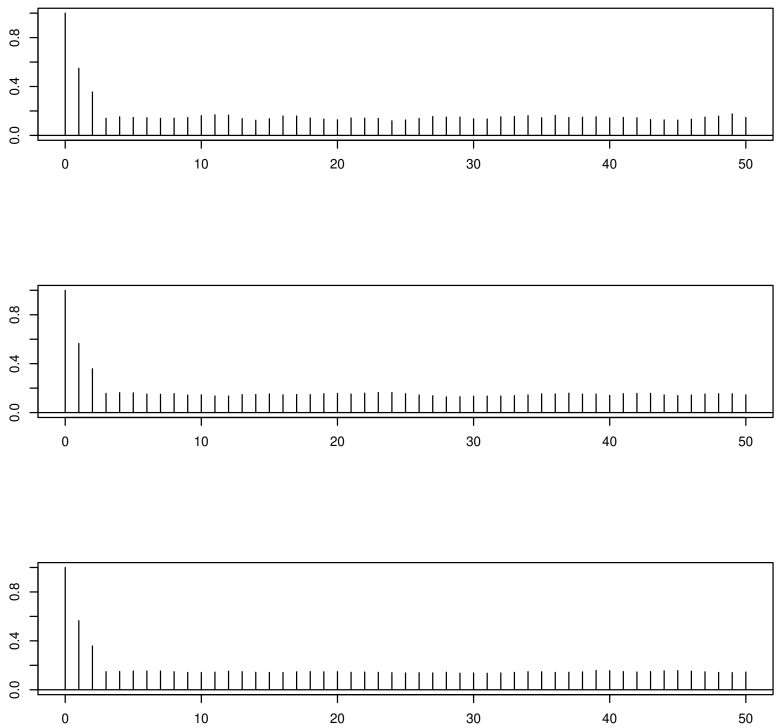

Figure 4

The R/S statistic as a function of the sample size