ELSEVIER

Agricultural Economics 27 (2002) 123–137

AGRICULTURAL ECONOMICS

www.elsevier.com/locate/agecon

Are commodity prices chaotic?

Arjun Chatrath*, Bahram Adrangi, Kanwalroop Kathy Dhanda

The Pamplin School of Business Administration, University of Portland, 5000 North Willamette Blvd., Portland, OR 97203, USA

Received 16 October 2000; received in revised form 27 August 2001; accepted 8 October 2001

Abstract

We conduct tests for the presence of low-dimensional chaotic structure in the futures prices of four important agricultural commodities. Though there is strong evidence of non-linear dependence, the evidence suggests that there is no long-lasting chaotic structure. The dimension estimates for the commodity futures series are generally much higher than would be for low dimension chaotic series. Our test results indicate that autoregressive conditional heteroskedasticity (ARCH)-type processes, with controls for seasonality and contract-maturity effects, explain much of the non-linearity in the data. We make a case that employing seasonally adjusted price series is important in obtaining robust results via some of the existing tests for chaotic structure. Finally, maximum likelihood methodologies, that are robust to the non-linear dynamics, lend strong support to the Samuelson hypothesis of maturity effects in futures price changes.

Keywords: Commodity prices; Chaos; Non-linearity

1. Introduction

It is well documented that a variety of deterministic non-linear relationships can produce highly complex time paths capable of passing most standard tests of randomness (see Brock, 1986, for a survey). Such random-looking but deterministic series have been dubbed 'chaotic' or 'strange' in the literature (e.g. Devaney, 1986; Guckenheimer and Holmes, 1986). Chaotic dynamics may explain a richer array of time series behaviour. For instance, sudden/large movements in commodity prices, exchange rates, equity prices and other financial or economic time series will not be properly captured by linear, or even most non-linear models, while chaotic models may be suitable in explaining such behaviour. Direct application of chaotic structures to economic theory has been initiated only in the last 20 years (Stutzer, 1980; Benhabib and Day, 1981, 1982), with researchers such as Brock and Sayers (1988) employing relatively new techniques to test the null hypothesis of chaos in a number of macroeconomic series (such as the US unemployment rate).[1] The evidence of chaos in economic time series such as GNP and unemployment has thus far been weak (Brock and Sayers, 1988).

On the other hand, the few studies on commodity price structure have generally found evidence consistent with low dimension chaos: Lichtenberg and Ujihara (1988) apply a non-linear cobweb model to US crude oil prices; Frank and Stengos (1989) estimate the correlation dimension and Kolmogorov entropy for gold and silver spot prices; Blank (1991) estimates the Lyapunov exponent for soybean futures; DeCoster et al. (1992) apply correlation dimension to

* Corresponding author. Tel.: +1-503-283-7465; fax: +1-503-978-8041.
E-mail address: chatrath@up.edu (A. Chatrath).

[1] For a more complete review of the applications of chaos to economic theory, see Baumol and Benhabib (1989).

0169-5150/02/$ – see front matter © 2002 Elsevier Science B.V. All rights reserved.
PII: S0169-5150(02)00029-4

daily sugar, silver, copper and coffee futures prices; Yang and Brorsen (1993) employ correlation dimension and the Brock, Dechert and Scheinkman (BDS) test on several futures markets, including soybean, corn and wheat.

Such evidence in favour of chaos in agricultural commodity markets is noteworthy. First, the search for predictable patterns in commodity markets is of obvious importance to farm policy. Given that many disparate domestic and international factors affect commodity prices, and that information on such factors is sadly lacking, the accurate structural modelling of such variables could be considered impossible. The evidence on chaos offers some scope for modelling price behaviour by simply employing the time series of prices, which are now readily available—even on an hourly basis. Furthermore, it has long been speculated that trading strategies based on technical analysis will be more successful if the time series are non-linear or even chaotic (Bohan, 1981; Brush, 1986; Pruitt and White, 1988, 1989; Clyde and Osler, 1997). Similarly, many studies demonstrate that common technical rules, inherently non-linear in nature, produce superior outcomes relative to trading rules based on linear models (LaBaron, 1991; Brock et al., 1992; Taylor, 1994; Blume et al., 1994; Chang and Osler, 1995). Clyde and Osler (1997) make a more direct link between chaotic systems and the viability of technical trading. They demonstrate that well-known technical patterns in prices (e.g. head-over-shoulder) will be more useful in short-term trading strategies when the underlying price series are chaotic.

Why do commodity prices exhibit stronger evidence of chaos? Baumol and Benhabib (1989) have suggested that disaggregated variables, such as commodity prices or production levels, that are inherently subject to resource constraints are generally better candidates for chaotic structure. Evidence of chaos in commodity prices may, in this case, explain why commodity markets, more than others, continue to attract large numbers of technical traders. However, do alternate explanations exist for the differences in the evidence across commodity prices and aggregated economic time series? The prior studies on the structure of commodity prices were conducted with fairly coarse tests for chaos and/or suffered from short data spans. In addition, most of these studies have also failed to control for seasonal variations in commodity prices. This failure may be especially important given the evidence that seasonality and non-linearity may be closely related. For instance, Deaton and Laroque (1992) demonstrate that commodity prices spend long periods in the 'doldrums', showing little movement but high autocorrelation, but frequently break out in violent fashion. The authors suggest that the standard rational-expectations model for commodity prices is capable of explaining such behaviour. Importantly, the authors suggest that the non-linearity in price patterns is related to the inherently seasonal inventory changes. Chambers and Bailey (1996) also show that 'periodic disturbances' in commodity prices (namely harvest-related shocks) can exist in rational-expectations equilibria. Their study further highlights the potential of (and the theoretical grounds for) incorporating controls for such shocks to enhance the capacity of models to account for observed fluctuations in commodity prices.

In this paper, we provide new evidence on the structure of commodity prices. We examine the non-linear dynamics and their explanations for four important commodity contracts: soybean, corn, wheat and cotton from the late 1960s to the mid-1990s. Testing for chaotic structure in commodity prices is a meaningful exercise for several reasons. Chaos would imply that while prices are deterministic, long-range prediction based on technical or statistical forecasting techniques becomes treacherous, as initial measurement errors will multiply exponentially.[2] We address this issue in greater detail in the next section.

The four commodities selected for this research play a dominant role in the US agriculture sector. The US is by far the world's largest producer of corn and soybean. Corn is the leading US crop in terms of dollar value and acreage under cultivation, and roughly 20% of US production is exported. The US is the fourth largest producer of wheat and the second largest producer and consumer of cotton, behind China, and the largest cotton exporter (for further details, see Foreign Agriculture Service (FAS) and the National Agricultural Statistical Service (NASS) of the US Department

[2] Definitions of 'chaos' structure are often based on the magnification of initial measurement errors. For instance, a popular definition requires that the largest Lyapunov exponent be positive (e.g. Brock et al., 1993), where the Lyapunov exponent is a measure of the speed of multiplication in forecast errors of initial measurement error.

of Agriculture). The prices of these commodities are subject to a great deal of seasonality. Generally, prices rise from the harvest lows to late spring, reflecting carrying costs and dwindling supplies. Of course, prices are also affected by weather patterns, the availability of substitute grains and products, and export demand—which may depend on factors varying from foreign production to the size of foreign livestock herds.

Our paper is distinguishable from the previous studies on chaos in commodity futures markets in that: (i) relatively long price histories are examined;[3] (ii) unlike prior papers, the data are subject to adjustments for seasonalities and maturity effects that may otherwise lead to an erroneous conclusion of deterministic structure; (iii) a wide range of autoregressive conditional heteroskedasticity (ARCH)-type models are considered as explanations for the non-linearities; (iv) alternate statistical techniques are employed to test the null hypothesis of chaotic structure. Like most prior studies, we present strong evidence that commodity futures prices display non-linear dependencies. Unlike these earlier studies, however, we find evidence that is clearly inconsistent with long-lasting chaotic structure. This difference may be attributed to differences in sample size and methodology. We make a case that employing seasonally adjusted price series may be critical to obtaining robust results with the existing tests for chaotic structure. We identify some commonly known ARCH-type processes that satisfactorily explain the non-linearities in the data. The exponential GARCH model of Nelson (1991) is found to generally perform the best in accounting for the non-linear dynamics in the commodities analysed.

The next section motivates the tests for chaos and further discusses the implications of chaotic structure in commodity prices. Simulated chaotic data is employed to highlight some important properties of chaos. Section 3 describes the procedures employed to test the null hypothesis of chaos. Section 4 presents the test results for the five commodities. Section 5 closes with a summary of the results.

[3] Yang and Brorsen (1993) examine the non-linear dynamics in daily futures prices for various commodity futures over the 1979–1988 interval. Blank (1991) examines only 2 years of data for soybean futures (the November 1986 contract). DeCoster et al. (1992) cover an interval more comparable to ours, from October 1972 to March 1989, for silver, copper, sugar and coffee contracts.

2. Chaos: concepts and implications for commodity markets

Since the concepts of chaos are well developed in the literature, our descriptions are brief relative to some papers that we refer to here. There are several definitions of chaos in use. A definition similar to the following is commonly found in the literature (Devaney, 1986; Brock, 1986; Deneckere and Pelikan, 1986; Brock and Dechert, 1988; Brock and Sayers, 1988; Brock et al., 1993): the series a_t has a chaotic explanation if there exists a system (h, F, x_0) where $a_t = h(x_t)$, $x_t + 1 = F(x_t)$, x_0 is the initial condition at $t = 0$, and where h maps the n-dimensional phase space, R^n to R^1 and F maps R^n to R^n. It is also required that all trajectories, x_t, lie on an attractor, A, and that nearby trajectories diverge so that the system never reaches an equilibrium or even exactly repeats its path.

The above definition restated: the time series a_t (e.g. daily returns for cotton futures) is said to have a chaotic explanation if there is some state vector x_t that evolves deterministically, $x_t + 1 = F(x_t)$, and there is some function $h(x)$ so that $a_t = h(x_t)$ for all t. If one knew (h, F) and could measure x_t without error, one could forecast x_{t+i} and, thus, a_{t+i} perfectly. In this respect, chaos is the opposite of the process that is instantaneously unpredictable. With respect to the divergence property and attractor, A: in order that F generates random-looking behaviour (which is deterministic), nearby trajectories must diverge exponentially. Moreover, in order that F generates deterministic behaviour, locally diverging trajectories must eventually fold back on themselves. The attractor A may be thought of as a subset of the phase space towards which sufficiently close trajectories are asymptotically attracted (Brock and Sayers, 1988).

Chaotic time paths will have the following properties that should be of special interest to commodity market observers:[4] (i) the universality of certain routes (such as the period doubling of trajectories) that are independent of the details of the map; (ii)

[4] See Brock et al. (1993) for a more complete description of the properties.

time paths that are extremely sensitive to microscopic changes in initial conditions, this property is often termed sensitive dependence upon initial condition or SDIC;[5] (iii) time series that appear stochastic even though they are generated by deterministic systems, i.e. the empirical spectrum and autocovariance functions of chaotic series are the same as those generated by random variables, implying that chaotic series will not be identified as such by most standard techniques (such as spectral analysis or autocovariance functions).

The above properties of chaos are probably better appreciated in the framework of a chaotic function. Here, we briefly illustrate some of these properties in the framework of the logistic equation, a function commonly employed to demonstrate the chaos phenomenon (Baumol and Benhabib, 1989; Hsieh, 1991). Consider the non-linear equation (logistic function) with a single parameter, w:

$$x_t = F(x_t) = wx_t(1 - x_t) \tag{1}$$

Fig. 1 graphs the relationship (x_{t+1}, x_t) for $w = 3.750, x_0 = 0.10$.[6] It is apparent that the (x_{t+1}, x_t) oscillations form a distinctive phase diagram (the bounding parabolic curve). As the oscillations expand, they encounter and bounce off the phase curve, moving closer to an apparent equilibrium on the negative slope of the phase curve. However, the convergence towards any equilibrium in that vicinity can only be temporary, since the slope of the phase curve $(\partial x_{t+1}/\partial x_t = w(1 - 2x_t))$ is less than -1. Fig. 1 also illustrates the property of period doubling of trajectories in chaotic systems and demonstrates the concept of low dimension: the chaotic map of x_{t+1} against x_t gives us a series of points in the phase curve. Even in the limit, these points would only form a one-dimensional set—a curve. If the x_{t+1} and x_t relationship was random, the points would be scattered about the two-dimensional phase space.

To illustrate the important property of SDIC, we graph in Figs. 2 and 3 the time paths $(x_t, t = 1, \dots, 60)$ for the logistic equation with $w = 3.750$, $x_0 = 0.10$, and $w = 3.750$, $x_0 = 0.103$, respectively. It is immediately apparent that the logistic equation has produced fairly complex time paths. Note that a change (an 'error') of only 0.003 introduced in x_0 has caused the time path to be vastly different after only a few time periods. For the first 11 periods, the time path in Fig. 2 looks almost identical to that in Fig. 3. However, the paths diverge substantially after $t = 11$. While we employ the logistic equation to demonstrate a chaotic time path here, the same sort of behaviour (where errors magnify exponentially) holds for a wide set of chaotic relations.

These illustrations allow us to suggest that the presence of chaos will hamper the success of technical traders and long-range forecasting models. Of course, one could forecast x_t perfectly if one could measure w and x_0 with infinite accuracy. As such measurement is not practical, both basic forecasting devices—extrapolation and estimation of structural forecasting models—become highly questionable in chaotic systems (see also Baumol and Benhabib, 1989).

A similar comment may be made with respect to the implications of chaos for policy makers (market regulators). If a price series is chaotic, regulators must have some knowledge of F and h to effect meaningful and non-transitory changes in price patterns. Then too, it is not obvious that regulators will succeed in promoting their agenda. Without highly accurate information of F and h, and the current state x_0, chaos would imply that regulators cannot extrapolate past behaviour to assess future movements. In effect, they would only be guessing as to the *need* for regulation. In other words, the sensible technical analyst and policy maker *ought* to be indifferent to whether or not the non-linear structure is chaotic, unless of course, she had detailed knowledge of the underlying chaotic structure.

It should be noted, however, that chaotic systems may provide some advantage to forecasting/technical analysis in the very short run (perhaps a few days when dealing with chaotic daily data). As indicated earlier, a deterministic chaotic system is, in some respects, polar to an instantaneously unpredictable system. For instance, Clyde and Osler (1997) simulate a chaotic series and demonstrate that the head-over-shoulder trading rule will be more consistent at generating profits

[5] This property follows from the requirement that local trajectories must diverge; if they were to converge, the system would be stable to disturbance and non-chaotic.

[6] The selection of $w > 3$ was not arbitrary. At $w < 3$, the series would converge to a single value. At $w = 3$, the series fluctuates between two values (or equilibria). The number of solutions continues to double (not infinitely) as w is increased beyond 3, producing a time path that is oscillatory (see Baumol and Benhabib, 1989).

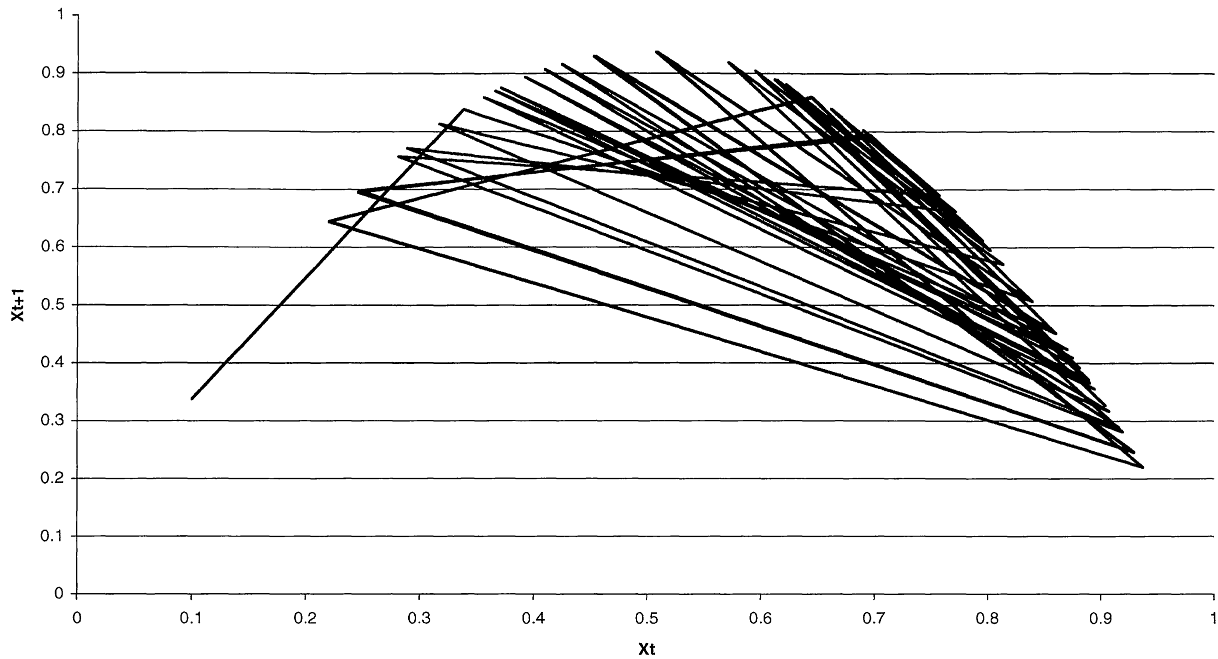

Fig. 1. Logistic map (x_{t+1}, x_t) for periods 1–60, $x_{t+1} = 3.75x(1-x)$, $x_0 = 0.10$.

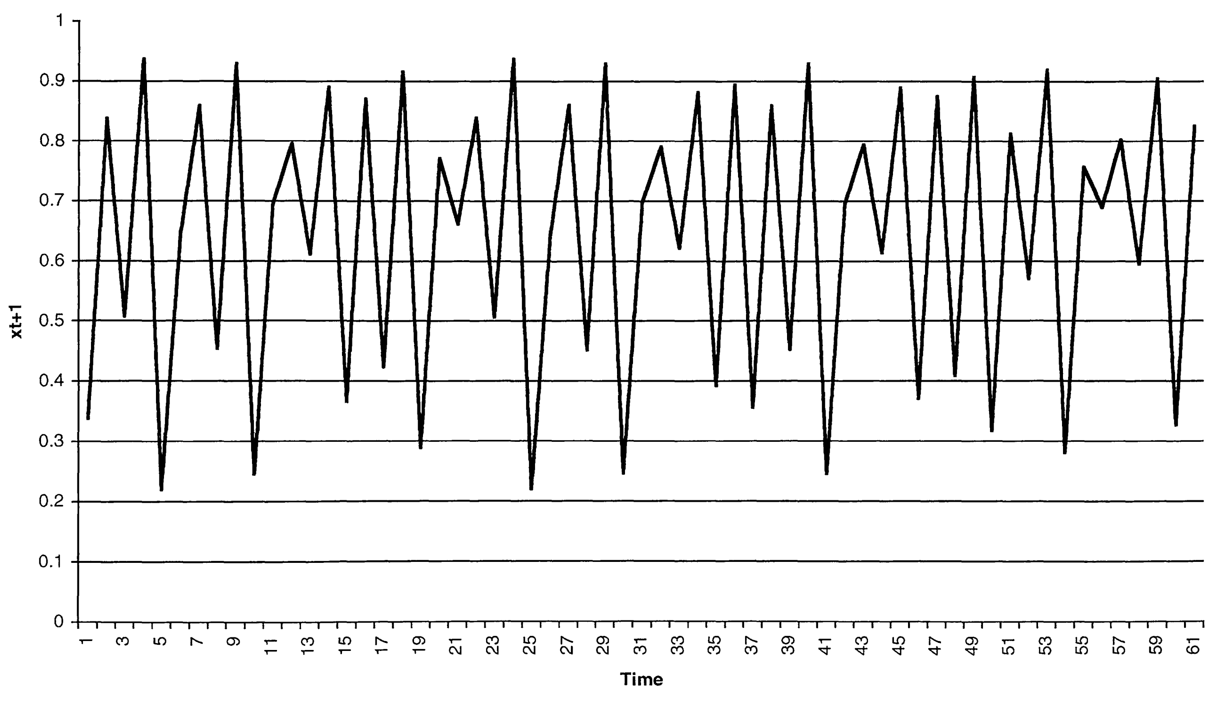

Fig. 2. Time path of logistic equation $x_{t+1} = 3.750x(1-x)$, $x_0 = 0.10$.

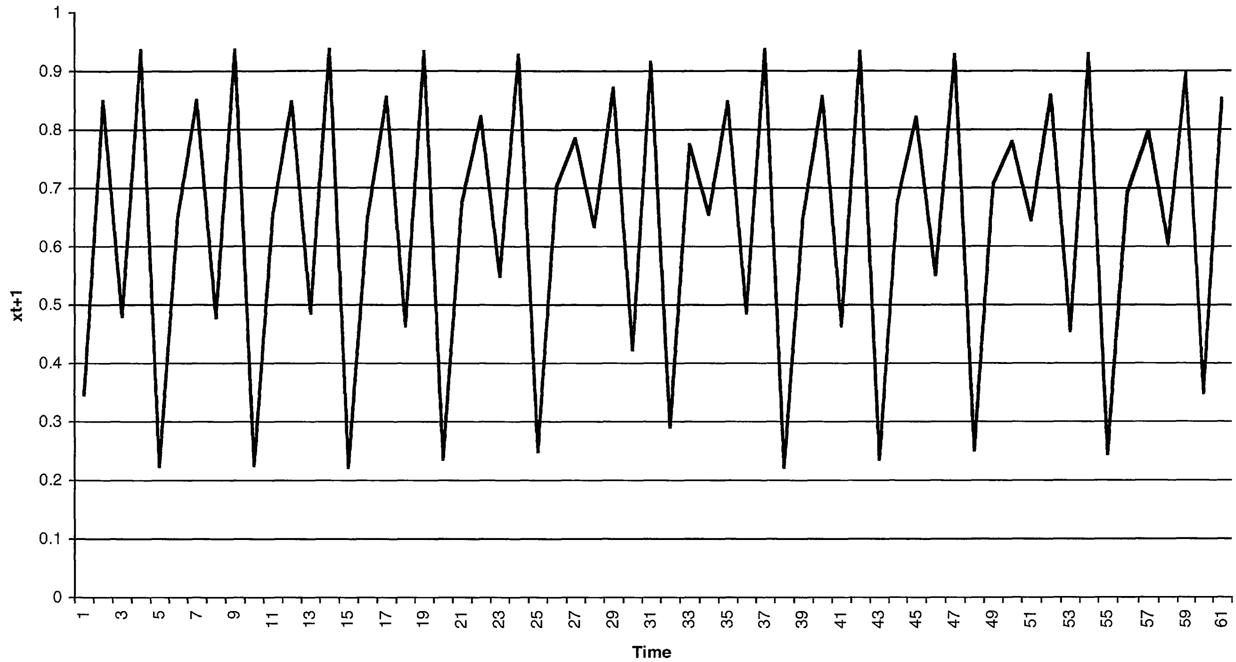

Fig. 3. Time path of logistic equation $x_{t+1} = 3.750x(1-x)$, $x_0 = 0.103$.

(relative to random trading) when applied to a known chaotic system. However, the results of Clyde and Osler also indicate that this consistency declines dramatically, so that the frequency of 'hits' employing the trading rule is not distinguishable from that of a random strategy after just a few trading periods (days).[7]

3. Testing for chaos

The known tests for chaos attempt to determine from observed time series data whether h and F are genuinely random. Three tests are employed here: the correlation dimension of Grassberger and Procaccia (1983) and Takens (1984), the BDS statistic of Brock et al. (1987) and a measure of entropy termed Kolmogorov–Sinai invariant, also known as Kolmogorov entropy. Among this group, Kolmogorov entropy probably represents the most direct test for chaos, measuring whether nearby trajectories separate as required by chaotic structure. However, this and other tests of SDIC (Lyapunov exponent) are known to provide relatively fragile results (Brock and Sayers, 1988), creating a need for alternate tests for chaos.[8] We briefly outline the construction of the tests, but do not address their properties at length, as they are well established (Brock et al., 1987, 1993).

3.1. Correlation dimension

Consider the stationary time series x_t, $t = 1, \dots, T$.[9] One imbeds x_t in an M-dimensional space by forming M-vectors, $x_t^M = \{x_t, x_{t+1}, \dots, x_{t+M-1}\}$, starting at each date. One employs the stack of these

[7] It is also noteworthy that short-term forecasting techniques, such as locally weighted regressions, are known to perform better for chaotic than for random data (Hsieh, 1991).

[8] Furthermore, there may be structural changes in the prices that lead to the failure to detect chaos despite the presence of short-run chaotic dynamics. For this reason, the failure to detect chaotic structure will be interpreted as a lack of long-lasting or long-term chaotic structure rather than as a complete lack of chaos.

[9] It is known that non-stationary processes can generate low dimensions even when not chaotic (Brock and Sayers, 1988). To rule out non-stationarity as a 'cause' for low dimension, one may difference the original series if it contains a unit root.

vectors to carry out the analysis. If the true system is n-dimensional, provided $M > 2n + 1$, the M-histories can help recreate the dynamics of the underlying system, if they exist (Takens, 1984). One can measure the spatial correlations among the M-vectors by calculating the correlation integral. For a given embedding dimension M and a distance ε, the correlation integral is given by:

$$C^M(\varepsilon) = \lim_{t \to \infty} \left\{ \frac{\text{the number of } (i, j) \text{ for which } \|x_t^M - x_j^M\| \le \varepsilon}{T^2} \right\} \tag{2}$$

where $\|\cdot\|$ is the distance induced by the norm.[10] For small values, one has $C^M(\varepsilon) \sim \varepsilon^D$ where D is the dimension of the system (see Grassberger and Procaccia, 1983). The correlation dimension in embedding dimension M is given by:

$$D^M(\varepsilon) = \lim_{\varepsilon \to 0} \lim_{T \to \infty} \left\{ \frac{\ln C^M(\varepsilon)}{\ln \varepsilon} \right\} \tag{3}$$

and the correlation dimension itself is given by:

$$D = \lim_{M \to \infty} D^M \tag{4}$$

If the value of D^M stabilises at some value (D) as M increases, then D is the correlation dimension. If D^M continues to increase with M, then the system can be regarded as stochastic since, for practical purposes, there is no difference between a high-dimensional system and randomness. Furthermore, one's computations can only be of finite resolution and data sets are of limited length, limiting the embedding level. On the other hand, if a stable low value of D^M is obtained (substantially lower than 10), there is evidence that the system is deterministic.[11]

A problem associated with the implementation of Eqs. (3) and (4) is that, with the limited length of the sample, it will almost always be possible to select a sufficiently small ε so that any two points will not lie within ε of each other (Ramsey and Yuan, 1987). A popular approach to overcome this difficulty is to instead estimate the statistic:

$$\mathrm{SC}^M = \frac{\ln C^M(\varepsilon_i) - \ln C^M(\varepsilon_{i-1})}{\ln(\varepsilon_i) - \ln(\varepsilon_{i-1})} \tag{5}$$

for various levels of M (Brock and Sayers, 1988). The SC^M statistic is a local estimate of the slope of the C^M versus ε function. Following Frank and Stengos (1989), we take the average of the three highest values of SC^M for each embedding dimension.

There are at least two ways to consider the SC^M estimates. First, the original data may be subjected to shuffling, thus destroying any chaotic structure that exists. If chaotic, the original series should provide markedly smaller SC^M estimates than its shuffled counterpart (Scheinkman and LeBaron, 1989).[12] Second, along with the requirement (for chaos) that SC^M stabilises at some low level as we increase M, we also require that linear transformations of the data leave its dimensionality unchanged (Brock, 1986). For instance, we would have evidence against chaos if AR errors provide SC^M levels that are dissimilar to those from the original series.

3.2. BDS statistic

Brock et al. (1987) employ the correlation integral to obtain a statistical test that has been shown to have strong power in detecting various types of non-linearity as well as deterministic chaos.

BDS show that if x_t is i.i.d. with a non-degenerate distribution:

$$C^M(\varepsilon) \to C^1(\varepsilon)^M, \quad \text{as } T \to \infty \tag{6}$$

for fixed M and ε. Employing this property, BDS show that the statistic:

$$W^M(\varepsilon) \to \sqrt{T}\, \frac{[C^M(\varepsilon) - C^1(\varepsilon)^M]}{\sigma^M(\varepsilon)} \tag{7}$$

where σ^M, the standard deviation of $[\cdot]$, has a limiting standard normal distribution under the null hypothesis of i.i.d. W^M is termed the BDS statistic.

[10] In practice, T is limited by the length of the data which in turn places limitations on the range of the values of ε and M to be considered.

[11] Grassberger and Procaccia (1983) determine the correlation dimension of the logistic map at 1.00 ± 0.02, the Henon map at 1.22 ± 0.01, and the Mackey Glass equation at 7.5 ± 0.15. For further discussion, see Brock et al. (1993).

[12] As discussed earlier, chaotic behaviour is associated with lower dimension than found in randomness.

Non-linearity is established if W^M is significant for a stationary series void of linear dependence. The absence of chaos is suggested if it is demonstrated that the non-linear structure arises from a known non-deterministic system. For instance, if one obtains significant BDS statistics for a stationary data series, but fails to obtain significant BDS statistics for the standardised residuals from an ARCH model, it can be said that the ARCH process explains the non-linearity in the data. This would preclude low dimension chaos.

Brock et al. (1993) examine the finite sample distribution of the BDS statistic and find the asymptotic distribution will well approximate the distribution of the statistic when: the sample has 500 or more observations; the embedding dimension is selected to be 5 or lower; ε is selected to be between 0.5 and 2 standard deviations of the data. The authors also find that the asymptotic distribution does not approximate the BDS statistic very well when applied to the standardised residuals of ARCH-type models (also see Brock et al., 1987). This is noteworthy as financial and commodity price movements are often found to have ARCH processes. The authors suggest bootstrapping the null distribution to obtain the critical values for the statistic when applying it to standardised residuals from these models.

3.3. Kolmogorov entropy

Kolmogorov entropy quantifies the concept of SDIC. Consider the two trajectories in Figs. 2 and 3. Initially, the two time paths are indistinguishable to the casual observer. As time passes, however, the trajectories diverge. Kolmogorov entropy (K) measures the speed with which this takes place.

Grassberger and Procaccia (1983) devise a measure for K which is more implementable than earlier measures of entropy. This measure is given by:

$$K_2 = \lim_{\varepsilon\to 0} \lim_{M\to\infty} \lim_{N\to\infty} \ln\left(\frac{C^M(\varepsilon)}{C^{M+1}(\varepsilon)}\right) \tag{8}$$

If a time series is non-complex and completely predictable, $K_2 \to 0$. If the time series is completely random, $K_2 \to \infty$. That is, the lower the value of K_2, the more predictable the system. For chaotic systems, one would expect $0 < K_2 < \infty$, at least in principle.

4. Evidence from commodity futures markets

We employ daily prices of nearby (expiring) futures contracts for soybeans (CBOT), corn (CBOT), wheat (CBOT), and cotton #2 (NYCE).[13] These commodities were selected for study because of their relatively long futures price histories. See Table 1 for the intervals studied. To obtain a spliced, continuous price series for each commodity, we follow common practice in tracking a particular contract until the last day of the pre-expiration month, at which point the series switch to the next nearby contract. Daily returns are obtained by taking the relative log of prices as in, $R_t = (\ln(P_t/P_{t-1})) \times 100$, where P_t represents the closing price (2:15 p.m. US Central Time for soybeans, corn, and wheat, and 3:00 p.m. for cotton) on day t.[14]

Table 1 presents diagnostics for the four R_t series. The series are stationary according to the augmented Dickey Fuller (ADF) statistics. The series are found to suffer from linear and non-linear dependencies as indicated by the $Q(12)$ and $Q^2(12)$ statistics. ARCH effects are strongly suggested by Engle's (1982) ARCH χ^2 statistic. Thus, as expected, there are clear indications that non-linear dynamics are generating the commodity price series. Whether these dynamics are chaotic in origin is the question that we turn to next.

To eliminate the possibility that linear structure or seasonalities may be responsible for the rejection of chaos by the tests employed below, we first estimate an autoregressive model for each of the four commodities with controls for possible seasonal effects, as in:

$$R_t = \sum_{i=1}^{p}\beta_i R_{t-i} + \sum_{j=1}^{12}\gamma_i M_{tj} + u_t \tag{9}$$

where the M_{jt} represent month-of-the-year dummy variables. The lag length for each series is selected based on the Akaike (1974) criterion.[15] The residual term (u_t) represents the price movements that are purged of linear relationships and seasonal influences.

[13] The data are obtained from the Futures Industry Institute, Washington, DC.

[14] We do not employ smoothing models to detrend the data, as we feel that the imposed trend reversions may erroneously be interpreted as structure (see Nelson and Plosser, 1982).

[15] The theoretical grounds for controlling for seasonality in explaining other aspects of commodity price behaviour are provided in Deaton and Laroque (1992) and Chambers and Bailey (1996), among others.

Table 1
Diagnostics on the returns (R_t) for soybean, corn, wheat and cotton

	Soybean	Corn	Wheat	Cotton
Interval	11 December 1969 to 30 March 1995	11 December 1969 to 30 March 1995	2 April 1968 to 30 March 1995	23 August 1972 to 30 March 1995
Observed	6395	6376	6802	5680
Mean	0.013	0.011	0.011	0.030
Standard deviation	1.576	1.373	1.554	1.451
ADF	−9.461[a]	−9.159[a]	−8.960[a]	−33.480[a]
ADF(*T*)	−9.514[a]	−9.197[a]	−8.989[a]	−33.504[a]
$Q(12)$	58.570[a]	52.520[a]	53.410[a]	45.540[a]
$Q^2(12)$	3066.330[a]	712.740[a]	773.570[a]	954.040[a]
ARCH(6)	864.630[a]	637.480[a]	312.080[a]	374.260[a]

ADF, ADF(*T*) represent the augmented Dickey Fuller tests (Dickey and Fuller, 1981) for unit roots with and without trend, respectively. The $Q(12)$ and $Q^2(12)$ statistics represent the Ljung–Box (Q) statistics for autocorrelation of the R_t and R_t^2 series, respectively. The ARCH(6) statistic is the Engle (1982) test for ARCH (of order 6).

[a] Represents significance at the 1% level.

Table 2
Linear structure and seasonality of the returns (R_t) for soybean, corn, wheat and cotton

RHS variable	Soybean	Corn	Wheat	Cotton
R_{t-1}	0.080[a] (6.43)	0.066[a] (5.24)	0.024[b] (1.99)	0.050[a] (3.64)
R_{t-2}	–	−0.045[a] (−3.61)	−0.063[a] (−5.20)	−0.048[a] (−3.81)
R_{t-3}	–	–	−0.017 (−1.40)	−0.003 (−0.27)
R_{t-4}	–	–	–	−0.015 (−1.13)
January	0.051 (0.76)	0.009 (0.15)	−0.063 (−0.98)	−0.031 (−0.47)
February	0.007 (0.10)	−0.067 (−1.09)	−0.104 (−1.53)	0.108 (1.56)
March	0.137[b] (2.10)	0.146[a] (2.56)	−0.118[c] (−1.85)	0.053 (0.76)
April	0.058 (0.84)	0.024 (0.48)	−0.062 (−0.94)	0.181[a] (2.66)
May	0.142[b] (2.07)	0.109[c] (1.83)	−0.125[c] (−1.92)	−0.052 (−0.77)
June	−0.031 (−0.45)	0.015 (0.26)	−0.007 (−0.11)	0.037 (0.56)
July	−0.043 (−0.64)	−0.186[a] (−3.12)	0.112[c] (1.73)	−0.230[a] (−2.80)
August	−0.113[c] (−1.69)	−0.062 (−1.08)	0.100 (1.58)	−0.028 (−0.44)
September	−0.044 (−0.64)	0.037 (0.60)	0.238[a] (3.61)	0.007 (0.11)
October	−0.037 (−0.55)	−0.031 (−0.54)	0.045 (0.71)	−0.008 (−0.13)
November	0.081 (1.18)	0.008 (0.14)	0.035 (0.52)	0.043 (0.63)
December	−0.064 (−0.95)	0.117[b] (1.99)	0.082 (1.25)	0.200[a] (2.94)
R^2	0.010	0.010	0.009	0.008
$Q(6)$	6.480	4.500	6.070	0.180
$Q(12)$	15.620	13.030	17.230	18.160

Coefficients and residual diagnostics from OLS regressions of returns on prior returns and 12 month-of-year dummies. The $Q(6)$ and $Q(12)$ statistics are Ljung–Box (Q) statistics for residual autocorrelation. Statistics in brackets are t-values.

[a] Represent significance at the 1% level.

[b] Represent significance at the 5% level.

[c] Represent significance at the 10% level.

Table 2 reports the results from the OLS regressions. There is evidence of seasonal effects in each of the four returns.[16] There is also significant linear structure in the returns, up to two lags for corn, wheat and cotton. The Durbin-h and Q statistics indicate that the residuals are free of linear structure.

4.1. Correlation dimension estimates

Table 3 reports the correlation dimension (SC^M) estimates for various components of the four returns

[16] To further examine the significance of the seasonality, we ran the above model with a constant and without the January dummy. The hypothesis that all the dummy variable coefficients are equal to zero is rejected at the 5% level for each of the commodities.

Table 3
Correlation dimension estimates

Series	Embedding dimension (m)			
	5	10	15	20
Logistic[a]	1.02	1.00	1.03	1.06
Logistic AR(1)	0.96	1.06	1.09	1.07
Logistic (AR(1), *S*)	0.97	1.06	1.08	1.06
Soybean returns	3.53	6.04	7.82	9.07
Soybean AR(1)	3.77	6.62	8.66	10.60
Soybean (AR(1), *S*)	3.98	7.07	9.34	12.80
Soybean shuffled	3.91	7.71	10.46	15.80
Corn returns	3.80	5.91	8.05	10.26
Corn AR(2)	4.03	7.28	10.56	14.87
Corn (AR(2), *S*)	4.12	7.35	10.77	17.33
Corn shuffled	3.70	7.32	11.89	18.10
Wheat returns	3.73	6.79	9.16	11.07
Wheat AR(3)	4.36	8.19	10.81	12.06
Wheat (AR(3), *S*)	4.55	8.27	11.38	15.14
Wheat shuffled	4.11	8.31	13.08	18.11
Cotton returns	4.08	7.82	11.80	13.00
Cotton AR(4)	4.29	8.07	12.84	22.92
Cotton (AR(4), *S*)	4.24	8.12	13.62	25.40
Cotton shuffled	4.18	8.82	13.83	27.16

AR(p) are autoregressive (order p) residuals, (AR(p), *S*) are residuals from autoregressive models that correct for seasonal (monthly) effects.
[a] $w = 3.750$, $n = 2000$.

series as well as the logistic series developed earlier. We report results for embedding dimensions up to 20 in order to check for saturation.[17] An absence of saturation provides evidence against chaotic structure. For instance, the SC^M estimates for the logistic map stay close to 1.00, even as we increase the embedding dimension. Moreover, the estimates for the logistic series do not change meaningfully after AR transformation or seasonal adjustment. Thus, as should be expected, the SC^M estimates are not inconsistent with chaos for the logistic series.

For the four-commodity series, on the other hand, the SC^M estimates provide evidence against chaotic structure. For instance, if one examines the estimates for the corn returns alone, one might (erroneously) make a case for low dimension chaos: the SC^M statistics seem to 'settle' close to 10. However, the estimates are substantially higher for the AR(2) series. Thus, the correlation dimension estimates suggest that there is no chaos in corn prices. Similar patterns are found for the other three commodities.

It is notable, however, that the SC^M estimates for the AR(p) series are generally smaller than the estimates of the series with seasonal correction (AR(p), *S*). For instance, the estimates for the AR(2) corn series are smaller than for the (AR(2), *S*) series. Moreover, note that the estimates for the (AR(2), *S*) corn series are not very different from the estimates from the random (shuffled) corn series. Thus, the correlation dimension estimates are found to be sensitive to controls for seasonal effects. This has important implications for future tests for chaos employing SC^M.

4.2. BDS test results

Table 4 reports the BDS statistics for the (AR(p), *S*) series corresponding to each of the four commodities and for standardised residuals ($\varepsilon/h^{0.5}$) from three types of ARCH model with their respective variance equations:

$$\text{GARCH(1, 1)}: h_t = \alpha + \alpha_1\varepsilon_{t-1}^2 + \beta_1 h_{t-1} + \beta_2\,\text{TTM}_t \tag{10}$$

$$\text{E-GARCH}: \log(h_t) = \alpha + \alpha_1\left|\frac{\varepsilon_{t-1}}{h_{t-1}}\right| + \alpha_2\left(\frac{\varepsilon_{t-1}}{h_{t-1}}\right) + \beta_1\log h_{t-1} + \beta_2\,\text{TTM}_t \tag{11}$$

Comp GARCH(1,1):

$$h_t = q_t + \alpha(\varepsilon_{t-1}^2 - q_{t-1}) + \beta_1(h_{t-1} - q_{t-1}) + \beta_2\,\text{TTM}_t,$$
$$q_t = \omega + \rho(q_{t-1} - \omega) + \Phi(\varepsilon_{t-1}^2 - h_{t-1}) \tag{12}$$

where the return equation which provides ε_t is the same as in Eq. (9), and TTM represents time-to-maturity (in days) of the futures contract.[18]

[17] Yang and Brorsen (1993), who also calculate correlation dimension for various commodity futures, compute SC^M only up to $M = 8$.

[18] The return equations from the ARCH systems provide coefficients similar to those in Table 2. We also estimated another familiar model, GARCH in mean (GARCHM). The BDS statistics from the GARCHM and GARCH(1, 1) models were found to be very similar so, in the interest of brevity, we do not provide the former. The GARCH model is due to Bollerslev (1986), the exponential model (EGARCH) is from Nelson (1991) and the asymmetric component ARCH model is a variation of the threshold GARCH model of Rabemananjara and Zakoian (1993).

Table 4
BDS statistics for (AR(p), S) residuals and standardised residuals from three ARCH models

Standard deviation	Embedding dimension			
	2	3	4	5
Panel A: soybean				
(AR(1), S) errors				
0.5	25.58[a]	33.50[a]	41.44[a]	52.80[a]
1.0	27.23[a]	34.03[a]	39.70[a]	46.10[a]
1.5	26.90[a]	32.92[a]	37.30[a]	42.13[a]
2.0	25.50[a]	30.87[a]	34.23[a]	36.80[a]
GARCH(1, 1) standard errors				
0.5	1.88[b]	2.05[a]	2.21[a]	2.78[a]
1.0	0.91	0.87	1.11	1.60[b]
1.5	0.17	0.03	0.36	0.80
2.0	0.01	−0.15	0.27	0.54
EGARCH(1, 1) standard errors				
0.5	0.79	0.76	0.78	1.32
1.0	−0.00	−0.11	0.04	0.47
1.5	−0.25	−0.37	−0.10	−0.35
2.0	0.01	0.01	0.40	0.78
AGARCH(1, 1) standard errors				
0.5	1.75[a]	2.56[a]	3.72[a]	5.62[a]
1.0	−0.02	0.42	1.38[b]	2.91[a]
1.5	−1.81	−1.90	−1.19	−0.09
2.0	−2.45	−2.97	−2.63	−2.00
Panel B: corn				
(AR(2), S) errors				
0.5	22.44[a]	30.07[a]	38.07[a]	49.18[a]
1.0	23.36[a]	29.50[a]	34.02[a]	38.98[a]
1.5	23.38[a]	28.75[a]	31.59[a]	34.16[a]
2.0	21.93[a]	25.73[a]	29.96[a]	33.57[a]
GARCH(1, 1) standard errors				
0.5	−0.08	0.68	0.72	1.18
1.0	−0.11	0.65	0.64	1.00
1.5	0.01	0.63	0.49	0.78
2.0	0.52	0.79	0.80	0.97
EGARCH(1, 1) standard errors				
0.5	−0.22	0.61	0.52	0.83
1.0	0.12	0.94	0.87	1.11
1.5	0.68	1.52	1.40	1.57
2.0	1.67	2.40	2.31	2.45[c]
AGARCH(1,1) standard errors				
0.5	0.36	0.69	1.01[c]	1.41[a]
1.0	0.48	1.10[c]	1.10[c]	1.34[a]
1.5	0.61	0.99[c]	1.01[c]	1.15[b]
2.0	0.91	1.23[c]	1.09[c]	1.23[b]
Panel C: wheat				
(AR(2), S) errors				
0.5	16.27[a]	20.32[a]	23.03[a]	26.16[a]
1.0	18.99[a]	22.88[a]	25.19[a]	27.77[a]
1.5	21.28[a]	25.33[a]	27.36[a]	29.39[a]

Table 4 (*Continued*)

Standard deviation	Embedding dimension			
	2	3	4	5
2.0	21.05[a]	25.07[a]	27.00[a]	28.60[a]
GARCH(1, 1) standard errors				
0.5	4.75[a]	5.30[a]	5.21[a]	5.72[a]
1.0	4.61[a]	4.73[a]	4.37[a]	4.66[a]
1.5	4.29[a]	4.14[a]	3.60[a]	3.68[a]
2.0	3.92[a]	3.64[a]	3.04[a]	2.89[a]
EGARCH(1, 1) standard errors				
0.5	2.83[b]	2.88[b]	2.29[b]	2.22[b]
1.0	2.89[c]	2.67[c]	1.89[c]	1.78[c]
1.5	2.94[c]	2.54	1.73	1.61
2.0	3.05[c]	2.64	1.86	1.66
AGARCH(1, 1) standard errors				
0.5	3.35[a]	3.98[a]	4.50[a]	5.62[a]
1.0	2.96[a]	3.25[a]	3.64[a]	4.84[a]
1.5	2.17[a]	2.26[a]	2.58[a]	3.78[a]
2.0	1.24	1.12[c]	1.19[c]	2.11[a]
Panel D: cotton				
(AR(4), S) errors				
0.5	16.60[a]	19.81[a]	21.86[a]	25.27[a]
1.0	16.58[a]	20.00[a]	22.36[a]	25.34[a]
1.5	15.29[a]	18.23[a]	20.65[a]	23.05[a]
2.0	14.58[a]	17.05[a]	19.22[a]	21.20[a]
GARCH(1, 1) standard errors				
0.5	3.87[a]	3.11[a]	2.22[a]	2.57[a]
1.0	2.60[a]	1.95[a]	1.39[c]	1.63[b]
1.5	1.64	0.93	0.52	0.66
2.0	1.00	0.18	−0.12	−0.06
EGARCH(1, 1) standard errors				
0.5	3.37[a]	2.45[c]	1.51	1.71[c]
1.0	2.41	1.62	0.92	1.00
1.5	1.74	0.92	0.42	0.50
2.0	1.43	0.65	0.34	0.44
AGARCH(1, 1) standard errors				
0.5	4.74[a]	7.06[a]	10.90[a]	17.10[a]
1.0	1.68[a]	2.05[a]	2.50[a]	3.32[a]
1.5	−0.20	−0.28	−0.15	0.07
2.0	0.16	−0.34	−0.56	−0.70

BDS statistics are evaluated against critical values obtained from Monte Carlo simulation (see Appendix A).
[a] Represent significance at the 1% level.
[b] Represent significance at the 5% level.
[c] Represent significance at the 10% level.

The time-to-maturity variable is intended to control for any maturity effects in the series (Samuelson, 1965). The BDS statistics are evaluated against critical values obtained by bootstrapping the null distribution for each of the GARCH models (see Appendix A).

The estimates from the above variance equations are discussed later in the paper.

The BDS statistics strongly reject the null hypothesis of no non-linearity in the (AR(p), S) errors for each of the commodity futures. This evidence, that the commodity futures have non-linear dependencies, is consistent with the findings in Table 1, and in Yang and Brorsen (1993), among others. The BDS statistics for the standardised residuals from the ARCH-type models, however, indicate that the source of the non-linearity in at least two of the four commodities is *not* chaos. For instance, for the soybean contract, the BDS statistics are dramatically lower (relative to those for the (AR(p), S) errors) for all the standardised residuals, and are consistently insignificant for the EGARCH model. For the corn contract, the GARCH(1, 1) model seems to satisfactorily capture the non-linear dependence in the data: the BDS statistics for the standardised residuals from this model are consistently insignificant. On the other hand, for the wheat and cotton contacts, the BDS statistics show persisting non-linearities even after the corrections for ARCH effects. Nonetheless, even for these two commodities, the BDS statistics for the ARCH residuals are much smaller (albeit significant) than those for the (AR(p), S) residuals.

On the whole, the BDS test results further support the notion that the non-linear dependence in commodity futures is explained by dynamics other than chaos. Certainly, for soybean and corn contracts the evidence is compelling that the non-linear dependencies in commodity futures arise from ARCH-type effects, rather than from complex, chaotic structures. Finally, it is also noteworthy that the EGARCH model performed reasonably well for each contract, even though it failed to completely explain the non-linearities in the majority of contracts.

4.3. Entropy estimates

Fig. 4 plots the Kolmogorov entropy estimates (embedding dimension 15–32) for the logistic map ($w = 3.750$, $x_0 = 0.10$) as well as the (AR(p), S) soybean, wheat, corn and cotton series. The entropy estimates for a twice-shuffled wheat return series are also presented for comparison. The estimates for the logistic map and the shuffled series provide the benchmarks for a known chaotic and a generally random series. The entropy estimates for the (AR(p), S) soybean, corn and wheat series show little signs of 'settling down' as those for the logistic map do. They behave much more like the entropy estimates for the shuffled series: a general rise in the K_2 statistic as one increases the embedding dimension. The entropy estimates for (AR(p), S) cotton series, however, are not as convincing. In general, the plots in Fig. 4 reaffirm the correlation dimension and BDS test results: there is little evidence of low dimension chaos in commodity futures prices.

4.4. ARCH and maturity effects in futures markets

It is apparent from the BDS statistics presented in Table 4 that the EGARCH model effectively explains the non-linearities in the soybean contract. In the corn contract, the GARCH(1, 1) model, along with the EGARCH model, performs well. Using these results, we can re-examine the Samuelson hypothesis on the relationship between contract maturity and variance employing the appropriately modelled variance structure. The Samuelson hypothesis implies that the volatility in futures price changes increases as a contract's delivery date approaches. If the Samuelson hypothesis is valid, proper valuation of futures and futures options would require that the term-structure of the volatility be estimated (see Bessembinder et al., 1996).

Table 5 reports the maximum likelihood results for soybean and cotton, the two contracts for which we have succeeded in isolating the appropriate non-linear model. In the interest of brevity, we do not present the results from the mean equations. The results indicate strong ARCH effects and, in the case of soybean, significant asymmetries in the variance structure. The Samuelson hypothesis is clearly supported for both the contracts: the time-to-maturity (TTM) variable is negative and significant in both equations. As we approach maturity (as TTM falls), the conditional variance ($\log(h_t)$ for soybean and h_t for corn) increases. However, it is notable that while TTM is found to be significant in the variance equation, this variable does not play a large role as a 'control variable' in the tests for chaos: the BDS statistics remained almost unchanged when we employed standardised residuals from models without TTM. In other words, the correlation-integration based tests for chaos are not as

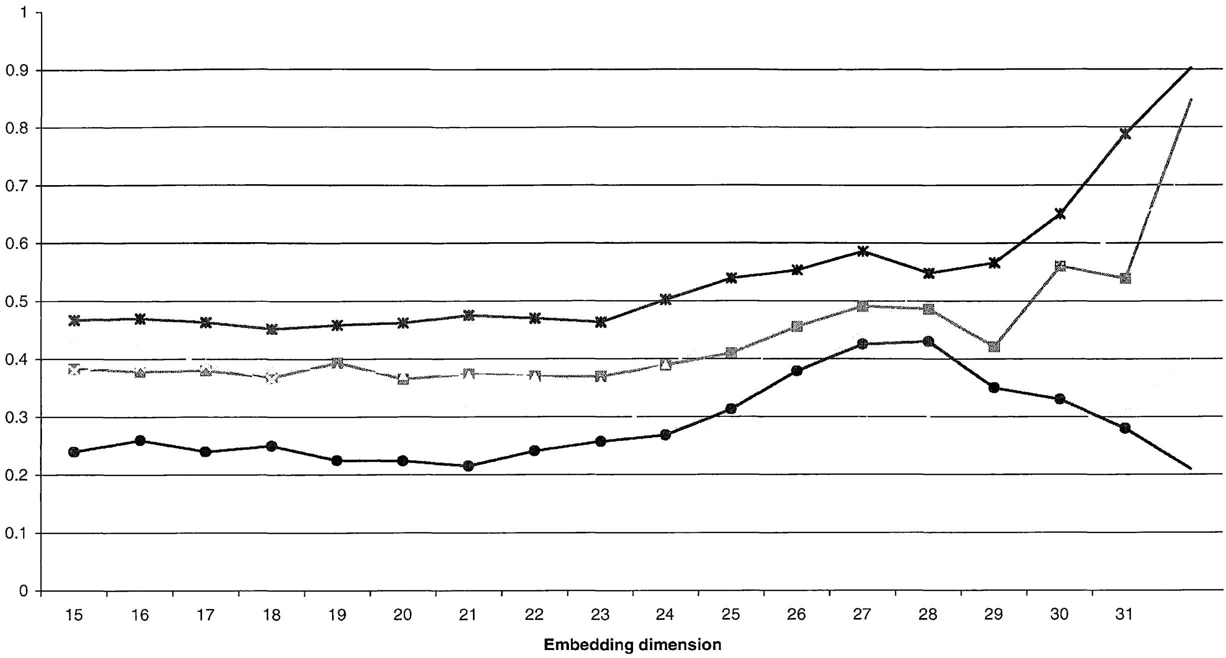

Fig. 4. Kolmogorov entropy estimates.

Table 5
Estimates of ARCH dynamics in soybean and corn futures

Soybean ($\log(h_t)$)		Corn (h_t)	
Constant	-0.0633^{a} (−8.98)	Constant	0.1072^{a} (22.47)
$\lvert\varepsilon/h\rvert_{t-1}$	0.0770^{a} (16.17)		
$(\varepsilon/h)_{t-1}$	0.0426^{a} (8.39)	ε_{t-1}	0.0617^{a} (19.81)
$\log(h_{t-1})$	0.8254^{a} (646.59)	h_{t-1}	0.9111^{a} (239.02)
TTM	-0.0014^{a} (−7.39)	TTM	-0.0026^{a} (26.91)
Log-likelihood	−10791.98	Log-likelihood	−10238.38
χ^2	2244.02^{a}	χ^2	1585.06^{a}

Statistics in brackets are t-values. TTM is time-to-maturity in days. The χ^2 statistics based on comparison with the corresponding OLS model.

[a] Represents significance at the 1% level.

sensitive to controls for TTM as they are to controls for seasonality.

5. Conclusion

The evidence of chaos in economic time series such as GNP and unemployment has thus far been weak. On the other hand, the few studies of commodity prices have generally found evidence consistent with low dimension chaos. Why is the evidence of chaos stronger in commodity prices? Could the relatively short data spans in earlier studies on commodities and the lack of controls for seasonal patterns account for the differences in the evidence between commodity prices and aggregated economic time series?

Employing over 25 years of data, we conduct a battery of tests for the presence of low-dimensional chaotic structure in four important commodity futures prices. Daily returns data from the nearby contracts are subjected to correlation dimension tests, BDS tests and tests for entropy. While we find strong evidence of non-linear dependence in the data, the evidence is not consistent with long-lasting chaotic structure. Our test results indicate that various ARCH processes explain the non-linearities in at least two of the contracts. We also make a case that employing seasonally adjusted price series is important to obtaining robust results via the existing tests for chaotic structure.

For the soybean and corn contract, we isolate appropriate ARCH models and examine the Samuelson hypothesis of a maturity effect in futures prices. The EGARCH results for soybean futures and the GARCH(1, 1) results for corn futures provide evidence in favour of the Samuelson hypothesis: volatility in futures returns increases as one approaches maturity. However, the tests for chaos were found to be less sensitive to controls for time-to-maturity than to controls for seasonality.

Acknowledgements

Pamplin Fellowship support is gratefully acknowledged. The paper has benefited from the comments of two anonymous reviewers. The authors are solely responsible for remaining errors.

Appendix A

Simulated critical values for the BDS test statistic

Embedding diversion	ε/σ			
	0.5	1.0	1.5	2.0
GARCH(1, 1) (97.5% critical values)				
2	1.62	1.53	1.42	1.25
3	1.76	1.63	1.45	1.44
4	2.35	2.21	2.16	1.97
5	2.42	2.28	2.25	2.10
Exponential GARCH (97.5% critical values)				
2	2.75	2.54	2.10	1.83
3	3.30	3.07	2.42	2.38
4	3.48	3.31	2.66	2.56
5	3.66	3.47	2.97	2.61

Appendix A. (*Continued*)

Embedding diversion	ε/σ			
	0.5	1.0	1.5	2.0
Asymmetric component GARCH (2.5% critical values)				
2	−2.86	−2.29	−1.78	−1.74
3	−3.51	−2.89	−2.49	−2.26
4	−3.64	−3.01	−2.81	−2.55
5	−3.67	−3.12	−3.08	−2.64
Asymmetric component GARCH (97.5% critical values)				
2	1.40	1.13	1.02	0.80
3	1.47	1.27	1.17	0.93
4	1.62	1.28	1.22	1.00
5	1.82	1.40	1.31	1.07

Simulated values based on Monte Carlo simulations of 2000 observations each. Two-hundred and fifty replications of the GARCH model ($\alpha_1 = 0.10$, $\beta_1 = 0.80$), the EGARCH model ($\alpha_1 = 0.05$, $\alpha_2 = 0.05$, $\beta_1 = 0.80$), and the asymmetric component model ($\alpha_1 = 0.05$, $\beta = 0.10$, $\rho = 0.80$, $\phi = 0.05$) were generated. BDS statistics for four embedding dimensions and $\varepsilon = 0.5$, 1, 1.5 and 2 standard deviations of the data were then computed for the 250×3 simulated series.

References

Akaike, H., 1974. A new look at statistical model identification. IEEE Trans. Automatic Control 19, 716–723.

Baumol, W.J., Benhabib, J., 1989. Chaos: significance, mechanism, and economic applications. J. Econ. Perspect. 3, 77–105.

Benhabib, W.J., Day, R.H., 1981. Rational choice and erratic behaviour. Rev. Econ. Stud. 48, 459–472.

Benhabib, W.J., Day, R.H., 1982. A characterisation of erratic dynamics in the overlapping generations model. J. Econ. Dyn. Control 4, 37–55.

Bessembinder, H., Coughenour, J.F., Seguin, P.J., Smoller, M.M., 1996. Is there a term structure of futures volatilities? Reevaluating the Samuelson hypothesis. J. Derivatives Winter, 45–58.

Blank, S.C., 1991. “Chaos” in futures markets? A non-linear dynamical analysis. J. Futures Markets 11, 711–728.

Blume, L., Easley, D., O’Hara, M., 1994. Market statistics and technical analysis: the role of volume. J. Finance 49, 153–181.

Bohan, J., 1981. Relative strength: further positive evidence. J. Portfolio Manage. Fall, 36–39.

Bollerslev, T., 1986. Generalised autoregressive conditional heteroskedasticity. J. Econometr. 31, 307–327.

Brock, W.A., 1986. Distinguishing random and deterministic systems. J. Econ. Theory 40, 168–195.

Brock, W.A., Dechert, W., 1988. Theorems on distinguishing deterministic and random systems. In: Barnett, W., Berndt, E., White, H. (Eds.), Proceedings of the Third Austin Symposium on Dynamic Econometric Modelling. Cambridge University Press, Cambridge.

Brock, W.A., Sayers, C.L., 1988. Is the business cycle characterised by deterministic chaos? J. Monetary Econ. 22, 71–90.

Brock, W.A., Dechert, W., Scheinkman, J., 1987. A Test of Independence Based on the Correlation Dimension, Unpublished Manuscript. University of Wisconsin, Madison/University of Houston/University of Chicago.

Brock, W.A., Lakonishok, J., LeBaron, B., 1992. Simple technical trading rules and the stochastic properties of stock returns. J. Finance 47, 1731–1764.

Brock, W.A., Hsieh, D.A., LeBaron, B., 1993. Non-Linear Dynamics, Chaos, and Instability: Statistical Theory and Economic Evidence. MIT Press, Cambridge, MA.

Brush, J., 1986. Eight relative strength methods compared. J. Portfolio Manage. Fall, 21–28.

Chambers, M.J., Bailey, R.E., 1996. A theory of commodity price fluctuations. J. Pol. Econ. 104, 924–957.

Chang, P.H.K., Osler, C.L., 1995. Head and Shoulder: Not Just a Flaky Pattern. Federal Reserve Bank of New York Staff Papers, No. 4.

Clyde, W.C., Osler, C.L., 1997. Charting: chaos theory in disguise? J. Futures Markets 17, 489–514.

Deaton, A., Laroque, G., 1992. On the behaviour of commodity prices. Rev. Econ. Stud. 59, 1–23.

DeCoster, G.P., Labys, W.C., Mitchell, D.W., 1992. Evidence of chaos in commodity futures prices. J. Futures Markets 12, 291–305.

Deneckere, R., Pelikan, S., 1986. Competitive chaos. J. Econ. Theory 40, 12–25.

Devaney, R.L., 1986. An Introduction to Chaotic Dynamical Systems. Benjamin/Cummings Publishing, Menlo Park, CA.

Dickey, D.A., Fuller, W.A., 1981. Likelihood ratio statistics for autoregressive time series with a unit root. Econometrica 49, 1057–1072.

Engle, R.F., 1982. Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation. Econometrica 50, 987–1007.

Frank, M., Stengos, T., 1989. Measuring the strangeness of gold and silver rates of return. Rev. Econ. Stud. 456, 553–567.

Grassberger, P., Procaccia, I., 1983. Measuring the strangeness of strange attractors. Physica 9, 189–208.

Guckenheimer, J., Holmes, P., 1986. Non-Linear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields. Springer, New York, NY.

Hsieh, D.A., 1991. Chaos and non-linear dynamics: applications to financial markets. J. Finance 46, 1839–1876.

LaBaron, B., 1991. Technical Trading Rules and Regimes Shifts in Foreign Exchange. University of Wisconsin, Social Sciences Research Institute Working Paper.

Lichtenberg, A.J., Ujihara, A., 1988. Application of non-linear mapping theory to commodity price fluctuations. J. Econ. Dyn. Control 13, 225–246.

Nelson, D., 1991. Conditional heteroskedasticity in asset returns: a new approach. Econometrica 59, 347–370.

Nelson, C., Plosser, C., 1982. Trends and random walks in macroeconomic time series. J. Monetary Econ. 10, 139–162.

Pruitt, S.W., White, R.E., 1988. The CRISMA trading system: who says technical analysis can't beat the market? J. Portfolio Manage. 14, 55–58.

Pruitt, S.W., White, R.E., 1989. Exchange-traded options and CRISMA trading: who says technical analysis can't beat the market? J. Portfolio Manage. 15, 55–56.

Rabemananjara, R., Zakoian, J.M., 1993. Threshold ARCH models and asymmetries in volatility. J. Appl. Econometr. 8, 31–49.

Ramsey, J., Yuan, H., 1987. The Statistical Properties of Dimension Calculations Using Small Data Sets. C.V. Starr Center for Applied Economics, New York University.

Samuelson, P.A., 1965. Proof that properly anticipated prices fluctuate randomly. Indus. Manage. Rev. 6, 41–63.

Scheinkman, J., LeBaron, B., 1989. Non-linear dynamics and stock returns. J. Business 62, 311–337.

Stutzer, M.J., 1980. Chaotic dynamics and bifurcation in a macro-model. J. Econ. Dyn. Control 2, 253–276.

Takens, F., 1984. On the numerical determination of the dimension of an attractor. In: Dynamical Systems and Bifurcations, Lecture Notes in Mathematics. Springer, Berlin.

Taylor, S.J., 1994. Trading futures using a channels rule: a study of the predictive power of technical analysis with currency examples. J. Futures Markets 14, 215–235.

Yang, S., Brorsen, B.W., 1993. Non-linear dynamics of daily futures prices: conditional heteroskedasticity or chaos? J. Futures Markets 13, 175–191.