Differences in process memory between high and medium frequency financial time series

Branko Pecar, PhD
University of Gloucestershire
The Park
Cheltenham GL50 2QF
Tel: ++44 (0)1242 536207
bpecar@glos.ac.uk

Abstract. Rescaled range analysis has been extensively used as a technique for characterising financial time series. In this paper we focus on the Hurst exponent characterisation of time series as a framework for describing some of the properties of financial time series. The objective is to establish, by estimating Hurst exponents, if stock returns scale differently for different horizons and for different return frequencies. The returns on stocks for ten real-life data sets taken from the NYSE and fifteen artificially generated time series were analysed. The frequency of data impacts only short memory processes. High frequency minute data indicate antipersistent behaviour, whilst medium frequency data are somewhat persistent. As the length of the return period increases, the properties of the time series begin to converge to their original pattern at zero lag. For long returns the frequency of data makes no impact on the process memory.

Keywords. Process memory, rescaled range analysis, the Hurst exponent, non-linear time series analysis, financial time series frequency

Journal of Economic Literature Classification System: C22, C81

1. Introduction

Two questions continue to keep on intriguing researchers and practitioners dealing with time series, and they are: do financial time series have embedded memory, and if so, does this memory exhibit any sensitivity to data frequency. Many studies, as shown below, have attempted to answer these two and numerous other closely related questions. Some of the answers are consistent, the others are very much contradictory. This paper will attempt to provide answers, although only partially, related to the process memory for only two different frequencies of stock return data.

One way to describe long memory processes is by analysing their autocorrelation functions. A long memory process is characterised by a hyperbolic decay in autocorrelation coefficients and a short memory process is characterised by an exponential decay in autocorrelation coefficients. However, rather than using autocorrelations, or corresponding frequency analysis, we'll focus on rescaled range analysis, or more specifically, the Hurst exponent. The technique pioneered by Hurst, and contextualised by Mandelbrot, is a very robust technique, insensitive to frequency distribution and potential non-stationarity of the data. Empirical evidence indicates that high frequency financial data are more likely to be stationary, whilst low frequency data are more likely to be non-stationary. This paper will examine both types of data. The use of appropriate techniques capable of handling both types of data is important.

2. Objectives

The objective of this paper is to establish whether the data frequency affects process memory. Specifically, it will be examined whether medium frequency data (daily stock returns) exhibit different process memory from very high frequency data (minute stock returns).

In order to achieve this, an experimental approach has been adopted. Ten time series of sufficient length (at least five thousand data points) were used. All the time series represent actual real-life stock closing values. From these closing values, the returns were calculated for different time horizons. Arbitrarily, the return horizons include 1, 100, 500 and 1,000 periods (lags)[1]. Equally, a selection of stocks from different industry sectors was made, representing: energy, IT, food and beverage, pharmaceutical and retail sectors. In addition to this, several artificial series were generated for benchmarking, as described further below. All the time series are described in greater detail in section 6.

As all the time series used in this paper are of reasonable length (depending on the series, between 5,400 and 10,400 observations), the use of classical rescaled range analysis for estimating the Hurst exponent seems appropriate[2]. To provide additional rationale for this decision, a brief overview of the most important literature in this domain is surveyed, followed by a short introduction of the rescaled range analysis and the Hurst exponent.

[1] Returns for 1,000 days correspond to 2.7 years and for 1,000 minutes to 16.7 hours, i.e. 2.5 days.
[2] For implications of the series length on calculating the Hurst exponent, see Ambrose, B. W., E. W. Ancel, et al. (1993). "Fractal Structure in the Capital Markets Revisited." Financial Analysis Journal(May-June): 73-77.

3. Current state of affairs

Generally, there seems to be a lack of consensus on how to characterise various financial time series. Partially this is due to the fact that different categories of series at different time frequencies are examined, and they do not necessarily behave in a uniform way. Some studies scrutinise the actual stock values or the trading volume, although most of the studies focus on stock returns, which is what this paper will do, and we'll define the returns shortly. Another often considered category is stock volatility, which could be expressed as the squared returns, or the absolute returns, or even as the logs of squared returns. A number of models are then used by different researchers to characterise financial time series, and some of the examples are: Gaussian distribution, Levy flight, various leptokurtic distributions, ARCH/GARCH models and an almost infinite number of their derivatives such as IGARCH, EGARCH, SWARCH, FIGARCH, etc., depending on what has been modelled. What follows is a brief overview of some of the work published in this domain.

Peters (1989; 1992) picked up on some early Mandelbrot hypotheses (to be described shortly) and used the so-called rescaled range (R/S) analysis to explore the question of persistence in time series. This was followed shortly by Ambrose, Ancel et al. (1993) who questioned some of the general claims made (those of Peters in particular) about persistent behaviour in monthly time series. They concluded that there is no long-term deterministic structure. The issue they challenged in particular was inadequate time series length as well as the inappropriate selection of the subset for the calculation of the Hurst exponent. However, the approaches they adopted to R/S analysis (referred to as F-Hurst and G-Hurst) give identical results for long time series (more than 2^{10} observations).

In order to add rigour to this area, Lo (1991) suggested a test, based on R/S analysis and the Hurst exponent, to establish the presence of long term dependencies in time series. According to Lo, classical R/S analysis is incapable of distinguishing between short-range and long-range dependencies in the time series. In computational terms, Lo's modified rescaled range method makes minor changes when compared to the original rescaled range method. The modified method divides the range with the standard deviation as well as with the weighted autocovariances up to some lag q. Pagan (1996) confirmed that the choice of q is critical for results. A small q favours the alternative hypothesis (presence of long memory) and a large q is more in favour of the null hypothesis, making the test somewhat biased.

These few examples represent a great body of effort made during the early nineties on this front. A good review of a number of early research attempts in this domain is given by Brock and deLima (1995). Since their monograph has been made available, a number of other researchers have published their findings. Some of the more recent and prominent sources are listed below.

Mantegna & Stanley (1995) analysed the S&P 500 index and found that, in support of Mandelbrot's early claims[3], the Levy distribution matches their data, although they

[3] Mandelbrot suggested a leptokurtic Levy distribution (which complies with the "fat tail" property of many stock price distributions).

found deviations of the tail of the distribution. Although the Levy distribution has an infinite second moment, in the case of their data, a finite second moment was present, explaining some of the deviations. They also examined a GARCH model and reported that it showed different properties from the observed data.

Evertsz (1995) analysed high frequency (intra-day) DEM-USD foreign exchange rates and German DAX daily closing data[4]. Unlike Mantegna & Stanley, he rejected the possibility that the Levy distribution models his data sets. He also found no evidence of change in distribution for various returns, in line with the changes in the Hurst exponent.

Hiemstra and Jones (1997) analysed daily returns of 1,952 common stocks and found that long memory is not a widespread characteristic of common stocks. Goodhart and O'Hara (1997), in a slightly different context, analysed whether high frequency data will reveal limitations to the efficiency of markets and concluded that, while asset market prices exhibit nonlinearities, they are not chaotic.

Bisaglia and Guegan (1998) compared a number of different methods (including the classical R/S method) for estimating the Hurst exponent. They used intra-day exchange rates DEM/FRF from 1 September 1994 to 1 September 1995, examined at different frequencies: (i) every hour (5904 observations), (ii) every 20 min (14819 observations) and (iii) every 10 min (24827 observations). They scrutinised log differenced data (although the absolute returns and square returns as volatility estimators were also examined). All three series show on average $H \approx 0.5$. In fact classical R/S method show $H \approx 0.46$, implying small anti persistence.

Harrison, Yu et al (1999) used exceptionally long time series of the daily S&P Composite Price Index and reported the presence of deterministic chaos.

Stanley (2000) quotes several papers, in particular Plerou, Gopikrishnan, et al who, after analysing 40 million data points from the US stock markets, found evidence that there is no Levy regime in their data, contrary to what Mandelbrot and later Mantegna hypothesised.

Lobato and Velasco (2000) examined trading volume and return volatility (not stock returns as in this paper), but in the frequency domain. Specifically, they used 30 stock data from DJIA (Dow Jones Industrial Average), each, with the exception of three, containing 8,180 daily observations. They discovered that the trading volume exhibits long memory. The volatility also exhibits the same long memory, but this component is different from the volume one.

Meade (2002), on the other hand, analysed predictability of an AR-GARCH model and found it more appropriate than non-linear models, implying that the high frequency exchange rate data he analysed follows the GARCH model.

[4] As an a-priori remark, note that for high frequency data he estimated H=0.45 and for daily data H=0.54, which resembles greatly the findings reported in this paper.

Strozzi (2002) used a variety of data sets, namely: intra-day foreign exchange spot rates, bid and ask quotes for precious metals, transaction prices from European future contracts and two stock indices (DJIA and S&P 500). All the data sets were high frequency (half-hourly). The reported Hurst exponents for all the series, without exception, were very high (almost 1) indicating long process memory.

Small and Tse (2003) explored three different daily data sets and rejected the random walk model altogether in favour of the hypothesis of deterministic nonlinear dynamics (although not necessarily that of chaotic behaviour). Deterministic structure cannot be modelled with conditional heteroscedastic processes, and they reject the usefulness of all GARCH models in modelling inter-day series. The conclusion is that financial series are deterministic but "swamped by stochastic behaviour"[5]. If this is the case, then process memory is by default embedded in such time series.

The brief overview that proceeded indicates that a number of papers, implicitly or explicitly, tried to establish the presence of long-term memory in financial time series. Not surprisingly, most of them reported contradictory findings, depending on circumstances. The causes of contradictory findings could be summarised as follows: the length of the data set being insufficient or inconsistent; the frequency of data sampling varying from study to study; some of the analyses covering stock returns, the others volatility and some others cross correlations between the volume and value; some reports focusing on stocks from only one industry sector, or just on aggregate series (indices); different time horizons being used to calculate the returns in different studies, etc. All in all, a number of issues could obscure findings and there is no consensus or a general model that could embrace all aspects of financial time series characterisation. Findings in one area, applicable to a specific data set, are not universal.

4. The background of rescaled range analysis

The origins of rescaled range analysis go back to H.E. Hurst's interest in the storage capacity of water reservoirs and his seminal paper from 1951 (Hurst 1951). Hurst devoted a substantial amount of time, energy and intellect to establishing how to compute the correct storage required in the Great Lakes of the Nile river basin. His objective was to determine the capacity of a reservoir guaranteeing a minimum discharge when the water intake is low. To determine the optimum capacity, he examined the maxima and minima of a number of reservoirs over a number of periods and concluded that there is an optimum range required to maintain the average discharge. When this range was measured against different time horizons, he discovered strong scaling properties. This scaling property is estimated as the slope of the line on a log plot showing the changes in the rescaled range[6] versus the number of cases used for the range. Contrary to his expectation that this slope should be equivalent to 0.5 (which would be normal for random events), he discovered that the number is around 0.7. In other words, the scaling of certain natural phenomena (reservoirs included) does not follow chance, but shows an increasing amount of long process memory and persistence, as in this particular case.

---

[5] Description of this stochastic behaviour is vague and implies either "apparently random" behaviour or possibly even extremely high dimensional dynamics.

[6] Range is a function of deviations from the mean. Rescaling the ranges means that they are divided by their appropriate standard deviation, i.e. R/σ.

Mandelbrot based some of his hypotheses on Hurst's work and made a breakthrough in 1965 (Mandelbrot 1983). He renamed the rescale analysis exponent as the Hurst exponent and changed its original notation from K, as used by Hurst, to H. More importantly, he linked it with the fractal dimension and established that they are related as D = 2 - H, where D is the fractal dimension. The real breakthrough came from his hypothesis that the assumption of normally distributed price variations should be rejected, although, as we saw above, there is no consensus as to which distribution it should be replaced with.

The Hurst exponent, as a technique, makes no assumption on the frequency distribution of the data and as such is a valuable tool. The Hurst exponent can takes values $0<H<1$. For $H=0.5$, the series is comparable to a white noise process. For H close to zero, the series is random and quite 'jumpy', showing a tendency for every observation to move dramatically either in the opposite direction from the previous observation or continue in the same direction in a very accentuated manner. For H close to unity, the series is still random but much smoother. It tends to jump up and down over much larger intervals. This kind of series is said to have an element of persistence built into it. The Hurst exponent, therefore, indicates both how rough the series is and the length of its memory.

5. Method of estimating the Hurst exponent

We indicated that there is a connection between the Hurst exponent and the fractal dimension. However, there is also a relationship between the Hurst exponent and the spectral density exponent β. Spectral densities S_x show the following relationship:

$$S_x \propto \frac{1}{f^{\beta}} \tag{1}$$

where, f are the frequencies and $\beta=2H+1$. This indicates that the Hurst exponent can also be estimated from the power spectrum of the series[7]. A number of other ways are known for estimating the Hurst exponent. Some are more conventional, such as a modified periodogram spectral estimator or an autoregressive spectral estimator, the others are more specialised, such as dispersion analysis or via the fractal dimension. Cannon, Percival et al. (1997), for example, used three scaled windowed variance (SWV) methods for estimating the Hurst exponent. In principle, an alternative to the R/S method for estimating the Hurst exponent is necessary if the time series is not of sufficient length[8]. As the time series selected in this paper are invariably greater than $N=2^{12}$ (some are $N>2^{13}$), the original Hurst R/S method was deemed to be appropriate. A software package called Fractan (release 4.3), freely available from the Internet, was used[9] and it calculates the Hurst exponent very much as prescribed by Hurst.

[7] For $H=0.5$, for example, we get ordinary Brownian motion, often also referred to as $1/f^2$ noise.
[8] Cannon, Percival et al. (1997) reported inconsistencies only with the time series of $N<2^9$
[9] Fractan was created by Vyacheslav Sychyov and is available from <http://fractan.boom.ru/soft.htm>.

If x_t are observations in a time series, $\bar{x}$ the mean, S the biased standard deviation and a time series fragment is of length τ, the time series can be rescaled for different lengths of the fragment in the following way:

$$\bar{x}_{\tau,t} = \frac{1}{\tau}\sum_{t=1}^{\tau} x_t \qquad (2)$$

$$S_{\tau,t} = \sqrt{\frac{1}{\tau}\sum_{t=1}^{\tau}(x_t - \bar{x}_{\tau,t})^2} \qquad (3)$$

For an interval τ, which is $1 \le \tau \le N$, $\bar{x}_{\tau,t}$ is a local mean and $S_{\tau,t}$ is a local standard deviation. The cumulative sum of all the deviations in an interval from the local mean is:

$$Z_{\tau,t} = \sum_{t=1}^{\tau}(x_t - \bar{x}_{\tau,t}) \qquad (4)$$

From this, a local range can be calculated as:

$$R_{\tau} = \max Z_{\tau,t} - \min Z_{\tau,t} \qquad (5)$$

In order to make them comparable, the ranges are rescaled (normalised) by the local standard deviation and averaged for all identical intervals to decrease statistical variability (Mandelbrot 2002). For all identical intervals with τ elements:

$$(R/S)_{\tau} = \frac{1}{\tau}\sum_{t=1}^{\tau}\frac{R_{\tau}}{S_{\tau}} \qquad (6)$$

As the interval τ changes, so does the expected rescaled value of R/S. Usually this scaling property is expressed as:

$$E[(R/S)_N] \cong cN^H \qquad (7)$$

Where c is a constant and H is the Hurst exponent. The slope of the straight line that connects all $\ln(R/S)_N$ and $\ln(N)$ is the estimate of the Hurst exponent.

For proper random process following the Gaussian distribution H=0.5. This is considered a special case of Brownian motion. For all other values of H, the process is potentially a fractional Brownian motion. More specifically, where $0 \le H < 0.5$ the process is considered antipersistent and, where $0.5 < H \le 1$ the process is persistent.

6. Time series description

In this paper the analysis was conducted using ten real-life financial time series. The data sets are stock movements taken from the New York Stock Exchange (NYSE). The series consist of daily and minute closing values for BP, IBM, Coca Cola, Pfeizer and Wall Mart (Tickers symbols: BP, IBM, KO, PFE and WMT). The details for all 10 series are given in Table 1. All minute data cover Monday to Friday intra-day trading from 9.30-16.00.

Stock values	BP	IBM	Coca Cola	Pfeizer	Wall Mart
Daily closing values					
Start	03-Jan-77	02-Jan-62	02-Jan-70	04-Jan-82	25-Aug-72
Finish	09-May-03	09-May-03	09-May-03	09-May-03	09-May-03
No of observations	6653	10405	8423	5391	7698
Minute closing values					
Start	22-Apr-03	22-Apr-03	22-Apr-03	22-Apr-03	22-Apr-03
Finish	05-May-03	05-May-03	05-May-03	05-May-03	05-May-03
No of observations	5458	5862	5473	5473	5473

Table 1: Five financial time series from the NYSE selected for analysis

As indicated, rather than analysing the closing values of the stocks, the return values were calculated. If a stock price, at any point in time, is defined as X_t and Δ_t is a change in time, then the price change is calculated as:

$$R_t = \frac{X_{t+\Delta t} - X_t}{X_t} = \frac{Z_t}{X_t} \qquad (8)$$

As $Z_t/X_t \approx \log X_{t+\Delta t} - \log X_t$, if Z_t/X_t is small, the value of R_t can be more efficiently expressed as:

$$R_t = \log X_{t+\Delta t} - \log X_t \qquad (9)$$

This was the method adopted for calculating returns. In this paper specific returns were calculated for periods (lagged values) of 1, 100, 500 and 1,000 observations.

To provide a benchmark and illustrate how the Hurst exponent behaves, three initial data sets, representing different classes of processes, were generated artificially and they are: a white noise process, a Wiener process[10] and the Lorenz attractor. For all three series 8,000 observations were generated. The white noise process was generated on Excel using the RAND() function. The Wiener process was rendered as integrated white noise process. The Lorenz attractor, representing chaotic processes, was generated using difference equations[11] with the following parameter values:

dx/dt= -10(x +y) (10)
dy/dt= 28x -y -xz (11)
dz/dt= -8/3z +xy. (12)

For these three simulated series, artificial ‘returns’ were also calculated for identical time periods as for the stock exchange data, i.e. for 1, 100, 500 and 1,000 lags.

[10] Brownian motion with Markov characteristics, earlier referred to as fractional Brownian motion (fBm)
[11] Only ‘x’ variable evolving in a time window has been used

7. The Hurst exponent estimates

The values of the Hurst exponent estimated for all three simulated time series are given in Table 2.

As expected, a white noise process shows the Hurst exponent value at around 0.5 (0.5153 ±0.14 in this case) and the Wiener process' Hurst exponent is close to 1 (0.9667 ±0.11). On the other hand, the Hurst exponent for the Lorenz attractor is 0.6282 ±0.17. These facts do not reveal anything new and are used only as a benchmark. The real question is: what happens to the Hurst exponent for the returns calculated at different lags for each of these series?

Hurst exponents	Actual	Lag 1	Lag 100	Lag 500	Lag 1000
White noise	0.5153	0.1064	0.2476	0.4284	0.5232
Wiener	0.9667	0.5186	0.7238	0.8748	0.9921
Lorenz	0.6282	0.1001	0.2802	0.5419	0.6174

Table 2: Hurst exponents for three artificially generated time series and their lagged log differences (returns)

In the case of the Lorenz attractor, the Hurst exponent drops dramatically for lag one, and it gradually recovers towards its original value, as the number of lags increases. A white noise process shows identical behavioural properties (a five-fold initial drop in the Hurst value). The Wiener process, which resembles most real-life stock exchange data, shows the same long-term tendency, i.e. for long lags it will reach its original Hurst exponent. However, for very short lags (lag of one) its Hurst value is only half its original value and it is on a comparable level to the white noise Hurst exponent for zero lag.

How should these results be interpreted from the perspective of process memory? Essentially, for very long lags (returns), all three processes resume their initial levels of process memory, i.e. a white noise process continues to show absence of any memory, the Weiner process shows great persistence and the Lorenz attractor shows some limited persistence. For very short lags (returns for just one period) the pattern is less uniform. The chaotic process (the Lorenz attractor), just like a white noise process, tends to become very antipersistent. The Wiener process, on the other hand, gets reduced to a process that looks like a white noise process.

What happens with the returns for real-life stock exchange time series? Estimated values of Hurst exponents are given in Table 3.

The Hurst exponent for all ten closing stock value series (five daily and five minute time series) shows very high values, resembling a Wiener process. We mentioned above that the short-term returns of a Wiener process behave like white noise with $H \approx 0.5$. The expectation is that the short-term returns for both minute and daily stock values for only one period will also behave just like a white noise process.

A majority of the daily returns for lag one show the Hurst exponent higher than 0.5. One return series (WMT) shows even more radical departure in the positive direction from 0.5. In this selection of stocks, the IBM daily returns are the only ones exhibiting antipersistent behaviour. Although there is not sufficient evidence to say, at least on the basis of this sample, that single day returns are fundamentally different from a white noise process, it remains a fact that the Hurst exponent for the majority of these series is on the upper side of the value of 0.5.

Hurst values	Actual	Lag 1	Lag 100	Lag 500	Lag 1000
Daily values					
BP	0.9652	0.50905	0.68770	0.89042	0.96000
IBM	0.9324	0.43335	0.87070	0.86155	0.89970
KO	0.9492	0.55218	0.84360	0.94931	1
PFE	1	0.52612	0.69930	0.87243	0.95040
WMT	0.9668	0.61411	0.82730	0.88602	0.93330
Average daily H	0.96272	0.52696	0.78572	0.89195	0.94868
Minute values					
BP	1	0.50853	0.68040	0.92881	0.97940
IBM	0.9193	0.43742	0.61900	0.82842	0.92960
KO	1	0.43901	0.64530	0.90459	0.91080
PFE	1	0.46158	0.74290	0.90050	0.97510
WMT	0.9327	0.49333	0.62020	0.82782	0.97040
Average minute H	0.97040	0.46797	0.65685	0.87803	0.95306

Table 3: Hurst exponent values for daily and minute returns separated 1, 100, 500 and 1,000 periods

A majority of the minute returns for lag one, on the other hand, have the Hurst exponent lower than 0.5 (with the exception of BP), implying some limited process memory, i.e. the antipersistent type of memory. Clearly there is a difference between minute and daily short-term returns. To put it very informally, today's return on stocks is in general faintly positively correlated with yesterday's return, whilst this minute's return on stocks is more likely to 'overreact' by either going in the opposite direction from the return recorded just a minute ago, or continuing at some larger order of magnitude in the same direction.

For very long returns, in our case 1,000 periods, both daily and minute return series resume the shape of their original time series, showing $0.9<H<1$. The recovery path for the average Hurst exponent from zero to 1,000 lags in Fig. 1 looks very much like the one for the Wiener process. The recovery rate of the Hurst exponents for daily returns is marginally different from the recovery rate of the Hurst exponents for minute returns. Daily returns seem to show a somewhat higher level of persistence than minute returns, although they are both recovering ultimately towards their original level of the Hurst exponent.

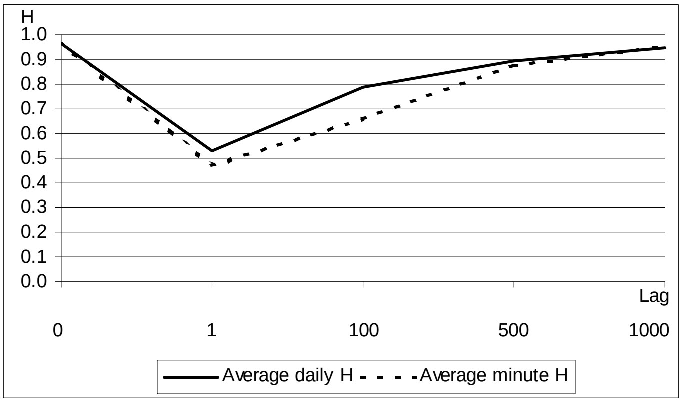

Fig 1. Average daily and minute Hurst exponents for corresponding five return series shown as a function of changes in return horizons (lags)

8. Conclusions

We set out to determine what kind of process memory is associated with daily and minute stock value returns. A limited sample of only five daily and five minute stock series from the NYSE (although varied from the industry point of view) was selected. The series were long enough to warrant the use of conventional R/S analysis to estimate Hurst exponents, which were used as a basis for the analysis. We concluded that experimental data indicate that there is a difference in memory properties between the medium and high frequency data.

Although the estimated Hurst exponents implied that both the single day stock returns and the single minute returns have some limited memory, the nature of this process memory is of the opposite character. Whilst the single day returns show some persistence, the single minute returns clearly show antipersistence. This was informally articulated as a rule according to which today's return tends to show an inclination to pick up yesterday's trend, whilst this minute's return is more likely to go in the opposite direction from the return recorded just a minute ago. For the long-term returns (calculated for lags of 1,000 observations) there is no difference between daily and minute returns and they both show similar properties to the original series from which they were derived. In other words, the frequency of data has no impact on long process memory, but has an impact on short process memory.

Using a very limited tool set (in this case the R/S analysis and estimating Hurst exponents) and only several experimental data sets some interesting findings are suggested. However, these findings, as interesting as they are, need to be validated in a more robust way using much larger sample of stocks.

References

Ambrose, B. W., E. W. Ancel, et al. (1993). "Fractal Structure in the Capital Markets Revisited." Financial Analysis Journal(May-June): 73-77.

Bisaglia, L. and D. Guegan (1998). "A comparison of techniques of estimation in long-memory processes." Computational Statistics & Data Analysis **27**: 61-81.

Brock, W. A. and P. J. F. Lima de (1995). Nonlinear time series, complexity theory and finance, http://212.67.202.199/~msewell/complexity/BrLi95.pdf.

Cannon, M. J., D. B. Percival, et al. (1997). "Evaluating scaled windowed variance methods for estimating the Hurst coefficient of time series." Physica **A**(241): 606-626.

Evertsz, C. J. G. (1995). "Fractal Geometry of Financial Time Series." Fractal Geometry and Analysis **3**(3): 609-616.

Goodhart, C. A. E. and M. O'Hara (1997). "High frequency data in financial markets: Issues and applications." Journal of Empirical Finance **4**: 73-114.

Harrison, R. G., D. Yu, et al. (1999). "Non-linear Noise Reduction and Detecting Chaos: Some evidence from the S&P Composite Price Index." Mathematics and Computer in Simulation **48**: 497-502.

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

Hurst, H. E. (1951). "Long-term storage capacity of reservoirs." Transactions of the American Societ of Civil Engineers **116**: 770-808.

Lo, A. W. (1991). "Long-Term Memory in Stock Market Prices." Econometrica **59**(5): 1279-1313.

Lobato, I. N. and C. Velasco (2000). "Long Memory in Stock-Market Trading Volume." Journal of Business & Economic Statistics **18**(4): 410-427.

Mandelbrot, B. (1983). The Fractal Geometry of Nature. New York, W H Freeman & Co.

Mandelbrot, B. B. (2002). Gaussian Self-Affinity and Fractals. New York, Springer Verlag.

Mantegna, R. N. and H. E. Stanley (1995). "Scaling behaviour in the dynamiccs of an economic index." Nature **376**: 46-49.

Meade, N. (2002). "A Comparison of the Accuracy of Short Term Foreign Exchange Forecasting Methods." International Journal of Forecasting **18**: 67-83.

Pagan, A. (1996). "The econometrics of financial markets." Journal of Empirical Finance **3**: 15-102.

Peters, E. (1989). "Fractal structures in the capital markets." Financial Analysis Journal(July/August).

Peters, E. (1992). "R/S Analysis using logarithmic returns." Financial Analysis Journal(November/December).

Small, M. and C. K. Tse (2003). "Determinism in financial time series." Studies in Nonlinear Dynamics and Econometrics **7**(3): Article 5.

Stanley, E. H. (2000). "Exotic statistical physics: Applications to biology, medicine and economics." Physica A **285**: 1-17.

Strozzi, F. (2002). Application of non-linear time series analysis techniques to high frequency currency exchange data, http://www.biblio.liuc.it:8080/biblio/liucpap/pdf/99.pdf.