

Uncovering Nonlinear Structure in Real-Time Stock-Market Indexes: The S&P 500, the DAX, the Nikkei 225, and the FTSE-100



A. Abhyankar; L. S. Copeland; W. Wong

Journal of Business & Economic Statistics, Vol. 15, No. 1. (Jan., 1997), pp. 1-14.

Stable URL:

<http://links.jstor.org/sici?sici=0735-0015%28199701%2915%3A1%3C1%3AUNSIIRS%3E2.0.CO%3B2-Q>

Journal of Business & Economic Statistics is currently published by American Statistical Association.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://uk.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://uk.jstor.org/journals/astata.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is an independent not-for-profit organization dedicated to creating and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact support@jstor.org.

Uncovering Nonlinear Structure In Real-Time Stock-Market Indexes: The S&P 500, the DAX, the Nikkei 225, and the FTSE-100

A. ABHYANKAR

Department of Accounting and Finance, University of Stirling, Stirling FK9 4LA, Scotland

L. S. COPELAND

Cardiff Business School, University of Wales, Cardiff CF1 3EU, Wales

W. WONG

Department of Accounting and Finance, University of Stirling, Stirling FK9 4LA, Scotland

This article tests for nonlinear dependence and chaos in real-time returns on the world's four most important stock-market indexes. Both the Brock–Dechert–Scheinkman and the Lee, White, and Granger neural-network-based tests indicate persistent nonlinear structure in the series. Estimates of the Lyapunov exponents using the Nychka, Ellner, Gallant, and McCaffrey neural-net method and the Zeng, Pielke, and Eyckholt nearest-neighbor algorithm confirm the presence of nonlinear dependence in the returns on all indexes but provide no evidence of low-dimensional chaotic processes. Given the sensitivity of the results to the estimation parameters, we conclude that the data are dominated by a stochastic component.

KEY WORDS: Brock–Dechert–Scheinkman test; Chaos; GARCH models; Lyapunov exponent; Nearest-neighbor method; Neural net; Nonparametric; Stock index futures; Stock returns.

1. INTRODUCTION

This article addresses two important questions that have been the focus of a substantial and still growing literature in recent years. Is there nonlinear dependence in stock-market returns? And, if so, is the nonlinear structure characterized by low-dimensional chaos? In other words, is the apparent randomness of the time series pattern of returns explicable, in part at least, by a deterministic process?

Until relatively recently, it was more or less taken for granted that movements in stock-market prices were overwhelmingly stochastic in nature, if not actually a random walk. The assertion seemed unchallengeable not only on empirical grounds but also for apparently sound theoretical reasons—namely, consistency with the ruling efficient-markets paradigm. Moreover, it seems improbable a priori that the pattern of returns could be explained to any substantial degree by a deterministic process, given that the major cause of market movements is normally assumed to be the random flow of news.

In the last few years, however, several developments have taken place that have led to serious questioning of the proposition that stock returns are inherently unforecastable. First, researchers using conventional econometric methods have uncovered several deviations from efficiency in the behavior of stock prices (Fama 1991). Second, the efforts of statisticians, econometricians, and physicists have resulted in the development of several tests capable of detecting nonlinear as well as linear patterns in the data (see Sec. 2). Third, the exciting progress made in the last 20 years in understanding the mathematics of nonlinear systems means

that we can now entertain the possibility of certain types of deterministic process in financial data. In particular, it has become clear that many low-dimension deterministic nonlinear systems are capable of generating output that is in most respects indistinguishable from white noise. It is important to note that, as far as financial series are concerned, this type of process *could* be consistent with market efficiency if it is only forecastable at horizons too short to allow for profitable exploitation by speculators. The unresolved issue addressed in this article relates, therefore, to whether stock-market index returns are best represented by a purely stochastic process or rather by a (nonlinear) deterministic structure, presumably with superimposed noise.

As can be seen from Table 1, there is already a substantial literature examining the questions addressed in this article. Most of the research so far has concentrated, as we do here, on stock-market indexes [especially the Standard & Poor (S&P)] or on exchange rates, though a few have looked elsewhere (futures markets, gold and silver prices). The most commonly deployed test is the Brock–Dechert–Scheinkman (BDS) test for independence (see Sec. 2), though several authors have relied on estimates of the correlation dimension itself.

As regards the main conclusions of the literature, there is a broad consensus of support for the proposition that the return process is characterized by a pattern of nonlinear dependence. In particular, BDS tests almost invariably

Table 1. Nonlinearity Testing on Financial Data: Summary of Published Results

Authors	Dataset	Sample info.	Tests	Results
Abhyankar, Copeland, and Wong (1995)	FTSE-100 cash	$N = 60,000$	(1) Bispectral linearity test (2) BDS (3) L.E.	(1) Nonlinear (3) No evidence of chaos
Eldridge and Coleman (1993)	FTSE-100 cash and futures	$N \approx 1000$, June '84 to Sept. '87	(1) Correlation dimension test (2) Wolf's L.E.	Not iid and consistent with chaos
Hsieh (1993)	Foreign currency spots and futures	$N = 1,275$ from 22/2/85 to 9/3/90 daily	Tests of linear and nonlinear predictabilities	No linear and nonlinear predicabilities
Philippatos, Pilarinu, and Malliaris (1993)	Ten major national stock indexes	$N = 833$, weekly levels and returns from Jan. '76 to Dec. '91	BDS tests	Nonlinear
Kräger and Kugler (1992)	Exchange rates (1) Japanese yen (2) German mark (3) French franc (4) Italian lira (5) Swiss franc	$N \approx 500$, weekly returns from June '80 to Jan. '90	BDS	Nonlinear
Vaidyanathan and Krehbiel (1992)	S&P 500 futures mispricings	$N = 1,500$	(1) BDS (2) Correlation dimension test	Nonlinear and low-dimensional chaos ($d = 6$)
Vassilicos, Demos, and Tata (1992)	(1) Deutsche mark (2) Swiss franc (3) NYSE (daily)	$N = 20,000$ to 30,000	(1) Wolf's L.E. (2) Correlation dimension test	No evidence of chaos
Brock, Hsieh, and LeBaron (1991)	(1) CRSP value-weighted index (2) S&P 500	$N = 2,510$, daily from 2/1/74 to 30/12/83	(1) BDS (2) Tsay (3) Dimension plots (4) Sign-scrambling plots (5) Recurrence plots	Nonlinearity; nonconstant variance; little evidence of nonlinear forecastability
Kodres and Papell (1991)	Daily futures (1) British pound (2) Canadian dollar (3) Deutsche mark (4) Japanese yen (5) Swiss franc	$N \approx 3,500$, from 1/7/73 to 17/3/87	BDS	Nonlinear

reject the null of an iid process. On the other hand, the evidence on chaos is more mixed, with some evidence of a low-dimensional structure in the U.S. stock-market index (Mayfield and Mizrach 1989; Vaidyanathan and Krehbiel 1992) but little or none in exchange-rate series (Hsieh 1989, 1993; Tata 1991). Note, however, that this conclusion is based for the most part on the results of correlation dimension tests rather than direct Lyapunov exponent (L.E.) estimates (but see Vassilicos, Demos, and Tata 1992; Eldridge and Coleman 1993).

A notable feature of the literature summarized in the table is that, because most of the published work in this area relies on relatively low-frequency (typically daily or even weekly) data, it invariably uses datasets of fewer than 5,000 observations and often much smaller. This is a serious drawback for several reasons, which are most apparent in testing for sensitivity to initial conditions, the hallmark of chaos. [See Devaney (1989) for a more formal definition.] In the first place, the scope for applying L.E. methods to datasets of only 1,000 or so observations is very limited. The prob-

lem is particularly acute in the light of the fact that there are no rigorous criteria for assessing the significance of L.E. estimates.

Perhaps most worrying of all, note that, to find sufficient observations to implement the tests, most researchers were forced to use data periods measured in years (up to 15 in some cases). The longer the data period, however, the less plausible is the assumption that the underlying process could have remained stationary from start to finish, a fact that makes the results in the table difficult to interpret.

By contrast, for each of our six stock-market indexes (four cash and two futures) we use high-frequency real-time datasets covering only three months but still involving a minimum of 10,000 observations, which means that we are able to implement several tests of nonlinear dependence. In particular, we apply the now-standard nonparametric test to the prewhitened series, as well as to the output of generalized autoregressive conditional heteroscedasticity (GARCH) filters. In addition, we make use of a recently developed neural-net-based test for nonlinear structure. In the light of the indications of nonlinear dependence uncovered

Table 1—Continued

Authors	Dataset	Sample info.	Tests	Results
Tata (1991)	Swiss franc	$N = 32,200$	(1) Correlation dimension test (2) BDS	No evidence of low-dimensional chaos but nonlinear
Hsieh (1991)	CRSP	$N = 1,297-2,017$, data from 1963 to 1987	(1) BDS (2) 3rd moment tests	Not iid, nonlinear
	S&P 500 (1) Weekly (2) Daily (3) Four 15-min. returns	(1) $N \approx 1,500$ (2) $N \approx 1,700$ (3) $N \approx 1,800$	BDS tests	Not iid
Vassilicos (1990)	Deutsche mark	$N = 20,408$, ask-quotes from 9/4/89 to 15/4/89	Correlation dimension test	No low-dimensional chaos
Frank and Stengos (1989)	Returns of (1) Gold (2) Silver prices	(1) $N \approx 2,900$ (2) $N \approx 3,100$	(1) Correlation dimension test (2) Kolmogorov entropy	(1) Dimension of 6–7 (2) Positive; low-dimensional chaos
Hinich and Patterson (1989)	Dow Jones Industrial Average	$N \approx 750$, from 1/9/78 to 31/8/81	Bispectral Gaussianity and linearity tests	Non-Gaussian and nonlinear; unalised data less nonlinear
Hsieh (1989)	Major foreign currencies daily closing bid prices	$N = 2,510$, from 2/1/74 to 30/12/83	(1) Box–Pierce (2) Ljung–Box (3) BDS	Nonlinear
Mayfield and Mizrach (1989)	S&P 500	$N = 20,088$, 20-sec. returns from Jan '87	Correlation dimension tests	Low-dimensional chaos
Scheinman and LeBaron (1989)	Daily returns on CRSP weighted index	$N = 5,200$	BDS tests on original and filtered	Evidence of nonlinearity
Brockett, Hinich, and Patterson (1988)	(1) 10 common U.S. stocks (2) U.S.\$–yen spot and forward rates	(1) N not given (2) $N \approx 400$	Bispectral Gaussianity and linearity tests	Non-Gaussian and nonlinear
Eckmann, Kamphorst, Ruelle, and Scheinkman (1988)	Daily returns on CRSP weighted index	$N = 5,200$	(1) Recurrence plots (2) Wolf's L.E.	Weak evidence of chaos

NOTE: L.E. is the Lyapunov exponent test.

by the tests, we proceed to estimate L.E.'s for each of the series, in an attempt to establish whether the underlying processes are characterized by extreme sensitivity to initial conditions.

The results reported here suffice to establish several points. First, we support the bulk of the literature in finding clear evidence of nonlinear dependence in all four series at all frequencies examined. Second, this dependence is largely, but not entirely, explained by volatility clustering, as specified in the class of conditional heteroscedasticity models already widely used in financial time series modeling. Third, if there is a low-dimensional deterministic structure generating the data, it is almost certainly not chaotic (i.e., not sensitive to initial conditions). In view of the instability of the L.E. estimates as the estimation parameters are varied, however, the most plausible explanation of the processes observed is that they are predominantly random.

This conclusion reinforces the results of Abhyankar, Copeland, and Wong (1995) (henceforth ACW) in several different respects. First, for reasons discussed in Section 2, our conclusions regarding nonlinear dependence are made

more robust by replacing the Hinich (1982) bispectrum test used by ACW with the Lee–White–Granger (1993) (LWG) test.

Second, where ACW examined the behavior only of the U.K. Financial Times Stock Exchange-100 (FTSE-100) index over the first six months of 1993, the conclusions of this article relate to all four of the world's most important stock-market indexes over a different data period. In the case of the United Kingdom and the United States, we also present results on the futures, as well as the cash index. Our dataset in fact consists of real-time observations for the period September 1 to November 31, 1991, at 1-minute frequency in the case of the FTSE-100, the Deutscher Aktienindex (DAX), and the Nikkei and 15-second frequency for the S&P 500, and transaction prices over the same period for the FTSE and S&P futures.

Third, and most important, we are concerned here with index data generated under a wide variety of differing market microstructures, ranging from the specialist system of the New York Stock Exchange to the auction markets of Tokyo and Frankfurt and the competitive dealership environment of London. In addition, each index is different, in

Table 2. Descriptive Statistics for Index Returns: September–November 1991

Frequency	Sample size N	Unique values NC	No. of zeroes NO	Minimum	Maximum	Mean	S.D.	Skewness	Kurtosis
S&P 500									
15-second	97,185	19,821	22,011	-.00108	.00084	-.00000	.00007	-.276	10.732
1-minute	24,504	14,884	2,579	-.00257	.00263	-.00000	.00018	-.548	21.486
5-minute	4,898	4,518	189	-.00560	.00450	-.00001	.00054	-.625	15.516
FTSE-100									
1-minute	31,200	9,723	14,123	-.00121	.00133	-.00000	.00009	-.911	17.134
5-minute	6,240	4,797	932	-.00324	.00332	-.00001	.00038	-1.122	13.429
DAX									
1-minute	11,340	9,796	1,446	-.01860	.00898	-.00001	.00036	-10.513	747.204
5-minute	2,268	2,232	33	-.01910	.00967	-.00003	.00100	-2.803	77.766
NIKKEI									
1-minute	16,348	16,164	159	-.00358	.00421	.00000	.00044	1.318	14.729
5-minute	3,172	3,169	3	-.00758	.00733	-.00000	.00113	.436	4.176
S&P 500 FUTURES									
1-minute	24,180	4,961	5,442	-.00420	.00359	-2.6E-06	.000365	-.341	8.540
5-minute	4,836	3,147	533	-.00774	.00479	-1.3E-05	.000814	-.573	7.159
FTSE-100 FUTURES									
1-minute	26,390	1,762	16,728	-.00334	.00375	-5.6E-06	.000303	-.107	7.776
5-minute	5,278	1,497	1,634	-.00606	.00562	-2.8E-05	.00065	-.305	5.820

NOTE: N is the number of observations in the sample. NC is N less the number of repeated observations in the sample. NO is the number of zero observations in the sample. Both skewness and kurtosis statistics are centered on 0.

terms of how its constituent stocks are selected, the type of price incorporated, and the way its weights are computed. The advantage of using such heterogeneous series is that we can be confident that any common patterns we succeed in uncovering must be independent of market structures and the details of index composition.

In Section 2, we outline our methodology, including two tests for nonlinearity and two different approaches to estimating Lyapunov exponents. Section 3 describes the main features of our datasets for each index, and our results are presented and discussed in Section 4.

2. METHODOLOGY

In this article we implement two tests for nonlinear dependence, the well-known Brock, Dechert, Scheinkman, and LeBaron (1987) test (henceforth BDS) and the more recent Lee, White, and Granger (1993) (LWG) test.

The BDS test relies on the limiting value of the correlation integral

$$C(m, \varepsilon, N) = I[(t, s): \|X_t^m - X_s^m\| < \varepsilon] / N^2, \quad (1)$$

where $X_t^m = (x(t), \dots, x(t - m + 1))$, $\|\cdot\|$ is the L_∞ norm on R^m , and $I[\cdot]$ denotes the number of elements. Subject only to modest regularity conditions, as $N \rightarrow \infty$, $C(m, \varepsilon, N)$ has a limit $C(m, \varepsilon)$ such that, if $\{x(t)\}$ is iid, it follows that

$$C(m, \varepsilon) = C(1, \varepsilon)^m. \quad (2)$$

This reasoning motivates the BDS test statistic

$$W(m, \varepsilon, N) = \sqrt{\frac{T}{V}} [C(m, \varepsilon, N) - C(1, \varepsilon, N)^m], \quad (3)$$

which converges in distribution to $N(0, 1)$ as $N \rightarrow \infty$. Moreover, Brock et al. (1987) derived the estimator variance, providing a basis for tests of the iid hypothesis.

Note, however, that the BDS test rejects iid for linear as well as nonlinear processes. Because we are concerned with nonlinearity, we apply the test here to data from which the autocorrelation has been removed by prior fitting of a Bayes information criterion (BIC)-minimizing autoregressive moving average model.

We also implement a newer test for nonlinearity introduced by Lee et al. (1993). The LWG test involves fitting a single hidden layer neural network to the residuals from a linear model, then testing its incremental contribution to explaining the movements in the dependent variable. More specifically, if the neural net modeling a series $x\{t\}$ can be represented as

$$o = x'\theta + \sum_{j=1}^q \beta_j \psi(x'\gamma_j), \quad (4)$$

where o is the network output, θ is a vector of (linear) weights, and $\psi(x'\gamma_j)$ is a given nonlinear mapping from \mathcal{R} to \mathcal{R} (the "activation function"), then linearity implies that the optimal network weights β_j , $j = 1, \dots, q$, are all 0. LWG suggested proceeding to implement the Lagrange multiplier test based on the statistic

$$nR^2 \rightarrow \chi^2(q^*), \quad (5)$$

Table 4. BDS Tests for Nonlinear Dependence (Sept.–Nov. 1991): Index Futures

m	ε/σ	S&P 500 Futures				FTSE-100 Futures			
		Linear		GARCH(p, q)		Linear		GARCH(p, q)	
		1-min	5-min	1-min	5-min	1-min	5-min	1-min	5-min
2	.50	30.15	12.79	8.30	1.33	19.70	17.68	8.30	3.37
3	.50	40.76	17.04	11.17	1.38	27.61	25.24	12.20	5.08
4	.50	51.05	20.53	14.58	1.56	33.96	31.29	15.04	6.27
5	.50	63.46	24.66	19.05	2.43	40.71	37.70	19.31	7.78
6	.50	80.04	30.11	25.62	3.46	48.78	45.39	23.93	9.56
7	.50	102.55	37.05	32.33	4.04	58.27	54.33	30.58	11.72
8	.50	138.91	47.14	44.57	4.73	70.19	65.52	40.34	14.69
9	.50	196.78	61.70	66.04	5.65	85.71	79.95	55.83	20.03
10	.50	293.15	81.80	105.66	6.49	105.82	98.98	83.08	31.36
2	.75	30.97	13.34	.70	1.17	28.68	11.97	16.32	1.10
3	.75	39.65	17.24	3.35	.81	37.60	18.06	20.56	1.11
4	.75	46.57	20.19	6.38	.80	45.56	22.59	24.63	1.97
5	.75	53.27	23.42	9.48	1.49	54.11	26.97	28.54	2.22
6	.75	60.57	27.48	12.66	2.09	64.33	31.62	32.94	2.27
7	.75	68.87	32.17	15.90	2.33	77.13	36.40	38.23	2.18
8	.75	78.99	38.50	20.33	2.87	93.93	42.07	44.96	2.07
9	.75	91.67	46.96	25.59	3.48	116.27	49.15	53.58	2.24
10	.75	107.60	57.85	33.33	4.07	145.02	57.72	64.29	2.32
2	1.00	30.76	14.25	1.14	.92	21.49	7.18	14.40	3.40
3	1.00	39.52	17.47	2.35	.45	30.76	11.52	18.71	4.46
4	1.00	46.50	19.61	3.69	.31	38.29	15.11	22.83	5.87
5	1.00	53.24	21.55	4.97	.73	45.85	18.15	26.86	7.01
6	1.00	60.55	23.79	6.05	1.13	54.37	20.57	30.85	7.96
7	1.00	68.84	26.08	6.95	1.21	64.70	22.71	35.46	8.82
8	1.00	78.89	28.74	7.89	1.52	77.90	24.91	41.04	9.88
9	1.00	91.48	32.02	8.74	1.94	95.04	27.12	48.11	11.24
10	1.00	107.24	35.79	9.84	2.35	116.52	29.38	56.73	12.59
2	1.25	31.85	15.15	.25	.54	9.35	3.77	17.71	1.58
3	1.25	39.55	18.20	1.29	-.07	17.16	6.69	21.24	1.90
4	1.25	44.80	20.07	2.70	-.35	23.11	9.38	24.72	2.46
5	1.25	49.29	21.66	4.12	-.17	28.81	11.30	27.64	3.04
6	1.25	53.51	23.42	5.32	.04	33.60	12.66	30.64	3.31
7	1.25	57.86	25.17	6.22	.03	38.26	13.69	33.53	3.56
8	1.25	62.48	27.12	7.12	.24	43.19	14.80	36.52	3.87
9	1.25	67.79	29.43	7.86	.56	48.73	15.73	39.82	4.35
10	1.25	73.58	32.00	8.64	.86	55.04	16.58	43.50	4.91
2	1.50	32.62	15.14	-.18	.24	44.88	2.17	13.78	1.47
3	1.50	39.96	17.65	.48	-.43	47.55	3.90	16.58	1.29
4	1.50	44.39	18.90	1.39	-.78	49.24	5.15	18.71	1.67
5	1.50	47.77	19.78	2.46	-.82	50.55	6.08	20.78	2.05
6	1.50	50.48	20.80	3.33	-.75	51.66	6.54	22.24	2.10
7	1.50	53.01	21.75	3.96	-.87	52.85	6.71	23.65	2.01
8	1.50	55.40	22.84	4.58	-.73	54.28	6.82	25.13	1.94
9	1.50	57.95	24.06	5.03	-.47	55.88	6.84	26.71	2.08
10	1.50	60.54	25.34	5.50	-.21	57.59	6.90	28.40	2.24

NOTE: m is the embedding dimension. GARCH residuals are from the BIC-minimizing model.

Our starting point is the familiar Takens (1981) phase space reconstruction, which allows us to write a noisy (scalar) time series $\{x(t), t = 1, 2, \dots\}$ in state-space form as

$$X_t = F(X_{t-1}) + \varepsilon_t, \quad (6)$$

where $X_t = (x(t), x(t-L), \dots, x(t-(d-1)L))$, d is the embedding dimension, L is the time delay, $\varepsilon_t = (\varepsilon_t, 0, \dots, 0)$ represents the stochastic component of the process, with $\{\varepsilon_t\}$ a sequence of iid random variables, and F is an $R^d \rightarrow R^d$ function that satisfies some general regularity conditions.

Given any two initial state vectors $X_0^{(1)}, X_0^{(2)}$ sufficiently close together, then after one time period has elapsed, the

following approximation will hold:

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

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

$$\lambda = \lim_{t \rightarrow \infty} \frac{1}{t} \ln \|J_{t-1} \cdot J_{t-2} \cdots J_0\|. \quad (8)$$

In practical terms, a bounded system with $\lambda > 0$ exhibits the sensitive dependence on initial conditions characteristic of chaos because, if this condition is satisfied, trajectories that start at two points arbitrarily close together will diverge exponentially as time passes.

Table 5. Lee-White-Granger Tests of Nonlinearity

Interval	AR(p)	p1	p2
<i>S&P 500</i>			
15-second	12	.0000	.0000
1-minute	3	.0000	.0000
5-minute	3	.0000	.0000
<i>FTSE-100</i>			
1-minute	7	.0000	.0000
5-minute	3	.0000	.0000
<i>DAX</i>			
1-minute	4	.0000	.0000
5-minute	5	.0005	.0000
<i>Nikkei</i>			
1-minute	11	.0000	.0000
5-minute	2	.0000	.0000
<i>S&P 500 Futures</i>			
1-minute	2	.0000	.0000
5-minute	3	.0000	.0001
<i>FTSE-100 Futures</i>			
1-minute	5	.0206	.0099
5-minute	1	.0001	.0000

NOTE: p1 and p2 are Hochberg-Bonferroni bounds on p value (Lee et al. 1993, sec. 6) with q* = 2 for p1, q* = 3 for p2 (Lee et al. 1993, sec. 4).

Most early work on L.E. estimation used the direct method of Wolf, Swift, and Vastano (1985). In essence, this approach involves averaging the observed divergence rates, which can be regarded as approximations to the left side of (7). If the series is chaotic, these divergences will tend to grow without limit. As McCaffrey, Ellner, Gallant, and Nychka (1992) showed, however, λ estimates derived in this fashion are liable to be biased upward when the process is contaminated by noise, as we must assume is the case here.

Rather than direct estimates of the rates of divergence, we prefer to rely on the Jacobian estimation methods of Briggs (1990), Nychka, Ellner, Gallant, and McCaffrey (1991), and Zeng, Pielke, and Eykholt (1992). This approach offers several advantages. First, it makes it possible to augment the approximation in (7) by the introduction of higher-order terms of the Taylor expansion. Moreover, the noise in the underlying process (6) can be smoothed out by using additional near neighbors in the estimation algorithm.

In the results reported here, we adopt the Zeng et al. (1992) estimation algorithm. This involves defining a "shell" (i.e., the zone between two spheres) rather than a ball from which near neighbors are to be selected, a modification intended to minimize the effect of noise on the estimates. In the present case, it was also preferred because it greatly reduced the difficulties presented by the large number of zero returns in the higher-frequency datasets (see Sec. 4).

As an alternative to this general line of approach, we also implement a neural-network algorithm, as set out by McCaffrey et al. (1992) and Nychka et al. (1992). This non-parametric regression procedure approximates the function F in (6) by a single hidden-layer feed-forward neural network with only a single output. More specifically, the esti-

mator takes the general form

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

where $G(u) = e^u / (1 + e^u)$ is the logistic distribution function, the γ_j are the weights modifying the inputs, the β_j are similarly applied to the outputs of the hidden units, and the μ_j are constant inputs equivalent to the column of ones in the standard econometric model [see McCaffrey et al. (1992) and Nychka et al. (1992) for details].

The neural-network algorithm outlined here offers two possible advantages over the nearest-neighbors approach. First, it avoids the so-called "curse of dimensionality," the increasing unreliability of estimates at higher dimensions. Second, because it is possible to obtain BIC values for each function approximation, we are able to derive a kind of numerical indication of the reliability of the L.E. estimates. In fact, the results of Nychka et al. (1992) suggest that this method works reasonably well on noisy systems even when the number of observations is far smaller than we have here.

3. THE DATA

We use a total of six series, four published cash indexes (the FTSE-100, the S&P 500, the DAX, and Nikkei) and two series constructed for futures on the FTSE-100 and S&P 500.

The main features of the FTSE-100 and the S&P 500 are well known. Both are value-weighted indexes compiled respectively at 1-minute and 15-second intervals. Neither index includes dividend payouts, and each represents a sizable proportion of its respective market, whether in terms of capitalization or turnover. In fact, although the FTSE accounts for about 70% of the value of the London market, the comparable figure for the S&P is 80%. Both indexes are arithmetic weighted means, with market capitalizations as weights. The most significant difference between the two is in the nature of the stock prices used. Although the S&P is based on the last transaction price of a constituent stock, the FTSE uses the midpoint of the best bid-and-ask prices taken from the London Stock Exchange's automated quotation system (SEAQ). As is well known, the U.S. index includes prices that may be "stale," in the sense that they are no longer up-to-date, in cases in which stocks trade less frequently than the 15-second interval at which the index is recomputed. On the other hand, the FTSE uses prices that are only notional quotes anyway because they apply neither to very small nor very large block sizes and because there may in some cases never be any trades at those prices within the minute. (See Sutcliffe 1993.)

The DAX is also a value-weighted arithmetic mean, but it includes only the 30 largest firms on the Frankfurt Stock Exchange, though it still represents 60% of the total market capitalization and over 65% of trading volume. It differs from the U.K. and U.S. indexes insofar as it incorporates dividend payments as reinvested income, with the weights being recomputed once a year to preserve a balance between high- and low-dividend stocks.

Table 6. Zeng, Pielke, and Eykholt (1992) L.E. Estimates for S&P

dim	15-second returns (N = 97,180)			1-minute returns (N = 24,500)			5-minute returns (N = 4,895)		
	max λ	$\sum \lambda$	K - Y dim	max λ	$\sum \lambda$	K - Y dim	max λ	$\sum \lambda$	K - Y dim
<i>nb = 20, q = 1</i>									
1	-1.403	-1.403		-.860	-.860		-.410	-.410	
2	-.507	-1.659		-.247	-1.168		.035	1.046	
3	-.253	-1.829		-.086	-1.385		.103	-.933	1.383
4	-.144	-1.900		-.022	-1.513		.135	-1.079	2.069
5	-.091	-1.993		.006	-1.615	1.034	.123	-1.183	2.446
6	-.061	-2.020		.019	-1.668	1.151	.129	-1.222	3.085
7	-.043	-2.075		.021	-1.765	1.229	.114	-1.304	3.502
8	-.032	-2.086		.026	-1.775	1.392	.100	-1.399	4.003
<i>nb = 40, q = 1</i>									
1	-1.463	-1.463		-.961	-.961		-.527	-.527	
2	-.542	-1.782		-.281	-1.271		-.027	-.848	
3	-.269	-1.978		-.098	-1.503		.072	-1.051	1.228
4	-.155	-2.090		-.037	-1.666		.100	-1.182	1.620
5	-.098	-2.187		-.005	-1.784		.099	-1.337	2.054
6	-.066	-2.236		.009	-1.868	1.053	.096	-1.405	2.365
7	-.047	-2.306		.011	-1.963	1.082	.096	-1.480	2.688
8	-.035	-2.341		.015	-2.006	1.149	.074	-1.613	2.944
<i>nb = 30, q = 2</i>									
1	-1.513	-1.513		-1.201	-1.201		-.682	-.682	
2	-.442	-1.428		-.258	-1.128		-.011	-.761	
3	-.127	-1.322		.025	-1.005	1.088	.232	-.654	2.098
4	.004	-1.251	1.020	.144	-.913	2.293	.333	-.516	3.247
5	.068	-1.175	2.050	.204	-.823	3.515	.399	-.455	4.349
6	.104	-1.109	3.274	.244	-.745	4.764	.424	-.365	5.468
7	.133	-1.054	4.579	.268	-.694	5.985	.456	-.307	6.558
8	.152	-.998	5.862	.294	-.630	7.074	.470	-.292	7.578
<i>nb = 60, q = 2</i>									
1	-1.717	-1.717		-1.440	-1.440		-.813	-.813	
2	-.553	-1.670		-.399	-1.370		-.169	-1.079	
3	-.202	-1.593		-.089	-1.297		.089	-.924	1.340
4	-.055	-1.548		.048	-1.215	1.295	.208	-.866	2.550
5	.017	-1.503	1.109	.118	-1.158	2.376	.267	-.766	3.802
6	.054	-1.448	1.661	.152	-1.078	3.575	.291	-.686	4.928
7	.076	-1.413	2.586	.175	-1.059	4.803	.318	-.650	6.071
8	.089	-1.366	3.716	.189	-.993	6.027	.351	-.589	7.136

NOTE: dim = dimensions estimated (1 to 8); q = order of estimation polynomial (1 or 2); nb = number of nearest neighbors used.

The Nikkei 225 Stock Average is the most widely quoted index of price movements on the Tokyo Stock Exchange (TSE). It is a price-weighted average of the 225 shares listed in the First Section of the TSE and is updated at 1-minute intervals throughout the trading day.

The two futures series were constructed following the established convention in the literature. Starting with raw transactions data, the price at any point of time, t_0 , was taken from the first recorded transaction after t_0 . In this fashion, we were able to generate futures series to match the cash indexes for the United Kingdom and the United States over the entire data period.

For each of the six series our data period runs from September 1 to November 30, 1991. We are concerned in this article with the return on the index, measured as the log change in the index level over a 15-second interval in the case of the S&P and also over a 1-minute and a 5-minute interval in all three cases. Descriptive statistics are presented in Table 2, page 4. There are several noteworthy characteristics.

First, the 1-minute datasets are of unequal size because the respective markets are not open for the same number of hours per day, or at least the proportion of the day over which their indexes are continuously updated varies—8 hours for the FTSE, $7\frac{1}{2}$ hours for the S&P, and only 4.5 hours for the DAX. Note that there are fewer observations on the futures than on the respective cash indexes, reflecting the fact that futures trading finishes well before the close of both London and New York stock markets. Thus, the U.K. dataset contains over 30,000 1-minute observations, compared to just under 25,000 for the New York market and only 11,000 for Frankfurt. Our largest dataset contains nearly 100,000 15-second returns on the S&P. The higher the data frequency, however, the greater the proportion of zero returns—well over 20% in the case of the 15-second S&P returns and over 10% of the 1-minute returns but fewer than 5% of the 5-minute returns.

Although all three series are approximately zero mean, all except the Nikkei are also characterized by negative

Table 7. L.E.'s for FTSE-100: Zeng et al. (1992) Estimates

dim	1-minute returns (N = 31,200)			5-minute returns (N = 6,230)		
	max λ	$\sum \lambda$	K - Y dim	max λ	$\sum \lambda$	K - Y dim
nb = 20, q = 1						
1	-.156	-.156		-.506	-.506	
2	.799	-.341	1.701	-.100	-.980	
3	.432	-.830	2.143	.027	-1.197	1.075
4	.253	-1.012	2.438	.071	-1.372	1.388
5	.152	-1.180	2.500	.074	-1.396	1.890
6	.155	-1.217	3.257	.089	-1.428	2.361
7	.152	-1.357	3.763	.099	-1.517	3.043
8	.147	-1.416	4.268	.094	-1.558	3.385
nb = 40, q = 1						
1	-.281	-.281		-.490	-.490	
2	.565	-.377	1.600	-.111	-1.095	
3	.244	-1.002	1.685	.004	-1.288	1.011
4	.160	-1.257	1.896	.049	-1.491	1.223
5	.104	-1.329	2.042	.055	-1.579	1.461
6	.134	-1.384	2.671	.060	-1.612	1.776
7	.135	-1.539	3.142	.064	-1.701	2.190
8	.130	-1.608	3.499	.057	-1.778	2.366
nb = 30, q = 2						
1	1.025	1.025		-.582	-.582	
2	1.204	.079		-.080	-.994	
3	1.409	.598		.153	-.838	1.697
4	1.060	-.158	3.871	.263	-.753	2.973
5	.690	-.459	4.563	.330	-.655	4.111
6	.460	-.639	5.140	.374	-.527	5.251
7	.380	-.687	6.047	.416	-.478	6.344
8	.284	-.824	6.612	.456	-.400	7.446
nb = 60, q = 2						
1	.471	.471		-.559	-.559	
2	.748	-.265	1.738	-.188	-1.291	
3	.712	-.373	2.596	.042	-1.148	1.122
4	.439	-.901	2.978	.128	-1.110	2.089
5	.499	-.847	4.181	.198	-1.006	3.203
6	.358	-.960	4.415	.238	-.907	4.431
7	.293	-.996	5.174	.276	-.858	5.660
8	.198	-1.153	5.415	.303	-.792	6.809

NOTE: See note to Table 6.

skewness, for reasons that are unclear. As expected, there is clear evidence both of leptokurtosis and strong autocorrelation in all cases, but most especially in the case of the DAX. It is somewhat surprising that the degree of leptokurtosis varies little over the different frequencies for the FTSE and the S&P, whereas the increase is extremely marked for the DAX.

Comparing the respective futures with the cash returns, it is noticeable that, although the futures distributions are less fat-tailed, they also exhibit greater volatility. This pattern is consistent with the argument of, for example, Whaley (1993) that infrequent trading is likely to increase the volatility of the futures relative to the spot, other things being equal. Specifically, Whaley (1993) showed that the less frequently stocks trade and the greater is the moving average component of the futures price induced by so-called "bid-ask bounce," the greater will be the volatility of the futures relative to the cash, other things being equal. In our

case, stocks often traded less frequently than our 1-minute observation interval, especially in the U.K. market. It is noteworthy that our futures series are volatile by comparison with the cash, and moreover that the disparity is greater at the higher frequency.

4. RESULTS

4.1 Nonlinearity

The results of the BDS tests are given in Table 3, page 5 (for the cash indexes), and Table 4, page 6 (for the futures). For each series, the computed values are given first for the residuals from a linear filter, followed by the residuals from BIC-minimizing GARCH(p, q) models. To save space, we present no results at frequencies lower than 5 minutes, though our results for observations at 15-minute, 30-minute, and 1-hour intervals are entirely consistent with the pattern of the high-frequency results given here (see ACW for the FTSE-100).

Table 8. L.E.'s for DAX: Zeng et al. (1992) Estimates

dim	1-minute returns (N = 31,200)			5-minute returns (N = 6,230)		
	max λ	$\sum \lambda$	K - Y dim	max λ	$\sum \lambda$	K - Y dim
nb = 20, q = 1						
1	-.468	-.468		-.306	-.306	
2	-.038	-.800		.018	-.780	1.023
3	.074	-1.005	1.248	.065	-1.066	1.216
4	.111	-1.139	1.827	.090	-1.137	1.752
5	.120	-1.220	2.339	.105	-1.351	2.322
6	.130	-1.301	2.910	.106	-1.287	3.040
7	.130	-1.356	3.521	.114	-1.371	3.767
8	.126	-1.383	4.164	.126	-1.333	4.535
nb = 40, q = 1						
1	-.509	-.509		-.441	-.441	
2	-.058	-.839		-.080	-.988	
3	.053	-1.067	1.164	.006	-1.215	1.015
4	.090	-1.275	1.538	.049	-1.346	1.291
5	.100	-1.345	2.079	.057	-1.529	1.544
6	.106	-1.439	2.477	.071	-1.495	2.102
7	.105	-1.492	3.020	.053	-1.580	2.412
8	.096	-1.553	3.415	.075	-1.664	3.102
nb = 30, q = 2						
1	-.707	-.707		-.612	-.612	
2	-.088	-.835		.010	-.679	1.014
3	.178	-.733	1.931	.214	-.685	2.058
4	.304	-.627	3.132	.356	-.544	3.225
5	.387	-.487	4.296	.430	-.379	4.455
6	.455	-.401	5.435	.506	-.266	5.639
7	.489	-.298	6.571	.618	-.095	6.867
8	.549	-.226	7.680	.705	.078	
nb = 60, q = 2						
1	-.818	-.818		-.785	-.785	
2	-.195	-1.040		-.182	-1.036	
3	.048	-1.014	1.168	.061	-1.031	1.222
4	.189	-.911	2.421	.184	-.858	2.453
5	.264	-.788	3.735	.275	-.737	3.841
6	.329	-.735	5.007	.330	-.680	5.076
7	.356	-.630	6.103	.402	-.570	6.236
8	.393	-.536	7.230	.453	-.382	7.452

NOTE: See note to Table 6.

Table 9. L.E.'s for Nikkei: Zeng et al. (1992) Estimates

dim	1-minute returns (N = 16,345)			5-minute returns (N = 6,230)		
	max λ	$\sum \lambda$	K - Y dim	max λ	$\sum \lambda$	K - Y dim
nb = 20, q = 1						
1	-.5603	-.5600		-.0984	-.0980	
2	-.0695	-.8380		.2435	-.3300	1.4250
3	.0681	-1.0230	1.2230	.2914	-.5100	2.2510
4	.1087	-1.1220	1.7930	.2736	-.6550	3.0600
5	.1173	-1.2340	2.3030	.2720	-.7720	3.8000
6	.1239	-1.2630	2.9710	.2465	-.8380	4.5580
7	.1184	-1.3610	3.5310	.2075	-1.0130	5.1070
8	.1128	-1.4020	4.1240	.1906	-1.1080	5.6560
nb = 40, q = 1						
1	-.6670	-.6670		-.2412	-.2410	
2	-.1096	-.9390		.1339	-.5150	1.2060
3	.0271	-1.1360	1.0770	.2363	-.6880	2.0540
4	.0706	-1.2770	1.3700	.2315	-.7790	2.7120
5	.0849	-1.3870	1.7900	.2354	-.8890	3.4450
6	.0899	-1.4350	2.2330	.2227	-.9690	4.1700
7	.0888	-1.5390	2.6550	.1741	-1.1860	4.3170
8	.0813	-1.5850	3.1160	.1629	-1.2570	4.7810
nb = 30, q = 2						
1	-.8028	-.8030		-.4618	-.4620	
2	-.0318	-.7610		.1071	-.5820	1.1550
3	.2109	-.6610	2.0730	.2905	-.5540	2.2230
4	.3193	-.5630	3.1900	.4085	-.4090	3.4250
5	.3795	-.4720	4.3280	.4685	-.3180	4.5570
6	.4273	-.3540	5.4870	.4880	-.2860	5.5940
7	.4719	-.2840	6.5960	.4755	-.3200	6.5640
8	.5079	-.2180	7.6930	.4899	-.1930	7.7250
nb = 60, q = 2						
1	-.9974	-.9970		-.7321	-.7320	
2	-.1917	-1.0300		-.0820	-.8510	
3	.0738	-.9580	1.2820	.1562	-.7990	1.8040
4	.1886	-.8640	2.5100	.2870	-.6230	3.1160
5	.2586	-.7900	3.7560	.3667	-.5910	4.1650
6	.3025	-.6990	5.0210	.3607	-.5840	5.1710
7	.3344	-.6200	6.1280	.3534	-.6070	6.1680
8	.3642	-.5590	7.2210	.3565	-.5620	7.2200

NOTE: See note to Table 6.

There are several features common to the results for all the series. First, the degree of departure from independence is invariably greatest at the highest frequencies. Second, in no case are the raw data independent, and moreover the dependence is not removed by linear filtering. Perhaps the only surprise is that the apparent departure from independence is smallest for the DAX and largest for the FTSE.

The evidence from the GARCH residuals is more mixed. On the one hand, in every case the BDS statistics are substantially reduced compared to the residuals from the linear filter. On the other hand, whether one concludes that GARCH explains all the structure or not depends very much on the embedding dimension, m , and the value of the neighborhood parameter, ϵ/σ , used in estimation. For example, the results for the 5-minute DAX returns are consistent with iid when $m \geq 2$ and $\epsilon/\sigma \geq .75$. On the other hand, iid is rejected for all parameter values at the 1-minute frequency. Similarly, for the other series there is a tendency for the BDS statistic to fall as the degrees of freedom are exhausted

by higher values of m and ϵ/σ .

Table 5, page 7, presents the results for the LWG neural-net-based test, in the form of Hochberg-Bonferroni bounds on the probability of the test statistic under the null of an iid process. As can be seen from the near-zero values in the table, the outcome is overwhelming for all series at all frequencies with the sole exception of the 1-minute observations on FTSE futures, in which the null hypothesis of independence is accepted at a probability bounded from above by 1% or 2%, depending on the network order. In all other cases, the maximum probability of independence is negligible. The implication is clear: The nonlinear component unambiguously improves forecast accuracy.

4.2 Sensitive Dependence

In Tables 6-11, we present estimates of the L.E.'s for the six series, derived from implementation of the Zeng et

Table 10. L.E.'s for S&P 500 Futures: Zeng et al. (1992) Estimates

dim	1-minute returns (N = 31,200)			5-minute returns (N = 6,230)		
	max λ	$\sum \lambda$	K - Y dim	max λ	$\sum \lambda$	K - Y dim
nb = 20, q = 1						
1	-1.324	-1.324		-.823	-.823	
2	-.486	-1.617		-.205	-1.119	
3	-.239	-1.772		-.048	-1.321	
4	-.139	-1.877		.007	-1.438	1.028
5	-.086	-1.947		.027	-1.561	1.147
6	-.059	-1.963		.032	-1.609	1.259
7	-.042	-2.031		.031	-1.688	1.347
8	-.032	-2.026		.029	-1.714	1.474
nb = 40, q = 1						
1	-1.420	-1.420		-.920	-.920	
2	-.530	-1.764		-.241	-1.245	
3	-.259	-1.943		-.071	-1.445	
4	-.151	-2.063		-.005	-1.596	
5	-.096	-2.160		.017	-1.755	1.069
6	-.071	-2.222		.026	-1.812	1.141
7	-.049	-2.260		.029	-1.872	1.210
8	-.037	-2.284		.027	-1.908	1.244
nb = 30, q = 2						
1	-1.568	-1.568		-1.279	-1.279	
2	-.440	-1.407		-.318	-1.190	
3	-.116	-1.276		-.011	-1.051	
4	.008	-1.211	1.048	.131	-.909	2.265
5	.075	-1.130	2.166	.190	-.834	3.504
6	.113	-1.068	3.428	.233	-.791	4.691
7	.136	-1.055	4.620	.251	-.759	5.838
8	.152	-1.013	5.833	.274	-.677	7.024
nb = 60, q = 2						
1	-1.815	-1.815		-1.599	-1.599	
2	-.581	-1.692		-.529	-1.527	
3	-.211	-1.574		-.161	-1.390	
4	-.058	-1.505		.013	-1.250	1.075
5	.017	-1.456	1.112	.088	-1.187	2.129
6	.052	-1.414	1.672	.134	-1.143	3.368
7	.074	-1.387	2.573	.153	-1.101	4.586
8	.086	-1.369	3.641	.173	-1.021	5.769

NOTE: See note to Table 6.

Table 11. L.E.'s for FTSE-100 Futures: Zeng et al. (1992) Estimates

dim	1-minute returns (N = 26,389)			5-minute returns (N = 5,277)		
	max λ	∑ λ	K - Y dim	max λ	∑ λ	K - Y dim
nb = 20, q = 1						
1	-2.272	-2.272		-1.303	-1.303	
2	-.816	-2.100		-.431	-1.509	
3	-.455	-2.200		-.217	-1.695	
4	-.293	-2.216		-.124	-1.793	
5	-.208	-2.217		-.083	-1.917	
6	-.154	-2.236		-.061	-1.925	
7	-.119	-2.271		-.047	-1.972	
8	-.106	-2.408		-.038	-2.036	
nb = 40, q = 1						
1	-2.310	-2.310		-1.424	-1.424	
2	-.905	-2.295		-.490	-1.668	
3	-.524	-2.477		-.243	-1.857	
4	-.334	-2.470		-.146	-2.012	
5	-.250	-2.518		-.097	-2.126	
6	-.179	-2.495		-.072	-2.179	
7	-.139	-2.542		-.052	-2.208	
8	-.122	-2.605		-.042	-2.320	
nb = 30, q = 2						
1	-2.052	-2.052		-1.616	-.162	
2	-.709	-1.839		-.437	-1.377	
3	-.335	-1.828		-.111	-1.246	
4	-.165	-1.679		.026	-1.150	1.182
5	-.094	-1.702		.087	-1.140	2.315
6	-.031	-1.552		.118	-1.040	3.507
7	.003	-1.529	1.050	.148	-.991	4.780
8	.027	-1.538	2.135	.179	-.972	6.068
nb = 60, q = 2						
1	-2.205	-2.205		-1.909	-1.909	
2	-.834	-2.081		-.587	-1.681	
3	-.452	-2.150		-.204	-1.547	
4	-.254	-2.016		-.053	-1.490	
5	-.176	-2.093		.001	-1.484	1.007
6	-.098	-1.907		.040	-1.411	1.657
7	-.067	-1.927		.065	-1.369	2.669
8	-.044	-1.943		.082	-1.379	3.906

NOTE: See note to Table 6.

al. (1992) algorithm. On the one hand, the estimated L.E.'s are, for the most part, negative so that the sum $\sum \lambda_i, \lambda > 0$, which is a measure of the Kolmogorov entropy, is negative in virtually every case. Our estimates of the maximal L.E. values vary in sign, however, with positive values occurring often at the higher dimensions, though it should be noted that positive L.E. estimates occur far less often for the futures than for the cash indexes. Where it is defined (i.e., where there are positive L.E.'s), we present the implied Kaplan-Yorke dimension in the final column of the tables. It is given by

$$D_{KY} = k + \frac{1}{|\lambda_{k+1}|} \sum_{i=1}^k \lambda_i, \quad \sum_{j=1}^k \lambda_j \geq 0 > \sum_{j=1}^{k+1} \lambda_j; \quad (10)$$

that is, k is the largest integer for which $\sum \lambda_i > 0$. Clearly, our estimates are highly sensitive to the choice of embedding dimension and polynomial order.

The results of applying the neural-net estimation methods of Nychka et al. (1992) are presented in Figures 1-5. As they recommend, in each case we examine the results from two perspectives.

The scatterplots (Figs. 1-3) show the spread of λ estimates for different (L, d, q) combinations, where L is the time delay used, $d = 1, \dots, 6$ is the embedding dimension, and q is the selection parameter equal to the number of units in the hidden layer of the net, chosen so as to minimize the BIC. This means that we have a total of $6 \times 50 = 300$ parameter combinations in all cases. We then plot the 10 BIC-minimizing estimates of the L.E. associated with each (L, d, q) triplet—that is, from 200 to 300 points in each figure.

The results are set out differently in the line graphs (Figs. 4 and 5). Here, for each embedding dimension $d = 1, \dots, 6$, the graphs show

1. The BIC-minimizing L.E. estimate (the "best fit")
2. The mean of the 10 BIC-minimizing L.E. estimates (the unbroken line)
3. Given the standard deviation of the 10 BIC-minimizing L.E. estimates, a range of one standard deviation about the mean ("Upper" and "Lower" marked by the two dotted lines)

For illustration purposes, Figure 1 shows the output from applying this algorithm to a dataset consisting of 500 points generated by simulating the Lorenz mapping. It can be seen from the scatterplot that as the computed BIC value falls, the L.E. estimate stabilizes in the neighborhood of the true value of .0745 per .05 time units (i.e., 1.49 per full time unit) (Wolf et al. 1985). The line graph for this system (not shown here) rises steeply from negative to positive as the embedding dimension goes from 1 to 3 and subsequently levels off as the estimates settle down in the region of the correct value.




Figure 1. Lorenz System (Delay = 2).




Figure 2. S&P 500 Cash 1-Minute Returns (Delay = 6).

In Figures 2–5, we show scatterplots for the S&P 500 cash and futures, and line graphs for the FTSE, at the 1-minute frequency only to save space. An earlier working-paper version of this article, available from us on request, presents both line and scatter graphs for each of our six series at both 5-minute and 1-minute frequencies, with results that are qualitatively very similar to those given here. By comparison with our simulations on the Lorenz equations, the results for the cash indexes are much less stable and therefore far harder to interpret, as can be seen from Figure 2. All that can be said with any confidence is that there is no evidence of a positive L.E. and no clear pattern of convergence toward a particular value. The line graphs are perhaps more informative, in most cases rising as the




Figure 3. S&P 500 Futures 1-Minute Returns (Delay = 2).




Figure 4. FTSE-100 Cash 1-Minute Returns (Delay = 9).

dimension increases without much sign of ever leveling off, though there is some evidence of an L.E. of about -0.07 in the FTSE 1-minute data (Fig. 4). It is noteworthy that the graphs for the FTSE bear a reassuring resemblance to those for the first six months of 1993, as reported by ACW. This may possibly be evidence that, for this particular index at least, the data-generating process has remained reasonably stable in recent years.

Looking at Figures 3 and 5, the L.E. estimates for the futures look slightly less erratic than those for the respective spot indexes. This is especially true in the case of the S&P futures. Nonetheless, even here it is hard to justify a conclusion more precise than simply that the dominant L.E. is almost certainly negative.

Two general conclusions appear to be justified. First, there is little sign of any sensitive dependence. Second, the true dimension of the systems generating our datasets is unclear, probably infinite.




Figure 5. FTSE-100 Futures 1-Minute Returns (Delay = 2).

5. CONCLUSIONS

In this article, we have provided what we believe to be reasonably robust answers to two questions regarding the behavior of stock-market index returns. In the first instance, we are able to reject the hypothesis of independence in favor of a nonlinear structure for all six data series. To some extent, it seems that this dependence can be attributed to volatility clustering, though this phenomenon appears unlikely to provide a complete explanation.

When we proceed to estimation of L.E.'s, however, we find no support for the view that the underlying processes are chaotic (i.e., exhibit sensitive dependence on initial conditions). Instead, we interpret the evidence as, for the most part, supporting the view that the data processes are dominated if not actually swamped by noise. Thus, although there might be an underlying deterministic nonlinear but nonchaotic process in the data, there is almost certainly also a stochastic component whose presence cannot be ignored. These conclusions are borne out by our results using both standard nonparametric methods and neural-network-based approaches.

Several important questions remain on the research agenda. Are the properties of stock-market index returns matched by those of their components? That is, are the characteristics of returns on individual stocks qualitatively similar to those described in this article? How general are these results? Last but by no means least, insofar as there are discernible patterns in returns, can they be forecast with enough accuracy and far enough into the future to generate trading profits?

ACKNOWLEDGMENTS

The second author is grateful for the support of the Nuffield Foundation. We record our thanks for data supplied by Graham Hart of the Quality of Markets Unit of the International Stock Exchange (London), Anthony Barnett at the CME, and Susanne Lotz at the Deutsche Börse. We also thank S. Ellner, D. W. Nychka, and A. R. Gallant for use of their LENNS software and the Manchester Computing Centre for providing time on their Fujitsu VPX 240/10 vector processor. Helpful comments on an earlier draft were received from Xubin Zeng, Craig Hiemstra, Stephen Taylor, Lawrence Kryzanowski, and participants in the 1995 conferences of the European Financial Management Association, European Finance Association, the European Institute for Advanced Studies in Management Conference on Empirical Research in Finance, the Asia-Pacific Finance Association, and seminars in the Statistics Departments of the Universities of Manchester, Dundee, and Glasgow and at Swansea Business School. As always, the responsibility for any remaining weaknesses is ours alone.

[Received February 1995. Revised January 1996.]

REFERENCES

Abhyankar, A., Copeland, L. S., and Wong, W. (1994), "Moment Condition Failure in High Frequency Financial Data: Evidence From the S&P 500," mimeo.

- (1995), "Nonlinear Dynamics in Real-Time Equity Market Indices: Evidence From the UK," *Economic Journal*, 105, 864–880.
- Briggs, K. (1990), "An Improved Method for Estimating Liapunov Exponents of Chaotic Time Series," *Physics Letters A*, 151, 27–32.
- Brock, W. A., Dechert, W., Scheinkman, J. A., and LeBaron, B. (1987), "A Test for Independence Based Upon the Correlation Dimension," working paper, University of Wisconsin-Madison, Dept. of Economics.
- Brock, W. A., Hsieh, D., and LeBaron, B. (1991), *Nonlinear Dynamics, Chaos and Instability: Statistical Theory and Empirical Evidence*, Cambridge, MA: MIT Press.
- Brockett, P. L., Hinich, M. J., and Patterson, D. (1988), "Bispectral-Based Tests for the Detection of Gaussianity and Nonlinearity in Time Series," *Journal of the American Statistical Association*, 83, 657–664.
- Devaney, R. L. (1989), *An Introduction to Chaotic Dynamical Systems* (2nd ed.), Redwood City, CA: Addison-Wesley.
- Eckmann, J. P. S., Kamphorst, S. O., Ruelle, D., and Scheinkman, J. A. (1988), "Lyapunov Exponents for Stock Returns," in *The Economy as an Evolving Complex System*, eds. P. W. Anderson, K. J. Arrow, and D. Pines, New York: Addison-Wesley, pp. 301–304.
- Eldridge, R. M., and Coleman, M. P. (1993), "The British FTSE-100 Index: Chaotically Deterministic or Random?" working paper, Fairfield University, School of Business.
- Fama, E. F. (1991), "Efficient Capital Markets: II," *Journal of Finance*, 46, 1575–1618.
- Frank, M., and Stengos, T. (1989), "Measuring the Strangeness of Gold and Silver Rates of Return," *Review of Economic Studies*, 56, 553–567.
- Hinich, M. (1982), "Testing for Gaussianity and Linearity of a Stationary Time Series," *Journal of Time Series Analysis*, 3, 169–176.
- Hinich, M., and Patterson, D. (1989), "Evidence of Nonlinearity in the Trade-by-Trade Stock Market Return Generating Process," in *Economic Complexity: Chaos, Transport, Bubbles and Nonlinearity*, eds. W. Barret, J. Geweke, and K. Shell, Cambridge, U.K.: Cambridge University Press, pp. 383–409.
- Hochberg, Y. (1988), "A Sharper Bonferroni Procedure for Multiple Tests of Significance," *Biometrika*, 75, 800–802.
- Hsieh, D. (1989), "Testing for Nonlinear Dependence in Daily Foreign Exchange Rates," *Journal of Business*, 62, 339–359.
- (1991), "Chaos and Nonlinear Dynamics: Applications to Financial Markets," *Journal of Finance*, 47, 1145–1189.
- (1993), "Using Non-linear Methods to Search for Risk Premia in Currency Futures," *Journal of International Economics*, 35, 113–132.
- Kodres, L. E., and Papell, D. H. (1991), "Nonlinear Dynamics in the Foreign Exchange Futures Market," working paper, University of Michigan, School of Business and Administration.
- Kröger, H., and Kügler, P. (1992), "Nonlinearity in Foreign Exchange Markets: A Different Perspective," working paper, Universität Mannheim, Germany.
- Lee, T.-H., White, H., and Granger, C. W. J. (1993), "Testing for Neglected Nonlinearity in Time Series Models: A Comparison of Neural Network Methods and Alternative Tests," *Journal of Econometrics*, 56, 269–290.
- Loretan, M., and Phillips, P. C. B. (1994), "Testing the Covariance Stationarity of Heavy-Tailed Time Series: An Overview of the Theory With Applications to Several Financial Datasets," *Journal of Empirical Finance*, 1, 211–248.
- Mayfield, E. S., and Mizrach, B. (1989), "On Determining the Dimension of Real Time Stock Price Data," working paper, Boston College, Dept. of Economics.
- 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*, 87, 682–695.
- Nychka, D., Ellner, S., Gallant, A. R., and McCaffrey, D. (1992), "Finding Chaos in Noisy Systems," *Journal of the Royal Statistical Society, Ser. B*, 52, 399–426.
- Phillipatos, G. C., Pilarina, 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, Dept. of Economics.
- Scheinkman, J. A., and LeBaron, B. (1989), "Nonlinear Dynamics and Stock Returns," *Journal of Business*, 62, 311–337.
- Sutcliffe, C. M. S. (1993), "Stock Index Futures: Theories and International Evidence," London: Chapman and Hall.
- Takens, F. (1981), *Detecting Strange Attractors in Turbulence, in Dynamical Systems and Turbulence* (Lecture Notes in Mathematics 898), Berlin:

- Springer-Verlag, pp. 366–381.
- Tata, F. (1991), “Is the Foreign Exchange Market Characterized by Nonlinearity?” Discussion Paper 118, Discussion Paper Series, London School of Economics, Financial Markets Group.
- Vaidyanathan, R., and Krehbiel, T. (1992), “Does the S&P 500 Futures Mispricing Series Exhibit Nonlinear Dependence Across Time?” *Journal of the Futures Market*, 12, 659–677.
- Vassilicos, J. A. (1990), “Are Financial Markets Chaotic? A Preliminary Study of the Foreign Exchange Market,” Discussion Paper 86, Financial Markets Group Discussion Paper Series, London School of Economics.
- 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 143, Discussion Paper Series, London School of Economics, Financial Markets Group.
- Wolf, A. J., Swift, J., and Vastano, J. (1985), “Determining Lyapunov Exponents From a Time Series,” *Physica*, 16D, 285–317.
- Whaley, R. E. (1993), “Predictability of Stock Index Basis Changes: Comments on ‘Behaviour of the FTSE 100 Basis’ by Chris Strickland and Xinzhong Xu,” *Review of Futures Markets*, 12, 503–508.
- Zeng, X., Pielke, R. A., and Eykholt, R. (1992), “Extracting Lyapunov Exponents From Short Time Series of Low Precision,” *Modern Physics Letters*, 6, 3229–3232.

LINKED CITATIONS

- Page 1 of 2 -



You have printed the following article:

Uncovering Nonlinear Structure in Real-Time Stock-Market Indexes: The S&P 500, the DAX, the Nikkei 225, and the FTSE-100

A. Abhyankar; L. S. Copeland; W. Wong

Journal of Business & Economic Statistics, Vol. 15, No. 1. (Jan., 1997), pp. 1-14.

Stable URL:

<http://links.jstor.org/sici?sici=0735-0015%28199701%2915%3A1%3C1%3AUNSIRS%3E2.0.CO%3B2-Q>

This article references the following linked citations:

References

Nonlinear Dynamics in Real-Time Equity Market Indices: Evidence from the United Kingdom

A. Abhyankar; L. S. Copeland; W. Wong

The Economic Journal, Vol. 105, No. 431. (Jul., 1995), pp. 864-880.

Stable URL:

<http://links.jstor.org/sici?sici=0013-0133%28199507%29105%3A431%3C864%3ANDIREM%3E2.0.CO%3B2-1>

Bispectral-Based Tests for the Detection of Gaussianity and Linearity in Time Series

Patrick L. Brockett; Melvin J. Hinich; Douglas Patterson

Journal of the American Statistical Association, Vol. 83, No. 403. (Sep., 1988), pp. 657-664.

Stable URL:

<http://links.jstor.org/sici?sici=0162-1459%28198809%2983%3A403%3C657%3ABTFTDO%3E2.0.CO%3B2-U>

Efficient Capital Markets: II

Eugene F. Fama

The Journal of Finance, Vol. 46, No. 5. (Dec., 1991), pp. 1575-1617.

Stable URL:

<http://links.jstor.org/sici?sici=0022-1082%28199112%2946%3A5%3C1575%3AECMI%3E2.0.CO%3B2-L>

Measuring the Strangeness of Gold and Silver Rates of Return

Murray Frank; Thanasis Stengos

The Review of Economic Studies, Vol. 56, No. 4. (Oct., 1989), pp. 553-567.

Stable URL:

<http://links.jstor.org/sici?sici=0034-6527%28198910%2956%3A4%3C553%3AMTSOGA%3E2.0.CO%3B2-I>

LINKED CITATIONS

- Page 2 of 2 -



A Sharper Bonferroni Procedure for Multiple Tests of Significance

Yosef Hochberg

Biometrika, Vol. 75, No. 4. (Dec., 1988), pp. 800-802.

Stable URL:

<http://links.jstor.org/sici?sici=0006-3444%28198812%2975%3A4%3C800%3AASBPFM%3E2.0.CO%3B2-N>

Testing for Nonlinear Dependence in Daily Foreign Exchange Rates

David A. Hsieh

The Journal of Business, Vol. 62, No. 3. (Jul., 1989), pp. 339-368.

Stable URL:

<http://links.jstor.org/sici?sici=0021-9398%28198907%2962%3A3%3C339%3ATFNDID%3E2.0.CO%3B2-8>

Estimating the Lyapunov Exponent of a Chaotic System With Nonparametric Regression

Daniel F. McCaffrey; Stephen Ellner; A. Ronald Gallant; Douglas W. Nychka

Journal of the American Statistical Association, Vol. 87, No. 419. (Sep., 1992), pp. 682-695.

Stable URL:

<http://links.jstor.org/sici?sici=0162-1459%28199209%2987%3A419%3C682%3AETLEOA%3E2.0.CO%3B2-N>

Nonlinear Dynamics and Stock Returns

Jose A. Scheinkman; Blake LeBaron

The Journal of Business, Vol. 62, No. 3. (Jul., 1989), pp. 311-337.

Stable URL:

<http://links.jstor.org/sici?sici=0021-9398%28198907%2962%3A3%3C311%3ANDASR%3E2.0.CO%3B2-U>