

MARTINGALES, NONLINEARITY, AND CHAOS

WILLIAM A. BARNETT

Department of Economics, Campus Box 1208
Washington University
One Brookings Drive
St. Louis, Missouri 63130-4810

APOSTOLOS SERLETIS

Department of Economics
The University of Calgary
Calgary, Alberta T2N 1N4

June 3, 1998

Abstract

In this article we provide a review of the literature with respect to the efficient markets hypothesis and chaos. In doing so, we contrast the martingale behavior of asset prices to nonlinear chaotic dynamics, discuss some recent techniques used in distinguishing between probabilistic and deterministic behavior in asset prices, and report some evidence. Moreover, we look at the controversies that have arisen about the available tests and results, and raise the issue of whether dynamical systems theory is practical in finance.

Keywords: Efficient markets hypothesis, Chaotic dynamics.

JEL classification: C22, G14

*We would like to thank Cars Hommes and Dee Dechert for useful comments.

1 Introduction

Recently the efficient markets hypothesis and the notions connected with it have provided the basis for a great deal of research in financial economics. A voluminous literature has developed supporting this hypothesis. Briefly stated, the hypothesis claims that asset prices are rationally related to economic realities and always incorporate all the information available to the market. This implies the absence of exploitable excess profit opportunities. However, despite the widespread allegiance to the notion of market efficiency, a number of studies have suggested that certain asset prices are not rationally related to economic realities. For example, Summers (1986) argues that market valuations differ substantially and persistently from rational valuations and that existing evidence (based on common techniques) does not establish that financial markets are efficient.

Although most of the empirical tests of the efficient markets hypothesis are based on linear models, interest in nonlinear chaotic processes has in the recent past experienced a tremendous rate of development. There are many reasons for this interest, one of which being the ability of such processes to generate output that mimics the output of stochastic systems, thereby offering an alternative explanation for the behavior of asset prices. In fact, the possible existence of chaos could be exploitable and even invaluable. If, for example, chaos can be shown to exist in asset prices, the implication would be that profitable, nonlinearity-based trading rules exist (at least in the short run and provided the actual generating mechanism is known). Prediction, however, over long periods is all but impossible, due to the sensitive dependence on initial conditions property of chaos.

In this paper, we survey the recent literature with respect to the efficient markets hypothesis and chaos. In doing so, in the next two sections we briefly discuss the efficient markets hypothesis and some of the more recent testing methodologies. In section 4, we provide a description of the key features of the available tests for independence, nonlinearity, and chaos, focusing explicit attention on each test's ability to detect chaos. In section 5, we present a discussion of the empirical evidence on macroeconomic and (mostly) financial data, and in section 6, we look at the controversies that have arisen about the available tests and address some important questions regarding the power of some of these tests. The final section concludes.

2 The Martingale Hypothesis

Standard asset pricing models typically imply the ‘martingale model’, according to which tomorrow’s price is expected to be the same as today’s price. Symbolically, a stochastic process x_t follows a martingale if

$$E_t(x_{t+1}|\Omega_t) = x_t \tag{1}$$

where Ω_t is the time t information set - assumed to include x_t . Equation (1) says that if x_t follows a martingale the best forecast of x_{t+1} that could be constructed based on current information Ω_t would just equal x_t .

Alternatively, the martingale model implies that $(x_{t+1} - x_t)$ is a ‘fair game’ (a game which is neither in your favor nor your opponent’s)¹

$$E_t[(x_{t+1} - x_t)|\Omega_t] = 0. \tag{2}$$

Clearly, x_t is a martingale if and only if $(x_{t+1} - x_t)$ is a fair game. It is for this reason that fair games are sometimes called ‘martingale differences’.² The fair game model (2) says that increments in value (changes in price adjusted for dividends) are unpredictable, conditional on the information set Ω_t . In this sense, information Ω_t is fully reflected in prices and hence useless in predicting rates of return. The hypothesis that prices fully reflect available information has come to be known as the ‘efficient markets hypothesis’.

In fact Fama (1970) defined three types of (informational) capital market efficiency (not to be confused with allocational or Pareto-efficiency), each of which is based on a different notion of exactly what type of information is understood to be relevant. In particular, markets are weak-form, semistrong-form, and strong-form efficient if the information set includes past prices and returns alone, all public information, and any information public as well as private, respectively. Clearly, strong-form efficiency implies semistrong-form efficiency, which in turn implies weak-form efficiency, but the reverse implications do not follow, since a market easily could be weak-form efficient

¹A stochastic process z_t is a fair game if z_t has the property $E_t(z_{t+1}|\Omega_t) = 0$.

²The martingale process is a special case of the more general submartingale process. In particular, x_t is a ‘submartingale’ if it has the property $E_t(x_{t+1}|\Omega_t) \geq x_t$. In terms of the $(x_{t+1} - x_t)$ process, the submartingale model implies that $E_t[(x_{t+1} - x_t)|\Omega_t] \geq 0$ and embodies the concept of a superfair game. LeRoy (1989, pp. 1593-4) also offers an example in which $E_t[(x_{t+1} - x_t)|\Omega_t] \leq 0$, in which case x_t will be a ‘supermartingale’, embodying the concept of a subfair game.

but not semistrong-form efficient or semistrong-form efficient but not strong-form efficient.

The martingale model given by (1) can be written equivalently as

$$x_{t+1} = x_t + \varepsilon_t$$

where ε_t is the martingale difference. When written in this form the martingale looks identical to the ‘random walk model’ - the forerunner of the theory of efficient capital markets. The martingale, however, is less restrictive than the random walk. In particular, the martingale difference requires only independence of the conditional expectation of price changes from the available information, as risk neutrality implies, whereas the (more restrictive) random walk model requires this and also independence involving the higher conditional moments (i.e., variance, skewness, and kurtosis) of the probability distribution of price changes.

In fact, Campbell, Lo, and MacKinlay (1997) distinguish between three versions of the random walk hypothesis — the ‘independently-and-identically-distributed-returns’ version, the ‘independent-returns’ version, and the version of ‘uncorrelated-returns’ — see Campbell, Lo, and MacKinlay (1997) for more details. The martingale difference model, by not requiring probabilistic independence between successive price changes, is entirely consistent with the fact that price changes, although uncorrelated, tend not to be independent over time but to have clusters of volatility and tranquility (i.e., dependence in the higher conditional moments) - a phenomenon originally noted for stock market prices by Mandelbrot (1963) and Fama (1965).

3 Tests of the Martingale Hypothesis

The random walk and martingale hypotheses imply a unit root in the level of the price or logarithm of the price series - notice that a unit root is a necessary but not sufficient condition for the random walk and martingale models to hold. Hence, these models can be tested using recent advances in the theory of integrated regressors. The literature on unit root testing is vast and, in what follows, we shall only briefly illustrate some of the issues that have arisen in the broader search for unit roots in financial asset prices.³

³It is to be noted that unit root tests have low power against relevant alternatives. Also, as Granger (1995) points out, nonlinear modelling of nonstationary variables is a

Nelson and Plosser (1982), using the augmented Dickey-Fuller (ADF) unit root testing procedure [see Dickey and Fuller (1981)] test the null hypothesis of ‘difference-stationarity’ against the ‘trend-stationarity’ alternative. In particular, in the context of financial asset prices, one would estimate the following regression

$$\Delta y_t = \alpha_0 + \alpha_1 y_{t-1} + \sum_{j=1}^{\ell} c_j \Delta y_{t-j} + \varepsilon_t$$

where y denotes the logarithm of the series. The null hypothesis of a single unit root is rejected if α_1 is negative and significantly different from zero. A trend variable should not be included, since the presence of a trend in financial asset prices is a clear violation of market efficiency, whether or not the asset price has a unit root. The optimal lag length, ℓ , can be chosen using data-dependent methods, that have desirable statistical properties when applied to unit root tests. Based on such ADF unit root tests, Nelson and Plosser (1982) argue that most macroeconomic and financial time series have a unit root.

Perron (1989), however, argues that most time series [and in particular those used by Nelson and Plosser (1982)] are trend stationary if one allows for a one-time change in the intercept or in the slope (or both) of the trend function. The postulate is that certain ‘big shocks’ do not represent a realization of the underlying data generation mechanism of the series under consideration and that the null should be tested against the trend-stationary alternative by allowing, under both the null and the alternative hypotheses, for the presence of a one-time break (at a known point in time) in the intercept or in the slope (or both) of the trend function.⁴ Hence, whether the unit root model is rejected or not depends on how big shocks are treated. If they are treated like any other shock, then ADF unit root testing procedures

new, complicated, and largely undeveloped area. We therefore ignore this issue in this paper, keeping in mind that this is an area for future research.

⁴Perron’s (1989) assumption that the break point is uncorrelated with the data has been criticized, on the basis that problems associated with ‘pre-testing’ are applicable to his methodology and that the structural break should instead be treated as being correlated with the data. More recently, a number of studies treat the selection of the break point as the outcome of an estimation procedure and transform Perron’s (1989) conditional (on structural change at a known point in time) unit root test into an unconditional unit root test.

are appropriate and the unit root null hypothesis cannot (in general) be rejected. If, however, they are treated differently, then Perron-type procedures are appropriate and the null hypothesis of a unit root will most likely be rejected.

Finally, given that integration tests are sensitive to the class of models considered (and may be misleading because of misspecification), ‘fractionally’-integrated representations, which nest the unit-root phenomenon in a more general model, have also been used - see Baillie (1996) for a survey. Fractional integration is a popular way to parameterize long-memory processes. If such processes are estimated with the usual autoregressive-moving average model, without considering fractional orders of integration, the estimated autoregressive process can exhibit spuriously high persistence close to a unit root. Since financial asset prices might depart from their means with long memory, one could condition the unit root tests on the alternative of a fractional integrated process, rather than the usual alternative of the series being stationary. In this case, if we fail to reject an autoregressive unit root, we know it is not a spurious finding due to neglect of the relevant alternative of fractional integration and long memory.

Despite the fact that the random walk and martingale hypotheses are contained in the null hypothesis of a unit root, unit root tests are not predictability tests. They are designed to reveal whether a series is difference-stationary or trend stationary and as such they are tests of the permanent/temporary nature of shocks. More recently a series of papers including those by Poterba and Summers (1988), and Lo and MacKinlay (1988) have argued that the efficient markets theory can be tested by comparing the relative variability of returns over different horizons using the variance ratio methodology of Cochrane (1988). They have shown that asset prices are mean reverting over long investment horizons - that is, a given price change tends to be reversed over the next several years by a predictable change in the opposite direction. Similar results have been obtained by Fama and French (1988), using an alternative but closely related test based on predictability of multiperiod returns. Of course, mean-reverting behavior in asset prices is consistent with transitory deviations from equilibrium which are both large and persistent, and implies positive autocorrelation in returns over short horizons and negative autocorrelation over longer horizons.

Predictability of financial asset returns is a broad and very active research topic and a complete survey of the vast literature is beyond the scope

of the present paper. We shall notice, however, that a general consensus has emerged that asset returns are predictable. As Campbell, Lo, and MacKinlay (1997, pp. 80) put it “[r]ecent econometric advances and empirical evidence seem to suggest that financial asset returns are predictable to some degree. Thirty years ago this would have been tantamount to an outright rejection of market efficiency. However, modern financial economics teaches us that other, perfectly rational, factors may account for such predictability. The fine structure of securities markets and frictions in the trading process can generate predictability. Time-varying expected returns due to changing business conditions can generate predictability. A certain degree of predictability may be necessary to reward investors for bearing certain dynamic risks”.

4 Tests of Nonlinearity and Chaos

Most of the empirical tests that we discussed so far are designed to detect ‘linear’ structure in financial data - that is, linear predictability is the focus. However, as Campbell, Lo, and MacKinlay (1997, pp. 467) argue “ ... many aspects of economic behavior may not be linear. Experimental evidence and casual introspection suggest that investors’ attitudes towards risk and expected return are nonlinear. The terms of many financial contracts such as options and other derivative securities are nonlinear. And the strategic interactions among market participants, the process by which information is incorporated into security prices, and the dynamics of economy-wide fluctuations are all inherently nonlinear. Therefore, a natural frontier for financial econometrics is the modeling of nonlinear phenomena”.

It is for such reasons that interest in deterministic nonlinear chaotic processes has in the recent past experienced a tremendous rate of development. Besides its obvious intellectual appeal, chaos is interesting because of its ability to generate output that mimics the output of stochastic systems, thereby offering an alternative explanation for the behavior of asset prices. Clearly then, an important area for potentially productive research is to test for chaos and (in the event that it exists) to identify the nonlinear deterministic system that generates it. In what follows, we turn to several univariate statistical tests for independence, nonlinearity and chaos, that have been recently motivated by the mathematics of deterministic nonlinear dynamical systems.

4.1 The Correlation Dimension Test

Grassberger and Procaccia (1983) suggested the ‘correlation dimension’ test for chaos. To briefly discuss this test, let us start with the 1-dimensional series, $\{x_t\}_{t=1}^n$, which can be embedded into a series of m -dimensional vectors $X_t = (x_t, x_{t-1}, \dots, x_{t-m+1})'$ giving the series $\{X_t\}_{t=m}^n$. The selected value of m is called the ‘embedding dimension’ and each X_t is known as an ‘ m -history’ of the series $\{x_t\}_{t=1}^n$. This converts the series of scalars into a slightly shorter series of (m -dimensional) vectors with overlapping entries - in particular, from the sample size n , $N = n - m + 1$ m -histories can be made. Assuming that the true, but unknown, system which generated $\{x_t\}_{t=1}^n$ is ϑ -dimensional and provided that $m \geq 2\vartheta + 1$, then the N m -histories recreate the dynamics of the data generation process and can be used to analyze the dynamics of the system — see Takens (1981).

The correlation dimension test is based on the ‘correlation function’ (or ‘correlation integral’), $C(N, m, \epsilon)$, which for a given embedding dimension m is given by:

$$C(N, m, \epsilon) = \frac{1}{N(N-1)} \sum_{m \leq t \neq s \leq n} H(\epsilon - \|X_t - X_s\|)$$

where ϵ is a sufficiently small number, $H(z)$ is the Heavside function (which maps positive arguments into 1 and nonpositive arguments into 0), and $\|\cdot\|$ denotes the distance induced by the selected norm (the ‘maximum norm’ being the type used most often). In other words, the correlation integral is the number of pairs (t, s) such that each corresponding component of X_t and X_s are near to each other, nearness being measured in terms of distance being less than ϵ . Intuitively, $C(N, m, \epsilon)$ measures the probability that the distance between any two m -histories is less than ϵ . If $C(N, m, \epsilon)$ is large (which means close to 1) for a very small ϵ , then the data is very well correlated.

The correlation dimension can be defined as

$$D_c^m = \lim_{\epsilon \rightarrow 0} \frac{\log C(N, m, \epsilon)}{\log \epsilon},$$

that is by the slope of the regression of $\log C(N, m, \epsilon)$ versus $\log \epsilon$ for small values of ϵ , and depends on the embedding dimension, m . As a practical matter one investigates the estimated value of D_c^m as m is increased. If as

m increases D_c^m continues to rise, then the system is stochastic. If, however, the data are generated by a deterministic process (consistent with chaotic behavior), then D_c^m reaches a finite saturation limit beyond some relatively small m .⁵ The correlation dimension can therefore be used to distinguish true stochastic processes from deterministic chaos (which may be low-dimensional or high-dimensional).

While the correlation dimension measure is therefore potentially very useful in testing for chaos, the sampling properties of the correlation dimension are, however, unknown. As Barnett, Gallant, Hinich, Jungeilges, Kaplan, and Jensen (1995, pp. 306) put it “[i]f the only source of stochasticity is [observational] noise in the data, and if that noise is slight, then it is possible to filter the noise out of the data and use the correlation dimension test deterministically. However, if the economic structure that generated the data contains a stochastic disturbance within its equations, the correlation dimension is stochastic and its derived distribution is important in producing reliable inference”.

Moreover, if the correlation dimension is very large as in the case of high-dimensional chaos, it will be very difficult to estimate it without an enormous amount of data. In this regard, Ruelle (1990) argues that a chaotic series can only be distinguished if it has a correlation dimension well below $2 \log_{10} N$, where N is the size of the data set, suggesting that with economic time series the correlation dimension can only distinguish low-dimensional chaos from high-dimensional stochastic processes - see also Grassberger and Procaccia (1983) for more details.

4.2 The BDS Test

To deal with the problems of using the correlation dimension test, Brock, Dechert, LeBaron, and Scheinkman (1996) devised a new statistical test which is known as the BDS test — see also Brock, Hsieh, and LeBaron

⁵Since the correlation dimension can be used to characterize both chaos and stochastic dynamics (i.e., the correlation dimension is a finite number in the case of chaos and equal to infinity in the case of an independent and identically distributed stochastic process), one often finds in the literature expressions like ‘deterministic chaos’ (meaning simply chaos) and ‘stochastic chaos’ (meaning standard stochastic dynamics). This terminology, however, is confusing in contexts other than that of the correlation dimension analysis and we shall not use it here.

(1991). The BDS tests the null hypothesis of whiteness (independent and identically distributed observations) against an unspecified alternative using a nonparametric technique.

The BDS test is based on the Grassberger and Procaccia (1983) correlation integral as the test statistic. In particular, under the null hypothesis of whiteness, the BDS statistic is

$$W(N, m, \epsilon) = \sqrt{N} \frac{C(N, m, \epsilon) - C(N, 1, \epsilon)^m}{\hat{\sigma}(N, m, \epsilon)}$$

where $\hat{\sigma}(N, m, \epsilon)$ is an estimate of the asymptotic standard deviation of $C(N, m, \epsilon) - C(N, 1, \epsilon)^m$ - the formula for $\hat{\sigma}(N, m, \epsilon)$ can be found in Brock et al. (1996). The BDS statistic is asymptotically standard normal under the whiteness null hypothesis - see Brock et al. (1996) for details.

The intuition behind the BDS statistic is as follows. $C(N, m, \epsilon)$ is an estimate of the probability that the distance between any two m -histories, X_t and X_s of the series $\{x_t\}$ is less than ϵ . If $\{x_t\}$ were independent then for $t \neq s$ the probability of this joint event equals the product of the individual probabilities. Moreover, if $\{x_t\}$ were also identically distributed then all of the m probabilities under the product sign are the same. The BDS statistic therefore tests the null hypothesis that $C(N, m, \epsilon) = C(N, 1, \epsilon)^m$ - the null hypothesis of whiteness.⁶

Since the asymptotic distribution of the BDS test statistic is known under the null hypothesis of whiteness, the BDS test provides a direct (formal) statistical test for whiteness against general dependence, which includes both nonwhite linear and nonwhite nonlinear dependence. Hence, the BDS test does not provide a direct test for nonlinearity or for chaos, since the sampling distribution of the test statistic is not known (either in finite samples or asymptotically) under the null hypothesis of nonlinearity, linearity, or chaos. It is, however, possible to use the BDS test to produce indirect evidence about nonlinear dependence [whether chaotic (i.e., nonlinear deterministic) or stochastic], which is necessary but not sufficient for chaos - see Barnett et al. (1997) and Barnett and Hinich (1992) for a discussion of these issues.

⁶Note that whiteness implies that $C(N, m, \epsilon) = C(N, 1, \epsilon)^m$ but the converse is not true.

4.3 The Hinich Bispectrum Test

The bispectrum in the frequency domain is easier to interpret than the multiplicity of third order moments $\{C_{xxx}(r, s) : s \leq r, r = 0, 1, 2, \dots\}$ in the time domain - see Hinich (1982). For frequencies ω_1 and ω_2 in the principal domain given by

$$\Omega = \{(\omega_1, \omega_2) : 0 < \omega_1 < 0.5, \omega_2 < \omega_1, 2\omega_1 + \omega_2 < 1\},$$

the bispectrum, $B_{xxx}(\omega_1, \omega_2)$, is defined by

$$B_{xxx}(\omega_1, \omega_2) = \sum_{r=-\infty}^{\infty} \sum_{s=-\infty}^{\infty} C_{xxx}(r, s) \exp[-i2\pi(\omega_1 r + \omega_2 s)].$$

The bispectrum is the double Fourier transformation of the third order moments function and is the third order polyspectrum. The regular power spectrum is the second order polyspectrum and is a function of only one frequency.

The skewness function $\Gamma(\omega_1, \omega_2)$ is defined in terms of the bispectrum as follows

$$\Gamma^2(\omega_1, \omega_2) = \frac{|B_{xxx}(\omega_1, \omega_2)|^2}{S_{xx}(\omega_1)S_{xx}(\omega_2)S_{xx}(\omega_1 + \omega_2)}, \quad (3)$$

where $S_{xx}(\omega)$ is the (ordinary power) spectrum of $x(t)$ at frequency ω . Since the bispectrum is complex valued, the absolute value (vertical lines) in Equation (3) designates modulus. Brillinger (1965) proves that the skewness function $\Gamma(\omega_1, \omega_2)$ is constant over all frequencies $(\omega_1, \omega_2) \in \Omega$ if $\{x(t)\}$ is linear; while $\Gamma(\omega_1, \omega_2)$ is flat at zero over all frequencies if $\{x(t)\}$ is Gaussian. Linearity and Gaussianity can be tested using a sample estimator of the skewness function. But observe that those flatness conditions are necessary but not sufficient for general linearity and Gaussianity, respectively. On the other hand, flatness of the skewness function is necessary and sufficient for third order nonlinear dependence. The Hinich (1982) 'linearity test' tests the null hypothesis that the skewness function is flat, and hence is a test of lack of third order nonlinear dependence. For details of the test, see Hinich (1982).

4.4 The NEGM Test

As it was argued earlier, the distinctive feature of chaotic systems is sensitive dependence on initial conditions - that is, exponential divergence of trajectories with similar initial conditions. The most important tool for diagnosing the presence of sensitive dependence on initial conditions (and thereby of chaoticity) is provided by the dominant Lyapunov exponent, λ . This exponent measures average exponential divergence or convergence between trajectories that differ only in having an ‘infinitesimally small’ difference in their initial conditions and remains well-defined for noisy systems. A bounded system with a positive Lyapunov exponent is one operational definition of chaotic behavior.

One early method for calculating the dominant Lyapunov exponent is that proposed by Wolf, Swift, Swinney, and Vastano (1985). This method, however, requires long data series and is sensitive to dynamic noise, so inflated estimates of the dominant Lyapunov exponent are obtained. Recently, Nychka, Ellner, Gallant, and McCaffrey (1992) have proposed a regression method, involving the use of neural network models, to test for positivity of the dominant Lyapunov exponent. The Nychka et al. (1992), hereafter NEGM, Lyapunov exponent estimator is a regression (or Jacobian) method, unlike the Wolf et al. (1985) direct method which [as Brock and Sayers (1988) have found] requires long data series and is sensitive to dynamic noise.

Assume that the data $\{x_t\}$ are real-valued and are generated by a non-linear autoregressive model of the form

$$x_t = f(x_{t-L}, x_{t-2L}, \dots, x_{t-mL}) + e_t \quad (4)$$

for $1 \leq t \leq N$, where L is the time-delay parameter and m is the length of the autoregression. Here f is a smooth unknown function, and $\{e_t\}$ is a sequence of independent random variables with zero mean and unknown constant variance. The Nychka et al. (1992) approach to estimation of the maximum Lyapunov exponent involves producing a state-space representation of (4)

$$X_t = F(X_{t-L}) + E_t, \quad F : \mathbb{R}^m \rightarrow \mathbb{R}^m$$

where $X_t = (x_t, x_{t-L}, \dots, x_{t-mL+L})'$, $F(X_{t-L}) = (f(x_{t-L}, \dots, x_{t-mL}), x_{t-L}, \dots, x_{t-mL+L})'$, and $E_t = (e_t, 0, \dots, 0)'$, and using a Jacobian-based method to es-

estimate λ through the intermediate step of estimating the individual Jacobian matrices

$$J_t = \frac{\partial F(X_t)}{\partial X'}.$$

After using several nonparametric methods, McCaffrey et al. (1992) recommend using either thin plate splines or neural nets to estimate J_t . Estimation based on neural nets involves the use of the a neural net with q units in the hidden layer

$$f(X_{t-L}, \theta) = \beta_0 + \sum_{j=1}^q \beta_j \psi(\gamma_{0j} + \sum_{i=1}^m \gamma_{ij} x_{t-iL})$$

where ψ is a known (hidden) nonlinear ‘activation function’ [usually the logistic distribution function $\psi(u) = 1/(1 + \exp(-u))$]. The parameter vector θ is then fit to the data by nonlinear least squares. That is, one computes the estimate $\hat{\theta}$ to minimize the sum of squares $S(\theta) = \sum_{t=1}^N [x_t - f(X_{t-1}, \theta)]^2$, and uses $\hat{F}(X_t) = (f(x_{t-L}, \dots, x_{t-mL}, \hat{\theta}), x_{t-L}, \dots, x_{t-mL+L})'$ to approximate $F(X_t)$.

As appropriate values of L , m , and q , are unknown, Nychka et al. (1992) recommend selecting that value of the triple (L, m, q) that minimizes the Bayesian Information Criterion (BIC) - see Schwartz (1978). As shown by Gallant and White (1992), we can use $\hat{J}_t = \partial \hat{F}(X_t) / \partial X'$ as a nonparametric estimator of J_t when (L, m, q) are selected to minimize BIC. The estimate of the dominant Lyapunov exponent then is

$$\hat{\lambda} = \frac{1}{2N} \log |\hat{v}_1(N)|$$

where $\hat{v}_1(N)$ is the largest eigenvalue of the matrix $\hat{T}_N' \hat{T}_N$ and where $\hat{T}_N = \hat{J}_N \hat{J}_{N-1}, \dots, \hat{J}_1$.

Another very promising approach to the estimation of Lyapunov exponents [that is similar in some respects to the Nychka et al. (1992) approach] has also been recently proposed by Gencay and Dechert (1992). This involves estimating all Lyapunov exponents of an unknown dynamical system. The estimation is carried out, as in Nychka et al. (1992), by a multivariate feedforward network estimation technique — see Gencay and Dechert (1992) for more details.

4.5 The White Test

In White's (1989) test, the time series is fitted by a single hidden-layer feed-forward neural network, which is used to determine whether any nonlinear structure remains in the residuals of an autoregressive (AR) process fitted to the same time series. The null hypothesis for the test is 'linearity in the mean' relative to an information set. A process that is linear in the mean has a conditional mean function that is a linear function of the elements of the information set, which usually contains lagged observations on the process.⁷

The rationale for White's test can be summarized as follows: under the null hypothesis of linearity in the mean, the residuals obtained by applying a linear filter to the process should not be correlated with any measurable function of the history of the process. White's test uses a fitted neural net to produce the measurable function of the process's history and an AR process as the linear filter. White's method then tests the hypothesis that the fitted function does not correlate with the residuals of the AR process. The resulting test statistic has an asymptotic χ^2 distribution under the null of linearity in the mean.⁸

4.6 The Kaplan Test

Kaplan (1994) used the fact that solution paths in phase space reveal deterministic structure that is not evident in a plot of x_t versus t , to produce a test statistic which has a strictly positive lower bound for a stochastic process, but not for a deterministic solution path. By computing the test statistic from an adequately large number of linear processes that plausibly might have produced the data, the approach can be used to test for linearity against the alternative of noisy nonlinear dynamics. The procedure involves producing linear stochastic process surrogates for the data and determining whether the surrogates or a noisy continuous nonlinear dynamical solution path better describe the data. Linearity is rejected, if the value of the test

⁷For a formal definition of linearity in the mean, see Lee, White, and Granger (1993, section 1). Note that a process that is not linear in the mean is said to exhibit 'neglected nonlinearity'. Also, a process that is linear is also linear in the mean, but the converse need not be true.

⁸See Lee, White, and Granger (1993, section 2) for a presentation of the test statistic's formula and computation method.

statistic from the surrogates is never small enough relative to the value of the statistic computed from the data - see Kaplan (1994) or Barnett et al. (1997) for more details about this procedure.

5 Evidence on Nonlinearity and Chaos

A number of researchers have recently focused on testing for nonlinearity in general and chaos in particular in macroeconomic time series. There are many reasons for this interest. Chaos, for example, represents a radical change of perspective on business cycles. Business cycles receive an endogenous explanation and are traced back to the strong nonlinear deterministic structure that can pervade the economic system. This is different from the (currently dominant) exogenous approach to economic fluctuations, based on the assumption that economic equilibria are determinate and intrinsically stable, so that in the absence of continuing exogenous shocks the economy tends towards a steady state, but because of stochastic shocks a stationary pattern of fluctuations is observed.⁹

There is a broad consensus of support for the proposition that the (macroeconomic) data generating processes are characterized by a pattern of nonlinear dependence, but there is no consensus at all on whether there is chaos in macroeconomic time series. For example, Brock and Sayers (1988), Frank and Stengos (1988), and Frank, Gencay, and Stengos (1988) find no evidence of chaos in U.S., Canadian, and international, respectively, macroeconomic time series. On the other hand, Barnett and Chen (1988), claimed successful detection of chaos in the (demand-side) U.S. Divisia monetary aggregates. Their conclusion was further confirmed by DeCoster and Mitchell (1991, 1994). This published claim of successful detection of chaos has generated considerable controversy, as in Ramsey, Sayers, and Rothman (1990) and Ramsey and Rothman (1994), who raised questions regarding virtually all published tests of chaos. Further results relevant to this controversy have recently been provided by Serletis (1995).

⁹Chaos could also help unify different approaches to structural macroeconomics. As Grandmont (1985) has shown, for different parameter values even the most classical of economic models can produce stable solutions (characterizing classical economics) or more complex solutions, such as cycles or even chaos (characterizing much of Keynesian economics)

Although the analysis of macroeconomic time series has not yet led to particularly encouraging results (mainly due to the small samples and high noise levels for most macroeconomic series), as can be seen from Table 1, there is also a substantial literature testing for nonlinear dynamics on financial data.¹⁰ This literature has led to results which are as a whole more interesting and more reliable than those of macroeconomic series, probably due to the much larger number of data available and their superior quality (measurement in most cases is more precise, at least when we do not have to make recourse to broad aggregation). As regards the main conclusions of this literature, there is clear evidence of nonlinear dependence and some evidence of chaos.

For example, Scheinkman and LeBaron (1989) studied United States weekly returns on the Center for Research in Security Prices (CRSP) value-weighted index, employing the BDS statistic, and found rather strong evidence of nonlinearity and some evidence of chaos.¹¹ Some very similar results have been obtained by Frank and Stengos (1989), investigating daily prices (from the mid 1970's to the mid 1980's) for gold and silver, using the correlation dimension and the Kolmogorov entropy. Their estimate of the correlation dimension was between 6 and 7 for the original series and much greater and non-converging for the reshuffled data.

More recently, Serletis and Gogas (1997) test for chaos in seven East European black market exchange rates, using the Koedijk and Kool (1992) monthly data (from January 1955 through May 1990). In doing so, they use three inference methods, the BDS test, the NEGM test, as well as the Lyapunov exponent estimator of Gencay and Dechert (1992). They find some consistency in inference across methods, and conclude, based on the NEGM

¹⁰For other unpublished work on testing nonlinearity and chaos on financial data, see Abhyankar, Copeland, and Wong (1997, Table 1).

¹¹In order to verify the presence of a nonlinear structure in the data, they also suggested employing the so-called 'shuffling diagnostic'. This procedure involves studying the residuals obtained by adapting an autoregressive model to a series and then reshuffling these residuals. If the residuals are totally random (i.e., if the series under scrutiny is not characterized by chaos), the dimension of the residuals and that of the shuffled residuals should be approximately equal. On the contrary, if the residuals are chaotic and have some structure, then the reshuffling must reduce or eliminate the structure and consequently increase the correlation dimension. The correlation dimension of their reshuffled residuals always appeared to be much greater than that of the original residuals, which was interpreted as being consistent with chaos.

test, that there is evidence consistent with a chaotic nonlinear generation process in two out of the seven series - the Russian ruble and East German mark. Altogether, these and similar results seem to suggest that financial series provide a more promising field of research for the methods in question.

A notable feature of the literature just summarized is that most researchers, in order to find sufficient observations to implement the tests, use data periods measured in years. The longer the data period, however, the less plausible is the assumption that the underlying data generation process has remained stationary, thereby making the results difficult to interpret. In fact, different conclusions have been reached by researchers using high-frequency data over short periods. For example, Abhyankar, Copeland, and Wong (1995) examine the behavior of the U.K. Financial Times Stock Exchange 100 (FTSE 100) index, over the first six months of 1993 (using 1-, 5-, 15-, 30-, and 60-minute returns). Using the Hinich (1982) bispectral linearity test, the BDS test, and the NEGM test, they find evidence of nonlinearity, but no evidence of chaos.

More recently, Abhyankar, Copeland, and Wong (1997) test for nonlinear dependence and chaos in real-time returns on the world's four most important stock-market indices - the FTSE 100, the Standard & Poor 500 (S&P 500) index, the Deutscher Aktienindex (DAX), and the Nikkei 225 Stock Average. Using the BDS and the NEGM tests, and 15-second, 1-minute, and 5-minute returns (from September 1 to November 30, 1991), they reject the hypothesis of independence in favor of a nonlinear structure for all data series, but find no evidence of low-dimensional chaotic processes.

Of course, there is other work, using high-frequency data over short periods, that finds order in the apparent chaos of financial markets. For example, Ghashghaie, Breymann, Peinke, Talkner, and Dodge (1996) analyze all worldwide 1,472,241 bid-ask quotes on U.S. dollar-German mark exchange rates between October 1, 1992 and September 30, 1993. They apply physical principles and provide a mathematical explanation of how one trading pattern led into and then influenced another. As the authors conclude, "... we have reason to believe that the qualitative picture of turbulence that has developed during the past 70 years will help our understanding of the apparently remote field of financial markets".

6 Controversies

Clearly, there is little agreement about the existence of chaos or even of nonlinearity in (economic and) financial data, and some economists continue to insist that linearity remains a good assumption for such data, despite the fact that theory provides very little support for that assumption. It should be noted, however, that the available tests search for evidence of nonlinearity or chaos in data without restricting the boundary of the system that could have produced that nonlinearity or chaos. Hence these tests should reject linearity, even if the structure of the economy is linear, but the economy is subject to shocks from a surrounding nonlinear or chaotic physical environment, as through nonlinear climatological or weather dynamics. Under such circumstances, linearity would seem an unlikely inference.¹²

Since the available tests are not structural and hence have no ability to identify the source of detected chaos, the alternative hypothesis of the available tests is that no natural deterministic explanation exists for the observed economic fluctuations anywhere in the universe. In other words, the alternative hypothesis is that economic fluctuations are produced by supernatural shocks or by inherent randomness in the sense of quantum physics. Considering the implausibility of the alternative hypothesis, one would think that findings of chaos in such nonparametric tests would produce little controversy, while any claims to the contrary would be subjected to careful examination. Yet, in fact the opposite seems to be the case.

We argued earlier that the controversies might stem from the high noise level that exists in most aggregated economic time series and the relatively low sample sizes that are available with economic data. However, it also appears that the controversies are produced by the nature of the tests themselves, rather than by the nature of the hypothesis, since linearity is a very strong null hypothesis, and hence should be easy to reject with any test and any economic or financial time series on which an adequate sample size is available. In particular, there may be very little robustness of such tests across variations in sample size, test method, and data aggregation method - see Barnett et al. (1995) on this issue.

It is also possible that none of the tests for chaos and nonlinear dynamics

¹²In other words, not only is there no reason in economic theory to expect linearity within the structure of the economy, but there is even less reason to expect to find linearity in nature, which produces shocks to the system.

that we have discussed completely dominates the others, since some tests may have higher power against certain alternatives than other tests, without any of the tests necessarily having higher power against all alternatives. If this is the case, each of the tests may have its own comparative advantages, and there may even be a gain from using more than one of the tests in a sequence designed to narrow down the alternatives.

To explore this possibility, Barnett with the assistance of Jensen designed and ran a single blind controlled experiment, in which they produced simulated data from various processes having linear, nonlinear chaotic, or nonlinear nonchaotic signal. They transmitted each simulated data set by email to experts in running each of the statistical tests that were entered into the competition. The emailed data included no identification of the generating process, so those individuals who ran the tests had no way of knowing the nature of the data generating process, other than the sample size, and there were two sample sizes: a ‘small sample’ size of 380 and a ‘large sample’ size of 2000 observations.

In fact five generating models were used to produce samples of the small and large size. The models were a fully deterministic, chaotic Feigenbaum recursion (Model I), a generalized autoregressive conditional heteroskedasticity (GARCH) process (Model II), a nonlinear moving average process (Model III), an autoregressive conditional heteroskedasticity (ARCH) process (Model IV), and an autoregressive moving average (ARMA) process (Model V). Details of the parameter settings and noise generation method can be found in Barnett et al. (1996). The tests entered into this competition were Hinich’s bispectrum test, the BDS test, White’s test, Kaplan’s test, and the NEGM test of chaos.

The results of the competition are available in Barnett et al. (1997) and are summarized in Table 2. They provide the most systematic available comparison of tests of nonlinearity and indeed do suggest differing powers of each test against certain alternative hypotheses. In comparing the results of the tests, however, one factor seemed to be especially important: subtle differences existed in the definition of the null hypothesis, with some of the tests being tests of the null of linearity, defined in three different manners in the derivation of the test’s properties, and one test being a test of the null of chaos. Hence there were four null hypotheses that had to be considered to be able to compare each test’s power relative to each test’s own definition of the null.

Since the tests do not all have the same null hypothesis, differences among them are not due solely to differences in power against alternatives. Hence one could consider using some of them sequentially in an attempt to narrow down the inference on the nature of the process. For example, the Hinich test and the White test could be used initially to find out whether the process lacks third order nonlinear dependence and is linear in the mean. If either test rejects its null, one could try to narrow down the nature of the nonlinearity further by running the NEGM test to see if there is evidence of chaos. Alternatively, if the Hinich and White tests both lead to acceptance of the null, one could run the BDS or Kaplan test to see if the process appears to be fully linear. If the data leads to rejection of full linearity but acceptance of linearity in the mean, then the data may exhibit stochastic volatility of the ARCH or GARCH type.

In short, the available tests provide useful information, and such comparisons of other tests could help further to narrow down alternatives. But ultimately we are left with the problem of isolating the nature of detected nonlinearity or chaos to be within the structure of the economy. This final challenge remains unsolved, especially in the case of chaos.

7 Conclusion

Recently there has been considerable criticism of the existing research on chaos, as for example in Granger's (1994) review of Benhabib's (1992) book. The presence of dynamic noise (i.e., noise added in each iteration step) makes it difficult and perhaps impossible to distinguish between (noisy) high-dimensional chaos and pure randomness. The estimates of the fractal dimension, the correlation integral, and Lyapunov exponents of an underlying unknown dynamical system are sensitive to dynamic noise, and the problem grows as the dimension of the chaos increases. The question of the 'impossibility' of distinguishing between high-dimensional chaos and randomness has recently attracted some attention, as for example in Radunskaja (1994), Bickel and Bühlmann (1996), and Takens (1997). Analogously, Bickel and Bühlmann (1996) argue that distinguishing between linearity and nonlinearity of a stochastic process may become impossible as the order of the linear filter increases. In a time series framework, it is prudent to limit such tests to the use of low order linear filters as approximations to nonlinear processes

when testing for general nonlinearity, and tests for low dimensional chaos, when chaotic nonlinearity is of interest — see also Barnett et al. (1997, footnote 11).

However, in the field of economics, it is especially unwise to take a strong opinion (either pro or con) in that area of research. Contrary to popular opinion within the profession, there have been no published tests of chaos ‘within the structure of the economic system’, and there is very little chance that any such tests will be available in this field for a very long time. Such tests are simply beyond the state of the art. Existing tests cannot tell whether the source of detected chaos comes from within the structure of the economy, or from chaotic external shocks, as from the weather. Thus, we do not have the slightest idea of whether or not asset prices exhibit chaotic nonlinear dynamics produced from the nonlinear structure of the economy (and hence we are not justified in excluding the possibility). Until the difficult problems of testing for chaos ‘within the structure of the economic system’ are solved, the best that we can do is to test for chaos in economic time series data, without being able to isolate its source. But even that objective has proven to be difficult. While there have been many published tests for chaotic nonlinear dynamics, little agreement exists among economists about the correct conclusions.

8 References

Abhyankar, A. H., L. S. Copeland, and W. Wong, 1995, Nonlinear Dynamics in Real-Time Equity Market Indices: Evidence From the UK, *Economic Journal* 105 (431), 864-880.

Abhyankar, A. H., L. S. Copeland, and W. Wong, 1997, Uncovering Nonlinear Structure in Real-Time Stock-Market Indexes: The S&P 500, the DAX, the Nikkei 225, and the FTSE-100, *Journal of Business and Economic Statistics* 15, 1-14.

Baillie, R. T., 1996, Long Memory Processes and Fractional Integration in Econometrics, *Journal of Econometrics* 73, 5-59.

Barnett, W. A. and P. Chen, 1988, The Aggregation-Theoretic Monetary Aggregates are Chaotic and Have Strange Attractors: An Econometric Application of Mathematical Chaos, in: W. A. Barnett, E. R. Berndt, and H. White, eds., *Dynamic Econometric Modeling* (Cambridge University Press, Cambridge, U.K.).

Barnett, W. A., A. R. Gallant, M.J. Hinich, J. Jungeilges, D. Kaplan, and M.J. Jensen, 1997, A Single-Blind Controlled Competition Among Tests for Nonlinearity and Chaos, *Journal of Econometrics* 82 (1997), 157-192.

Barnett, W. A., A. R. Gallant, M.J. Hinich, J. Jungeilges, D. Kaplan, and M.J. Jensen, 1995, Robustness of Nonlinearity and Chaos Tests to Measurement Error, Inference Method, and Sample Size, *Journal of Economic Behavior and Organization* 27, 301-320.

Barnett, W. A., A. R. Gallant, M.J. Hinich, J. Jungeilges, D. Kaplan, and M.J. Jensen, 1996, An Experimental Design to Compare Tests of Nonlinearity and Chaos, in: W. A. Barnett, A. Kirman, and M. Salmon, eds., *Nonlinear Dynamics in Economics* (Cambridge University Press, Cambridge, U.K.).

Barnett, W. A. and M. J. Hinich, 1992, Empirical Chaotic Dynamics in Economics, *Annals of Operations Research* 37,1-15.

Benhabib, J., 1992, *Cycles and Chaos in Economic Equilibrium* (Princeton University Press, Princeton, NJ).

Bickel, P.J. and P. Bühlmann, 1996, What is a Linear Process?, *Proceedings*

of the National Academy of Science, USA, Statistics Section 93, 12,128-12,131.

Brillinger, D. R., 1965, An Introduction to the Polyspectrum, *Annals of Mathematical Statistics* 36, 1351-1374.

Brock, W. A., W. D. Dechert, B. LeBaron, and J. A. Scheinkman, 1996, A Test for Independence Based on the Correlation Dimension, *Econometric Reviews* 15, 197-235.

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

Brock, W. A. and C. Sayers, 1988, Is the Business Cycle Characterized by Deterministic Chaos?, *Journal of Monetary Economics* 22, 71-90.

Brockett, P. L., M. J. Hinich, and D. Patterson, 1988, Bispectral-Based Tests for the Detection of Gaussianity and Nonlinearity in Time Series, *Journal of the American Statistical Association* 83, 657-664.

Campbell, J. Y., A. W. Lo, and A. C. MacKinlay, 1997, *The Econometrics of Financial Markets* (Princeton University Press, Princeton, NJ).

Cochrane, J. H., 1988, How Big is the Random Walk in GNP?, *Journal of Political Economy* 96, 893-920.

DeCoster, G. P. and D. W. Mitchell, 1991, Nonlinear Monetary Dynamics, *Journal of Business and Economic Statistics* 9, 455-462.

DeCoster, G. P. and D. W. Mitchell, 1994, Reply, *Journal of Business and Economic Statistics* 12, 136-137.

Dickey, D. A. and W. A. Fuller, 1981, Likelihood Ratio Tests for Autoregressive Time Series With a Unit Root, *Econometrica* 49, 1057-1072.

Fama, E. F., 1970, Efficient Capital Markets: A Review of Theory and Empirical Work, *Journal of Finance* 25, 383-417.

Fama, E. F., 1965, The Behavior of Stock Market Prices, *Journal of Business* 38, 34-105.

Fama, E. F. and K. R. French, 1988, Permanent and Temporary Components

- of Stock Prices, *Journal of Political Economy* 96, 246-273.
- Frank, M., R. Gencay, and T. Stengos, 1988, International Chaos, *European Economic Review* 32, 569-1584.
- Frank, M. and T. Stengos, 1988, Some Evidence Concerning Macroeconomic Chaos, *Journal of Monetary Economics* 22, 423-438.
- Frank, M. and T. Stengos, 1989, Measuring the Strangeness of Gold and Silver Rates of Return, *Review of Economic Studies* 56, 553-567.
- Gallant, R. and H. White, 1992, On Learning the Derivatives of an Unknown Mapping With Multilayer Feedforward Networks, *Neural Networks* 5, 129-138.
- Gencay, R. and W. D. Dechert, 1992, An Algorithm for the n -Lyapunov Exponents of an n -Dimensional Unknown Dynamical System, *Physica D* 59, 142-157.
- Ghashghaie, S., W. Breymann, J. Peinke, P. Talkner, and Y. Dodge, 1996, Turbulent Cascades in Foreign Exchange Markets, *Nature* 381, 767-770.
- Grandmont, J.-M., 1985, On Endogenous Competitive Business Cycles, *Econometrica* 53, 995-1045.
- Granger, C. W. J., 1994, Is Chaotic Economic Theory Relevant for Economics? A Review Article of: Jess Benhabib: Cycles and Chaos in Economic Equilibrium, *Journal of International and Comparative Economics*, 3, 139-145.
- Granger, C. W. J., 1995, Modelling Nonlinear Relationships Between Extended-Memory Variables, *Econometrica* 63, 265-279.
- Grassberger, P. and I. Procaccia, 1983, Characterization of Strange Attractors, *Physical Review Letters* 50, 346-349.
- Hinich, M. J., 1982, Testing for Gaussianity and Linearity of a Stationary Time Series, *Journal of Time Series Analysis* 3, 169-176.
- Hinich, M. J. and D. Patterson, 1989, Evidence of Nonlinearity in the Trade-by-Trade Stock Market Return Generating Process, in: W. A. Barnett, J. Geweke, and K. Shell, eds., *Economic Complexity: Chaos, Bubbles, and Nonlinearity* (Cambridge University Press, Cambridge, U.K.).

- Hsieh, D. A., 1991, Chaos and Nonlinear Dynamics: Applications to Financial Markets, *Journal of Finance* 46, 1839-1877.
- Kaplan, D. T., 1994, Exceptional Events as Evidence for Determinism, *Physica D* 73, 38-48.
- Koedijk, K. G., and C. J. M. Kool, 1992, Tail Estimates of East European Exchange Rates, *Journal of Business and Economic Statistics* 10, 83-96.
- Lee, T.-H., H. White, and C. W. J. Granger, 1993, Testing for Neglected Nonlinearities in Time Series Models, *Journal of Econometrics* 56, 269-290.
- LeRoy, S. F., 1989, Efficient Capital Markets and Martingales, *Journal of Economic Literature* 27, 1583-1621.
- Lo, A. W. and A. C. MacKinlay, 1988, Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test, *Review of Financial Studies* 1, 41-66.
- Mandelbrot, B. B., 1963, The Variation of Certain Speculative Stock Prices, *Journal of Business* 36, 394-419.
- McCaffrey, D., S. Ellner, R. Gallant, and D. Nychka, 1992, Estimating the Lyapunov Exponent of a Chaotic System with Nonparametric Regression, *Journal of the American Statistical Association* 87, 682-695.
- Nelson, C. R. and C. I. Plosser, 1982, Trends and Random Walks in Macroeconomic Time Series, *Journal of Monetary Economics* 10, 139-162.
- Nychka, D., S. Ellner, R. Gallant, and D. McCaffrey, 1992, Finding Chaos in Noisy Systems, *Journal of the Royal Statistical Society B* 54, 399-426.
- Perron, P., 1989, The Great Crash, the Oil Price Shock, and the Unit Root Hypothesis, *Econometrica* 57, 1361-1401.
- Poterba, J. M. and L. H. Summers, 1988, Mean Reversion in Stock Prices: Evidence and Implications, *Journal of Financial Economics* 22, 27-59.
- Radunskaya, A., 1994, Comparing Random and Deterministic Series, *Economic Theory* 4, 765-776.
- Ramsey, J. B. and P. Rothman, 1994, Comment on "Nonlinear Monetary Dynamics" by DeCoster and Mitchell, *Journal of Business and Economic*

Statistics 12, 135-136.

Ramsey, J. B., C. L. Sayers, and P. Rothman, 1990, The Statistical Properties of Dimension Calculations Using Small Data Sets: Some Economic Applications, *International Economic Review* 31, 991-1020.

Ruelle, D., 1990, Deterministic Chaos: The Science and the Fiction, *Proc. R. Soc. London A* 427 (1873), 241-248.

Scheinkman, J. A., 1990, Nonlinearities in Economic Dynamics, *Economic Journal* 100, 33-48.

Scheinkman, J. A. and B. LeBaron, 1989, Nonlinear Dynamics and Stock Returns, *Journal of Business* 62, 311-337.

Schwartz, G., 1978, Estimating the Dimension of a Model, *The Annals of Statistics* 6, 461-464.

Serletis, A., 1995, Random Walks, Breaking Trend Functions, and the Chaotic Structure of the Velocity of Money, *Journal of Business and Economic Statistics* 13, 453-458.

Serletis, A. and P. Gogas, 1997, Chaos in East European Black-Market Exchange Rates, *Research in Economics* 51 (1997), 359-385.

Summers, L. H., 1986, Does the Stock Market Rationally Reflect Fundamental Values, *Journal of Finance* 41, 591-601.

Takens, F., 1981, Detecting Strange Attractors in Turbulence, in D. Rand and L. Young, eds., *Dynamical Systems and Turbulence* (Springer-Verlag, Berlin).

Takens, F., 1997, The Effect of Small Noise on Systems with Chaotic Dynamics, in S. J. van Strien and S. M. Verduyn Lunel, eds., *Stochastic and Spatial Structures of Dynamical Systems* (North-Holland, Amsterdam).

White, H., 1989, Some Asymptotic Results for Learning in Single Hidden-Layer Feedforward Network Models, *Journal of the American Statistical Association* 84, 1003-1013.

Wolf, A., J. B. Swift, H. L. Swinney, and J. A. Vastano, 1985, Determining Lyapunov Exponents from a Time Series, *Physica* 16D 16, 285-317.