The Economic Journal, **105** (*July*), 864–880. © Royal Economic Society 1995. Published by Blackwell Publishers, 108 Cowley Road, Oxford OX4 1JF, UK and 238 Main Street, Cambridge, MA 02142, USA.

NONLINEAR DYNAMICS IN REAL-TIME EQUITY MARKET INDICES: EVIDENCE FROM THE UNITED KINGDOM*

A. Abhyankar, L. S. Copeland and W. Wong

This paper tests for the presence of nonlinear dependence and chaos in real-time returns on the U.K. FTSE-100 Index, using a six month sample of about 60,000 observations. Since there is clear evidence of nonlinearity, we follow other researchers in this field by applying the same tests to the residuals from a GARCH process fitted to the data, in order to find out whether or not the nonlinearity can be explained by this type of model. In the event, our results suggest that GARCH can explain some but not all of the observed nonlinear dependence.

In the past few years, a large literature has appeared on nonlinearity in finance and economics. At a theoretical level, it has been shown that even very simple economic models often involve a rich variety of dynamic processes, including in some cases the possibility of nonlinear or complex chaotic behaviour for some range of parameter values (see survey in Boldrin (1988)). More recently, the published empirical literature has concentrated on testing economic and financial time series for the presence of nonlinear dependencies using various measures indicative of complex dynamics.

The issues involved in this area are of critical importance, not least in their implications for market efficiency. For example, the presence of a well-behaved nonlinear structure would be inconsistent with market efficiency, at least if accompanied by risk-neutrality and negligible transaction costs. On the other hand, a chaotic process, defined for our purposes simply as one characterised by sensitive dependence on initial conditions,[1] need not necessarily imply the existence of exploitable profit opportunities. In the first place, the complexity of the process may make it impossible for agents to identify, though if researchers can uncover the true model, it could be argued that the market is equally capable of doing so. More importantly, however, sensitive dependence on initial conditions means that knowing the function driving the market price may be insufficient to guarantee a profit, because forecast accuracy may degenerate too rapidly to leave time for profitable trades to be executed.

In order to investigate these empirical questions, researchers in economics have for the most part had to use methods originally developed in the physical

* The authors are grateful to the International Stock Exchange, London and its Quality of Markets Unit for providing the data used here. We would like to thank S. Ellner, D. W. Nychka and A. R. Gallant for use of their LENNS software and K. Briggs for his Lyapunov exponent estimation code. The Manchester Computing Centre generously provided time on their Fujitsu VPX 240/10 vector processor. The paper has benefited from comments by Professors Maurice Priestley, Subba Rao and Xubin Zeng, participants in the Annual Conference of the Society of Economic Dynamics and Control, UCLA, June 1994, the European Institute for Advanced Studies in Management Conference on Empirical Methods in Finance, Brussels, May 1994 and seminars in a number of departments in the United Kingdom. The second author acknowledges with gratitude the support of the Nuffield Foundation.

[1] See Devaney (1989) for a more formal definition. Baumol and Benhabib (1989) discuss the implications of chaos in an economic context.

sciences for analysing the relatively large and less noisy data sets available in those areas. Applying the same methods to financial time series has inevitably involved the use of relatively low frequency (e.g. daily) data obtained over a number of years so as to provide sufficient observations (Frank and Stengos, 1989; Hsieh, 1989; Scheinkman and LeBaron, 1989). So far, very few studies have used the high frequency data which are nowadays becoming available from the equity, foreign exchange and derivatives markets. This is unfortunate, since higher frequency datasets offer several important advantages: larger sample sizes, more potential for observing microstructural effects and, given the shorter time period involved, other things being equal, a greater likelihood that the underlying process has remained stationary over the sample period.

The present paper has the advantage in this regard in relying on a dataset consisting of the 60,000 minute-by-minute real time returns on the UK FTSE-100 Index for the first 6 months of 1993. Moreover, the dataset has another attractive feature, insofar as it provides a relatively clean measure of the index return, since it is based on the mid-quote prices of market-makers who are obliged, on the London Stock Exchange, to quote firm two-way prices. In each of our tests we are able to examine evidence of nonlinearity at a range of different frequencies (1-, 5-, 15-, 30-, and 60-min).

There are a number of tests for nonlinearity in the literature, including non-parametric tests in the frequency and time domain and parametric and semi-parametric tests in the time domain. In this paper we implement two well-known tests, the Hinich (1982) bispectrum and the BDSL (Brock *et al.* 1987). We then proceed to test for chaos using two different methods for estimating the Lyapunov exponent, the neural nets approach of Nychka *et al.* (1992), and the method of higher-order local neighbourhood-to-neighbourhood mappings (Briggs, 1990; Brown *et al.* 1991; Zeng *et al.* 1992).

Our results clearly indicate the presence of nonlinear dependence in high frequency FTSE returns. However, we find very little evidence to support the view that returns could be characterised by a low-dimension chaotic process. Specifically, while both the Hinich and BDSL tests reveal significant nonlinearities at all frequencies, the estimates of the Lyapunov exponents are in some cases unstable (i.e. highly sensitive to the embedding dimension) and in other cases consistently negative, depending on the estimation methodology employed. We conclude that either the process is low-dimensional chaos contaminated by large amounts of noise, or the attractor is of very high or, indeed, infinite dimension.

The paper is organised as follows. In Section I, we briefly review earlier empirical work. In the next section, we describe the tests used in this paper. In Section III, we give some details of our dataset and present our results and Section IV contains our conclusions.

I. REVIEW OF PREVIOUS WORK ON FINANCIAL MARKETS

From the substantial literature which has appeared in the last ten years or so, there is broad agreement that a nonlinear structure is to be found in financial series. This statement applies equally to exchange rates (e.g. Kodres and Papell, 1991; Kräger and Kügler, 1992), to stock market series (Hsieh, 1991; Philippatos *et al.* 1993), and to gold and silver prices (Frank and Stengos, 1989).[2] In particular, BDSL tests almost invariably reject independence.

The published evidence on chaos is more mixed. By and large, there appears to be more evidence to support a low dimensional structure in the US stock market index than in exchange rate series (compare, on the one hand, Mayfield and Mizrach (1989), Vaidyanathan and Krehbiel (1992), with Hsieh (1989; 1993), Tata (1991)), though for the most part this conclusion is based on the results of correlation dimension tests rather than direct Lyapunov exponent (L.E.) estimates (but see Vassilicos *et al.* 1992; Eldridge and Coleman, 1993).

Even where Lyapunov exponent estimates are presented, there is no rigorous criterion for assessing their significance. For the most part, moreover, the estimates were derived by direct fitting of the state space mapping (Wolf *et al.* 1985), rather than by working with the Jacobian, which is preferable in the context of noisy time series. In general, the results suggest that L.E. estimation may be highly sensitive to the algorithm used, a suspicion which is borne out by the results presented in Section III below.

With this caveat in mind, it is unclear how much importance should be attached to estimates of positive L.E.'s. Certainly, it is unsurprising that a number of researchers report dimension estimates of 6–7, at or near the upper limit of the dimension which can be estimated with datasets of the size used in most of the work reported in the table. It should also be borne in mind that L.E. estimates are known to be subject to upward bias in cases where the dataset is contaminated by noise (McCaffrey *et al.* 1992) or equally where the number of observations is inadequate.

There remain a number of questions outstanding.

(1) Is the apparent nonlinearity capable of being explained by any of the time series models prevalent in the finance literature, in particular the ARCH family: GARCH, IGARCH, EGARCH etc.? Although there is evidence that GARCH or its variants can account for much, though probably not all of the nonlinear structure in exchange rates (Hsieh, 1989; Kodres and Papell, 1991), the situation as regards stock markets is less clear.

(2) Are the answers to these and other questions dependent on the frequency with which the data are observed? Or is the degree of nonlinearity and/or apparent sensitivity to initial conditions greater at higher frequencies? The results to be reported in this paper will confront this issue directly. It is noticeable that while early work relied on relatively small datasets made up of no more than a few hundred or so daily or weekly observations, researchers in the last few years have increasingly used very high frequency data (see especially Vassilicos *et al.* 1992). In part, this trend simply exploits the fact that high frequency datasets have become available. But, more importantly, it also reflects a widespread recognition that only with very large amounts of data is

[2] Note that the literature includes a number of studies of future prices. However, there is no clear difference between the results reported for futures and spot prices.

it ever going to be possible to arrive at reliable estimates of the sensitivity of an observed stochastic process to disturbances, as measured by the Lyapunov exponent (see next section).[3] Moreover, given the structural changes in financial markets during the 1980s, it is unlikely that stationarity could have been expected to hold over anything more than a few months, or a year or two at most. In practice, therefore, a large dataset must involve high-frequency observations.[4]

At the same time, it is noticeable that, at least as far as exchange rates are concerned, there appears to be little qualitative difference between the results achieved with the real time dataset of Vassilicos *et al.* (1992), and those published by researchers using weekly or daily returns.

II. TESTS FOR NONLINEARITY AND CHAOS

In this paper, we implement two nonlinearity tests: the Hinich (1982) bispectrum and BDSL tests. We then proceed to estimate the Lyapunov exponent, following the neural nets approach of Nychka *et al.* (1992), and the nearest-neighbour methods of Zeng *et al.* (1992) and Briggs (1990).

II.1. *The Hinich (1982) Bispectrum Test*

In heuristic terms, the bispectrum is a generalisation of the spectral density, with covariance terms $\mathrm{E}(x_t x_{t+\tau})$ replaced by third-order cumulants $\mathrm{E}(x_t x_{t+\tau} x_{t+\kappa})$. Brillinger (1965) shows that if a 3rd-order stationary time series $\{x(t)\}$ is IID-linear, then the skewness function:

$$\Gamma^2(\omega_1, \omega_2) = \frac{B(\omega_1, \omega_2)^2}{f(\omega_1) f(\omega_2) f(\omega_1 + \omega_2)}, \tag{1}$$

where $f(\omega_1)$ and $B(\omega_1, \omega_2)$ are respectively the 2nd- and 3rd-order spectra of $\{x(t)\}$, is constant over all frequencies (ω_1, ω_2) in the principal domain given by:

$$\Omega = (\omega_1, \omega_2) : 0 < \omega_1 < 0{\cdot}5, \omega_2 < \omega_1, 2\omega_1 + \omega_2 < 1. \tag{2}$$

Furthermore, if $\{x(t)\}$ is Gaussian, then the expression in (1) is zero in Ω.

These results were first applied to tests for nonnormality and nonlinearity by Subba Rao and Gabr (1980). Using results on the asymptotic distribution of the sample bispectrum, Hinich (1982) developed a method based on robust test statistics which is both simpler and more stable, and has been widely studied and applied (Ashley *et al.* 1986; Brockett *et al.* 1988; Hinich and Patterson 1985, 1989). We follow the procedures in Ashley *et al.* (1986). In particular, we base our tests on their interquartile range and 80% quantile statistics, with their recommended window size of $M = N^{\frac{1}{2}}$, where N is the number of observations available.

Finally, we note that the fact that bispectrum-based tests have been shown

[3] Thus, Vassilicos *et al.* (1992) argue that a minimum size of dataset required for work of this type is around 5,000 observations.

[4] This is not to say that high-frequency data are necessarily stationary. For example, we are grateful to a referee for pointing out that microstructural effects are likely to induce nonstationarities into intra-daily returns (Brock and Kleidon, 1992). Our contention is simply that these nonstationarities are likely to be far less of a problem than those present in longer period datasets.

to be invariant to linear filtering (Ashley *et al.* 1986) means that they nicely complement the BDSL test, which is of course sensitive to linear structure.[5]

II.2. *The BDSL Test*

Brock *et al.* (1987) have developed a powerful test for independence and identical distribution based on the correlation integral. Given a time series of N observations $\{x(t), t = 1, 2, 3, \ldots N\}$, the correlation integral is defined as:

$$C(m, \epsilon, N) = I[(t, s) : \|X_t^m - X_s^m\| < \epsilon]/N^2, \tag{3}$$

where $X_t^m = [x(t), \ldots, x(t-m+1)]$, $\|.\|$ is the L_∞ norm on R^m and $I[.]$ denotes the number of elements in the set. Intuitively, the correlation integral measures the proportion of embedded vectors of dimension m lying within the ϵ-neighbourhood of an initial embedding, X_t. Under modest regularity conditions $C(m, \epsilon, N)$ has a limit $C(m, \epsilon)$ as $N \to \infty$. Now if $\{x(t)\}$ is IID, then we have

$$C(m, \epsilon) = C(1, \epsilon)^m, \tag{4}$$

from which we define the BDSL test statistic:

$$(m, \epsilon, N) = \sqrt{\frac{T}{V}}[C(m, \epsilon, N) - C(1, \epsilon, N)^m], \tag{5}$$

which converges in distribution to N(0, 1) as $N \to \infty$. The variance V can be consistently estimated from the data, as detailed in Brock *et al.* (1987).

II.3. *Lyapunov Exponent Tests*

It was shown in Takens (1981) that the underlying dynamics of a process can be understood in terms of the phase space reconstruction. Thus a noisy (scalar) time series $\{x(t), t = 1, 2 \ldots\}$ can be rewritten in state space form as:

$$X_t = F(X_{t-1}) + \boldsymbol{\varepsilon}_t, \tag{6}$$

where $X_t = \{x(t), x(t-L), \ldots x[t-(d-1)L]\}$, d is the embedding dimension, L the time delay, $\boldsymbol{\varepsilon}_t = (e_t, 0 \ldots 0)$ represents the stochastic component of the process, with $\{e_t\}$ a sequence of IID random variables, and F a $R^d \to R^d$ function satisfying some general regularity conditions.

Now, given two initial state vectors $X_0^{(1)}, X_0^{(2)}$ sufficiently close together, then after one time point:

$$\|X_1^{(2)} - X_1^{(1)}\| \approx \|J_0(X_0^{(2)} - X_0^{(1)})\| \tag{7}$$

where J_0 is the $d \times d$ Jacobian matrix of partial derivatives of F evaluated at $X_0^{(2)}$. Then the Lyapunov exponent (L.E.) of the system can be defined as:

$$\lambda = \lim_{t \to \infty} \frac{1}{t} \ln \|J_{t-1} \cdot J_{t-2} \cdot \ldots \cdot J_0\|. \tag{8}$$

One operational definition of chaos is a bounded system with $\lambda > 0$, since if this condition is satisfied, trajectories which start at two almost identical state vectors will diverge exponentially as time passes (the property known as 'sensitive dependence on initial conditions').

[5] Ashley *et al.* (1986) also use Monte Carlo methods to show that the distribution of the bispectral skewness statistic is unlikely to be greatly affected by estimation error in the linear filter.

The direct L.E. estimation method of Wolf *et al.* (1985) involves averaging observed divergence rates which, as approximations of the LHS of (7), will tend to grow without limit if the series is chaotic. However, this direct approach is less appropriate when the process is assumed to be contaminated with noise, in which case λ estimates derived in this fashion are liable to be biassed upward (McCaffrey *et al.* 1992).

Instead, we use Jacobian methods based on estimates of (7) and (8). One advantage of this approach is that the approximation in (7) can be improved by the introduction of higher-order terms in the Taylor expansion. Moreover, the noise in the underlying process (6) can be smoothed out by using additional near-neighbours in the estimation algorithm, an approach which has been successfully implemented by Briggs (1990), Brown *et al.* (1991) and Zeng *et al.* (1992).

In the work reported here, we first follow Zeng *et al.* (1992) in choosing near neighbours from within a 'shell' (i.e. the zone between two spheres) rather than from a ball. This method is claimed to be an improvement in so far as it minimises the effect of noise on the estimates, though in the present case, it was especially attractive because it greatly reduced the difficulties presented by the large number of zero returns in the higher frequency datasets (see Section III).

An alternative non-parametric regression approach which we implement here involves approximating the function F in (6) by using a single hidden-layer feed-forward neural network with a single output, of the general form:

$$\hat{f}(\boldsymbol{X}_t) = \beta_0 + \sum_{j=1}^{q} \beta_j G(\gamma_j' \boldsymbol{X}_t + \mu_j), \tag{9}$$

where we have introduced the logistic distribution function:[6] $G(u) = e^u/(1+e^u)$ (see McCaffrey *et al.* 1992 and Nychka *et al.* 1992). Unlike the nearest-neighbour algorithms, this approach avoids the 'curse of dimensionality'. It is also an advantage to be able to obtain Bayesian Information Criterion values for each function approximation, giving some kind of numerical indication of the reliability of the L.E. estimates. In fact, the results in Nychka *et al.* (1992) suggest this method works reasonably well on noisy systems even when the number of observations is far smaller than is the case here.

III. THE RESULTS

The data set used in this paper consisted of real-time observations of the return on the FTSE-100 Index of stocks quoted on the London International Stock Exchange over the period from 4 January 1993 to 30 June 1993. The index is

[6] In practice, Ellner *et al.* (1991) use the approximation:

$$\Psi(x) = \frac{x\left(1+\dfrac{|x|}{2}\right)}{2+|x|+\dfrac{x^2}{2}}.$$

However, in our case, the optimisation routine converged faster using (9).

recomputed every minute, so that at the highest (i.e. 1-min) frequency, our dataset includes a total of about 60,000 points.[7] However, our tests were also conducted at lower frequencies. Table 1 gives a selection of descriptive statistics

Table 1

Descriptive Statistics for FTSE Returns

Frequency	Sample size N	Unique values NC	No. of zeroes NO	Mean	s.d.	Skewness	Kurtosis
60-min	992	981	10	2·860E-06	1·980E-03	0·843	8·702
30-min	1,984	1,942	34	1·430E-06	1·320E-03	0·201	4·537
15-min	3,968	3,737	152	7·150E-07	8·720E-04	0·194	8·247
5-min	11,904	8,975	1,292	2·380E-07	3·840E-04	−0·226	16·355
1-min	59,520	14,722	22,491	4·770E-08	1·110E-04	−0·993	293·082

Notes: N is number of observations in sample; NC is N less the number of repeated observations in sample; NO is number of zero observations in sample. Both skewness and kurtosis statistics centred on zero.

on the data set sample at a range of frequencies from hourly to minute-by-minute. A number of points are worth noting.

First, looking at the columns headed N (sample size), NC (number of unique i.e. non-repeated values) and NO (number of zero values), it is clear that there is a potential problem with the maximal (i.e. 1-min) data set. Our concern here is with the return on the FTSE-100, defined as the log difference in the level of the index. It follows that on every occasion when the index remained unchanged over the observation interval, the return was zero. Obviously, the shorter the observation interval, the more likely it is that the FTSE will remain unchanged. In fact, as can be seen from the table, the return was zero at well over one third of the points in the 1-min data set, and about 10% of the points at the 5-min frequency.

For nearest-neighbour estimation methods, it is also relevant to consider the number of unique (i.e. excluding repeated) values. In fact, over 75% of the maximal data set consisted of non-unique values (zero or otherwise), leaving fewer than 15,000 unique values. By contrast, of the 992 hourly observations, only 11 were repeated, of which all but one were zeroes.

Looking at the remaining columns of Table 1, there is very little obvious sign of mean reversion, insofar as the standard deviation increases (though not necessarily in strict proportion) as the frequency falls. Although the dataset is only mildly skewed, the degree of leptokurtosis is overwhelming, as is usually the case with financial data. Moreover, while it is significantly non-Gaussian in this respect at all frequencies, the extent of the leptokurtosis is vastly greater for the 1-min data set, a finding which is consistent with the well-established conclusion in the literature that deviations from normality are more marked at higher than lower frequencies.

[7] For obvious reasons, the 'overnight' data point (i.e. the first return of the day) was excluded from the data set at each frequency.

Table 2

BDSL Tests for Independence of FTSE Returns

(January–June 1993)

dim	1-hour	30-min	15-min	5-min	1-min
ARMA residuals					
2	3·941	4·912	12·332	51·995	136·453
3	4·627	6·790	14·814	53·585	143·908
4	4·817	8·089	16·177	54·642	147·126
5	4·749	8·726	17·570	56·452	150·426
6	4·188	9·166	19·224	59·197	154·567
7	3·724	9·125	20·994	62·701	159·733
8	3·642	8·810	22·985	66·953	166·244
9	3·853	8·305	24·961	72·247	174·222
10	4·187	7·386	27·066	78·757	183·628
GARCH residuals					
2	2·702	2·450	10·421	9·384	5·771
3	2·896	3·914	12·755	9·969	8·410
4	2·595	4·817	14·089	11·292	13·057
5	1·961	5·092	15·433	13·004	17·899
6	1·092	5·113	16·856	14·847	22·630
7	0·419	4·556	18·338	16·859	27·480
8	0·187	3·853	19·822	18·955	32·203
9	0·226	2·967	21·216	21·036	37·154
10	0·446	1·866	22·560	23·290	42·505

Notes: BDSL statistics are distributed N(0,1) under null hypothesis of IID residuals. $\epsilon/\sigma = 1{\cdot}0$ for results given above. Results for $\epsilon/\sigma = 0{\cdot}75, 1{\cdot}25, 1{\cdot}50$ available from authors.

Nonlinearity (I): *the BDSL Test* (*Table 2*)

Table 2 shows estimated values of the BDSL test for embedding dimensions from 2 to 10.[8] The computed statistics are asymptotically standard normal variates. In fact, in the light of the Monte Carlo studies reported in Brock *et al.* (1991), normality holds well for sample sizes of 1,000 and upwards, and for values of ϵ of between 0·5 and 2·0 standard deviations. It should be noted that the same authors also warn against relying on asymptotic normality for values of N/m of less than 200. Thus, for our hourly data set, this would indicate embedding dimensions no greater than 5. We note also their conclusion that asymptotic normality is robust both to skewness and leptokurtosis.

It is well known, however, that the BDSL test rejects IID for a wide range of processes, including linear ARMA and nonlinear GARCH. Since our concern here is with nonlinearity, an ARMA model was fitted to the data prior to testing, so that the results in the top half of Table 2 refer to the residuals from an ARMA model selected by the Bayesian Information Criterion.

One conclusion can be drawn immediately. In virtually every case, our results are inconsistent with an IID process. This is true at all dimensions and for all frequencies.[9] Moreover, as expected, the higher the frequency, the

[8] The results given in Table 2 are for $\epsilon/\sigma = 1{\cdot}0$. The computations were performed for values of ϵ ranging from 0·5 to 1·5 times the standard deviation, with qualitatively similar results (available from authors).

[9] And at all values of ϵ/σ (not presented in the table). As noted by Brock *et al.* (1991), at high dimensions the results are potentially very sensitive to the chosen value of ϵ/σ.

Table 3

Bispectral Normality and Linearity Tests (Hinich, 1982)

	60-min	30-min	15-min	5-min	1-min
(*a*) Raw data					
n	990	1,980	3,960	11,900	59,500
m	31	45	63	109	243
Normality tests					
csq	335·1	609·0	1,322·0	6884·0	2·06E+06
D.F.	140	294	600	1,908	9,760
p	0	0	0	0	0
Hinich linearity tests					
λ	2·79	2·14	2·41	5·22	420·15
D.F.	2	2	2	2	2
z1	2·328	0·852	0·179	3·318	125·696
p1	0·01	0·197	0·4288	0·0005	0
z2	2·264	1·799	2·303	6·525	661·44
p2	0·0118	0·036	0·0106	0	0
(*b*) GARCH residuals					
n	980	1,980	3,960	11,900	59,500
m	31	45	63	109	243
Normality tests					
csq	285·5	489·3	972·9	3,439	38,180
D.F.	140	294	600	1,908	9,760
p	0	0	0	0	0
Hinich linearity tests					
λ	2·08	1·33	1·24	1·6	5·82
D.F.	2	2	2	2	2
z1	0·423	−0·205	0·096	−0·801	16·306
p1	0·3363	0·581	0·4619	0·7885	0
z2	156	1·587	1·384	1·388	16·489
p2	0·4379	0·0563	0·0832	0·0825	0

Notes: n, Number of data points used; *m*, width of spectral window (= SQRT(n) + 1); csq, chi-sq statistic for Gaussianity tests; D.F., degrees of freedom; p, p-value of Gaussianity test statistics; λ, non-centrality parameter; z1, N(0,1) test statistic for inter-quartile-range (IQR) linearity test; p1, p-value of IQR linearity test; z2, N(0,1) test statistic for 80%-quantile-range (R80) linearity test; p2, p-value of R80 linearity test.

further into the tail of the distribution are the computed z-statistics. But there is little reason to doubt the basic result that the data are almost certainly characterised by nonlinear dependence at high frequencies, and probably at lower frequencies too.

The bottom half of the table repeats the same tests on the residuals from a GARCH(1,1) process. While the null hypothesis of IID residuals is unequivocally rejected in virtually every case, the values of the BDSL statistic are very noticeably reduced at all dimensions. The improvement is particularly dramatic in the 1-min case. To that extent, our results endorse the conclusions of Hsieh (1990) and others.

Nonlinearity (II): the Hinich (1982) Test (Table 3)

Our data set is well suited to spectral tests, not least because the use of a wide range of frequencies reduces the aliasing problem discussed in Hinich and Patterson (1989).

Table 3 presents the results of tests based on the bispectrum of the raw data, at each sampling frequency. Lines 3 to 5 give the results of the test for Gaussianity, exploiting the fact that under the null hypothesis of normally distributed observations, the bispectrum is zero. Plainly, this condition is violated here at all frequencies, with vanishingly small associated p-values of the chi-square statistic. The evidence of skewness implies that the data cannot be adequately described by any symmetric model. In this light, it is not surprising that GARCH residuals fail the BDSL tests.

For the nonlinearity test, we present values of the standard normal test statistics both for the interquartile range (z_1) and for the 80%-quantile range (z_2), with their associated p-values, p_1 and p_2 respectively. The advantage in restricting attention to the quantile range statistics is that they exclude outliers in the extreme tails of the distribution.

Perhaps the superiority of the 80% range is reflected in the fact that it yields results which are more in line with our *a priori* expectations, as well as with the BDSL results already reported. What is surprising about the values of z_1 in Table 3 is that they are actually lower (i.e. less extreme) at the 30-min than the 60-min frequency. Moreover, the value falls even further at 15-min, where the p-value is actually over 40%. This flies in the face of our expectation that the degree of nonlinearity would increase, not decrease, as the frequency increases. The results for the 80% range are vastly more plausible. The IID hypothesis is unambiguously rejected, at least at the standard 5% level, although the results for 60, 30 and 15 min are very similar. At the 5-min frequency, and especially at the 1-min, the rejection of linear independence is overwhelming.

The bottom half of the table shows the results of the same tests applied to the residuals from GARCH processes fitted to the raw returns. The difference between the two halves is dramatic. It is plain that on the basis of this type of test, GARCH removes almost every trace of any structure in the series at all frequencies except 1-min, where once more we see an improvement, but nowhere near sufficient to eliminate the dependence entirely. On the evidence here, one cannot reject the proposition that the data are completely characterised by a GARCH process. However, a caveat is in order here. The fitting of a GARCH model to what are clearly skewed data involves an obvious misspecification and may well have distorted the underlying structure (Brock *et al.* 1991), a point which is probably relevant to many of the recently published empirical studies of financial markets.

Nearest-neighbour Estimates of Lyapunov Exponents (Table 4)

Table 4 presents estimates of the maximum Lyapunov exponents, λ_i, of our data series, derived from fitting the Jacobian matrix by the Zeng *et al.* (1992) algorithm.[10] The estimates are arranged in ascending order of embedding dimension (from 1 to 8), order of the fitted polynomial ($q = 1, 2$) and number of extra neighbours used (nb).

Although the results are mixed, the overall conclusion seems to point to a

[10] Computations of the other (smaller) L.E. values are available from the authors on request.

Table 4

Lyapunov Exponent (LE) Modified Estimates (Zeng, 1992)

Minimum required number of neighbours:

Order	1	2	3	4	5	6	7	8
1	1	2	3	4	5	6	7	8
2	2	5	9	14	20	27	35	44

	60-min returns ($N = 992$)			30-min returns ($N = 1,984$)			15-min returns ($N = 3,968$)			5-min returns ($N = 11,904$)			1-min returns ($N = 59,520$)		
dim	max λ	$\Sigma\lambda$	K-Y DIM	max λ	$\Sigma\lambda$	K-Y DIM	max λ	$\Sigma\lambda$	K-Y DIM	max λ	$\Sigma\lambda$	K-Y DIM	max λ	$\Sigma\lambda$	K-Y DIM
	$nb = 10$	$q = 1$		$nb = 10$	$q = 1$		$nb = 10$	$q = 1$		$nb = 10$	$q = 1$		$nb = 10$	$q = 1$	
1	−0·497	−0·497		−0·518	−0·518		−0·616	−0·616		−0·680	−0·680		−1·051	−1·051	
2	−0·056	−0·782		−0·084	−0·800		−0·184	−0·982		−0·243	−1·234		−0·389	−1·443	
3	0·055	−0·908	1·205	0·046	−0·977	1·168	−0·050	−1·193		−0·102	−1·361		−0·194	−1·590	
4	0·093	−1·067	1·777	0·063	−1·121	1·534	0·001	−1·228	1·007	−0·043	−1·434		−0·111	−1·681	
5	0·109	−1·165	2·319	0·074	−1·135	2·173	0·017	−1·387	1·193	−0·013	−1·525		−0·066	−1·726	
6	0·098	−1·238	2·890	0·078	−1·308	2·724	0·028	−1·366	1·603	0·008	−1·545	1·102	−0·037	−1·735	
7	0·112	−1·196	3·605	0·071	−1·272	3·428	0·036	−1·443	2·318	0·016	−1·621	1·303	−0·021	−1·732	
8	0·091	−1·268	4·205	0·079	−1·372	4·206	0·045	−1·462	3·019	0·023	−1·618	1·674	−0·007	−1·735	
	$nb = 20$	$q = 1$		$nb = 20$	$q = 1$		$nb = 20$	$q = 1$		$nb = 20$	$q = 1$		$nb = 20$	$q = 1$	
1	−0·759	−0·759		−0·682	−0·682		−0·630	−0·630		−0·616	−0·616		−0·906	−0·906	
2	−0·140	−0·929		−0·144	−0·948		−0·230	−1·145		−0·268	−1·398		−0·378	−1·486	
3	0·050	−1·093	1·137	−0·009	−1·177		−0·102	−1·338		−0·137	−1·570		−0·196	−1·729	
4	0·059	−1·250	1·305	0·034	−1·344	1·170	−0·045	−1·444		−0·075	−1·653		−0·126	−1·880	
5	0·097	−1·350	1·700	0·045	−1·447	1·380	−0·020	−1·529		−0·044	−1·726		−0·085	−1·956	
6	0·088	−1·388	2·100	0·056	−1·546	1·716	−0·006	−1·624		−0·023	−1·800		−0·057	−1·989	
7	0·087	−1·410	2·602	0·049	−1·563	2·108	0·003	−1·704	1·036	−0·015	−1·844		−0·045	−1·992	
8	0·080	−1·515	2·985	0·053	−1·685	2·384	0·006	−1·727	1·146	−0·009	−1·844		−0·033	−2·006	

	nb = 15	*q* = 2		*nb* = 15	*q* = 2		*nb* = 15	*q* = 2		*nb* = 30	*q* = 2		*nb* = 30	*q* = 2	
1	−0·917	−0·917		−0·859	−0·859		−0·728	−0·728		−0·660	−0·660		0·215	0·215	
2	−0·103	−0·834		−0·031	−0·711		−0·030	−0·734		−0·114	−1·031		−0·244	−1·207	
3	0·199	−0·600	2·087	0·200	−0·600	2·114	0·218	−0·587	2·147	0·120	−0·856	1·532	−0·007	−1·152	
4	0·302	−0·490	3·255	0·341	−0·466	3·323	0·337	−0·459	3·309	0·243	−0·695	3·003	0·105	−1·047	2·027
5	0·388	−0·480	4·368	0·382	−0·365	4·478	0·415	−0·350	4·488	0·336	−0·565	4·193	0·168	−0·962	3·200
6	0·448	−0·279	5·560	0·431	−0·326	5·527	0·468	−0·241	5·649	0·416	−0·449	5·369	0·215	−0·880	4·421
7	0·393	−0·367	6·419	0·466	−0·230	6·672	0·528	−0·147	6·780	0·466	−0·346	6·507	0·247	−0·812	5·618
8	0·403	−0·258	7·603	0·482	−0·196	7·717	0·571	−0·082	7·883	0·509	−0·275	7·617	0·273	−0·757	6·783
	nb = 30	*q* = 2		*nb* = 30	*q* = 2		*nb* = 30	*q* = 2		*nb* = 40	*q* = 2		*nb* = 40	*q* = 2	
1	−1·267	−1·267		−1·160	−1·160		−0·926	−0·926		−0·657	−0·656		−0·197	−0·197	
2	−0·333	−1·196		−0·285	−1·104		−0·260	−1·146		−0·243	−1·345		−0·311	−1·395	
3	−0·008	−1·044		0·003	−1·004	1·010	0·033	−0·999	1·125	−0·024	−1·223		−0·109	−1·483	
4	0·154	−0·832	2·331	0·145	−0·896	2·367	0·143	−0·872	2·329	0·092	−1·100	1·888	−0·015	−1·454	
5	0·227	−0·797	3·645	0·225	−0·756	3·705	0·222	−0·778	3·707	0·170	−0·997	3·160	0·043	−1·377	1·488
6	0·282	−0·647	5·031	0·267	−0·722	4·943	0·281	−0·654	5·055	0·227	−0·897	4·459	0·089	−1·301	2·757
7	0·273	−0·618	6·066	0·299	−0·626	6·091	0·315	−0·584	6·155	0·273	−0·797	5·757	0·116	−1·234	4·018
8	0·294	−0·595	7·106	0·322	−0·596	7·123	0·363	−0·531	7·258	0·309	−0·723	6·972	0·141	−1·181	5·238

N, sample size; *nb*, number of additional neighbours used; *q*, order of estimation polynomial (1 or 2); dim, embedding dimension for estimation; max λ, maximum estimated value of L.E.; $\Sigma\lambda$, sum of L.E.'s; K-Y DIM, Kaplan-Yorke dimension implied by sum of L.E.'s.

high, probably infinite dimension. It should be noted that a large element of judgement is involved in assessing the estimates, simply because nothing is known of the asymptotic distribution of the estimators.

The maximal λ_i for the 1-min returns are sometimes positive (especially for $q = 2$), and at lower frequencies the sign also varies. All that can be said is that it is more often positive for higher dimensions when more neighbouring points are used and the second-order polynomial is fitted. In absolute terms, the λ_i are invariably small, usually less than 0·25. Moreover, the sum $\Sigma\lambda_i$ over all $\lambda_i > 0$, which is a measure of the Kolmogorov entropy, is negative in virtually every case, and the Kaplan–Yorke dimension, which is given by:

$$D_{KY} = k + \frac{1}{|\lambda_{k+1}|}\sum_{i=1}^{k}\lambda_i \quad \text{for:} \quad \sum_{i=1}^{k}\lambda_i \geqslant 0 > \sum_{i=1}^{k+1}\lambda_i$$

is unstable across dimensions and polynomial order. It is worth noting that, even in the case of the estimate most favourable to low dimensional chaos, for $nb = 30$, $q = 2$ and dimension 8, the implied maximum forecast horizon is less than $1\frac{1}{2}$ min. It is also noteworthy that the L.E. estimates tend to diminish as we increase the number of extra neighbours used in estimation, but tend to increase as the order of the polynomial is increased.[11]

Neural Net Estimates of Lyapunov Exponents[12]

One of the attractive features of the Nychka *et al.* (1992) approach is that, in order to remedy the lack of any rigorous significance testing procedure, they suggest the use of scatter plots that are informative about precision.

Figs. 1*a* and 1*b* show the scatter of λ estimates for different (L, d, q) combinations, where L is the time delay used ($L = 1$ for hourly returns, $L = 2$ for 1-min returns), $d = 1, \ldots, 10$ is the embedding dimension and $q = 1, \ldots, 5$ is the selection parameter equal to the number of units in the hidden layer of the net.[13] This means that we have a total of $1 \times 10 \times 5 = 50$ parameter combinations in each case. We then plot the 20 BIC-minimising estimates of the L.E. associated with each (L, d, q) triplet i.e. 1,000 points in each figure.

As can clearly be seen, although the estimates are quite sensitive to the estimation parameters, the conclusion that the λ_i are all negative is extremely robust. Moreover, this is obvious even though the estimates are poorly defined, with a number of different L.E. values associated with the same minimum BIC, especially at the half- and one-hour frequencies. This result in itself may be taken as a sign that the data are characterised by a low signal-to-noise ratio.

[11] We also experimented with Briggs (1990) estimates, but the results were not very plausible. Though they indicated the presence of at least one positive Lyapunov exponent in every case, the estimates were extremely large in absolute terms, especially at the 1-min interval. Moreover, the estimates of all the λ_i appeared to be very sensitive to the values of m and q, so that the entropy measure fluctuated wildly, indicating that the series is virtually unforecastable at any horizon. The difficulty with the Briggs (1990) algorithm in the present case seems to lie in the large number of zero returns in the high frequency data, many of which are excluded by the Zeng *et al.* (1992) 'shell' technique.

[12] The results in this section were achieved by vectorising the LENNS program of Ellner, Nychka and Gallant to run on a supercomputer.

[13] To save space, we only present graphs for the hourly and 1-min returns. Our conclusions were qualitatively similar at the other frequencies.

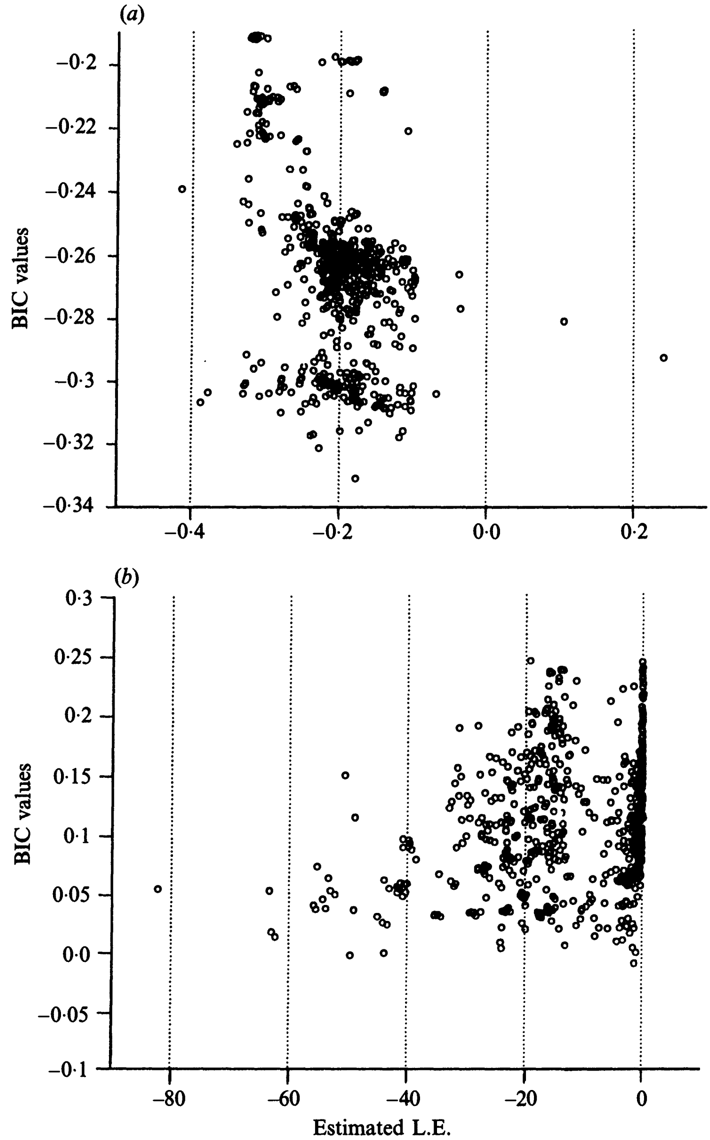

Fig. 1. Scatter of λ estimates for different combinations. (*a*), 1-min returns; (*b*), hourly returns.

Figs. 2*a* and 2*b* give a somewhat different perspective on the results. At each embedding dimension $d = 1, \dots, 10$, the graphs show:
(*a*) the BIC-minimising L.E. estimate (the 'best fit')
(*b*) the mean of the 10 BIC-minimising L.E. estimates (the unbroken line)
(*c*) given the standard deviation of the 10 BIC-minimising L.E. estimates, a range of one standard deviation about the mean ('Upper' and 'Lower', marked by the two dotted lines).

Again, it is plain that even the upper bounds on the estimates are below zero in almost every case. In other respects, however, the results are ambiguous. For

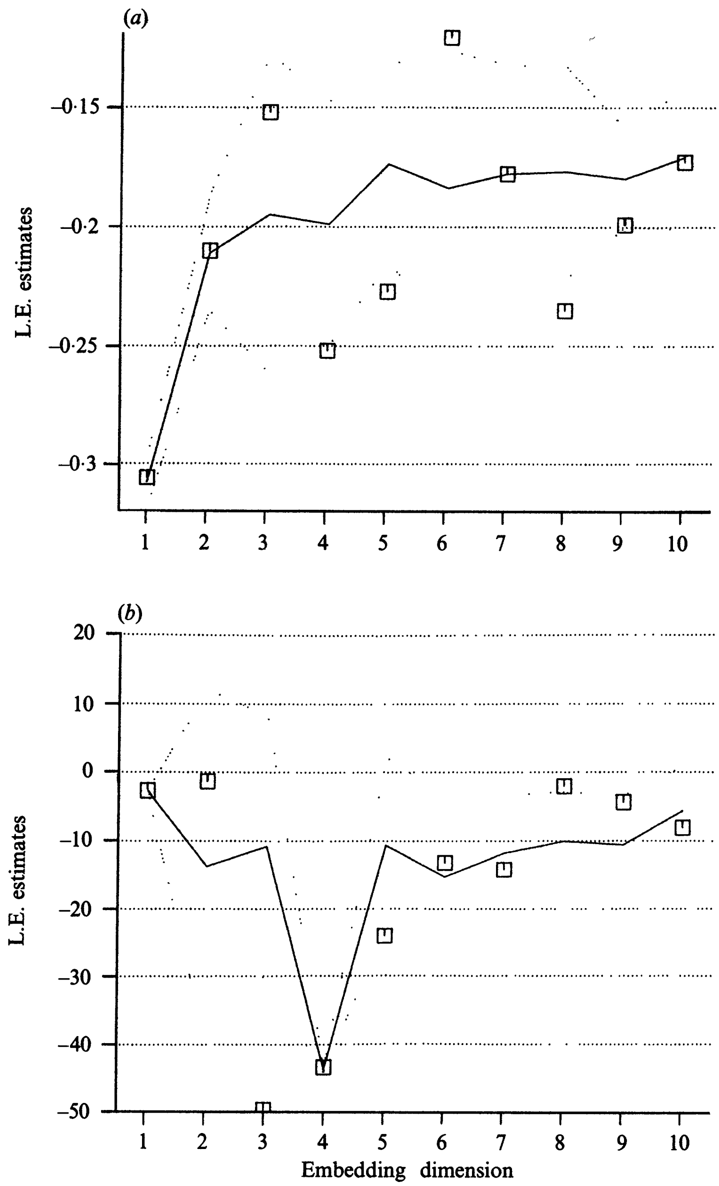

Fig. 2. Estimates and embedding dimensions. (*a*), 1-min returns; (*b*), hourly returns. ☐, best fit; ——, mean; ······, upper/lower.

example, in the case of the 1-min returns the mean L.E. estimate appears to be levelling off at around $-0{\cdot}17$. But the 'best' estimate fluctuates wildly, falling outside the one-standard deviation bounds in several cases, in spite of the fact that the bounds themselves are very wide. The estimates with hourly returns (Fig. 2*b*) are even more erratic and show no clear levelling off in the λ estimates, though they are still negative at $d = 10$.

IV. CONCLUSIONS

This paper has addressed two major questions. Is there a nonlinear structure in the process generating returns on the FTSE 100? And, if so, is there any evidence that the process may be chaotic?

Our answers are plain. There is clear evidence of nonlinear dependence at all frequencies, especially the 1- and 5-min. However, it seems that some of the nonlinearity can be explained by a simple GARCH process. We find little to support the view that the process is chaotic at any frequency. Our estimates of the maximal Lyapunov exponent seem very sensitive both to the estimation methodology and to the chosen embedding dimension. In the case of the nearest-neighbour estimates, they are very much dependent on the choice of embedding dimension, as well as on the number of additional neighbours used and the order of the fitted polynomials. As far as the neural nets are concerned, the result that the maximum L.E. is negative seems highly robust.

We are very conscious that the research agenda in this area is still very full indeed. The obvious next step would be to proceed to forecasting the index using locally weighted regression. As far as other datasets are concerned, not only is it of interest to know whether other countries' stock market indices have the same dynamic characteristics as the FTSE, it would also be intriguing to see whether the associated futures prices follow a similar pattern. It is also essential to proceed at an early stage to examining the patterns in individual stock returns. Not only is the subject important in its own right, it may also help to explain the time series pattern of index returns, though the relationship between the statistical properties of the index and its components is potentially complex in the extreme.

University of Stirling

Date of receipt of final typescript: November 1994

References

Ashley, A., Patterson, D. M. and Hinich, M. (1986). 'A diagnostic test for nonlinear serial dependence in time series fitting errors.' *Journal of Time Series Analysis*, vol. 7, no. 3, pp. 165–78.

Baumol, W. J. and Benhabib, J. (1989). 'Chaos: significance, mechanism, and economic applications.' *Journal of Economic Perspectives*, vol. 3(1), pp. 77–105.

Boldrin, M. (1988). 'Persistent oscillations and chaos in economic models: notes for a survey.' In *The Economy as an Evolving Complex System* (ed. P. W. Anderson, K. J. Arrow and D. Pines). Redwood City, CA: Addison-Wesley.

Briggs, K. (1990). 'An improved method for estimating Liapunov exponents of chaotic time series', vol. 151, 1, 2, *Physics Letters* A, pp. 27–32.

Brillinger, D. R. (1965). 'An introduction to the polyspectrum.' *Annals of Mathematical Statistics*, vol. 36, pp. 1351–74.

Brock, W. A., Dechert, W., Scheinkman, J. A. and LeBaron, B. (1987). 'A test for independence based upon the correlation dimension.' Working Paper, Department of Economics, University of Wisconsin-Madison.

Brock, W. A., Hsieh, D. and LeBaron, B. (1991). *Nonlinear Dynamics, Chaos and Instability: Statistical Theory and Empirical Evidence*, Cambridge, MA: MIT Press.

Brock, W. A. and Kleidon, A. (1992). 'Periodic market closure and trading volume: a model of intraday bids and asks.' *Journal of Economic Dynamics and Control*, vol. 16, pp. 451–89.

Brockett, P. L., Hinich, M. J. and Patterson, D. (1988). 'Bispectral-based tests for the detection of

Gaussianity and linearity in time series.' *Journal of the American Statistical Association*, vol. 83, 403, pp. 657–64.

Brown, R., Bryant, P. and Abarbanel, H. D. I. (1991). 'Computing the Lyapunov spectrum of a dynamical system from an observed time series.' *Physical Review*, A43, pp. 2787–806.

Devaney, R. L. (1989). *An Introduction to Chaotic Dynamical Systems*, 2nd ed. Redwood City, CA: Addison-Wesley.

Eldridge, R. M. and Coleman, M. P. (1993). 'The British FTSE-100 index: chaotically deterministic or random? Working Paper, School of Business, Fairfield University.

Ellner, S., Gallant, A. R., McCaffrey, D. and Nychka, D. (1991). 'Convergence rates and data requirements for Jacobian-based estimates of Liapunov exponents from data.' *Physics Letters* A, vol. 153, pp. 357–63.

Ellner, S., Nychka, D. W. and Gallant, A. R. (1992). 'LENNS, a program to estimate the dominant Lyapunov exponent of noisy nonlinear systems from time series data.' Institute of Statistics Mimeo Series no. 2235 (BMA Series no. 39), Statistics Department, North Carolina State Univerity.

Frank, M. and Stengos, T. (1989). 'Measuring the strangeness of gold and silver rates of return.' *Review of Economic Studies*, vol. 56, pp. 553–67.

Hinich, M. (1982). 'Testing for Gaussianity and linearity of a stationary time series.' *Journal of Time Series Analysis*, vol. 3, pp. 169–76.

Hinich, M. J. and Patterson, D. M. (1985). 'Evidence of nonlinearity in daily stock returns.' *Journal of Business and Economic Statistics*, vol. 3, 1, pp. 69–77.

Hinich, M. J. and Patterson, D. (1989). 'Evidence of nonlinearity in the trade-by-trade stock market return generating process.' In *Economic Complexity: Chaos, Sunspots, Bubbles and Nonlinearity* (ed. W. Barnett, J. Guweke and K. Shell), pp. 383–409. Cambridge: Cambridge University Press.

Hsieh, D. (1989). 'Testing for nonlinear dependence in daily foreign exchange rates.' *Journal of Business*, vol. 62, 3, pp. 339–59.

Hsieh, D. (1991). 'Chaos and nonlinear dynamics: applications to financial markets.' *Journal of Finance*, vol. 47, 5, pp. 1145–89.

Hsieh, D. (1993). 'Using non-linear methods to search for risk premia in currency futures.' *Journal of International Economics*, vol. 35, pp. 113–32.

Kodres, L. E. and Papell, D. H. (1991). 'Nonlinear dynamics in the foreign exchange futures market.' Working Paper, School of Business and Administration, University of Michigan.

Kräger, H. and Kügler, P. (1992). 'Non-linearities in foreign exchange markets: a different perspective.' Working Paper, Universität Mannheim, Germany.

Mayfield, E. S. and Mizrach, B. (1989). 'On determining the dimension of real time stock price data.' Working Paper, Boston College, MA.

McCaffrey, D. F., Ellner, S., Gallant, A. R. and Nychka, D. W. (1992). 'Estimating the Lyapunov exponent of a chaotic system with nonparametric regression.' *Journal of the American Statistical Association*, vol. 87, 419, pp. 682–95.

Nychka, D., Ellner, S., Gallant, A. R. and McCaffrey, D. (1992). 'Finding chaos in noisy systems.' *Journal of the Royal Statistical Society*, series B, vol. 52, 2, pp. 399–426.

Phillipatos, G. C., Pilarinu, E. and Malliaris, A. G. (1993). 'Chaotic behaviour in stock prices of European stock markets: a comparative analysis of major economic regions.' Working Paper, University of Tennessee.

Scheinkman, J. A. and LeBaron, B. (1989). 'Nonlinear dynamics and stock returns.' *Journal of Business*, vol. 62, pp. 311–37.

Subba Rao, T. and Gabr, M. (1980). 'A test for linearity of stationary time series.' *Journal of the Royal Statistical Society*, series B, vol. 54, pp. 329–51.

Tata, F. (1991). 'Is the foreign exchange market characterized by nonlinearity?' Discussion Paper No. 118, LSE Financial Markets Group Discussion Paper Series, London.

Vaidyanathan, R. and Krehbiel, T. (1992). 'Does the S & P 500 futures mispricing series exhibit nonlinear dependence across time?' *Journal of Futures Markets*, vol. 12, 6, pp. 659–77.

Vassilicos, J. C., Demos, A. and Tata, F. (1992). 'No evidence of chaos but some evidence of multifractals in the foreign exchange and the stock markets.' Discussion Paper No. 143, LSE Financial Markets Group Discussion Paper Series, London.

Wolf, A. J., Swift, J. and Vastano, J. (1985). 'Determining Lyapunov exponents from a time series.' *Physica*, vol. 16D, pp. 285–317.

Zeng, X., Pielke, R. A. and Eykholt, R. (1992). 'Extracting Lyapunov exponents from short time series of low precision.' *Modern Physics Letters*, vol. 6, 2, pp. 3229–32.