Finance Stochast. 3, 1–13 (1999)

Finance and Stochastics

Stock market prices and long-range dependence*

Walter Willinger[1], **Murad S. Taqqu**[2], **Vadim Teverovsky**[2]

[1] AT&T Labs-Research, 180 Park Avenue, C284, Florham Park, NJ 07932, USA
(e-mail: walter@research.att.com)
[2] Boston University, Department of Mathematics, Boston, MA 02215, USA

Abstract. Using the CRSP (Center for Research in Security Prices) daily stock return data, we revisit the question of whether or not actual stock market prices exhibit long-range dependence. Our study is based on an empirical investigation reported in Teverovsky, Taqqu and Willinger [33] of the modified *rescaled adjusted range* or R/S statistic that was proposed by Lo [17] as a test for long-range dependence with good robustness properties under "extra" short-range dependence. Our main conclusion is that because the modified R/S statistic shows a strong preference for accepting the null hypothesis of no long-range dependence, irrespective of whether long-range dependence is present in the data or not, Lo's acceptance of the hypothesis for the CRSP data (i.e., no long-range dependence in stock market prices) is less conclusive than is usually regarded in the econometrics literature. In fact, upon further analysis of the data, we find empirical evidence of long-range dependence in stock price returns, but because the corresponding degree of long-range dependence (measured via the Hurst parameter H) is typically very low (i.e., H-values around 0.60), the evidence is not absolutely conclusive.

Key words: Long-range dependence, fractional Gaussian noise, fractional ARIMA, long memory, R/S, stock prices

JEL classification: C13, C15, C52, G10

Mathematics Subject Classification (1991): 60G18, 62-09

---

* This research was partially supported by the NSF Grants NCR-9404931 and DMS-9404093 at Boston University.
Manuscript received: May 1997; final version received: September 1997

1 Introduction

Long-range dependence is widespread in nature (e.g., see Mandelbrot [22] for details) and has been extensively documented in hydrology, meteorology and geophysics (see for example Mandelbrot and Wallis [24, 25, 27]). More recently, long-range dependence has also started to play an important role in the analysis and performance modeling of traffic measurements from modern high-speed communications networks (for a recent bibliographical survey of this area, see Willinger, Taqqu and Erramilli [34]).

In economics and finance, long-range dependence has a long history and has remained a topic of active research is the study of economic and financial time series (e.g., see Lo [17], and Cutland, Kopp and Willinger [6], and references therein). Historical records of economic and financial data typically exhibit distinct nonperiodic cyclical patterns that are indicative of the presence of significant power at low frequencies (i.e., long-range dependence). This led Granger [12], for example, to talk about the "typical spectral shape of an economic variable". However, the statistical investigations that have been performed to test for the presence or absence of long-range dependence in economic data have offered a much less coherent picture. In the case of economic time series representing returns on common stocks, these investigations have often become a source of major controversies, mainly because of the important implications that the presence of the long-range dependence phenomenon in asset returns has on many of the paradigms used in modern financial economics: it is inconsistent with the efficient market hypothesis, and plays havoc with stochastic analysis techniques that have formed the basis of a big part of modern finance theory and its applications (for more details, see the discussions in Mandelbrot [21], Lo [17], Rogers [30] and Cutland, Kopp and Willinger [6]).

Historically, the importance of long-range dependent processes as stochastic models lies in the fact that they provide an elegant explanation and interpretation of an empirical law that is commonly referred to as *Hurst's law* or the *Hurst effect*. In short, for a given set of observations $(X_i, i \geq 1)$, with partial sum $Y(n) = \sum_{i=1}^{n} X_i, n \geq 1$, and sample variance $S^2(n) = n^{-1}\sum_{i=1}^{n}(X_i - n^{-1}Y(n))^2, n \geq 1$, the *rescaled adjusted range statistic* or R/S*-statistic* is defined by

$$\frac{R}{S}(n) = \frac{1}{S(n)}\left[\max_{0\leq t\leq n}\left(Y(t) - \frac{t}{n}Y(n)\right) - \min_{0\leq t\leq n}\left(Y(t) - \frac{t}{n}Y(n)\right)\right], \quad n \geq 1. \tag{1}$$

Hurst [16] found that many naturally occurring empirical records appear to be well represented by the relation $E\left[R/S(n)\right] \sim c_1 n^H, \quad \text{as } n \to \infty$, with typical values of the *Hurst parameter* H in the interval $(0.5, 1.0)$, and c_1 a finite positive constant that does not depend on n. On the other hand, if the observations X_i come from a short-range dependent model, then it is known [11, 2] that $E\left[R/S(n)\right] \sim c_2 n^{0.5}, \quad \text{as } n \to \infty$, where c_2 is independent of n, and finite and positive. The discrepancy between these two relations is generally referred to as the *Hurst effect* or the *Hurst phenomenon*.

Classical R/S-analysis (originally due to Mandelbrot and Wallis [25]; see also Mandelbrot and Taqqu [23]) aims at inferring from an empirical record the value of the Hurst parameter H for the long-range dependent process that presumably generated the record at hand. While classical R/S-analysis is not a very reliable method in the presence of small samples, it can be highly effective and useful as a graphical or “eyeballing” method for reasonably large samples, where it often provides a rather accurate picture of the presence or absence of long-range dependence in a given empirical record and, in the former case, about the intensity of long-range dependence as measured by the Hurst parameter. For practical purposes, the most useful feature of the classical R/S-analysis is its relative robustness under changes in the marginal distribution of the data, even if the marginals exhibit heavy tails with infinite variance (see for example Mandelbrot and Wallis [26] and also Mandelbrot and Taqqu [23]). In the context of asset return, Mandelbrot [20] was one of the first who considered the possibility of long-range dependence in common stock returns. His empirical studies were largely based on the classical R/S analysis and led him to suggest that representative values of the Hurst parameter for asset returns might be around $H = 0.55$. Mandelbrot’s empirical findings were essentially confirmed by a later large-scale study by Greene and Fielitz [13], who used the classical R/S-method on 200 daily stock return series of securities listed on the New York Stock Exchange and reported that many of the series are characterized by long-range dependence.

The less attractive features of the classical R/S-analysis are its sensitivity to the presence of explicit short-range dependence structures, its bias, and a lack of a distribution theory for the underlying statistic (1). These characteristics stand in the way of using classical R/S as a rigorous statistical inference method. To overcome some of these shortcomings, Lo [17] proposed a *modified* R/S-*statistic* that is obtained by replacing the denominator S in (1), i.e., the sample standard deviation, by a consistent estimator of the square root of the variance of the partial sum $Y(n)$. The motivation for this modification is that in the case of dependent random variables, the variance of the partial sum is not simply the sum of the variances of the individual X_i’s but also includes their autocovariances up to some lag, for a judicious choice of the truncation lag. Lo derives the limiting distribution of his modified R/S-statistic under both short-range and long-range dependence, claims that it is robust with respect to short-range dependence, and illustrates through Monte Carlo simulations that it has reasonable power against certain long-range dependence alternatives. When applied to the data sets of daily stock return indices from the Center for Research in Security Prices (CRSP), Lo finds no evidence of long-range dependence in the data and concludes that the dynamic behavior of asset returns may be adequately described by traditional, short-range dependent models.

In contrast to the prevailing view in the econometrics literature (Hauser, Kunst, Reschenhofer [14], Huang and Yang [15], Campbell, Lo and MacKinlay [5]), Lo’s modified R/S method does not appear to provide the “ultimate” test for long-range dependence. Indeed, Teverovsky, Taqqu and Willinger [33] identify a number of problems associated with Lo’s method and it’s use in practice. Among the most relevant findings, they show that Lo's method has a strong preference for accepting the null hypothesis of no long-range dependence, irrespective of whether long-range dependence is present in the data or not. In practical terms, this result implies that when Lo's method indicates that there is no evidence of long-range dependence in a given data set, it is necessary to re-examine the data to confirm or refute this finding. We illustrate this point with the CRSP data sets that Lo used in his empirical study and where his use of the modified R/S-statistic led him to conclude that there is no evidence of long-memory in the data. By performing an in-depth analysis of these same data sets (and certain transformations thereof), we identify the causes that led to the acceptance of the null hypothesis (no long-range dependence) by Lo [17] and arrive at a much less conclusive picture. Based on this experience as well as on additional evidence, we do not recommend to use Lo's method in isolation, i.e., as the sole technique to test for long-range dependence, but to always rely on a diverse portfolio of graphical and statistical methods for checking for long-range dependence, as described for example in Beran [3], Taqqu, Teverovsky and Willinger [32], Taqqu and Teverovsky [31] and Abry and Veitch [1].

The remaining part of the paper is organized as follows. In Sect. 2, we contrast the classical R/S-method with Lo's modified R/S-statistic and summarize the main findings of and new insights gained from the Monte Carlo simulations performed in Teverovsky, Taqqu and Willinger [33]. In view of this new evidence, we (i) revisit in Sect. 3 the CRSP data sets of daily stock returns that Lo used in his empirical study of long-memory in stock returns, (ii) illustrate the results of our in-depth analysis of the CRSP data sets, and (iii) discuss their practical implications. In Sect. 4, we provide some constructive suggestions for dealing with the long-range dependence issue in financial data.

2 Lo's modified R/S statistic and its properties

In practice, classical R/S-analysis is based on a heuristic graphical approach, originally developed by Mandelbrot and Wallis [25], that attempts to exploit as fully as possible the information in a given historical record. In short, a given sample of N observations is subdivided into K *blocks*, each of size N/K. Then, for each *lag* n, $n \le N$, estimates $R(k_m, n)/S(k_m, n)$ of the R/S statistic given in (1) are computed by starting at the points, $k_m = (m-1)N/K + 1$, $m = 1, 2, \dots, K$, and such that $k_m + n \le N$. Thus, for any given m, all the data points before $k_m = mN/K + 1$ are ignored. For values of n smaller than N/K, there are K different estimates of $R(n)/S(n)$; for values of n approaching N, there are fewer values, as few as 1 when $n \ge N - N/K$. Also note that the $R(k_m, n)/S(k_m, n)$ values corresponding to neighboring values of k_m and n are strongly interdependent; for a given $n > N/K$, the various estimates of $R(n)/S(n)$ involve overlapping observations, and so do the estimates when evaluated at different lags but for a fixed starting point k_m. The graphical R/S approach consists then of calculating the estimates $R(k_m, n)/S(k_m, n)$ for logarithmically spaced values of n, and plotting $\log(R(k_m, n)/S(k_m, n))$ vs. $\log(n)$, for all starting points k_m. This results in the *rescaled adjusted range plot*, also known as the *pox plot of* R/S. This type of plot is displayed in Fig. 1.

In contrast, Lo [17], instead of considering multiple lags, focuses only on lag $n = N$, the length of the series. Furthermore, instead of simply using the sample standard deviation, S, to normalize R, he uses a weighted sum of autocovariances, namely,

$$S_q(N) := \left(\frac{1}{N} \sum_{j=1}^{N} (X_j - \bar{X}_N)^2 + \frac{2}{N} \sum_{j=1}^{q} \omega_j(q) \left[\sum_{i=j+1}^{N} (X_i - \bar{X}_N)(X_{i-j} - \bar{X}_N) \right] \right)^{1/2}, \tag{2}$$

where $\bar{X}_N$ denotes the sample mean of the time series, and the weights $\omega_j(q)$ are given by

$$\omega_j(q) = 1 - \frac{j}{q+1}, \quad q < N.$$

Lo then defines the modified R/S-statistic, $V_q(N)$, by setting

$$V_q(N) = N^{-1/2} R(N)/S_q(N), \tag{3}$$

with R given in (1). Since

$$\lim_{N \to \infty} P\{V_q(N) \in [.809, 1.862]\} = 0.95, \tag{4}$$

(see Lo [17] and Teverovsky, Taqqu and Willinger [33]), Lo uses the interval $[0.809, 1.862]$ as the 95% (asymptotic) acceptance region for testing the null hypothesis

$$H_0 = \{\text{no long-range dependence, i.e., } H = 0.5\}$$

against the composite alternative

$$H_1 = \{\text{there is long-range dependence, i.e. } 1/2 < H < 1\}.$$

The main innovation here is, of course, using $S_q(N)$ in the denominator instead of $S_0(N) = S$ as in the classical R/S-statistic. Lo's motivation for using $S_q(N)$ is to normalize R by a quantity that compensates for the "extra" short-range dependencies that may be present in a given data set. If a time series has both short- and long-range dependence, Lo's statistic V_q should typically fall outside the confidence interval (4), implying correctly that long-range dependence is present. Notice that unlike the graphical R/S-method, which usually provides a rough estimate of the Hurst parameter H, Lo's method only indicates whether long-range dependence is present or not.

Lo's results are asymptotic, in that they assume $N \to \infty$ and $q = q(N) \to \infty$. In practice, the sample size N is finite, and the question naturally arises as to whether there is a "right" choice of $q(N)$. As discussed in detail in Teverovsky, Taqqu and Willinger [33], the right choice of q in Lo's method is essential, because it influences both the actual size of the test, $P(\text{reject } H_0 \mid H_0)$, and its power, $P(\text{accept } H_0 \mid H_1)$. In particular, we show in [33] that for large q-values,

$$V_q(N) \simeq q^{1/2-H}. \tag{5}$$

That is, as q increases, the test statistic V_q decreases (for $H > 1/2$) and, for large enough q, it will be well within the confidence interval for the null hypothesis, i.e., $V_q \in [.809, 1.862]$ (see Eq. (4)). This finding points to a serious problem when relying on V_q as the sole indicator for whether or not a given data set is consistent with long-range dependence and led us to perform a detailed Monte Carlo simulation study, the results of which are described in [33]. Using synthetically generated "purely" long-range dependent time series, i.e., fractional Gaussian noise (FGN) with Hurst parameter $.5 < H < 1$, "purely" short-range dependent sequences, i.e., ARMA(p, q) processes, as well as hybrid short-range and long-range dependent processes, i.e., fractional autoregressive-moving average stationary time series FARIMA(p, d, q) with $d = H - 1/2, 0 < d < .5$, we conclude in [33], that while Lo's modified R/S-statistic represents a theoretical improvement over the classical rescaled adjusted range statistic, its practical application requires care and has a number of problematic features. Most importantly, we found a strong dependence between the outcome of the test (based on the test-statistic V_q) and the choice of the truncation lag q, with a pronounced bias toward accepting the null hypothesis of no long-range dependence for large q's. This happens even in ideal scenarios of "purely" long-range dependent data with high values of the Hurst parameter. Moreover, when interpreting a region of "stability", where the test statistic V_q remains fairly constant as q changes, as a sign that the statistic can be trusted, then the conclusion will almost always be to accept the null hypothesis of no long-range dependence – the only time that V_q enters a stable region is when it is already well within the acceptance region of the null hypothesis. Based on these findings, our recommendations are to always rely on a wide range of different q-values and associated V_q-values. Moreover, one ought never to use Lo's method in isolation, but always in conjunction with other graphical and statistical techniques for checking for long memory, especially when Lo's method results in accepting the null hypothesis of no long-range dependence.

3 Common stock returns and long-range dependence

We start with a brief overview of past and more recent studies of the long-range dependence phenomenon in common stock returns. Then we revisit the data from the Center for Research in Security Prices (CRSP) *daily stock returns* files that were used by Lo [17] and present our own in-depth analysis that was motivated by the new evidence reported in [33] and summarized in Sect. 2.

3.1 Empirical evidence

The Black-Scholes model introduced by Black and Scholes [4] and Merton [28] has become synonymous with modern finance theory. It assumes that the dynamics of stock prices is well described by exponential Brownian motions; that is, stock price returns behave like a sequence of i.i.d. Gaussian random variables. While it is commonly accepted that the Gaussian nature of the marginal distribution of asset returns is a mathematical convenience that is not consistent with empirical stock price returns (e.g., [19, 8, 29]), the dependence structure of stock price returns has been at the center of intense scrutiny for the last 30 or more years. Early investigations into the dependence structure of asset returns (e.g., by Fama [9, 10], who used simple significant tests for checking whether or not the first few autocorrelation coefficients are close to zero) concluded that successive returns can be assumed to be independent; i.e., stock price returns follow a random walk. Later on, Lo and MacKinlay [18] revisited this random walk hypothesis; after a careful analysis of market returns from a 25-year period (1962-1987), Lo and MacKinlay found substantial short-range dependence in the data and strongly rejected the hypothesis that asset returns are i.i.d. .

More recently, motivated by the empirical findings of of long-term memory in common stock returns by Mandelbrot [20] and Greene and Fielitz [13] (see Sect. 1), Lo [17] re-examined the question of long-run memory in asset returns – using his modified R/S-statistic $V_q(N)$ considered in Sect. 2 and the CRSP daily stock returns from 1962 to 1987, a time series of length approximately $N = 6400$. While there is overwhelming evidence that the series of *absolute values* of the CRSP daily stock returns exhibit long-range dependence (R/S gives $H = 0.85$), Lo was mainly interested in the question whether the series itself is consistent with long-range dependence. Analyzing the entire series, as well as fractions of it ($1/2$ and $1/4$ of the original series), and choosing truncation lags $q = 90, 180, 270, 360$ (note that these are relatively large values of q for such a short series; see [33]), Lo's main finding is that the daily stock returns do not exhibit long-range dependence because the values of $V_q(N)$ are within the 95% confidence interval (4). Moreover, Lo observed that the values of the test-statistic V_q do not change much for $q = 90, 180, 270, 360$ and are relatively "stable", which he takes as strong supporting evidence that the test statistic can be trusted, and that – contrary to the findings in [13] – the data are not consistent with long-range dependence. In fact, Lo attributes the findings in [13] to the fact that the classical R/S-method is sensitive to the presence of short-range dependence and concludes that traditional short-range dependent models are adequate to describe actual stock returns.

3.2 The CRSP data revisited

Based on our experience with the modified R/S-statistic V_q reported in [33] and summarized in Sect. 2, the claim made by Lo [17] that the CRSP stock return indices show no evidence of long-range dependence must be reexamined. On the one hand, Lo relies in his analysis of the CRSP data exclusively on the modified R/S-statistic, and we have seen earlier that in isolation, V_q has a strong preference for accepting the null hypothesis of no long-run memory, even if the data at hand are "truly" long-range dependent. On the other hand, we have also mentioned in Sect. 2 that Lo's argument in support of the "correctness" of the null hypothesis, i.e., the existence of a region of "stability" where the test-statistic V_q does not change much as the truncation lag q varies, has its definite downside: The only time V_q is relatively stable as q changes is when q is large enough for the cross-over from the alternative to the null hypothesis to occur and when V_q is well within the acceptance region of the null. Following the recommendations given in [33], further investigations of the CRSP data are required in order to minimize the possibility of a Type II error, i.e., of wrongly accepting the null hypothesis of no long-range dependence.

In the following, we re-examine some of CRSP data used by Lo in his study, identify the causes that led to a "clear" acceptance of the null by Lo, and provide empirical evidence which suggests that some of the CRSP data are, in fact, consistent with long-range dependence; i.e., Lo's claim of no evidence of long-run memory in the data may well be due to the fact that the modified R/S-statistic cannot, in general, guarantee a low Type II error probability. In particular, we focus here on the daily CRSP data for the equal-weighted (EW) returns indices (1962-1987), a time series of length $N = 6409$. (We also analyzed the daily CRSP data for the value-weighted (VW) returns indices; performing a similar analysis as illustrated below for the EW indices, we found no evidence of long-run memory in the VW data, in full agreement with Lo's findings in [17].) A plot of the EW time series is shown in Fig. 1 a, and a preliminary analysis of the data set (not shown here) indicates that, for all practical purposes, the time series can be considered stationary, with a marginal distribution that has heavier than Gaussian tails, but still with finite variance – a quick check (based on qq-plots) of the aggregated series reveals that the marginals appear to be in the domain of attraction of a Gaussian. Fig. 1 b depicts the pox plot of R/S corresponding to the EW series and results in an estimate of the Hurst parameter H of about 0.62 (using the shaded region).

Returning to the modified R/S-statistic V_q, Fig. 2 a shows a plot of V_q vs. q, for a range of q-values (solid line), with the approximate 95% confidence intervals for the hypothesis of no long-range dependence (dashed lines). As can be seen, for small q's, the values of the test-statistics fall outside the 95% confidence region (e.g., $V_0 = 2.63$), which can be taken as evidence for long-run memory in the data. However, for q-values around 10, V_q crosses over into the acceptance region of the null hypothesis and remains fairly stable and well within this region for q-values larger than about 50. This very observation was used by Lo (he chose q-values of $90, 180, 270, 360$) as convincing evidence that the data are not consistent with long-range dependence and that the test-statistic V_q can be trusted. To question this conclusion of Lo, Fig. 2 (a) also depicts the V_q-vs-q-plots for two random realizations of FGN with $H = 0.6$ (dotted lines). Clearly, the EW index is fully consistent with the simulated curves which, at the very least, gives rise to a much less conclusive picture concerning the presence or absence of long-range dependence in the EW index than is portrayed in [17]. Indeed, from the plot on the right side of Fig. 2 which shows the $\log - \log$ version of the V_q-vs-q- plot for the EW index, we observe a linear behavior over a significant

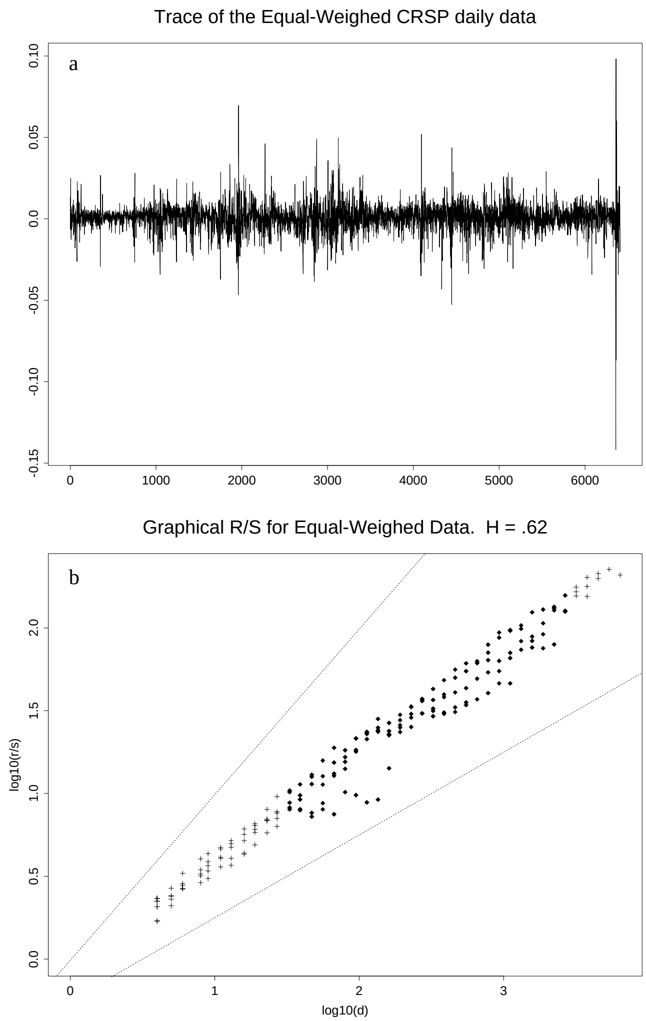

Fig. 1. Time series of the daily CRSP data for the equal-weighted returns indices, from 1962–1987, $N = 6409$ **a**, and corresponding pox plot of R/S **b**. The straight lines serve as reference: their slopes are 0.5 and 1.0 respectively

portion of the range, resulting, using (5), in an H-estimate of about 0.60. Similar results (not shown here) hold for subsets of the EW series.

On the log-log plot of V_q, we can see that at approximately q=50, the plot flattens out. This phenomenon can also be seen in some of the sample FARIMA series in Teverovsky, Taqqu and Willinger [33]. In an ideal long-range dependent series of infinite length, this would not occur. However, in real life, the correlations at very large lags are so small that they are very sensitive to slight

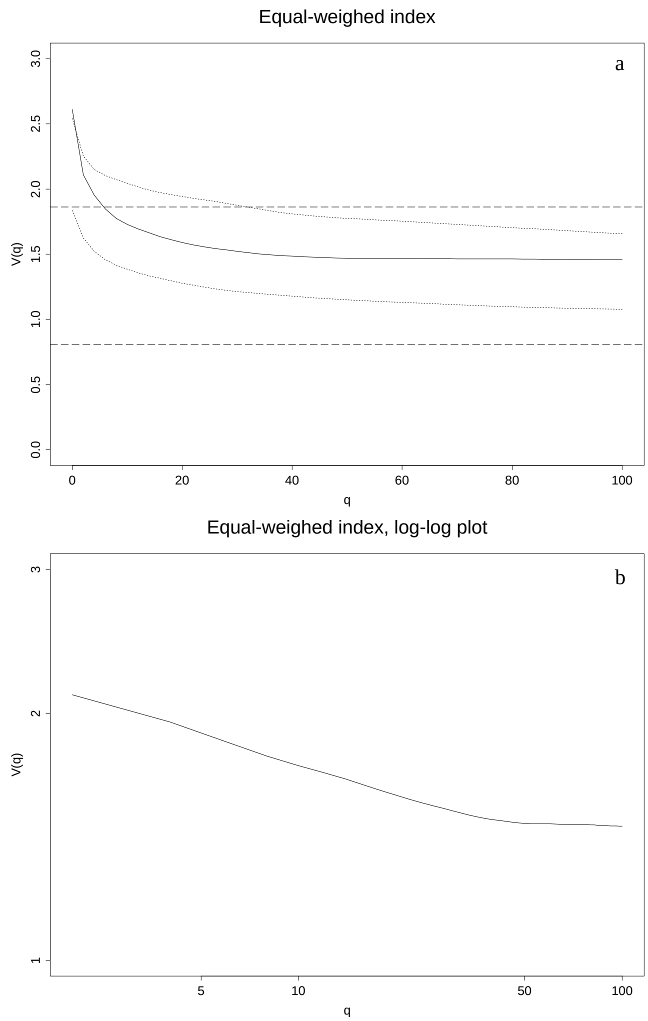

Fig. 2. a. V_q for the EW series (solid line) and 2 realizations of FGN($H = 0.6$) (dotted lines). **b.** V_q for the EW series on log – log scale

deviations. Thus, there will always be a q such that the log-log plot will have negligible fluctuations beyond it.

Next we used several other graphical and statistical methods for estimating H (for a detailed account of the different methods, see Taqqu, Teverovsky and Willinger [32]) and obtained values ranging consistently from 0.55 to 0.65. Especially noteworthy are the results obtained from using Whittle's approximate MLE method [3, Chapter 5]. When assuming an underlying FARIMA$(1, d, 1)$ model, the resulting H-estimate is about 0.63, with AR and MA coefficients $[(-0.16, -0.40)]$ that are significantly different from zero. When aggregating the EW series over blocks of size 10 and 20, respectively, and assuming now an underlying FGN model, we obtain approximately the same H-estimate: for blocks of size 10, the estimated H-value is about 0.63, with approximate 95% confidence intervals of $(0.58, 0.68)$; for blocks of 20, $H \approx 0.58$, with approximate 95% confidence intervals of $(0.51, 0.65)$. Note that by aggregating the data, the resulting marginals become more Gaussian (see above) and any "extra" short-range dependence tends to become insignificant.

Another way of removing "extra" short-range dependence and isolating hidden "pure" long-range dependence effects that may be present in a given data set is to partition the time series into non-overlapping blocks of size m, e.g. $m = 10, 20, 50$, and "shuffle" the observations *within* each block; that is, pick a random permutation of the time indices (shuffling in the context of long-range dependent data was used, for example, in Erramilli, Narayan and Willinger [7]). Intuitively, the effect of such a shuffling experiment is to destroy any particular structure of the autocorrelation function below lag m but to leave the high lags (i.e., low frequencies) essentially unchanged. Analyzing these shuffled data (using the above-mentioned methods, including the graphical R/S-method, the Whittle's estimate and the modified R/S-statistic), we obtain results that are very similar to those found for the original, un-shuffled series. This observation provides, by itself, strong empirical evidence that our findings for the EW series are not likely due to the presence of any particular short-range dependence structure but are the result of some "genuine" long-range dependence effect in the data. To emphasize this point even more, we then performed a shuffling experiment, where we shuffled the blocks but left the observations within each block intact. Such a shuffle has the effect of maintaining the short-range dependence but destroying any potential long-run memory in the data. For block sizes of 40 and above, the results consistently indicate the absence of any long-run memory, with H-estimates that are consistently lower than those obtained for the original series and are typically close to $H = 0.5$, regardless of the method that was used.

To summarize the results of our analysis of the CRSP EW data, we have obtained strong empirical evidence that Lo's acceptance of the null hypothesis of no long-range dependence in this time series is due to the strong preference of the modified R/S-method for accepting the null, even if there is "genuine" long-range dependence present in the data. Moreover, when applying a variety of different time domain-based as well as frequency domain-based graphical and statistical methods for checking for long-run memory to the EW series and certain transformations thereof, the results consistently indicate the presence of some – though rather weak (i.e., H-values of around 0.60) – long-range dependence. We agree with Lo, however, that it is typically difficult to distinguish between long-range dependence with low H and some types of short-range dependence.

4 Conclusion

Based on our findings described in Teverovsky, Taqqu and Willinger [33] about the properties of Lo's modified R/S-statistic, we strongly advise against its use as the sole technique for testing for long-memory in a given data set and advocate instead the use of a diverse portfolio of time domain-based and frequency domain-based graphical and statistical methods described in Taqqu, Teverovsky and Willinger [32] and Taqqu and Teverovsky [31]. These include the graphical R/S-method, the modified R/S-statistic with corresponding V_q-vs-q-plots, and Whittle's approach. Although Lo' modified R/S-statistic is a conceptual improvement over the classical R/S-statistic, it should not be used blindly nor in isolation. In particular, an acceptance of the null hypothesis of no long-range dependence based on the modified R/S-statistic should never be viewed as the "final word", mainly because of the serious difficulties that the test-statistic V_q has in identifying "genuine" long-range dependence (e.g., FGN with high H-values). Instead, an acceptance of the null based on the test-statistic V_q should always be accompanied and supported by further analysis of the data. To illustrate this, we revisited the CRSP daily equal-weighted (EW) returns indices that motivated Lo [17] to develop and use the modified R/S-statistic and led him to believe that there is no evidence of long-range dependence in the data. Upon further analysis of this data set, we noted that the EW time series seem to display long-range dependence, but because the corresponding H-values are typically very low, around 0.60, the evidence is not absolutely conclusive. We also found that Lo's conclusion is basically a result of the overly conservative nature of his proposed test-statistic in rejecting the null hypothesis of no long-range dependence.

Finally, while statistical analyses cannot be expected to provide a definitive answer concerning the presence or absence of long-range dependence in asset price returns, a more revealing but also much more challenging approach to tackle this problem consists of providing a mathematically rigorous physical "explanation" for the presence or absence of the long-range dependence phenomenon in stock returns. However, such a phenomenological approach will require a deeper understanding of the nature of the micro/macro-economic market forces that determine the price movements.

Acknowledgements. We thank Prof. J. Patel and John Neumann for providing us with the CRSP data used in this study.

References

1. P. Abry, D. Veitch: Wavelet analysis of long range dependent traffic. IEEE Transactions on Information Theory **44** (No. 1), 2–15 (1998)
2. A. A. Annis, E. H. Lloyd: The expected value of the adjusted rescaled Hurst range of independent normal summands. Biometrika **63**, 111–116 (1976)
3. J. Beran: Statistics for Long-Memory Processes. New York: Chapman & Hall 1994
4. F. Black, M. Scholes: The pricing of options and corporate liabilities. J. Polit. Econ. **3**, 637–654 (1973)
5. J. Y. Campbell, A. W. Lo, A. C. MacKinlay: The Econometrics of Financial Markets. Princeton, NJ: Princeton University Press 1997
6. N. J. Cutland, P. E. Kopp, W. Willinger: Stock price returns and the Joseph effect: a fractional version of the Black-Scholes model. In: E. Bolthausen, M. Dozzi, F. Russo (eds.) Seminar on Stochastic Analysis, Random Fields and Applications. Boston: Birkhäuser 1995, pp 327–351

7. A. Erramilli, O. Narayan, W. Willinger: Experimental queueing analysis with long-range dependent packet traffic. IEEE/ACM Transactions on Networking **4**(2), 209–223 (1996)
8. E. Fama: Mandelbrot and the stable Paretian hypothesis. J. Business **36**, 420–429 (1963) Reprinted in: P. Cootner (ed.) The Random Character of Stock Market Prices. Cambridge, MA: MIT Press 1964, pp 297-306
9. E. Fama: The behavior of stock market prices. J. Business **38**, 34–105 (1965)
10. E. Fama: Efficient capital markets: A review of theory and empirical work. J. Finance **25**, 383–417 (1970)
11. W. Feller: The asymptotic distributions of the range of sums of independent random variables. Ann. Math. Statist. **22**, 427–432 (1951)
12. C. W. J. Granger: The typical spectral shape of an economic variable. Econometrica **34**, 150–161 (1966)
13. M. T. Greene, B. D. Fielitz: Long-term dependence in common stock returns. J. Financial Econ. **4**, 339–349 (1977)
14. M. A. Hauser, R. M. Kunst, E. Reschenhofer: Modelling exchange rates: Long-run dependence versus conditional heterscedasticity. Appl. Financial Econ. **4**, 233–239 (1994)
15. B. N. Huang, C. W. Yang: The fractal structure in multinational stock returns. Appl. Econ. Lett. **2**, 67–71 (1995)
16. H. E. Hurst: Long-term storage capacity of reservoirs. Trans. Amer. Soc. Civil Eng. **116**, 770–808 (1951)
17. A. W. Lo: Long-term memory in stock market prices. Econometrica **59**, 1279–1313 (1991)
18. A. W. Lo, A. C. MacKinlay: Stock market prices do not follow random walks: Evidence from a simple specification test. Rev. Financial Stud. **1**, 41–66 (1988)
19. B. B. Mandelbrot: The Pareto-Lévy law and the distribution of income. Int. Econ. Rev. **1**, 79–106 (1960)
20. B. B. Mandelbrot: Forecasts of future prices, unbiased markets and martingale models. J. Business **39**, 242–255 (1967)
21. B. B. Mandelbrot: A fast fractional Gaussian noise generator. Water Resources Res. **7**, 543–553 (1971)
22. B. B. Mandelbrot: The Fractal Geometry of Nature. San Francisco: W.H. Freeman 1982
23. B. B. Mandelbrot, M. S. Taqqu: Robust R/S analysis of long-run serial correlation. In Proceedings of the 42nd Session of the International Statistical Institute, Manila (1979) Bulletin of the International Statistical Institute. Vol.48, Book 2, pp. 69-104
24. B. B. Mandelbrot, J. R. Wallis: Noah, Joseph and operational hydrology. Water Resources Res. **4**, 909–918 (1968)
25. B. B. Mandelbrot, J. R. Wallis: Computer experiments with fractional Gaussian noises, Parts 1,2,3. Water Resources Res. **5**, 228–267 (1969)
26. B. B. Mandelbrot, J. R. Wallis: Robustness of the rescaled range R/S in the measurement of noncyclic long-run statistical dependence. Water Resources Res. **5**, 967–988 (1969)
27. B. B. Mandelbrot, J. R. Wallis: Some long-run properties of geophysical records. Water Resources Res. **5**, 321–340 (1969)
28. R. C. Merton: An intertemporal capital asset pricing model. Econometrica **41**, 867–887 (1973)
29. S. Mittnik, S. T. Rachev: Modeling asset returns with alternative stable distributions. Economet. Rev. **12**, 261–330 (1993)
30. L. C. G. Rogers: Arbitrage with fractional Brownian motion. Math. Finance **7**, 95–105 (1997)
31. M. S. Taqqu, V. Teverovsky: Estimating long-range dependence in finite and infinite variance series. In: R. Adler, R. Feldman, M. S. Taqqu (eds.) A Practical Guide to Heavy Tails: Statistical Techniques for Analyzing Heavy-Tailed Distributions. Boston: Birkhäuser 1997
32. M. S. Taqqu, V. Teverovsky, W. Willinger: Estimators for long-range dependence: an empirical study. Fractals **3**(4) 785–798 (1995) Reprinted in C.J.G. Evertsz, H-O Peitgen, R.F. Voss (eds.) Fractal Geometry and Analysis. Singapore: World Scientific Publishing Co. 1996
33. V. Teverovsky, M. S. Taqqu, W. Willinger: On Lo's modified R/S statistic. Preprint, 1997
34. W. Willinger, M. S. Taqqu, A. Erramilli: A bibliographical guide to self-similar traffic and performance modeling for modern high-speed networks. In: F. P. Kelly, S. Zachary, I. Ziedins (eds.) Stochastic Networks: Theory and Applications. Oxford: Clarendon Press 1996, pp. 339–366