Available online at www.sciencedirect.com
SCIENCE DIRECT®
ELSEVIER
Pacific-Basin Finance Journal 11 (2003) 175–195
PACIFIC-BASIN FINANCE JOURNAL
www.elsevier.com/locate/econbase

The cross-sectional and cross-temporal universality of nonlinear serial dependencies: Evidence from world stock indices and the Taiwan Stock Exchange

Peter A. Ammermann[a,*], Douglas M. Patterson[b]

[a] *Department of Finance, Real Estate, and Law, College of Business Administration, California State University, 1250 Bellflower Boulevard, Long Beach, CA 90840-8505, USA*

[b] *Department of Finance, Insurance, and Law, Pamplin College of Business, Virginia Polytechnic Institute and State University, 1016 Pamplin Hall, Blacksburg, VA 24061-0221, USA*

Received 20 August 2001; accepted 5 June 2002

Abstract

This paper studies a variety of world financial market indices to determine how widespread the phenomenon of nonlinear serial dependency is, and then, by studying a relatively financially isolated market, the Taiwan Stock Exchange of the 1980s, examines more closely the extent to which nonlinearity appears to be an inherent feature of financial trading behavior. Nonlinearity is found to be a cross-sectionally universal phenomenon, existing within all the markets studied and within the vast majority of individual stocks traded on the Taiwan Stock Exchange. However, closer examination of the nonlinearity via a windowed testing procedure reveals that such dependencies do not appear to be cross-temporally universal; rather, the data seem to be characterized by relatively few brief episodes of extremely strong dependencies that are followed by longer stretches of relatively quiet behavior. Thus, the modeling of the extant nonlinearity appears to be problematic at best.

JEL classification: G15; C49
Keywords: Nonlinearity; Taiwan Stock Exchange; Bispectrum; Windowed testing

1. Introduction

Over the past decade and a half, numerous studies have documented the existence of nonlinear serial dependence in financial markets (e.g., Hinich and Patterson, 1985; Hsieh,

* Corresponding author. Tel.: +1-562-985-7526; fax: +1-562-985-1754.
E-mail address: pammerma@csulb.edu (P.A. Ammermann).

0927-538X/03/$ - see front matter
doi:10.1016/S0927-538X(02)00113-0

1989; Scheinkman and LeBaron, 1989). Such findings have important implications throughout the field of finance. For finance practitioners, the existence of nonlinear serial dependence raises the possibility of predictability in financial returns, even in the absence of return autocorrelation, or, alternatively, an exposure to greater levels of risk than would be expected under a random walk. For finance academics and econometricians, these findings raise concerns about the statistical adequacy of statistical models used to examine financial time series (for a detailed exposition along these lines, see Spanos (1999); for a specific application involving the weak-form efficient markets hypothesis, see Spanos (1995)). Unfortunately, however, many of the previous studies in this area have focused on a relatively small number of financial time series, such as foreign exchange rates and the returns for a small number of NYSE common stocks, which are heavily and directly influenced by U.S. financial activity.

In addition to using data from a number of national stock exchanges to document the prevalence of nonlinear dynamics across national borders, this paper seeks to extend the existing literature by using the Taiwan Stock Exchange of the 1980s as a case study to examine the prevalence of nonlinearity within a given market. The Taiwan Stock Exchange of this era is chosen because it was the most heavily traded, but among the most financially isolated, of the increasingly important emerging markets. Thus, empirical tests of this market are less likely to be affected by problems associated with illiquidity, while any positive nonlinearity results are more likely to have arisen from within the Taiwan stock market itself, rather than from interactions with U.S. financial transactions. Furthermore, rather than using an open outcry system for executing trades such as the NYSE uses, the Taiwan Stock Exchange uses a computerized trade matching system (for more information on the Taiwan Stock Exchange and its trading system, see, Aaron (1990), Chou (1989), and Rhee and Wang (1997)). Thus, any significant nonlinearity results for this market would provide a stronger indication that nonlinearity is an inherent feature of financial time series and is a result neither of interactions with options or futures markets (Taiwan's stock market had neither) nor of market microstructure effects nor of some other anomalous feature of U.S. financial markets.

Finally, in addition to studying the pervasiveness of nonlinearity across a given stock market, this paper goes a step further to examine the persistence of nonlinear serial dependencies across time. A test developed by Hinich and Patterson (1996), which breaks the stock returns down into short "windows" of time, is used to determine the stability of both the linear and the nonlinear serial dependency structures for the return generating processes for the above financial time series. Similar to the results of Hinich and Patterson (1996), it is found that neither the linear nor the nonlinear dependencies are persistent across time. Rather, the overall results appear to be driven by the activity within a small number of subperiods during which serial dependencies are highly significant, while the remaining majority of subperiods exhibit no significant serial dependencies. In other words, most of the time the markets move along at a close approximation to a random walk, but occasionally, as during October 1987, they suddenly "fall out of bed" and become highly autocorrelated and/or nonlinear. For the Taiwan stock market data, such changes in the nonlinear dependency structures do not appear to have any clearcut driving factor. Some of the changes in the linear dependency structures, however, appear to be directly attributable to changes in the Exchange's daily price limits that were made during 1987 and 1988.

To reiterate, this paper will examine three questions. First, is nonlinearity an inherent feature of financial time series that is pervasive across all markets? Second, is it also pervasive across all the stocks within a given market? And third, given that nonlinearity does exist within all of these time series, does it appear to be stable and of constant strength across time? The first two questions deal with the existence of nonlinearity cross-sectionally, while the third deals with its existence cross-temporally.

2. Testing for nonlinearity

Any time series model that cannot be written in the form of a linear ARMA or ARIMA model, i.e., any type of model that exhibits some form of serial dependency other than simple correlation or autocorrelation, is, by definition, a nonlinear model. Thus, there is literally an infinite number of potential nonlinear models. Concomitantly, there is a wide variety of tests designed to detect nonlinearity, each designed to search for a different feature of nonlinearity (see e.g., Ashley and Patterson (2000) and Hsieh (1991) for more details).

The most popular test for nonlinearity, and among the most general, is the BDS test of Brock et al. (1987) (see Brock et al. (1991) for more detailed information about this test). This test has been used successfully by a number of researchers (see, e.g., Hsieh, 1989, 1991; Scheinkman and LeBaron, 1989) for the detection and analysis of nonlinearity within a variety of financial time series, including both stock market data and foreign exchange rate data. As these papers find, the BDS test is quite good at detecting nonlinear dependencies, in terms of being very sensitive to departures from the null hypothesis in various directions. However, the null hypothesis for the BDS test is that the observations are independent and identically distributed, so on the way to detecting nonlinearity, other possible departures from i.i.d., such as linear dependencies (i.e., autocorrelation) or nonstationarity, must first be eliminated as possible causes of significant test statistics. So, in terms of a test that could be applied to raw returns, the BDS test is too sensitive; it has power against too broad a range of alternative hypotheses.

In an effort to narrow down the range of possible alternatives and enable modeling of the existing nonlinearity, Hsieh (1989) divides the realm of nonlinear dependencies into two broad categories—additive nonlinearity and multiplicative nonlinearity. Additive nonlinearity, or nonlinearity-in-mean, enters a process through its expected value, so that each element in the sequence can be expressed as the *sum* of a zero-mean random element and a nonlinear function of past elements:

$$y_t = \varepsilon_t + f(y_{t-1}, \ldots, y_{t-k}, \varepsilon_{t-1}, \ldots, \varepsilon_{t-k}). \tag{1}$$

With multiplicative nonlinearity, or nonlinearity-in-variance, each element can be expressed as the *product* of a zero-mean random element and a nonlinear function of past elements, so that the nonlinearity affects the process through its variance:

$$y_t = \varepsilon_t \times f(y_{t-1}, \ldots, y_{t-k}, \varepsilon_{t-1}, \ldots, \varepsilon_{t-k}). \tag{2}$$

One of the simplest examples of a nonlinear model that would exhibit additive nonlinearity is the Nonlinear Moving Average model of Robinson (1977), given by

$$y_t = \varepsilon_t + \beta\varepsilon_{t-1}\varepsilon_{t-2}. \tag{3}$$

A similar, but more general, type of additively nonlinear model is the bilinear model (see e.g., Granger and Anderson, 1978; Subba Rao and Gabr, 1980). The unconditional mean for this type of model is equal to zero, and its realizations would not exhibit any autocorrelation. However, in contrast to a normal linear process, such a process would exhibit nonzero *bicorrelations*, where a bicorrelation, assuming $\varepsilon_t \sim \mathrm{IID}(\mu=0, \sigma^2=1)$, is defined as $E(y_t y_{t-m} y_{t-n})$, the third-order moment. The existence of nonzero bicorrelations could lead a process to be predictable, thereby explaining part of the interest in the use of nonlinear models in finance. One way to test for the existence of nonzero bicorrelations is through the use of a bispectrum test (see Hinich, 1982), as was used in Hinich and Patterson (1985) to detect the existence of nonlinearity in NYSE common stock returns.

However, the most commonly used nonlinear model in finance and economics exhibits nonlinearity in variance, rather than in mean. It is the GARCH(1,1) model of Bollerslev (1986):

$$y_t = \varepsilon_t \times h_t,$$

$$\varepsilon_t \sim \mathrm{NIID}(0,1),$$

$$h_t^2 = \omega + \alpha\varepsilon_{t-1}^2 + \beta h_{t-1}^2. \tag{4}$$

This is the most frequently seen application of GARCH-type models, which were first developed by Engle (1982) and for which many variations are described in Bollerslev et al. (1992). The type of nonlinear dependency seen in these models can be detected via the McLeod and Li (1983) test, which was originally developed as a diagnostic for fitted ARIMA models. McLeod and Li's $Q_{xx}(p)$ statistic is a modification of the Ljung–Box $Q_x(p)$-statistic, which, rather than testing for autocorrelation among the residuals from a fitted ARIMA model, is instead fitted to the *squared* residuals. Thus, the McLeod and Li test is designed to detect autocorrelation among the squared residuals, such as would be generated by a GARCH process.

However, such an interpretation of McLeod and Li test results brings up a potential problem with using Hsieh's dichotomy—the two types of nonlinearity can mimic each other to some extent (see, e.g., Weiss, 1986). For example, both additive and multiplicative nonlinear processes could generate squared returns that are correlated with their own lags. Thus, either type of process could trigger significant McLeod and Li test results. Moreover, some types of nonlinearity could even resemble nonstationary linear processes. Data generated by an ARCH(q) process, for example, are observationally equivalent to data generated by a time-varying-parameter MA(q) process (see, e.g., Bollerslev et al., 1992; Grillenzoni, 1993).

Hence, rather than try to distinguish between the two categories of nonlinear processes, this paper will instead focus on using a less sensitive and more narrowly focused test for nonlinearity, so that significant results would more likely be an indication of nonlinearity rather than either of some other type of dependency or of nonstationarity. The test chosen for this purpose is the bispectrum test of Hinich (1982). As found in Barnett et al. (1994), the bispectrum test is relatively insensitive to many possible forms of nonlinearity, including ARCH or GARCH. This is due to the fact that, as discussed below, bispectral tests are based on bicorrelations, or third-order moments, which are zero-valued for GARCH processes. Thus, asymptotically, the bispectrum test will have the proper size, even in the presence of GARCH effects. Consequently, any significant results with this test, as were obtained by Hinich and Patterson (1985), are a very strong indication of an underlying nonlinear serial dependency structure of a form that is likely to be more complex than simple ARCH or GARCH dependence. Moreover, as Ashley et al. (1986) show, the bispectrum test can be used to test for nonlinearity even in the presence of linear dependencies, with no loss of power.

As noted above, the bispectrum test is based on the fact that linear processes have no nonzero bicorrelations. Assuming that $\{y_t\}$ is a third-order stationary, zero mean time series, then its bispectrum, $B(f_1,f_2)$, is defined as the double Fourier transform of its bicorrelation function:

$$B_y(f_1,f_2) = \sum_{n=-\infty}^{\infty}\sum_{m=-\infty}^{\infty} E(y_t y_{t-m} y_{t-n})\mathrm{e}^{-2\pi i(f_1 m + f_2 n)}. \tag{5}$$

It is a complex valued, spatially periodic function with a principal domain in the triangular set:

$$\Omega = \left\{0 < f_1 < \frac{1}{2}, 0 < f_2 < f_1, 2f_1 + f_2 < 1\right\}. \tag{6}$$

Conversely, the bicorrelations can be expressed as the double inverse Fourier transform of the bispectrum:

$$E(y_t y_{t-m} y_{t-n}) = \int_{\Omega}\int B_y(f_1,f_2)\mathrm{e}^{2\pi i(f_1 m + f_2 n)}\mathrm{d}f_1\mathrm{d}f_2. \tag{7}$$

If $\{y_t\}$ is a stationary *linear* process, then it can be expressed as

$$y_t = \sum_{n=0}^{\infty} a_n \varepsilon_{t-n}, \tag{8}$$

where $\{\varepsilon_t\}$ is a purely random zero mean process and the a_n are constants. In this case, the bispectrum can be written as

$$B_y(f_1,f_2) = \mu_3 A(f_1)A(f_2)A^*(f_1+f_2), \tag{9}$$

where $\mu_3 = E(\varepsilon_t^3)$, $A(f) = \sum_{n=0}^{\infty} a_n \mathrm{e}^{-2\pi i f_n}$, and $A^*(f)$ is its complex conjugate. Since the spectrum of $\{y_t\}$ is

$$S_y(f) = \sigma_\varepsilon^2 |A(f)|^2, \tag{10}$$

then the "squared skewness function" for $\{y_t\}$ is

$$\Psi^2(f_1,f_2) \equiv \frac{|B_y(f_1,f_2)|^2}{S_y(f_1)S_y(f_2)S_y(f_1+f_2)} \equiv \frac{\mu_3^2}{\sigma_\varepsilon^6}, \tag{11}$$

for all f_1 and f_2 in Ω whenever $\{y_t\}$ is linear. Thus, a stationary linear process will have a constant valued skewness function. Furthermore, if the $\{\varepsilon_t\}$ are also normally distributed, then μ_3 will equal zero, and the constant value to which the skewness function will be equal is zero. In other words, the bispectrum test actually comprises two tests—a test of linearity and a test of Gaussianity or normality. The linearity test is a test of the flatness of the bispectrum, while the normality test is a test of both the flatness of the bispectrum and its equality to zero across all frequency pairs. (Note: a process can be linear but not normal, but if it is a normal process, then it cannot be nonlinear.)

Hinich and Patterson (1996) develop an alternative test statistic for nonlinearity that is a time-domain analog to the bispectrum test (i.e., it is the bicorrelations themselves, rather than the double Fourier frequency transformation of the bicorrelation function, which are directly estimated for the test statistic). Their test statistic, for which the null hypothesis is that the observations are i.i.d., is given as follows:

$$H_N = \frac{1}{L}\sum_{s=2}^{L}\sum_{r=1}^{s-1}(G^2(r,s)-1) \xrightarrow{D} N(0,1), \tag{12}$$

where

$$L = N^c,\ 0 < c < \frac{1}{2}, \tag{13}$$

$$G(r,s) = \frac{1}{\sqrt{(N-s)}}\sum_{k=1}^{N-s} u(t_k)u(t_k+r)u(t_k+s). \tag{14}$$

As with the bispectrum test, because this test is based on the bicorrelations of a process, it will have the proper size, asymptotically, even in the presence of GARCH effects.

Hinich and Patterson use this test in conjunction with a procedure of dividing the overall sample period into shorter windows of time to allow a closer examination of the precise time periods during which nonlinear (or linear) dependencies are occurring. Such a test procedure will also be utilized in a later section of this paper.

3. Nonlinearity and stock market indices

The first question this paper will examine using the types of tests described above is the question of the pervasiveness of nonlinearity across different stock markets. The data set that will be used to examine this question consists of the daily closing values for stock market indices from six different stock markets across the world: the Dow Jones Industrial Average (DJIA) from the New York Stock Exchange, the Taiwan Stock Exchange Weighted Stock Index (Taiex) from the Taiwan Stock Exchange, the Nikkei 225 Stock

Table 1
Summary statistics

	DJIA	Taiex	Nikkei	Hang Seng	STI	FT-30
No. of observations	2810	3142	2991	2750	2755	2800
Mean	0.00048	0.00066	0.00026	0.00056	0.00027	0.00052
Median	0.00052	0.00093	0.00061	0.00110	0.00039	0.00092
Standard deviation	0.01130	0.01920	0.01148	0.01861	0.01297	0.01070
Skewness	− 4.2018	− 0.33142	− 0.57600	− 5.2124	− 3.2892	− 1.0184
Excess kurtosis	99.538	2.1788	22.051	97.149	54.444	11.988
Bera–Jarque test	1,168,308	679	60,764	1,093,882	345,227	17,250
(p-value)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)

Composite Index from the Tokyo Stock Exchange, the Hang Seng Index from the Hong Kong Stock Exchange, the Singapore Straits Times Industrials Index from the Singapore Stock Exchange, and the London Stock Index from the London Stock Exchange. Observations were taken from January 1982 through February 1993 for each index, but due to differences in holidays across countries, as well as the existence of Saturday trading on the Taiwan and Tokyo Stock Exchanges, the total number of observations range from 2750 for the Hang Seng Index to 3142 for the Taiex. Rates of return are calculated by taking the logarithmic differences in closing values between trading days.

Table 1 provides summary statistics for the data. Notably, the three emerging markets, Taiwan, Hong Kong, and Singapore, exhibit the highest levels of volatility. Of these, Taiwan's is the most volatile market, despite the existence of daily price limits that are imposed on this market. Also, largely as a consequence of the events of October 1987, all of the indices exhibit leptokurtosis and left-skewness. Interestingly, despite its high volatility, the Taiex's price limits do appear to have given this market some insulation from the Crash of 1987, so that its returns exhibit the lowest levels of kurtosis and skewness among the six markets. Nonetheless, the Bera–Jarque skewness–kurtosis test for normality results in a rejection of normality for each of the six indices, including the Taiex.

Table 2 contains the results of the nonlinearity test statistics for the six markets. The first column shows the Ljung–Box Q-statistics for autocorrelation of up to six lags ($Q_x(6)$). These are highly significant for each of the six markets, with p-values of less than 0.002 in each case. Thus, linear dependencies play a statistically significant role in the return generating process for each of these indices. Column 2 contains the McLeod and Li

Table 2
Nonlinearity statistics

	Ljung–Box test $Q_x(6)$	McLeod and Li test $Q_{xx}(6)$	Bispectrum test (Z-statistics)[a]	
			Normality	Linearity
DJIA	20.44	124.03	22.43	20.03
Taiex	115.28	2938.25	62.28	22.67
Nikkei	49.26	337.00	57.04	18.47
Hang Seng	31.09	32.22	42.16	16.00
STI	99.49	870.34	44.78	14.92
FT-30	23.07	1100.86	11.63	9.21

[a] For the bispectrum test results, a Z-statistic greater than 3.49 implies a p-value less than 0.0002.

test statistics for autocorrelation among the squared index returns ($Q_{xx}(6)$). As can readily be seen, these are dramatically significant for each of the indices save the Hang Seng Index, and the concomitant p-value for even the Hang Seng is still less than 0.001. The high values of these test statistics relative to the $Q_x(6)$ test statistics gives a strong indication of the presence of nonlinear dependencies in the data.

This possibility is tested formally via the Hinich bispectrum test, for which the results are shown in the third and fourth columns of Table 2. Column 3 shows the Z-statistics from the test of normality. Consistent with the results from the skewness–kurtosis test, normality is strongly rejected for each of the indices. The results of the test of linearity are shown in Column 4. Given that a Z-statistic greater than 3.49 implies a p-value of less than 0.0002, linearity is very strongly rejected for each set of index returns. Thus, by implication, nonlinearity plays a very significant role in the return dynamics for each of the indices.

As these results indicate, nonlinearity is evident in the returns for stock market indices across the world. Thus, nonlinearity is widely prevalent across countries. The next question to ask is, given that nonlinear dependencies are evident in the returns for the market as a whole, do they play an important role for all of the securities that trade on that market, or is nonlinearity only important for a select few; i.e., how pervasive is nonlinearity across the stocks within a given stock market. To examine this question, the stocks trading on the Taiwan Stock Exchange (whose index, the Taiex, happens to exhibit the most significant levels of nonlinearity of all the indices studied) will be tested individually for the presence of nonlinear serial dependencies.

4. The Taiwan Stock Exchange

Among world stock markets, the Taiwan Stock Exchange (TSE) is relatively new, with formal operations having begun during 1962. As of the end of its first year of operations, shares of stock in 18 firms were listed on the TSE. By December 1992, the number of firms listed on the TSE had grown to 246. These are divided by Taiwan's Securities and Exchange Commission (S.E.C.) into three categories. "Category A" firms are the largest and most profitable companies with the most diverse shareholders, roughly analogous to NYSE stocks in the U.S. "Category B" firms are somewhat smaller, more analogous to AMEX stocks. The third category, "Stocks Requiring Full Delivery," includes the shares of financially distressed firms, including bankrupt firms and even firms that have no ongoing business and may have been liquidated. Traders in these stocks must make full advance payment with no margin lending and must take full delivery of the share certificates for these firms.

The TSE is the only exchange in Taiwan for trading securities. Unlike the NYSE, these trades do not take place according to an open outcry trading system on a central trading floor under the supervision of individual market makers; rather, trades are executed via a computerized trading system under which trades are matched electronically. The hours during which these trades can occur are also more limited than for the NYSE. From Monday through Friday, the exchange is only open for trading during the hours from 9:00 a.m. to 12:00 noon. Unlike most exchanges, the TSE is also open for trading on Saturday, but only for the two hours from 9:00 a.m. to 11:00 a.m.

Both stocks and bonds are traded on the Taiwan Stock Exchange, but, because the S.E.C. wanted to discourage speculation, there were no listed options or futures traded on the market during the period under study. Similarly, short selling of listed shares is prohibited. Furthermore, the S.E.C. has imposed daily price limits on stock price movements, although these are in the process of slowly being eased out. Before October 1987, this limit was 5%. The limit was 3% during the period from October 1987 through November 1988, after which it was restored to the 5% level. In October 1989, the limit was raised to 7%, from which level further increases are anticipated, until the limits are eventually eliminated.

In addition to the above restrictions on the trading activities of Taiwanese investors, there are strict controls on foreign investment in the TSE. The TSE was totally closed to foreign firms and individuals before December 1983. Since then, foreigners have been allowed to invest in Taiwan through one of four mutual funds, the Taiwan (ROC) Fund, the Formosa Fund, the Taipei Fund, and the Taiwan Fund. On January 1, 1991, the market was opened for more general foreign investment, although the government still sets fairly strict controls on the levels of foreign investment, as well as on the foreign exchange transactions involving the New Taiwan Dollar that would need to be made for foreign investors to make investments in the TSE or to withdraw their profits back to their home countries. Although these restrictions are slowly being relaxed, they are still an important consideration.

All of these factors combined to lead the Taiwan Stock Exchange to be among the most financially isolated stock exchanges in the world. Not only was it isolated internationally, but it was also relatively isolated within Taiwan itself, in terms of there having been no derivative securities contingent on the stocks of the Taiwan Stock Exchange. However, despite this relative isolation, the Taiwan Stock Exchange does not suffer from the problems of illiquidity that mark the stock exchanges of some other emerging markets. On the contrary, despite its short trading hours and relatively small number of issues trading, the Taiwan stock market was often among the top three stock exchanges in the world, by volume, during the late 1980s. During January 1990, in fact, the Taiwan Stock Exchange was the most heavily traded stock exchange in the world, with US$6 billion average daily trading volumes. As a result of this heavy trading activity, the Taiwan Stock Exchange should be more immune to empirical problems related to illiquidity, such as nonsynchronous trading, than would other markets that could be studied.

Thus, aside from the effects of the exchange's price limits, which turn out to be rather direct and obvious, there are relatively few complicating factors which could confuse the interpretation of empirical results from the Taiwan Stock Exchange. All of the above factors combine to make the Taiwan Stock Exchange the ideal market to study to examine the pervasiveness of nonlinear serial dependencies among financial time series. When related to the results of previous studies, significant nonlinearity results for a large proportion of the stocks trading on the Taiwan Stock Exchange would indicate that nonlinear serial dependencies are an inherent, fundamental feature of financial time series.

5. Nonlinearity and Taiwan stock returns

The data available for studying the existence of nonlinearity on the Taiwan Stock Exchange include the daily closing prices for the total of 258 common and preferred stocks that traded on the exchange at some time between January 1984 and December 1992. These data were obtained from the E.P.S. database that is maintained by the Department of Education of the Republic of China on Taiwan. For these individual stocks, the longest price series is 2581 observations and the shortest is 2 observations. However, due to testing constraints, no price series with fewer than 55 observations were actually tested for nonlinearity. Only about 20% of the stocks have observations over the full sample period from January 1984 through December 1992. However, the economic growth of Taiwan and the increase in stock market prices on the Taiwan Stock Exchange throughout the 1980s led many more firms to become listed at some point after the sample period began. From these price series, returns were calculated as

$$R_t = \ln\left(\frac{P_t + D_t}{P_{t-1}}\right), \tag{15}$$

where P_t is the closing price of the stock on day t, and D_t is the dollar value of a cash or stock dividend for which day t is the ex-dividend day.

Before proceeding to look at the nonlinear dynamics in the stock returns, one potential confounding effect is first examined—linear dependencies, specifically autocorrelation, within the data. An AR(1) model is fitted to each of the stock return series to remove the largest component of linear serial dependencies, then the Ljung–Box ($Q_x(5)$) and the McLeod and Li ($Q_{xx}(6)$) test statistics are calculated for the residuals of the AR(1) model. The results for these models, averaged across the stocks within each industry, are presented in Tables 3 and 4. It is notable that the AR(1) parameter estimates are quite large, with values above 0.2 not uncommon for individual stocks, and a few of the estimates are even above 0.3. These estimated parameters are significant for a large majority of the stocks. Furthermore, the Ljung–Box $Q_x(5)$ statistics for the AR(1) model residuals are also significant for a large number of the stocks, indicating that although the first-degree autocorrelations are significant for most stocks, this autocorrelation does not account for all of the linear dependencies for many of the stocks. Interestingly, a substantial minority of the stocks also exhibit significant autocorrelation at lag 3 (as is also seen in the Taiex index daily returns). The factors underlying these linear dependencies will be examined later.

In addition to significant linear dependencies, the McLeod and Li test statistics ($Q_{xx}(6)$) are also highly significant (p-value < 0.001) for most of the stocks, giving a preliminary indication that significant nonlinear dependencies do seem to be pervasive throughout the Taiwan Stock Exchange. However, there also seem to be some differences in significance levels according to the industry in which a company operates. For example, the $Q_{xx}(6)$ statistics are all significant at the 0.001 level in the cement industry (1100), but only 4 of the 16 Banking and Insurance Industry stocks (Industry Number 2800) exhibit $Q_{xx}(6)$ statistics that are significant at this level.

To obtain a more complete picture of the characteristics of the nonlinear dynamics underlying the Taiwan stock returns, the bispectrum test is next applied to the stock returns. A summary of the results for this test, again averaged across the stocks within each industry, is also shown in Table 4. Among the stocks studied, for the 171 Category A stocks, the results are similar to those obtained by Hinich and Patterson (1985) for NYSE

Table 3
Average results, by industry, of tests for linear serial dependencies

Industry no.	Industry name	No. of stocks in industry	Ave. no. of observations	Ave. Phi(1)[a]	Ave. t-ratio	% of Stocks with Sig. Phi(1) statistics[b]	Ave. $Q_x(5)$	Ave. p-value	% of Stocks with Sig. $Q_x(5)$ results[b]
1100	Cement	9	1962.78	0.107	4.950	88.9	19.890	0.055	77.8
1200	Food	25	1331.92	0.157	4.780	84.0	11.618	0.180	48.0
1300	Plastics	18	1496.78	0.138	5.147	83.3	13.224	0.181	55.6
1400	Textiles	42	1503.07	0.135	5.140	83.3	14.126	0.145	52.4
1500	Electrical Machinery	12	1270.67	0.152	3.958	75.0	14.928	0.218	66.7
1600	Electrical Appliances, Wire, and Cable	11	2003.09	0.143	6.133	100.0	17.715	0.043	63.6
1700	Chemicals	14	1885.57	0.151	6.546	100.0	18.493	0.048	78.6
1800	Glass	6	1236.33	0.176	5.025	50.0	10.898	0.194	50.0
1900	Pulp and Paper	10	1999.90	0.116	5.014	100.0	16.493	0.079	80.0
2000	Iron and Steel	20	853.25	0.175	4.481	80.0	10.457	0.176	40.0
2100	Rubber	8	1485.75	0.118	4.393	75.0	23.731	0.171	62.5
2200	Electronics	4	1345.75	0.063	2.573	75.0	15.690	0.210	50.0
2300	Automobile	26	882.00	0.138	3.781	76.9	10.238	0.264	38.5
2500	Construction	13	1531.62	0.160	6.006	84.6	14.952	0.079	76.9
2600	Shipping	8	1004.38	0.124	3.943	75.0	13.906	0.073	62.5
2700	Hospitality	6	1709.67	0.120	5.058	83.3	14.328	0.046	66.7
2800	Banking and Insurance	16	2028.25	0.115	5.448	87.5	19.174	0.118	75.0
2900	Department Stores	7	1637.86	0.159	6.124	100.0	17.099	0.063	71.4
Overall		255	1456.93	0.140	4.952	83.9	14.549	0.144	58.4
Max[c]		1	2580	0.536	18.570		50.950	0.992	
Median[c]		1	1262	0.138	4.720		12.900	0.024	
Min[c]		1	8	− 0.086	− 0.840		0.500	0.000	

[a] Phi(1) refers to the coefficient of first-order autocorrelation.

[b] The "% of Stocks with Sig..." figures provide the percentages of stocks within a given industry, or the percentage of stocks overall, whose relevant coefficients or test statistics, under the appropriate assumptions, would be considered significant at a 5% level.

[c] The Max, Median, and Min figures give the maximum, the median, and the minimum values, respectively, across all of the stocks considered, of the statistic examined in a given column. E.g., 0.138 is the median value, across all of the stocks examined, of Phi(1), the coefficient of first-order autocorrelation; 4.720 is the median value, across all of the stocks, of the t-statistics obtained from dividing Phi(1) by its respective standard error; note that 4.720 is *not* necessarily the t-statistic for the stock that had the median Phi(1) value of 0.138.

stocks. Of these 171 stocks, normality is rejected in all but one case, and linearity is rejected at a size of 0.05 for all but two of the stocks. The two exceptions are Chu Wa Mill (stock 1439, in the Textile industry) and Yieh Loong (stock 2014, of the Iron and Steel industry). For the former, normality is rejected, but no significant nonlinearity is detected; for the latter, neither normality nor linearity can be rejected. Of the 75 remaining (non-Category A) stocks that were tested via the bispectrum test, both normality and linearity are rejected, at a 0.05 level, for 60 of them.

Table 4
Average results, by industry, of McLeod and Li and bispectrum tests for nonlinear serial dependencies

Industry no.	McLeod and Li test results			Bispectrum test results					
	Average $Q_{xx}(6)$	Average p-value	% Sig. Stocks[a]	No. of stocks tested	Average no. of observations	Normality	% Sig. Stocks[a]	Linearity	% Sig. Stocks[a]
1100	885.123	0.000	100.0	9	1962.778	26.349	100.0	11.623	100.0
1200	446.647	0.089	88.0	23	1443.565	17.073	100.0	8.956	87.0
1300	466.007	0.109	72.2	16	1584.353	19.744	100.0	11.389	100.0
1400	452.625	0.140	78.6	41	1503.071	13.781	100.0	9.182	90.2
1500	428.552	0.141	83.3	11	1383.636	17.279	100.0	9.927	100.0
1600	555.568	0.000	100.0	11	2003.091	23.495	100.0	12.870	100.0
1700	685.589	0.055	92.9	14	1885.571	19.797	100.0	10.415	100.0
1800	462.730	0.000	100.0	4	1236.333	15.360	100.0	8.153	100.0
1900	553.320	0.100	90.0	10	1999.900	20.720	100.0	12.552	100.0
2000	264.555	0.111	75.0	17	995.941	9.851	94.1	8.064	82.4
2100	747.449	0.118	87.5	7	1691.000	17.971	100.0	11.219	100.0
2200	563.800	0.013	100.0	4	1345.750	12.973	100.0	9.095	100.0
2300	185.839	0.158	73.1	25	915.440	9.866	100.0	6.800	96.0
2500	244.822	0.208	76.9	12	1531.615	19.458	100.0	10.111	91.7
2600	160.679	0.015	75.0	8	1004.375	12.193	100.0	6.728	100.0
2700	621.265	0.026	83.3	6	1709.667	14.780	100.0	11.925	100.0
2800	128.711	0.286	43.8	16	2028.250	23.759	100.0	13.006	100.0
2900	456.226	0.000	100.0	7	1637.857	16.850	100.0	10.293	100.0
Overall	421.564	0.111	80.8	241	1508.768	16.658	99.6	9.846	95.0
Max[b]	2214.200	1.000			2580	47.63		26.08	
Median[b]	193.100	0.000			1304	13.45		9.72	
Min[b]	0.230	0.000			58	− 0.55		− 0.63	

[a] The "% Sig. Stocks" figures provide the percentages of stocks within a given industry, or the percentage of stocks overall, whose relevant test statistics, under the appropriate assumptions, would be considered significant at a 5% level.

[b] The Max, Median, and Min figures give the maximum, the median, and the minimum values, respectively, across all of the stocks considered, of the statistic examined in a given column. E.g., 2580 is the maximum number of observations of any stock to which the bispectrum test was applied; similarly, 47.63 is the highest normality test statistic, among all the stocks examined, that was obtained from the bispectrum test for normality; it cannot be inferred, however, that the stock whose normality test statistic was 47.63 had 2580 daily return observations.

The sensitivity of the nonlinearity results to the test threshold chosen is shown in Table 5. As is readily apparent from this table, nonlinearity clearly plays a role in the return dynamics for the vast majority of stocks trading on the Taiwan Stock Exchange, with the

Table 5
Test threshold sensitivity analysis: percentage of stocks with significant test results

Threshold	McLeod and Li test ($Q_{xx}(6)$)	Bispectrum linearity test
0.100	83.9%	93.9%
0.050	80.8%	93.1%
0.010	77.3%	90.2%
0.001	72.6%	87.4%
No. of stocks tested	255	246

choice of test threshold making little difference to this conclusion. This is even more evident with the bispectrum test results than with the McLeod and Li test results.

It is noteworthy that nonlinearity is detected by the bispectrum test for a number of stocks for which the $Q_{xx}(6)$ statistic does not indicate the presence of any significant nonlinear dependencies. In the banking industry, especially, linearity is rejected at a level of less than 0.005 for all 15 of the banks tested, while the $Q_{xx}(6)$ test statistic rejects linearity at a size of less than 0.005 for only four of the banks, and at a size of less than 0.05 for only three more, or a total of seven stocks.

The key result, however, is that for the more narrowly targeted bispectrum test. As stated above, all but 2 of the 171 Category A common stocks exhibit significant nonlinear dependencies, and the test results for most of these lead to a very strong rejection of both normality and linearity. Thus, these results provide strong evidence that the returns for the vast majority of stocks on the Taiwan Stock Exchange are generated by some nonlinear process that (as the results reported in Barnett et al. (1994) would imply) is more complicated than a simple ARCH or GARCH process. In this regard, Taiwan's stocks appear to be very similar to those of the U.S.

6. Are serial dependency structures stable and constant over time?

The above results provide strong evidence that nonlinearity is quite widespread across economic sectors. However, for this finding to have much benefit to investors, the existing nonlinear dependencies must also be persistent across time. Thus, the next important question to ask is, for a given data set that exhibits significant nonlinear serial dependencies, is the strength and form of these dependencies stable across time, or do these dependencies appear suddenly during some time frames only to disappear during others?

To examine this possibility, the test procedure developed by Hinich and Patterson (1996) is employed. Once a window length has been chosen for the number of observations in each individual subsample, the overall sample is then broken down into half-overlapping windows of the desired window size. For example, if a 60-day window length is chosen, then the first window will run from day 1 through day 60. Window two will run from day 31 through day 90. The third window will start on day 61 and end on day 120. Subsequent windows will follow in a similar manner until the end of the data series is reached.

After the data have been divided into overlapping windows in this manner, two test statistics are calculated for each window. The first statistic, the C or autocorrelation test statistic, is a variation of the Ljung–Box test statistic used for detecting autocorrelation, or linear dependencies. The second statistic, the H or bicorrelation test statistic developed by Hinich and Patterson (1996) and described in Section 2, examines the data in the window for bicorrelation, or nonlinear dependencies. The significance levels for each of these statistics are then plotted against the dates at which the windows are centered, thereby allowing the researcher to easily examine how the strength of linear and nonlinear dependencies varies over time.

One question that has been raised regarding this procedure concerns the use of overlapping rather than non-overlapping windows. However, the results using overlapping windows are reported in this paper, since they can easily be seen to subsume the results that would be obtained from using non-overlapping windows. Thus, if the results using non-overlapping windows wish to be known, the researcher can simply take the results for the overlapping windows, and then just disregard the results for every other window. In other words, rather than focus on all of the windows, the researcher would just look at the results for windows 1, 3, 5,... or else windows 2, 4, 6,... However, keeping the results for the overlapping windows could provide more complete information about the "fleetingness" of dependencies within the data, without facing as great a loss of power as from using more narrowly focused non-overlapping windows. For example, if one window exhibits significant serial dependencies (whether linear or nonlinear) but neither of the half-overlapping windows on either side of this window exhibits such dependencies, then this would provide stronger evidence that linear or nonlinear dependencies are strong but highly localized in time, rather than simply being realizations from a persistent, but low-level, dependency structure whose results are seen only occasionally due to the results of random selection.

6.1. Index results

For the index data, this test procedure was performed for three separate window lengths—250, 125, and 62 days. Also, the most extreme 1% of the observations were clipped from each data set before the test statistics were calculated, in an effort to ensure that the results are not being driven by outliers. A summary of the results from this procedure for the stock index data is shown in Table 6. The "# Sig." and "% Sig." columns show the number and percentage, respectively, of data windows that exhibit any significant serial dependencies, whether linear or nonlinear or both (where significance is defined by a test statistic p-value of less than or equal to 0.01).

One result is readily apparent from this table—as the number of windows is increased and the width of the individual windows is decreased, the proportion of windows exhibiting significant linear or nonlinear dependencies declines sharply. This result implies that the significant full sample results for autocorrelation and, especially, nonlinearity are being triggered by the activity within a few relatively short "pockets" of highly autocorrelated and/or nonlinear data. In other words, the serial dependency structures

Table 6
Stability results for stock index returns

Window length: 250 Days				125 Days			62 Days		
No. of windows	Total	# Sig.	% Sig.	Total	# Sig.	% Sig.	Total	# Sig.	% Sig.
DJIA	21	5	23.8	44	4	9.1	89	5	5.6
Taiex	24	9	37.5	49	10	20.4	100	10	10.0
Nikkei	22	9	40.9	46	10	21.7	95	11	11.6
Hang Seng	21	8	38.1	43	11	25.6	88	7	8.0
STI	21	15	71.4	43	19	44.2	88	16	18.2
FT-30	21	8	38.1	44	8	18.2	89	8	9.0
Average	21.7		41.6	44.8		23.2	91.5		10.4

The threshold used for determining test significance is 0.01; also, to reduce the influence of outliers on the test results, the most extreme 1% of the observations are deleted from the data set.

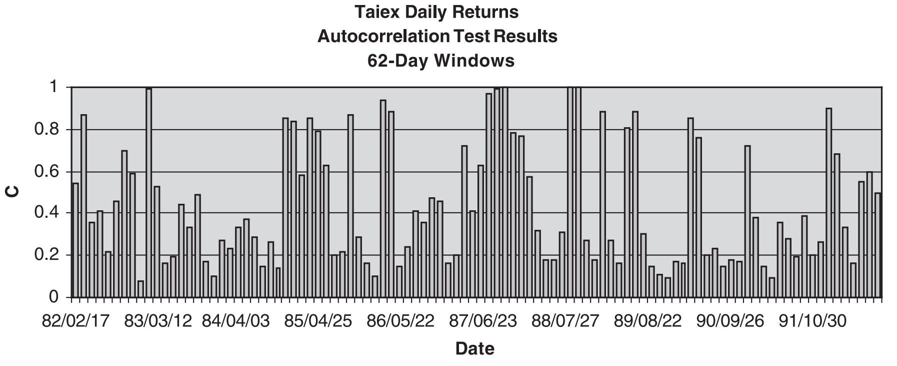

Fig. 1. Windowed serial dependency test results for daily index returns. This plot shows the results over time of the windowed test procedure for the Taiex daily returns for a 62-day window. The vertical axis shows the significance level, or one minus the p-value, for the serial dependency test statistic (e.g., if the significance level for a given window is above 0.95, then the serial dependency test statistic for that window has a p-value of less than 0.05 and would generally be considered to be significant). This figure gives the significance levels over time for the autocorrelation or C statistic.

for the financial time series examined are not stable, but rather vary over time, with the returns during most time periods rather closely approximating a random walk, while the returns during the relatively few remaining time periods are the only ones characterized by highly significant autocorrelation and/or nonlinear serial dependence. Moreover, except for one case, these periods of strong serial dependencies appear to be rather randomly distributed across the overall sample.

This is illustrated in Figs. 1 and 2, which show the results of this procedure for the Taiex daily returns for the most narrow window length, 62 days. The first plot, Fig. 1,

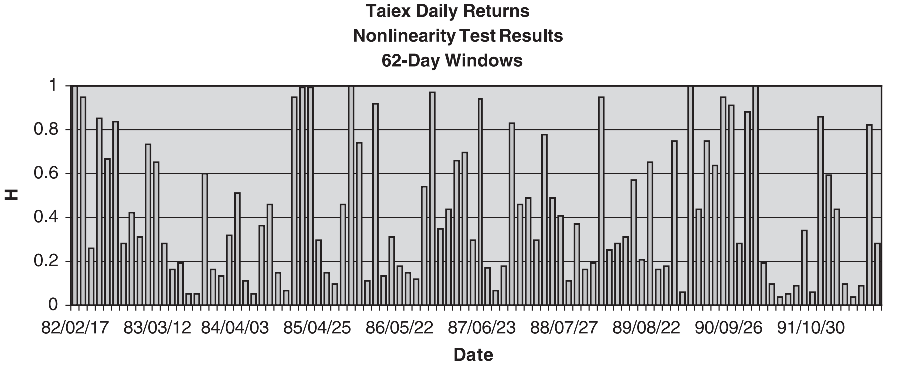

Fig. 2. Windowed serial dependency test results for daily index returns. This plot shows the results over time of the windowed test procedure for the Taiex daily returns for a 62-day window. The vertical axis shows the significance level, or one minus the p-value, for the serial dependency test statistic (e.g., if the significance level for a given window is above 0.95, then the serial dependency test statistic for that window has a p-value of less than 0.05 and would generally be considered to be significant). This figure shows the significance levels for the bicorrelation or H statistic, for nonlinearity.

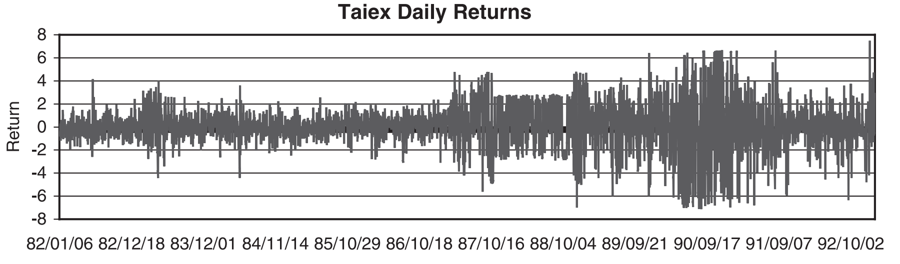

Fig. 3. Time plot of Taiex daily returns.

gives the significance levels over time for the autocorrelation or C statistic, while the second plot, Fig. 2, shows the nonlinearity significance levels from the bicorrelation or H statistic. For each plot, the vertical axis shows the significance level, or one minus the p-value, for the serial dependency test statistic. Thus, for example, if the significance level for a given window is above 0.95, then the serial dependency test statistic for that window has a p-value of less than 0.05 and would generally be considered to be significant. Across the horizontal axis for each plot are the dates at the centerpoints of each window.

One interesting result that can be seen in the figures that are shown concerns the distribution of the significant windows across time. For Fig. 2, the plot of the nonlinearity test significance levels over time, the windows with significant test results appear to be fairly randomly distributed over time, with no apparent pattern to their appearance or disappearance, which is a result that is typical for the indices studied. On the other hand, there does seem to be a pattern underlying the significant autocorrelation windows for the Taiex. In this case, some of the changes in autocorrelation levels appear to be driven by changes in the price limits that were imposed on this market. As can be seen in Fig. 1, the significant autocorrelations for this market are heavily clustered around the period of 1987 and 1988. This happens to be the time period during which Taiwan imposed its strictest price limits relative to contemporary levels of volatility. This can be seen in Figs. 3–5, which show the time series of Taiex returns, along with the concomitant levels of stock market variance as estimated via an AR(3)–GARCH(1,1)–t model that was fit to the Taiex returns. The price limits were set at 5% during the first part of this time series, until October 1987. During the initial portion of this interval, the

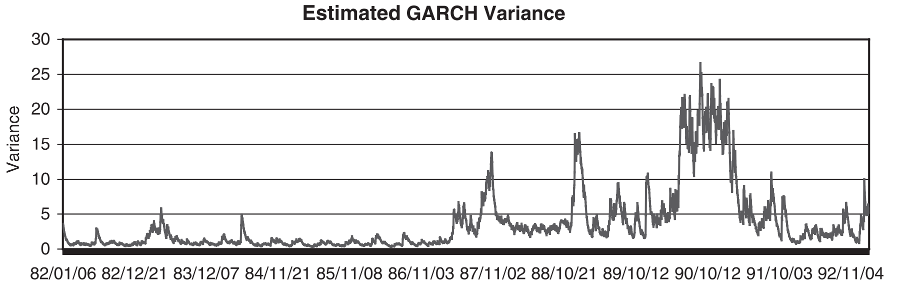

Fig. 4. Time plot of variances.

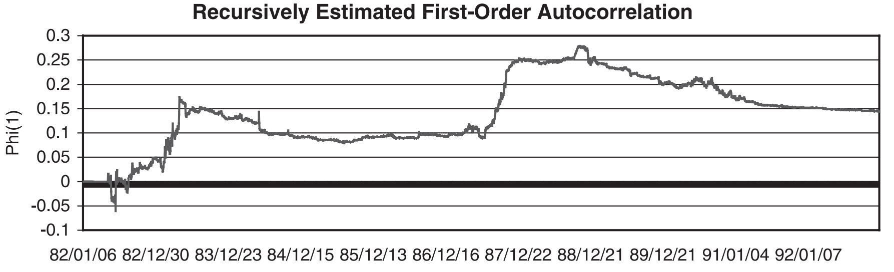

Fig. 5. Time plot of autocorrelations.

Taiex returns rarely hit against these price limits. However, starting around early 1987, the daily returns started to run into the price limits on a regular basis. To try to control this higher volatility, the Taiwan Stock Exchange tightened the price limits to 3%, where they remained for a year, during which time the Taiex's stock returns regularly ran up against the price limits, so that the outline of these limits can readily be seen from the plot of daily index returns.

Before the price limits were reduced, the volatility on the Taiwan Stock Exchange had steadily increased relative to the range of the price limits, as can be seen from the GARCH variance estimates. This variance spiraled upward until October 1987, after which the price limits were narrowed to 3% and the GARCH-estimated variance quickly fell to reflect this barrier. However, this did not represent a genuine reduction in market volatility. Rather, as Figs. 4 and 5 clearly illustrate, the underlying volatility still remained within the data, but its form was transferred by the price limits from variance within daily returns into autocovariance across daily returns. This is due to the fact that price movements that hit the price limits were forced to wait until the following day to continue.

Fig. 5 shows the time series of recursive least squares estimates for the first order autocorrelation within the Taiex returns. (This series is estimated by taking the first hundred observations for the Taiex daily returns, estimating the first-order autocorrelation among these daily returns, then adding the 101st observation, re-estimating the first-order autocorrelation among these 101 daily returns, then adding the 102nd observation, and so forth, until, for the last first-order autocorrelation estimate, the full sample of 3142 Taiex daily return observations is included in the estimation. Note that this means that each additional observation will have a much greater impact toward the beginning of the series than toward the end.) As is clearly visible, the levels of measured autocorrelation are relatively low (around 0.10) until the returns from the 1980s start to be included in the estimates. The estimated autocorrelation then quickly escalates to a level around 0.25 as the daily returns are limited to 3% moves. After the price limits are relaxed again to 5%, the level of autocorrelation starts to gradually decline. This decline is briefly reversed during January 1990, when daily returns were once again large enough to run into the 5% price limits on a regular basis, but afterward the decline in the recursive autocorrelation estimates continues unabated until the end of the sample period.

Thus, the linear serial dependencies for the Taiex returns are driven, at least in part, by the price limits that are imposed on the market, which cause the instances of significant autocorrelation to be heavily clustered. The windows containing significant nonlinear serial dependencies, however, do not seem to bear any relation to the changes in price limits and appear to be more randomly distributed across time.

6.2. Individual stock results

After testing the stock index data, the windowed test procedure was next repeated for the individual stocks trading on the Taiwan Stock Exchange. Given the large number of

Table 7
Average results, by industry, of tests for serial dependency stability

Industry no.	No. of stocks tested	Windowed stability test results		
		Ave. # of Windows[a]	Ave. # Sig. Win.[b]	Ave. % Sig. Win.[c]
1100	9	71.222	11.111	17.0
1200	24	52.043	8.000	17.0
1300	17	57.235	9.824	0.219
1400	42	54.190	8.286	0.144
1500	11	49.636	9.364	0.220
1600	11	72.636	8.818	11.4
1700	14	68.286	10.000	0.143
1800	6	44.333	7.000	41.9
1900	10	72.500	9.600	13.6
2000	17	35.412	5.529	14.2
2100	7	61.143	9.286	21.3
2200	4	48.500	6.500	13.4
2300	25	32.480	4.280	12.7
2500	13	55.231	11.154	29.1
2600	8	35.750	4.375	11.3
2700	6	61.833	9.000	13.1
2800	16	73.625	11.438	15.2
2900	7	59.286	8.714	15.7
Overall	247	55.402	8.321	16.8
Max[d]		94	36	100.0
Median[d]		47	7	13.8
Min[d]		1	0	0.0

[a] The "Ave. # of Windows" figures provide the averages, across all of the stocks within a given industry, or across all of the stocks overall, of the numbers of windows into which the stocks' sets of returns are subdivided.

[b] The "Ave. # Sig. Win." figures provide the averages, across all of the stocks within a given industry, or across all of the stocks overall, of the numbers of windows for the stocks' returns that exhibit significant serial dependencies, whether significant autocorrelation, significant bicorrelation, or both, where significance is determined by a p-value for the relevant test statistic or statistics of 1% or less.

[c] The "Ave. % Sig. Win." figures provide the averages, across all of the stocks within a given industry, or across all of the stocks overall, of the percentages of the windows into which the stocks' returns are subdivided that exhibit significant serial dependencies, whether significant autocorrelation, significant bicorrelation, or both, where significance is determined by a p-value for the relevant test statistic or statistics of 1% or less.

[d] The Max, Median, and Min figures give the maximum, the median, and the minimum values, respectively, across all of the stocks considered, of the statistic examined in a given column. E.g., one or more stocks had only one window's worth of data available, so the Min number of windows for the stocks is one (1); the minimum number of significant windows is zero (0), meaning that one or more stocks had returns for which none of the windows exhibited significant serial dependencies; it is not necessarily the case, however, that the stock with only one window's worth of data exhibited no significant serial dependencies.

stocks to examine, one standard window length was chosen that was applied to all of the stocks, regardless of the number of daily return observations available for a given stock. In order to more precisely pinpoint the exact time periods during which autocorrelation and/or nonlinearity enter into the return generating processes, a relatively narrow window length needed to be chosen. Somewhat arbitrarily, a window length of 54 days (9 trading weeks or approximately 2 months of returns) is chosen. This is equal to the square root of the longest stock return data series, rounded upward to the nearest number of full weeks.

A summary of the stability test results for the stocks trading on the Taiwan Stock Exchange, averaged across the stocks within each industry, is provided in Table 7. In general, the test results for the individual TSE stocks are similar to those for the daily Taiex index returns. There is a relatively small proportion (about 17%, on average) of significant windows of highly linear and/or nonlinear activity that are underlying the significant full-sample test results for these stocks. As was also seen with the Taiwan stock index data, moreover, changes in the return autocorrelation structure for many of the longer stock return time series appear to be driven by changes in the price limits imposed by the Taiwan Stock Exchange, so that autocorrelation is a significant factor in Taiwan stock returns almost exclusively during the 1987–1988 time period, when price limits were very strict relative to the volatility of the market. However, the appearance of significant nonlinear dependencies seems to be much more random, with no apparent connection to changes in price limits by the Taiwan Stock Exchange.

7. Conclusions

The results from this paper indicate that nonlinear serial dependencies do play a significant role in the returns for a broad range of financial time series, including returns from six different stock market indices from across the world, as well as the stock returns for the vast majority of stocks trading on the Taiwan Stock Exchange. Given the relative financial isolation of the Taiwan Stock Exchange, along with its operational differences from the New York Stock Exchange, the fact that nonlinearity is as important a factor in the returns for this market as for those of the New York Stock Exchange or the foreign currency exchanges indicates that nonlinear serial dependencies are a very fundamental aspect of financial time series.

Unfortunately, contrary to some initial hopes regarding such a finding, these nonlinear dependencies do not seem to be persistent enough to allow improvements over the assumption of a random walk for predicting securities' returns. Rather, these dependencies show up at random intervals for a brief period of time but then seem to disappear again before they can be exploited. At the same time, however, the nonlinear dependencies do seem to arise sufficiently frequently to prevent the random walk (i.e., Brownian motion) from being an adequate model for the purposes of risk assessment and management.

The results for linear dependencies for the most of the markets are the same as for the nonlinear dependencies, with windows of significant autocorrelation appearing only sporadically throughout the full sample. For the Taiwan stock market, some of the changes in the autocorrelation function of the return generating process appear to be driven by changes in the price limits imposed on the market. Variations in the nonlinear dependency structure for this market, on the other hand, do not appear to have any clear relation with the changes in the price limits.

This paper has not really addressed the question of statistical adequacy, and it is still conceivable that some stable nonlinear model could be developed that could adequately describe the nonlinear statistical dependencies in the data. However, the important question that this paper raises concerns the *economic* activity underlying these statistical results: What type of economic behavior would lead to long periods of time during which stock and index returns follow a random walk, interspersed with brief periods of highly autocorrelated and/or highly nonlinear activity? Given that such results have been found for markets across the world, this question would seem to be of fundamental importance to the field of finance.

Acknowledgements

The authors would like to thank Dr. Tai Ma of National Sun Yat-Sen University in Taiwan for data used in this study.

References

Aaron, B.D.L., 1999. Liberalizations in the Taiwan Stock Exchange. Unpublished thesis, University of Virginia.
Ashley, R.A., Patterson, D.M., 2000. A Nonlinear Time Series Workshop: A Toolkit for Detecting and Identifying Nonlinear Serial Dependence. Kluwer Academic Publishing, New York.
Ashley, R.A., Patterson, D.M., Hinich, M.J., 1986. A diagnostic test for nonlinear serial dependence in time series fitting errors. Journal of Time Series Analysis 7, 165–178.
Barnett, W.A., Gallant, A.R., Hinich, M.J., Jungeilges, J.A., Kaplan, D.T., Jensen, M.J., 1994. A single-blind controlled competition among tests of nonlinearity and chaos. Journal of Econometrics 82, 157–192.
Bollerslev, T., 1986. Generalized autoregressive conditional heteroskedasticity. Journal of Econometrics 31, 307–327.
Bollerslev, T., Chou, R.Y., Kroner, K.F., 1992. ARCH modeling in finance: A review of the theory and empirical evidence. Journal of Econometrics 52, 5–60.
Brock, W., Dechert, W.D., Scheinkman, J.A., LeBaron, B., 1987. A test for independence based on the correlation dimension. Econometric Reviews 15, 197–235.
Brock, W., Hsieh, D., LeBaron, B., 1991. Nonlinear Dynamics, Chaos, and Instability: Statistical Theory and Economic Evidence. MIT Press, Cambridge, MA.
Chou, S.-R., 1989. An empirical analysis of stock market anomalies: Evidence from the Republic of China in Taiwan. Unpublished dissertation, University of Kentucky.
Engle, R.F., 1982. Autoregressive conditional heteroskedasticity with estimates of the variance of U.K. inflation. Econometrica 50, 987–1008.
Granger, C., Anderson, A., 1978. An Introduction to Bilinear Time Series Models Vandenhoeck und Ruprecht, Gottingen.
Grillenzoni, C., 1993. ARIMA processes with ARIMA parameters. Journal of Business & Economic Statistics 11, 235–250.
Hinich, M.J., 1982. Testing for Gaussianity and linearity of a stationary time series. Journal of Time Series Analysis 3, 169–176.
Hinich, M.J., Patterson, D.M., 1985. Evidence of nonlinearity in daily stock returns. Journal of Business & Economic Statistics 3, 69–77.
Hinich, M.J., Patterson, D.M., 1996. Detecting epochs of transient dependence in intra day stock returns. Unpublished manuscript, Virginia Polytechnic Institute and State University.

Hsieh, D., 1989. Testing for nonlinearity in daily foreign exchange rate changes. Journal of Business 62, 339–368.

Hsieh, D., 1991. Chaos and nonlinear dynamics: Application to financial markets. Journal of Finance 46, 1839–1877.

McLeod, A.I., Li, W.K., 1983. Diagnostic checking ARMA time series models using squared-residual autocorrelations. Journal of Time Series Analysis 4, 269–273.

Robinson, P., 1977. The estimation of a non-linear moving average model. Stochastic Processes and their Applications 5, 81–90.

Rhee, S.G., Wang, C.-J., 1997. The bid–ask bounce effect and the spread size effect: Evidence from the Taiwan stock market. Pacific-Basin Finance Journal 5, 231–258.

Scheinkman, J., LeBaron, B., 1989. Nonlinear dynamics and stock returns. Journal of Business 62, 311–337.

Spanos, A., 1995. On theory testing in econometrics: Modeling with nonexperimental data. Journal of Econometrics 67, 189–226.

Spanos, A., 1999. Probability Theory and Statistical Interference: Econometric Modeling with Observational Data. Cambridge University Press, Cambridge.

Subba Rao, T., Gabr, M.M., 1980. An Introduction to Bispectral Analysis and Bilinear Time Series Models. Springer-Verlag, Berlin.

Weiss, A.A., 1986. ARCH and bilinear time series models: Comparison and combination. Journal of Business & Economic Statistics 4, 59–70.