Journal of Business & Economic Statistics, July 1992, Vol. 10, No. 3

On Determining the Dimension of Real-Time Stock-Price Data

E. Scott Mayfield
Department of Economics, Boston College, Chestnut Hill, MA 02167

Bruce Mizrach
Department of Finance, The Wharton School, Philadelphia, PA 19104

We estimate the dimension of high-frequency stock-price data using the correlation integral of Grassberger and Procaccia. The data, even after filtering, appear to be of low dimension. To control for dependence in higher moments, we use a new technique known as the method of delays in our reconstruction. Delaying the data leads dimension estimates similar to random processes. We conclude that the data are either of low dimension with high entropy or nonlinear but of high dimension.

KEY WORDS: Correlation integral; Entropy; Method of time delays; Nonlinear dynamics.

The study of nonlinear dynamics has generated conceptual breakthroughs in areas as diverse as fluid dynamics, cardiology, and biology. Seemingly random phenomena have been modeled successfully as low-dimensional nonlinear maps. Several authors (Barnett and Chen 1988; Brock 1986; Brock and Sayers 1988; Frank and Stengos 1988, 1989, Scheinkman and LeBaron 1989) have applied nonlinear analysis to economic and financial data.

This article is an attempt to uncover evidence of complex dynamics in U.S. equity markets. We determine the dimension and entropy of real-time stock-price data using the correlation integral of Grassberger and Procaccia (1983a). Our data track the intradaily movements in the Standard and Poor's (S&P) 500 cash index, sampled at 20-second intervals, approximately 20,000 observations in all. No previous analysis of economic or financial data has used sample sizes approaching those used in the physical sciences, typically 15,000 to 40,000 observations. This allows us to compute dimension estimates with a much higher level of statistical reliability. As Ramsey and Yuan (1990) noted, dimension estimates can be significantly biased in samples under 5,000 observations.

The work of Takens (1980, 1983) underlies the empirical analysis of nonlinear systems. Takens showed that one can determine the dimension of a dynamical system from a univariate time series. Proper reconstruction of the dynamical system is critical though and requires eliminating temporal dependence from the data; failure to do so often results in dimension estimates that are biased downward. Brock (1986) noted that dimension estimates should be invariant to smooth transformations of the data and advocated reconstruction using the residuals from a linear time series model. Subsequently, economists have adopted the Brock residual test as the standard diagnostic.

We argue that filtering techniques can be misleading. Our concern is that the Brock residual test has little power against the alternative of dependence in higher moments, a property often found in financial data. We show that a smooth nonlinear data-generating mechanism for stock returns will generate temporal dependence in *all* of the data's moments. Under these circumstances, no finite set of filters is adequate to remove temporal dependence. Even after autoregressive moving average (ARMA) and generalized autoregressive conditional heteroscedasticity (GARCH) filtering, dimension estimates will still be biased.

This article proposes that the economics literature follow the physical sciences in using the method of time delays. Rather than trying to filter out the dependence in the time series, one uses lags of the data at which the series has become approximately independent. Choosing this lag is a crucial free parameter. If we choose too short a delay, we fail to eliminate temporal dependence. Alternatively, if the delay is too long, it is difficult to detect underlying low-dimensional structure. With sensitive dependence, any noise on the system will cause the attractor's structure to be obliterated after a limited number of iterations. Data sampled at a daily frequency stand little chance of uncovering the dynamic properties of a chaotic attractor operating at an hourly frequency.

There is an optimal frequency, in a statistical sense, at which to reconstruct the system's dynamics. This occurs at some fraction of the mean orbital time at which the data are nearly independent. In our analysis, a minimum of mutual information occurs at a sampling interval of approximately five minutes.

After filtering with both ARMA and GARCH models, we find very strong evidence of low dimension. From the shift in the correlation integral, we also find a very high degree of entropy. By the methods currently employed in the economics literature, one could construe this as very strong evidence for chaotic dynamics.

When we reconstruct using the method of time delays, however, we find very different results. Our dimension estimates resemble those of random processes. We conclude that the data are either not of low dimension or that entropy renders the market nearly random after only five minutes.

The organization of the article is as follows: Section 1 describes the properties of the data we are looking for, dimension and entropy. Section 2 describes the correlation integral and how to construct estimates of dimension and entropy. Section 3.1 details the standard practice in the economics literature of filtering with ARMA and GARCH models. The bulk of our contribution is in the empirical analysis of Section 3.2, where we compute estimates of dimension and entropy using filtered data. Section 4 argues why filtering is likely to be inadequate. Section 5 repeats the dimension calculations using the method of delays and comes to very different conclusions. Section 6 includes a summary of our results.

1. MATHEMATICAL CHAOS

In this article, we study discrete dynamical systems of the form

$$x_{t+1} = F(x_t), \tag{1}$$

where $F : R^n \rightarrow R^n$. Much of the empirical work in nonlinear dynamics is concerned with uncovering the dimension of the attracting set of (1). An attractor is a compact set, Λ, with the property that there is a neighborhood of Λ such that for almost every initial condition the limit set of iterates of (1) as $t \rightarrow \infty$ is Λ. As a practical matter, we focus attention on finding systems of low dimension; these are the only systems that can be reliably distinguished from random ones.

1.1 Reconstructing Complex Dynamics

In our application, x_t might be thought of as *the market*. We receive only a scalar signal of its "heartbeat" in the form of a unvariate time series of the S&P 500 index of stocks,

$$p_t = h(x_t), \tag{2}$$

where $h : R^n \rightarrow R$ is an observer function of the market. What hope, if any, might we have of recovering the market dynamics from p_t? Takens (1980, 1983) showed that much of the system's dynamics is preserved as long as F and h are at least C^2 functions. Define an m-dimensional vector constructed from our univariate time series,

$$p_t^m = (p_t, \ldots p_{t+m-1}) = (h(x_t), \ldots, h(F^{m-1}(x_t))) \equiv J_m, \tag{3}$$

where F^{m-1} is the composition of F with itself $m - 1$ times. For example, $F^2(x_t) = F(F(x_t))$. We can now state the following result.

Proposition 1.1 (*Takens 1980*). For smooth pairs (h, F), the map $J_m : R^n \rightarrow R^m$ will be an embedding for $m \geq 2n + 1$.

Takens' theorem is really quite remarkable and has motivated nearly all of the empirical research on chaos. As Brock (1986) noted, the theorem implies that the dynamical behavior of the m vectors of stock-price data will resemble the unobservable dynamical behavior of the market process. Most important for our purposes is that the embedding preserves both the dimension and entropy of the dynamical system.

In summary, the Takens embedding theorem allows the degree of complexity of an underlying system to be recovered from a scalar time series that is smoothly related to the state variables of the system. In practice, however, proper reconstruction of the attractor is crucial; it is extremely important to remove the temporal dependence of nearby points on the reconstructed attractor. In implementing the Grassberger and Procaccia (GP) algorithm, the choice of time delay in the m vectors is crucial in properly constructing the embedding. If the data are sampled at very fine or coarse intervals, the dynamics can remain hidden from the analyst. We develop the GP procedure in Section 2 and discuss in greater detail various experimental considerations in Section 3.

1.2 Lyapunov Exponents and Entropy

A distinguishing property of chaotic processes is that of sensitive dependence; points that are initially close together tend to spread apart eventually. This property may lead the analyst to mistake a chaotic system for a random one. Combined with measurement limitations of the current state, sensitive dependence places an upper bound on the ability to forecast chaotic processes, even if the model F is known with perfect certainty.

To formalize the notion of sensitive dependence, we use the concept of the Lyapunov exponent. Let D_xF^N be the $n \times n$ Jacobian matrix evaluated at $x \in R^n$, and let $(D_xF^N)^*$ be the transpose of D_xF^N.

Definition 1.1 (*Guckenheimer and Holmes 1985*). Consider subspaces $V_t^{(1)} \supset V_t^{(2)} \supset \cdots \supset V_t^{(n)}$ in the tangent space at $F^N(x)$ and numbers $\mu_1 \geq \mu_2 \geq \cdots \geq \mu_n$ with the properties that (a) $D_xF(V^{(j)t}) = V_{t+1}^{(j)}$, (b) dimension $(V_t^{(j)}) = n + 1 - j$, and (c) $\lim_{N\rightarrow\infty}(1/N) \times \ln\sqrt{(D_xF^N)^*(D_xF^N) \cdot v} = \mu_j$ for all $v \in V_0^{(j)} - V_0^{(j+1)}$. The μ_j are then called the *Lyapunov exponents*.

As is standard in dynamic models, the largest exponent is crucial. If the map F is chaotic, at least one of the exponents must be positive. We treat the existence of a positive Lyapunov exponent as the definition of a chaotic system.

The Lyapunov exponents tell us about the average rate of expansion or contraction along the entire trajectory. In a chaotic system, points are being separated continually from one another in at least one dimension. Small discrepancies in the initial state become magnified and eventually become distinct trajectories.

In the empirical work that follows, we will not estimate the exponents directly. Rather, we will estimate the system's entropy. Pesin (1977) showed that the metric entropy equals the sum of the positive Lyapunov exponents. If we find evidence of low dimension and positive entropy, this will be strong evidence for the existence of chaotic dynamics.

2 THE CORRELATION INTEGRAL

This section details the workhorse of the empirical literature on nonlinear dynamics, the correlation integral. The first part is devoted to dimension estimation and the second part to entropy.

2.1 Correlation Dimension

Consider vector m histories of the S&P 500 index,

$$p_j^m = (p_j, \ldots, p_{j+m-1}). \tag{4}$$

The correlation integral measures the number of m vectors within an ε neighborhood of one another (Grassberger and Procaccia 1983a, 1984). In our notation, the correlation integral is defined as

$$\tilde{C}_m(\varepsilon) = \lim_{N\to\infty} \frac{1}{N^2} \times \#\{(j, k) | \; \|p_j^m - p_k^m\| < \varepsilon\}, \quad m = 2, 3, \ldots, \tag{5}$$

where $\|\cdot\|$ is some norm, N is the number of m histories, m is the embedding dimension, and # denotes the cardinality of the set. As $\varepsilon \to 0$, $\tilde{C}_m(\varepsilon) \sim \varepsilon^\nu$, where ν, the correlation exponent, is a lower-bound estimate of the Hausdorff dimension. Thus, for small ε,

$$\ln_2 \tilde{C}_m(\varepsilon) = \ln_2 k + \nu \ln_2 \varepsilon, \tag{6}$$

where k is a constant. In practice, (6) is calculated over a range of ε's. Brock (1986) showed that the correlation exponent is independent of any two norms and independent of m for $m \geq 2n + 1$.

2.2 Kolmogorov Entropy

Direct numerical computation of Lyapunov exponents has proven to be quite difficult in experimental systems. The most popular algorithm has been the one proposed by Wolf, Swift, Swinney, and Vastano (1985). Eckmann, Kamphurst, Ruelle, and Scheinkman (1988) and Barnett and Chen (1988) applied this algorithm to stock-market data and monetary indexes, respectively. Recently, McCaffrey, Ellner, Gallant, and Nychka (1991) proposed a new approach that works directly with the Jacobian matrix using nonparametric regression.

Estimates of Lyapunov exponents, however, have proven to be quite sensitive to embedding dimension and initial conditions, so we motivate an alternative procedure that can be implemented using the correlation integral. Grassberger and Procaccia (1983b, 1984) showed that the vertical change in the position of the invariant portion of the correlation integral (i.e., where the slope is unchanging) is a lower-bound estimate of Kolmogorov entropy. Specifically, they defined

$$K_{2,d}(\varepsilon) = \frac{1}{\tau} \ln \frac{C_d(\varepsilon)}{C_{d+1}(\varepsilon)}, \tag{7}$$

where $C_d(\varepsilon)$ is the value of the correlation integral for embedding dimension d and delay time between observations τ. They showed that

$$\lim_{\substack{d\to\infty \\ \varepsilon\to 0}} K_{2,d}(\varepsilon) \sim K_2, \tag{8}$$

where K_2 is order-2 Renyi entropy, which is a lower-bound estimate of Kolmogorov entropy. As embedding dimension increases, the average vertical distance between the integrals is the GP lower-bound estimate of Kolmogorov entropy.

3. DIMENSION AND ENTROPY ESTIMATES USING FILTERED DATA

In this section, we use the correlation integral to estimate both the dimension and entropy of our real-time stock-price data. In seeking to replicate previous work in the economics literature, we first look at transformations of the data that rely only on filtering. In Section 5, we repeat our calculations using the method of delays.

3.1 Filtering

Brock and Sayers (1988) studied the effect of temporal dependence on dimension estimates. They reported low-dimension estimates for a number of quarterly economic time series, but they believed these estimates to be spurious because of near unit roots. In phase space, data that are highly correlated will lie nearly along a line. Thus the reconstructed attractor will be stretched along a ray, leading the data analyst to underestimate the true dimension. Grassberger and Procaccia (1983a) noted the same phenomena in continuous time processes sampled at very close intervals.

Addressing this problem, Brock (1986) proposed using the residuals from a linear time series model to estimate dimension. Under the chaotic null, the residuals will preserve the dimension of the attracting set.

Proposition 3.1 (*Brock* [*1986*] *Residual Test Theorem*). Consider the model (1), (3) with F possessing a chaotic attractor. The residuals from a finite dimensional autoregressive (AR) process fit to p_t will have the same dimension as p_t.

After filtering the data with various ARMA models, Brock and Sayers [1988] rejected the hypothesis that the true data-generating process is of low dimension. All subsequent work in the economics literature has followed Brock's procedure of whitening the data with filters.

Scheinkman and LeBaron (1989) proposed another diagnostic tool, shuffling the data. By randomizing the original series, one creates a series without temporal structure. For an iid process, randomizing will not effect the dimension, since the shuffled series will also be iid. For data generated by a low-dimensional chaotic attractor, however, the loss of structure will cause the data to become more space filling. Thus dimension calculations based on the shuffled data are a useful benchmark against which to compare actual dimension estimates.

Few studies have passed the shuffle diagnostic after ARMA filtering. Frank and Stengos (1989) and Scheinkman and LeBaron (1989) reported dimension estimates, in the range of 6 to 7 after filtering, that pass the shuffle diagnostic. If filtering has not removed all of the temporal dependence, however, these results may still be biased. Only if the dimension estimates are robust to a delay time reconstruction can we be confident that filtering has removed all of the temporal dependence.

3.2 Data

We analyze the S&P 500 stock index, sampled at approximately 20-second intervals, as a measure of real-time market-wide price fluctuations. The S&P 500 index is a weighted average of stock prices, providing a smooth aggregator of the underlying market dynamics. The principal problem in using an index is the possibility of introducing noise into the system through aggregation bias. The quality of the index can be crucial for estimates of dimension and entropy. Barnett and Chen (1988) found evidence for chaotic dynamics in five Divisia but only one simple-sum monetary aggregate. Since the S&P 500 is a widely traded asset and can be replicated with positions in the underlying stocks, arbitrage possibilities are likely to keep the noise level low.

The data are from January 1987 and were obtained from the Chicago Mercantile Exchange, which monitors the S&P 500 for the trading of index futures. There are 19,027 observations in this trading month, comparable to the sample sizes of experimental data used in the physical sciences.

Previous studies employing relatively low-frequency economic data, such as those of Brock and Sayers (1988) and Frank and Stengos (1988), have only been able to analyze data sets of 200 observations because they are limited to post-World War II quarterly observations. In a previous analysis of the S&P 500, Scheinkman and LaBaron (1989) examined a daily data set of approximately 5,000 observations. The actual dimension calculations, however, are computed for weekly return series of one-fifth that size. Frank and Stengos (1989) looked at a precious-metal series at a daily frequency but have only 12 years of data. In contrast, our real-time data set provides a virtually limitless number of time series observations reflecting market conditions.

We log-difference the original data to create a series of real-time returns. In Table 1, we show descriptive statistics for the original and log-differenced series. References to the S&P 500 are to this transformed series. There are significant departures from normality in the third and fourth cumulants. Although excess kurtosis is a common feature of high-frequency asset returns, the skewness seems to be a unique aspect of the real-time data.

Following Brock (1986), we filter the log-differenced series with linear ARMA models. On the basis of the Akaike information criterion, an ARMA (12, 0, 0) model is used to filter the data. Since the stock exchange does not trade continuously, dummy variables for the first and last hours of the day are also included to account for nontrading effects.

Our next concern is dependence in the second moments, particularly because of the excess kurtosis. The squared residuals from the ARMA model are tested, and evidence of Engle's (1982) autoregressive conditional heteroscedasticity is detected. A GARCH (1, 1) model is then fit by maximum likelihood, and the data are filtered again. Since we are taking a smooth transformation of the data, the Brock residual theorem applies to this series as well. Under the GARCH null, the data will be a random process after filtering. If we denote the conditional mean as $\bar{p}_t$ and the conditional variance as $\bar{\sigma}_t$, the standardized residuals,

$$v_t = \frac{(p_t - \bar{p}_t)}{\sqrt{\bar{\sigma}_t}}, \tag{9}$$

are distributed $N(0, 1)$.

After prewhitening the data, we have three time series: (1) SP500, the log difference of the original S&P 500 cash index; (2) ARMA, the residuals from passing SP500 through an ARMA (12, 0, 0) filter; and (3) GARCH, the standardized ARMA residuals.

Table 1. Descriptive Statistics: S&P 500 Index

Statistic	Levels	Log-differences
Mean	263.35	.65257*E*-05
Standard deviation	8.8518	.13975*E*-03
Skewness	−.34592	−3.0980
Kurtosis	−.82171	192.47
Minimum	242.22	−.57800*E*-02
Maximum	280.96	.25684*E*-02

Table 2. Correlation Exponent Estimates

	S&P500		ARMA		GARCH		DELAY 16	
Embedding	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)
2	1.39	1.42	1.70	1.76	1.76	1.71	1.41	1.42
3	2.33	2.49	2.39	2.55	2.45	2.48	2.40	2.49
4	2.95	3.32	2.94	3.40	3.02	3.32	3.29	3.29
5	3.37	4.16	3.38	4.16	3.51	4.02	3.98	4.15
6	3.53	4.91	3.60	4.93	3.88	4.78	4.74	4.89
7	3.53	5.67	3.63	5.82	3.95	5.72	5.43	5.64
8	3.45	6.34	3.53	6.67	3.98	6.38	6.30	6.32
9	3.37	7.25	3.44	7.38	3.88	7.15	6.91	7.25
10	3.21	7.78	3.26	8.40	3.81	8.10	7.69	7.74
15	2.83	10.77	2.85	10.89	3.17	11.44	10.33	9.90
20	2.59	12.30	2.60	12.73	3.12	14.09	11.80	12.46

NOTE: Columns (1) report estimates for original series and columns (2) report estimates for the shuffled series.

3.3 Dimension Estimates

We calculate estimates of the correlation exponent over the range of embedding dimensions $M = 2, 3, \ldots, 10, 15, 20$. Theoretically, dimension estimates are made as $\varepsilon \to 0$. In practice, estimates are made over a range of values. The smallest value of ε is determined by the precision of the raw data; the original S&P 500 cash index is reported in dollars and cents. Consequently, the smallest nonzero change in the index that can be recorded is .01. This determines the smallest nonzero distance between any two m vectors. For the SP500 time series, the lower bound of the meaningful portion of the correlation integral is $\ln_2 \bar{\varepsilon} = -14.7$. We set an upper bound for ε such that 50% of the calculated norms are eliminated. This rule for estimating the correlation dimension roughly coincides with estimating the slope of the steepest segment of the correlation integral as identified by the nonparametric procedure developed by Mayfield and Mizrach (1991).

Results are reported in Table 2. Columns (1) report estimates for the ordered time series, and columns (2) report estimates for the shuffled series. Using a uniform pseudorandom number generator, each shuffled series is constructed by random draws without replacement from the associated original series. For the three series SP500, ARMA, and GARCH, comparison of adjacent embedding spaces indicates that marginal increases in embedding beyond 6 and 8 fail to reveal additional structure in the attractor. By comparing dimension estimates for the original series to those for the corresponding shuffled series, it is clear that there is low-dimensional structure present in the data.

In summary, across all three times series, dimension estimates become invariant to embedding and plateau at about 3.5 to 4.0. These estimates are in striking contrast to those based on the shuffled series. For each of the shuffled series, estimates continue to increase with embedding. It is clear that, by shuffling the data, the resulting series are much more space-filling than the original ones. In Figures 1, 2, and 3, we compare dimension estimates of the ordered series with those of the shuffled series. All three series pass the shuffle diagnostic.

In addition, by comparing the largest estimated dimension for each series (SP500: 3.53, ARMA: 3.63, GARCH: 3.98) it is clear that the low-dimension estimates reported in Table 2 are not due to the effects of a near unit-root process or GARCH process. In addition, these estimates are well below the estimates of 6 to 7 found by Scheinkman and LaBaron (1989) using daily closes.

3.4 Entropy Calculations

The results in Section 3.3 indicate that a low-dimensional process for stock-market prices exists. Furthermore, our evidence indicates that the nonlinearities are beyond GARCH. These estimates do not, by themselves, however, indicate that the underlying attractor is chaotic; the attractor must also be shown to have positive entropy.

Under the premise that the reconstruction is correct, we calculate a lower-bound estimate of its entropy based on the procedure described by Grassberger and Pro-

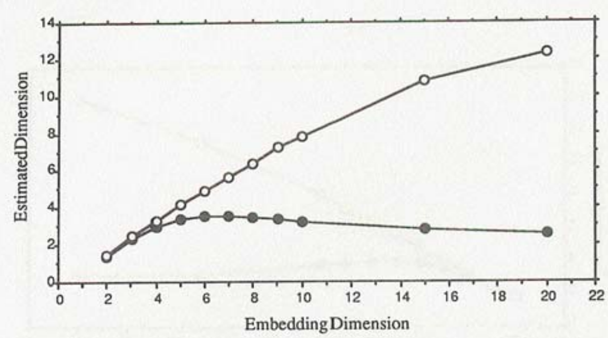

Figure 1. Dimension Estimates SP500: Ordered Versus Shuffled Data. This figure graphs the estimated GP correlation dimension for a range of embedding dimensions. The dark points (●) represent estimates for the ordered data and plateau around 3.5 as the embedding dimension rises. The dimension estimates for the shuffled data, as represented by the light points (○), continue to rise with embedding dimension.

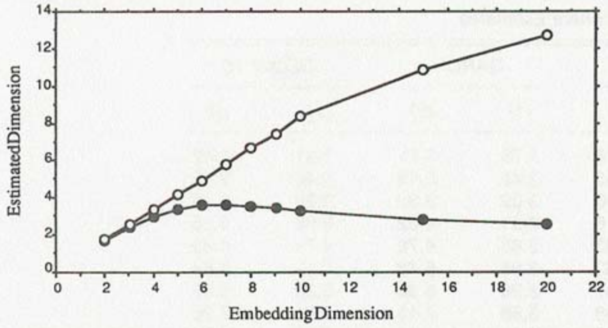

Figure 2. Dimension Estimates After ARMA Filtering: Ordered Versus Shuffled Data. The estimated correlation dimensions are graphed against their associated embedding dimensions. Using the ordered data, the estimates approach 3.5, while the estimates based on shuffled data continue to increase with embedding: ●, ordered data; ○, shuffled data.

caccia (1983b, 1984) and described in Section 2.2. Table 3 displays the vertical change in the correlation integral and the implied entropy estimate per one minute for the series. For small ε, the entropy estimate is proportional to the negative of the change in the intercept term as the embedding dimension increases. Figure 4 shows a graph of K_2 versus embedding dimension.

Based on these calculations, the K_2 lower-bound estimate of Kolmogorov entropy is approximately one-third of a bit of information per minute. Since we reconstruct the attractor from the difference in the S&P 500 cash index, which is reported in dollars and cents, our accuracy of the current state of the system is only 1 part in 50. Thus we have no more than six bits of information on the current state of the system. Given an observation on the current state of the system and knowledge of the true underlying system, an investor would have no knowledge of the system's state after 15 to 18 minutes. We define this duration as the implied forecast horizon and use it to support our choice of delay time in Section 4.

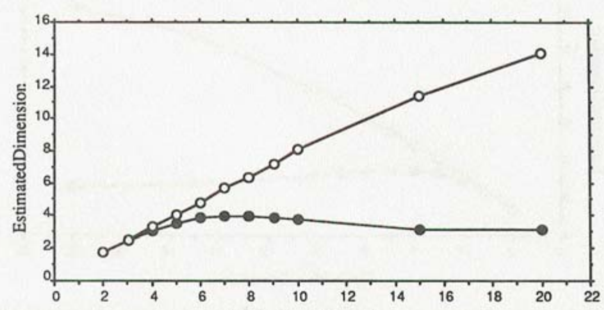

Figure 3. Dimension Estimates of Standardized GARCH Residuals: Ordered Versus Shuffled Data. After applying the GARCH filter, estimated correlation exponents peak at slightly less than 4, while the estimates using shuffled data continue to be space filling: ●, ordered data; ○, shuffled data.

Table 3. Entropy Calculations

Embedding dimension	*Δ intercept*	*K_2 (bits/min.)*
3	−1.884	5.652
4	−1.688	5.064
5	−1.256	3.768
6	−.687	2.061
7	−.319	.957
8	−.177	.531
9	−.118	.354
10	−.095	.285
15	−.486	.291
20	−.579	.348

On finding positive entropy, a researcher might conclude that the stock-return series is chaotic; however, if filtering has not removed temporal dependence in the data, this conclusion will be incorrect. In Section 5 of the article, we address the issue of time delays. If the reconstruction is correct, dimension calculations will be robust to delays within the region of sensitive dependence implied by our entropy calculations. In fact, we show that our dimension estimates are not invariant to delay time.

4. HIGHER ORDER TEMPORAL DEPENDENCE

In this section, we motivate why the data may still be temporally dependent after filtering. Assume for expository purposes that F in (1) is a scalar analytic function. Consider a series expansion of F around 0.

$$F(x_t) = \sum_{j=0}^{\infty} c_j x_t^j. \tag{10}$$

Suppose that one fits an AR(1) model to the data-generating mechanism (10). Define the residuals

$$\xi_{t+1} = x_{t+1} - \hat{\beta} x_t. \tag{11}$$

White (1980) showed that $\hat{\beta}$ will not coincide with the c_j's of the power series expansion if one estimates β

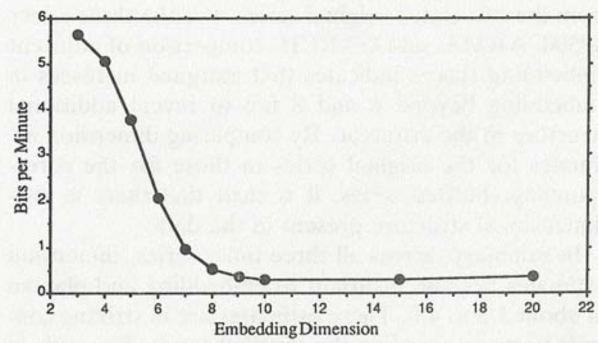

Figure 4. K_2 Lower Bound Entropy Estimate (bits per minute). This figure plots the vertical shift in the invariant segment of the correlation integral against the corresponding embedding dimension. For large embeddings, this is a lower-bound estimate of Kolmogorov entropy. At an embedding dimension of 10, the values plateau at approximately one-third of a bit of information per minute.

using ordinary least squares. More important for our purposes is that the residuals will be uncorrelated with lagged x's, but *not* independent; that is, $E[\xi_{t+1}|x_t] = 0$, but $E[\xi_{t+1}|x_t^2] \neq 0$. This result will obtain for any filter of degree k. Define the kth order residuals,

$$\xi_{t+1}^k = x_{t+1} - \sum_{j=0}^{k} \hat{\beta}_j x_t^j. \quad (12)$$

These residuals will still be correlated with the $k + 1$th power of x, $E[\xi_{t+1}^k|x_t^{k+1}] \neq 0$.

The data seem to show precisely this type of high-order dependence. Although GARCH filtering may account for the excess kurtosis, it is not equipped to remove the skewness. Only by delaying the data can we be sure that some higher order dependence does not remain.

5. THE METHOD OF TIME DELAYS

Takens's theorem allows us, in principle, to choose m histories with any delay time τ in reconstructing the attractor:

$$p_j^m = (p_j, p_{j+\tau}, p_{j+2\tau}, \ldots, p_{j+(m-1)\tau}). \quad (13)$$

In applied work, however, the quality of the reconstruction depends critically on the choice of τ. Empirically, it is desirable to reconstruct with data that are widely dispersed on the attractor. Filtered data, however, are no more likely to satisfy this criteria than the unfiltered data. In short, ARMA and GARCH filtering will rule out only very specific types of temporal dependence.

To ensure a wide dispersion of points on the attractor, the delay time τ is chosen so as to minimize the information on p_t contained in $p_{t-\tau}$. Holzfuss and Mayer-Kress (1986) noted that, in highly periodic data, the first zero-crossing of the autocorrelation function corresponds to the first minimum of mutual information function. In the log-differenced data, the first zero-crossing point of the autocorrelation function occurs at the 16th lag.

Using a delay of 16 observations between coordinates, we recalculate our dimension estimates. This delay corresponds to 5.33 minutes, which is well within our implied forecast horizon of 15 to 18 minutes. The fourth column of Table 2 reports these estimates as DELAY16. With this reconstruction, our dimension estimates are no longer invariant to embedding. In addition, by comparing the estimates to those from the corresponding shuffled series, as in Table 2, it is clear that the data are as space-filling as a stochastic process. With a conservative delay, the apparent structure of the system is removed.

The dramatic rise in dimension estimates is evidence of dependence in higher moments. It is not possible, however, to uniquely identify the source. With data sampled at such high frequency, nontraded stocks might be the cause. If a security does not trade within a given interval, then its most recent value is used in the calculation of the index. Thus nontraded securities could be a source of spurious dependence.

Two empirical observations lead us to the conclusion that nontraded securities are not the fundamental cause of the biased dimension estimates. First, the filters include opening and closing dummies for the first and last hours of the day. This should remove some nontrading effects, yet dimension estimates remain essentially unchanged using the filtered data. Second, based on the Wood transactions data, we compute a rough gauge on how often the stocks in the S&P500 trade. During January 1987, the lowest volume New York Stock Exchange stock in our sample is Brown and Sharpe (B&S). From a random selection of days, we estimate B&S has 35 quotes and trades per day, on average. At the other extreme, Ford Motor Company is a representative high-volume stock, with an average of 1,320 trades and quotes per day. For Ford, information arrives every 19.1 seconds, while for B&S it takes 722.9 seconds. With no delay, at an embedding dimension of 10, each m vector spans approximately 199.8 seconds (3.33 minutes). By interpolating between the rate at which Ford and B&S trade, we estimate that new information arrives on at least 371 of the S&P 500 stocks ($\approx$75%) within the time spanned by a given m vector.

Given that large stocks trade more frequently, over 75% of the S&P's market value is updated in this time interval. As embedding increases from 3 to 10 (or even 20), however, the dimension estimates remain very low, even though most stocks in the index have traded. Consequently, we do not believe that the dramatic change in dimension estimates using the method of delays is due to nontrading effects.

These results demonstrate that, with real-time data, filtering may not remove temporal dependence. Since filtering techniques remove only very specific types of nonlinear dependence, there is great potential for incorrectly detecting low dimension in real-time data. For this reason, we argue that the method of time delays is the appropriate reconstruction technique.

6. CONCLUSION

To gain deeper insight into the determination of stock-market prices, we apply nonlinear analysis to real-time data on the S&P 500 cash index. The use of high-frequency data enables us to examine the precise evolution of the market. Using the correlation integral, we find evidence of a low-dimensional attractor with positive entropy. These calculations are robust to ARMA and GARCH filtering but are not robust to changes in delay time.

The analysis of real-time data requires special considerations. Filtering is unable to remove temporal dependence in the stock-market data. In this instance, the choice of delay time is crucial for the proper reconstruction of the attractor; a delay long enough to eliminate the possible stretching of the attractor must be chosen. We choose a conservative delay of five minutes. With this delay time, the apparent structure of the system is removed.

Our inference concerning these results is, unfortunately, only heuristic. As Barnett and Hinich (1991) noted, no formal asymptotic theory for estimates of the correlation dimension and Kolgomorov entropy exists under the null hypothesis of chaos. Until such a theory is devised, we can only buttress evidence for nonlinear structure using existing statistical tools. Barnett and Hinich, for example, used the bispectrum.

We conclude that either the underlying system is of very high dimension or, if the true system is, in fact, of low dimension, its entropy is so high that it cannot be predicted beyond five minutes. From the standpoint of financial-market participants, these conclusions are essentially equivalent. Accepting the first conclusion, that there is a chaotic attractor, does not make the data in any sense more predictable. For very near-term forecasts, accurate prediction of the system's future state is possible; however, entropy causes these forecasts to deteriorate quickly. Our findings indicate that stock prices, even accounting for the nonlinearities, cannot be predicted over horizons of significant length.

ACKNOWLEDGMENTS

We thank Lisa Calise for her research assistance. The IBM Cambridge Research Facility generously provided time on their vector processor. Ken Kroner supplied the GARCH program used in Section 3. William Brock made helpful comments on an earlier draft that was presented at the International Symposium on Evolutionary Dynamics and Nonlinear Economics in Austin, Texas, April 1989. We also thank two anonymous referees and an associate editor for their comments. We retain responsibility for any remaining errors.

[*Received August 1990. Revised February 1992.*]

REFERENCES

Barnett, W. A., and Chen, P. (1988), "The Aggregation-Theoretic Monetary Aggregates Are Chaotic And Have Strange Attractors: An Econometric Application of Mathematical Chaos," in *Dynamic Econometric Modeling*, eds. W. Barnett, E. Berndt, and H. White, Cambridge, U.K.: Cambridge University Press, pp. 199–246.

Barnett, W. A., and Hinich, M. J. (1991), "Empirical Chaotic Dynamics in Economics," working paper, Washington University and University of Texas at Austin, Dept. of Economics, submitted to *Annals of Operations Research*.

Brock, W. A. (1986), "Distinguished Random and Deterministic Systems: An Expanded Version," *Journal of Economic Theory*, 90, 168–195.

Brock, W. A., and Sayers, C. (1988), "Is The Business Cycle Characterized by Deterministic Chaos?" *Journal of Monetary Economics*, 22, 71–90.

Eckmann, J. P., Kamphurst, S. O., Ruelle, D., and Scheinkman, J. (1988), "Lyapunov Exponents for Stock Returns," in *The Economy as an Evolving Complex System*, SFI Studies in the Sciences of Complexity, Reading, MA: Addison-Wesley, pp. 301–304.

Engle, R. (1982), "Autoregressive Conditional Heteroscedasticity With Estimates of the Variance of U.K. Inflation," *Econometrica*, 50, 987–1008.

Frank, M., and Stengos, T. (1988), "Some Evidence Concerning Macroeconomic Chaos," *Journal of Monetary Economics*, 22, 423–438.

——— (1989), "Measuring the Strangeness of Gold and Silver Rates of Return," *Review of Economic Studies*, 56, 553–568.

Grassberger, P., and Procaccia, I. (1983a), "Measuring the Strangeness of Strange Attractors," *Physica*, 9D, 189–208.

——— (1983b), "Estimation of the Kolmogorov Entropy From a Chaotic Signal," *Physical Review A*, 28, 2591–2593.

——— (1984), "Dimensions and Entropies of Strange Attractors From a Fluctuating Dynamics Approach," *Physica*, 13 D, 34–54.

Guckenheimer, J., and Holmes, P. (1985), *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*, New York: Springer-Verlag.

Holzfuss, J., and Mayer-Kress, G. (1986), "An Approach to Error-Estimation in the Application of Dimension Algorithms," in *Dimensions and Entropies in Chaotic Systems*, ed. G. Mayer Kress, New York: Springer-Verlag, pp. 114–22.

Mayfield, E. S., and Mizrach, B. (1991), "Nonparametric Estimation of the Correlation Exponent," *Physical Review A*, 88, 5298–5301.

McCaffrey, D. F., Ellner, S., Gallant, A. R., and Nychka, D. (1991), "Estimating the Lyapunov Exponent of a Chaotic System With Nonparametric Regression," Working Paper 1977R, North Carolina State University, Institute of Statistics, submitted to *Journal of the American Statistical Association*.

Pesin, J. B. (1977), "Characteristic Liapunov Exponents and Smooth Ergodic Theory," *Russian Mathematical Survey*, 32, 55–114.

Ramsey, J. B., and Yuan, H. J. (1990), "Statistical Properties of Dimension Calculations Using Small Data Sets," *Nonlinearity*, 3, 155–176.

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

Takens, F. (1980), "Detecting Strange Attractors in Turbulence," in *Dynamical Systems and Turbulence* (Lecture Notes in Mathematics 898), eds. D. Rand and L. Young, Berlin: Springer-Verlag, pp. 366–382.

——— (1983), "Distinguishing Deterministic and Random Systems," in *Nonlinear Dynamics and Turbulence*, eds. G. Borenblatt, G. Iooss, and D. Joseph, Boston: Pittman, pp. 315–333.

White, H. (1980), "Using Least Squares to Approximate Unknown Regression Functions," *International Economic Review*, 21, 149–170.

Wolf, A., Swift, J., Swinney, H., and Vastano, J. (1985), "Determining Lyapunov Exponents From a Time Series," *Physica*, 16 D, 285–317.