Cyclical Behavior in Eighteenth Century Stock Market Prices

E. A. Olszewski[1]
Department of Physics
University of North Carolina at Wilmington
Wilmington, North Carolina 28403-3297
email: moof@pootie.phy.uncwil.edu

K. L. Ross[2]
European 1 Department
International Monetary Fund
Washington, D.C. 20431
email: kross@imf.org

[1]Supported by the North Carolina Supercomputing Program.
[2]Supported by a research grant from United Carolina Bank. The opinions expressed in the paper are those of the author and do not necessarily represent the views of the IMF.

Contents

List of Tables

List of Figures

Summary

We have analyzed a bivariate time series consisting of eighteenth century stock prices for the South Sea Trading Company listed on both the London and Amsterdam exchanges. The stock prices listed on the Amsterdam exchange are for forward settlement of three weeks, and thus more closely resemble futures than stock prices. Peters (1992) has described a technique for exposing long-term periodicities using the normalized rescaled range or V statistic. After adjusting the data for the introduction of Barnard's Act, which restricted forward settlement in the London market, we apply this technique to the London data using a V statistic modified by Lo (1989) to account for biasing which results from short memory effects in samples of finite size. We find strong evidence for a long-term periodicity of approximately 20 years.

In order to test more rigorously for this periodicity we have modeled the bivariate series using the seminonparametric (SNP) method of Gallant and Tauchen (1990). The model includes the effect of the cointegration between the price series, but not the apparent periodicity. The model indicates that the data are markedly conditionally heteroscedastic and that the conditional probability distribution of the data is extremely non-normal. In addition, the model prescribes a value for the truncation lag, a free parameter required for calculating the V statistic. Generating 1000 simulations of data from the model we have calculated the V statistic for each simulation and constructed fractiles for the statistic. Comparing the fractiles derived from the SNP model to those derived by Lo for a general class of models, we conclude that the fractiles of Lo are only slightly biased toward a type 1 error for this particular model. Specifically for the London data, the value of the calculated V statistic occurs with approximately 7% probability according to the SNP model and approximately 6% probability according to Lo's general class of models. In spite of this slight descrepency, we conclude that an approximate 20 year cycle has been present in the London data during the eighteenth century.

I Introduction

In an impressive study on the origins of financial capitalism, Larry Neal collected financial data comprising more than a century of stock prices for companies traded simultaneously in the London and Amsterdam markets [18]. These data were used to examine diverse questions ranging from how well capital markets were integrated during the eighteenth century to why the original financial bubbles such as the South Sea and Mississippi bubbles began and transpired as they did. Furthermore, these markets were analyzed over periods of war and peace as well as over the mass migrations of people that these wars precipitated. Applying modern day statistical and financial analysis to these historical data Neal found evidence of very efficient European capital markets which were well integrated and among which, information flowed relatively quickly. Tests for weak form market efficiency using Box-Jenkins ARIMA models were performed, which indicated that capital markets were efficient throughout the 18th century. These results are especially noteworthy since following such tumultuous periods one may have expected the evolution of inefficient, segmented markets. Furthermore, estimated stock betas, which reflect an individual stock's volatility as compared to a market index, were reasonable. In summary, these emerging capital markets appear to have been efficient and well integrated from their inception, not as a continuing function of time, which is in contrast to conventional stock market wisdom concerning emerging markets and their information processing capabilities. The purpose of this paper is to reexamine Neal's original cross listed stock market data using modern time series techniques, while focusing on a single listing, the South Sea Trading Company.[1] These techniques, which take account of cointegration, fractional integration, and various types of nonlinearities such as conditional heterogeneity, allow one to construct better fitting models as well as to detect long-term cyclical patterns.

The first question addressed in the paper is whether stock prices listed on both the London and Amsterdam exchanges are cointegrated.[2] Cointegration rests on the premise that a long run equilibrium relationship exists between two time series. Specifically, there exists a contemporaneous linear combination of the two series which is stationary. If two cointegrated series deviate excessively from each other they must at some time in the future approach each other in order to restore equilibrium.

A second question which has recently been addressed is whether modern financial

[1]Companies traded on the Amsterdam exchange have been priced for forward settlement of three weeks. Thus, their prices more closely resemble futures rather than cash prices. Consequently, a reexamination of Neal's findings may be relevant to modern futures markets.

[2]This is akin to whether modern day futures and cash prices are cointegrated. In the case of futures and cash prices the situation is complicated because settlement of the futures contract occurs on a specific day. Thus, the cointegration will be concealed by a time varying cost of carry. In this case, however, the situation is less complicated since the settlement for a stock listed on the Amsterdam exchange is always three weeks in the future. This is a relatively short time so that if the price series are cointegrated it may be readily apparent.

data contain long-term cycles. The conclusion based on spectral analysis is that, in general, they do not. Recently, a technique based on the rescaled range analysis of Hurst has been used to test for long-term memory (fractional integration) and cycles in financial data. Will the application of these techniques uncover cycles in these old data? Clearly, both of these questions relate to important issues about market behavior including market efficiency. The rest of the paper is organized as follows. In Section II we describe the bivariate data series of stock prices to be analyzed and the methodology used in the paper. In Section III we test for stationarity in the data. Finding each data series to be integrated we test for cointegration. Then, using spectral analysis we search for cyclical behavior in the stock prices on the London exchange. Finding no cycles we test for long-term periodicities using rescaled range analysis. We next construct a bivariate model. Because the data exhibit anomalous behavior during the introduction of Barndard's Act from August 31, 1733 to February 1, 1734, we divide the data into three intervals: prior to, during, and after this period. We test for systematic differences between the three periods and adjust the data accordingly. Using the seminonparametric methodology of Gallant and Tauchen we construct a model which takes account of the cointegration between the data series, but does not take account of any long-term cyclical behavior. We test the adequacy of the model using both short-term and long-term diagnostics. In Section IV, we generate 1000 simulations of the original data. We apply the rescaled range analysis to the simulations in order to construct confidence bands for the rescaled range test. In Section V we summarize our results and discuss areas of future investigation.

II Data and Methodology

Bi-monthly stock price data of the the South Sea Trading Company have been compiled with quotes from the Course of the Exchange in London and with matching quotes from the Amsterdamsche Courant in Amsterdam. This process resulted in 1,675 observations quoted simultaneously in London and Amsterdam markets, dating from August 23, 1723 to December 19, 1794. Barnard's Act of 1734 forbade options or forward transactions on any stock traded in the London market. Stocks quoted on the Amsterdam market were time prices, given the Dutch tradition of not legally transferring title until company books were open for payment of dividends. Thus, the data consist of pure cash market prices in London from 1734 onward and futures market prices in Amsterdam.

First, we analyze each data series for non-stationarity using the Phillips and Perron (PP) test [19]. The PP test is based on several statistics: $Z(\phi_3)$, $Z(\phi_2)$, $Z(\tilde{\alpha})$, $Z(\phi_1)$, and $Z(\alpha^*)$. These statistics are used in testing whether the data series possesses either a deterministic trend or a unit root with or without drift.

After showing that each time series is non-stationary because of a unit root, we then test for cointegration between the two time series. If cointegration is present, it is inapproriate in constructing a model to make each series stationary by first-differencing. Rather one must construct the model using not only the individual series, first-differenced, but also a so called error correction mechanism. Alternatively, one can construct a model using either one of the individual time series, first differenced, and the cointegrating relationship. Since the bivariate data series consists of prices of the same stock on different exchanges, a reasonable guess is that the difference between the two time series is a cointegrating relationship. We test this conjecture using the PP test.

Next, we investigate the stock prices from the London exchange for possible long-term cycles. It is often speculated that long-term cycles exist in modern financial data such as stocks and commodities. If such cycles exist they may be a consequence of the four year business cycle, long-term weather cycles, etc. Traditionally, spectral analysis has been the principal tool used in uncovering cycles in time series. However, the application of spectral analysis to modern financial data has been relatively unsuccessful in finding or confirming the presence of any long-term cycles. Recently, an alternative technique has been used to expose long-term trends and cycles in financial data. Applying this technique Peters has analyzed approximately 100 years of daily Dow Jones Industrial Average data and concluded that the Dow has exhibited, first, a long-term persistence which cannot be explained by a model derived from a random walk and, second, an approximate four year cycle which is not apparent from spectral analysis [20]. His conclusions are based on a technique derived from rescaled range analysis. Rescaled range analysis has been pioneered by the hydrologist H. E. Hurst in his study of the long-term storage capacity of reservoirs [20, 8]. Extending Hurst's work Mandelbrot and others have introduced a family of Gaussian random functions, designated fractional Brownian motions (fBm's), which exhibit long-term persistence similar to that studied by Hurst and Peters [13, 14, 16, 15, 1]. A fractional Brownian motion, $B(t)$, of exponent H is defined as a moving average process in which past increments, $dB(s)$, are weighted by $(t-s)^{H-.5}$. Specifically,

$$B(t,w) - B(0,w) = \frac{1}{\Gamma(H+\frac{1}{2})} \int_0^t (t-s)^{H-.5} dB(s,w) \, . \tag{1}$$

Here w is the set of all values of a random function, and $0 < H < 1$. If $H = .5$, Eq. 1 reduces to Brownian motion. FBm's possess a number of interesting properties. The spectral densities of fBm's are proportional to ω^{1-2H}, ω being the frequency. They exhibit the following scaling behavior in the variance $\mathrm{VAR}(B(t,w) - B(0,\omega)) = t^{2H}\mathrm{VAR}(B(1,w) - B(0,w))$. As detailed in Mandelbrot and Wallis for $H > .5$ the future and past are positively correlated, with correlations approaching 1 as H approaches 1. For $H < .5$ the future and past are negatively correlated, with correlations decreasing to -.5. For $H = .5$ the future and past are uncorrelated.

There are several techniques for identifying fBm's. We shall now describe them briefly. In one technique the spectral density, $h(\omega)$, of a time series is estimated. Then, the $\log h(\omega)$ is plotted versus $\log \omega$. If persistence or anti-persistence is present, then for $\omega \to 0$ the plot should be a straight line whose slope is $1 - 2H$. For values of H, $0 < H < 1$, the series is stationary.

Another technique for identifying persistence utilizes the rescaled range, R/S, as first proposed by Hurst [8]. The basic premise of R/S analysis is that the variance of an integrated white noise process expands directly proportional to time. Consequently, any process which deviates from this behavior may be derived from an fBm. The quantity employed in the analysis is the rescaled range which is defined as follows. Given n consecutive points from a discrete time series $\{x_t\}$. Let

$$\bar{x} = \frac{1}{n} \sum_{k=1}^{n} x_k \,, \tag{2}$$

and

$$s_n = \sqrt{\frac{\sum_{k=1}^{n} (x_k - \bar{x})^2}{n}} \,. \tag{3}$$

The sample sequential range r_n is defined as

$$r_n = \max_{1 \le k \le n} (x_1 + \ldots + x_k - k\bar{x}) - \min_{1 \le k \le n} (x_1 + \ldots + x_k - k\bar{x}) \,, \tag{4}$$

and the rescaled range, R_n, is then given as

$$R_n = r_n / s_n \,. \tag{5}$$

A method proposed by Mandelbrot for assessing long run dependence in time series entails plotting the $\log r_n$ against $\log n$ [16, 12]. Asymptotically for large n, the slope of such a graph yields an estimate of H. Slopes greater than $\frac{1}{2}$, ($H > \frac{1}{2}$), indicate long run persistence, and slopes less than $\frac{1}{2}$, ($H < \frac{1}{2}$), indicate anti-persistence. As has been noted by various authors, the finite sample properties of H are modified by the short-range behavior of the time series [12, 3, 11]. In particular Davies and Harte have considered a stationary AR(1) process and show how regression estimates of H based on the slope method can be significantly biased in favor of rejecting $H = \frac{1}{2}$ even though, in fact, $H = \frac{1}{2}$ [3].

In order to reduce the effect of short-range biasing in the determination of H using the slope method Peters has proposed the following. First fit the data with an AR(1) model, and then perform the R/S analysis on the residuals generated from the model [20]. Peter's version of R/S analysis differs in some respects from that described by Mandelbrot and Wallis [16]. For a data series with T points Peters subdivides the series into smaller series of length n_j such that T/n_j equals an integer N_j. There will be J such divisions depending on the number of integers which divide into T exactly. For each j, ($j = 1, 2, \cdots, J$) Peters calculates N_j different rescaled ranges $R^k_{n_j}$, ($k = 1, 2, \cdots, N_j$), one for each subdivision. Peters then defines the rescaled range corresponding to each n_j as

$$R_{n_j} = \frac{\sum_{k=1}^{N_j} R^k_{n_j}}{N_j} \,. \tag{6}$$

The analysis proceeds as described previously, plotting $\log(R_{n_j})$ against $\log(n_j)$ and determining H from an estimate of the slope of the plot. Another statistic used originally by Hurst and more recently by Peters is V_{n_j},

$$V_{n_j} \equiv \frac{R_{n_j}}{\sqrt{n_j}} . \tag{7}$$

If the data series is derived from a white noise process, a plot of V_{n_j} against $\log(n_j)$ will be horizontal, since $R_{n_j} \approx \sqrt{n_j}$ for such data. Peter's has found this statistic useful in testing for cycles which may not be apparent from spectral analysis.

To assess the effect of biasing due to short-memory processes on his version of R/S analysis Peters has applied this analysis to data simulated from a variety of models including autoregressive, moving average, ARCH, and GARCH. Although the filtering method proposed by Peters appears to remove the short-range biasing effects of ARIMA models, the method seems not to be totally satisfactory in the case of ARCH and GARCH models. Specifically, when applied to data derived from ARCH or GARCH processes this method yields results which are consistent with data derived from processes exhibiting long-range dependence. The ARCH derived data show properties similar to a persistent process, and the GARCH derived data show properties similar to an anti-persistent process.

Lo has constructed a test statistic for long-range dependence similar to the rescaled range, R_n [11]. However, this test statistic takes account of short-range dependence by adjusting the variance used in R_n. The test statistic Q_n is defined as

$$Q_n \equiv \frac{r_n}{\sigma_n(q)} , \tag{8}$$

where

$$\sigma_n^2(q) = s_n^2 + \frac{2}{n}\sum_{j=1}^{q} \omega_j(q) \left\{ \sum_{i=j+1}^{n} (x_i - \bar{x})(x_{i-j} - \bar{x}) \right\} , \tag{9}$$

and

$$\omega_j(q) = 1 - \frac{j}{q+1} . \tag{10}$$

The value of the truncation lag, q, depends on the data being considered. Its value must be large enough to account for short-term dependencies in the data but not so large as to alter the finite sample distribution of Q_n radically (See reference [11] for Lo's discussion about how to determine an appropriate value for q). Lo defines the statistic $V_n(q)$, which is based on the test statistic Q_n,

$$V_n(q) \equiv \frac{Q_n}{\sqrt{n}} . \tag{11}$$

This statistic is identical to the statistic V_{n_j} used by Hurst and Peters, except that the variance has been adjusted to account for short range memory effects. Lo derives

Prob($V < v$)	.005	.0250	.050	.100	.200	.300	.400	.500
v	0.721	0.809	0.861	0.927	1.018	1.090	1.157	1.223
Prob($V < v$)	.543	.600	.700	.800	.900	.950	.975	.995
v	$\sqrt{\frac{\pi}{2}}$	1.294	1.374	1.473	1.620	1.747	1.862	2.098

Table 1: Fractiles of the limiting distribution of the V statistic under the assumption of no long memory.

the limiting distribution of V_n for which $V_n(q)$ is an estimator, under the assumption of no long memory and a general set of conditions which include strong mixing (short memory) and conditional heteroskedasticity.

We now discuss how rescaled range analysis can be used to uncover long-term cycles. In our analysis we consider only the South Sea Trading Company data listed on the London exchange.[3] First, we calculate $V_n(q)$ for various values of q, where $n = n_j$ for all j (See Eqs. 7 and 11). Next we plot a graph of $V_{n_j}(q)$ against n_j. We then inspect the plot for the value of n_j where the slope of the graph changes from greater than or equal to zero to a value significantly less than zero. This is the time scale corresponding to the onset of anti-persistence, i. e. the variance of the data no longer continues to expand, and is the length of the cycle. Peters discusses this point extensively [20]. If there is is a single break in the slope so that the anti-persistence remains at all larger time scales, we can test the statistical significance of the slope change as follows. For various values of q, we evaluate $V_{n_J}(q)$, the estimate of the V statistic for the entire time series. We then compare the values of $V_{n_J}(q)$ to the values given in Table 1. The values in the table compiled by Lo are derived from the limiting distribution of V_n. Small values of V correspond to anti-persistence, and large values correspond to persistence. If $V_{n_J}(q)$ is less than .861, we would accept the hypothesis of anti-persistence at the 5% confidence level, and therefore conclude that a long-term cycle is present.

It should be possible to increase the power of the V statistic for a given set of data, by first constructing a model which accounts for many of the characteristics typically found in modern financial data. The model, however, should not account for any long-term periodicity, thus, providing a specific null hypothesis against which to test for the presence of a long-term periodicity using the V statistic. The procedure involves simulating data based on the model and constructing a probability table based on the simulation data.

Since Fama's [4] examination of financial asset prices and his finding of skewed and leptokurtic data, researchers in finance have been attempting to incorporate

[3]As is discussed later the data listed on the Amsterdam exchange are more influenced by various systematic effects. When the data are adjusted for these effect, both the London and Amsterdam data yield similar results.

this finding into their statistical hypothesis testing design. Fama's work and later extensions have demonstrated that financial asset prices are characterized by unconditional distributions which are generally leptokurtic. In addition, their conditional distributions have been found to be extremely conditionally heterogenous. If these old data exhibit similar properties, this would confirm that conditional heterogeneity has been a characteristic of financial data for more than 250 years. In order to take account of possible nonlinearities in the data we have adopted the seminonparametric (SNP) methodology developed by Gallant and Tauchen for modelling the data [7]. The primary benefit of the SNP methodology is that very few assumptions are made about the underlying process generating the multivariate time series being modeled. Other methodologies such as ARCH or GARCH [2] are much more restrictive in their assumptions about the underlying process. The main disadvantages of the SNP methodology are that it can be much more computer intensive and estimated models tend to have a relatively large number of parameters.

In order to take account of possible cointegration we have modified the SNP methodology as follows. When the individual series possess a unit root and cointegration is not present, the bivariate series $\mathbf{v}_t$ to be modeled consists of univariate time series which are obtained by first–differencing each of the original time series. When cointegration is present, the bivariate time series comprises one of the original series which has been first–differenced and the other which is the cointegration relationship. Specifically, given a bivariate time series of financial data $(s_{1,t}, s_{2,t})$[4], the bivariate series, $\mathbf{v}_t$, to be modeled is given by

$$\mathbf{v}_t = \begin{pmatrix} v_{1,t} \\ v_{2,t} \end{pmatrix}, \tag{12}$$

where $v_{1,t} = s_{1,t} - s_{1,t-1}$, and $v_{2,t}$ is the cointegrating relationship, i. e. it is some linear combination of $s_{1,t}$ and $s_{2,t}$ which is stationary. Note that $v_{1,t}$ depends on t and $t-1$, whereas $v_{2,t}$ depends only on t. Thus, any model building strategy in which lagged values of $\mathbf{v}_t$ are required must take this into account in order to be consistent. For example, any model which contains lagged values of the data up to lag K must include $v_{1,t-i}$, $(i = 1, \ldots, K)$ and $v_{2,t-i}$, $(i = 1, \ldots, K+1)$, i. e. the term $v_{2,t-(K+1)}$ must also be included in the model.

III Data Analysis

In Figure 1 we show stock prices of the South Sea Trading Company as listed on the London exchange. The corresponding prices on the Amsterdam exchange are shown in Figure 2. The analysis of these time series is complicated by the fact that the time interval between successive data points is not constant, ranging from 3 to 79

[4] We assume that the series $s_{1,t}$ and $s_{2,t}$ have already been log transformed to homogenize their data.

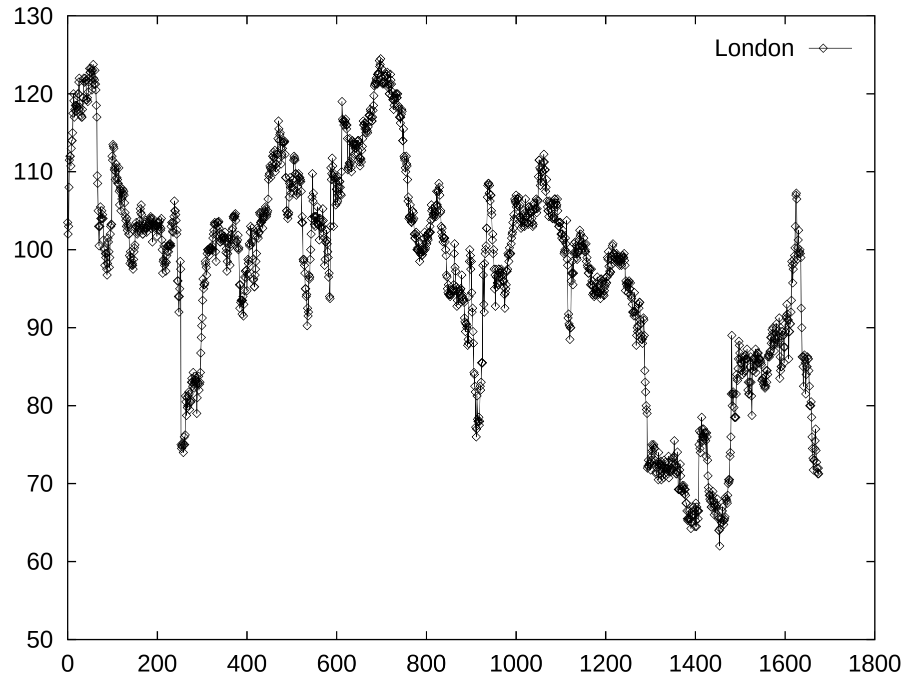

Figure 1: Stock prices of the South Sea Trading Company quoted on the London exchange.

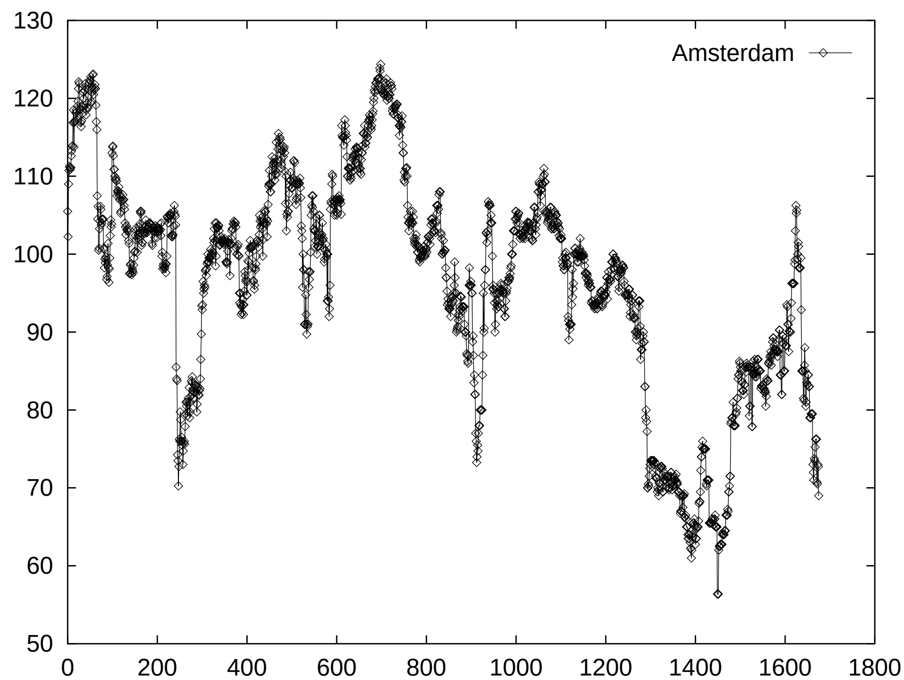

Figure 2: Stock prices of the South Sea Trading Company quoted on the Amsterdam exchange.

Time Interval	number of Occurrences	Time Interval	number of Occurrences	Time Interval	number of Occurrences
14	544	13	15	51	2
16	232	26	11	3	1
12	156	25	8	6	1
21	114	5	7	29	1
17	109	8	5	36	1
19	94	22	4	37	1
11	73	30	4	40	1
9	63	35	4	44	1
18	48	4	3	49	1
7	43	20	3	58	1
10	33	42	3	61	1
23	23	27	2	63	1
28	20	31	2	65	1
24	17	32	2	79	1
15	16	38	2		

Table 2: The time intervals between successive data points and their number of occurrences.

days. In Table 2 we list the various time intervals and the number of occurrences of that interval. The most common time interval between consecutive data points is 14 days; however, there are a large number of other intervals present. In Figures 3 and 4 we have plotted the log first-differences of the prices as a function of the number of days between successive data points. In order to assess the biasing effect of the non-constant time intervals between consecutive data points, we have tested for differences in the mean and variance of the various time intervals. We first have assigned dummy variables to each time interval. Then, we have applied the procedure described in Appendix A.1 to the log first difference of prices ($s_{i,t} - s_{i,t-1}, i = 1, 2$) on the London and Amsterdam exchanges. In Table 3 we have summarized the results of this analysis.

	Mean				Variance			
	F Stat.	Signif.	Pairwise Compar.	Range Test	F Stat.	Signif.	Pairwise Compar.	Range Test
London	1.5	.021	36,51	5.861	1.15	.229		
Amster.	.76	.872			.79	.917		

Table 3: Regression and Studentized Range tests. Only one pairwise comparison was significant according to the Tukey Studentized Range Test. The 5% confidence level of the range test is 5.573.

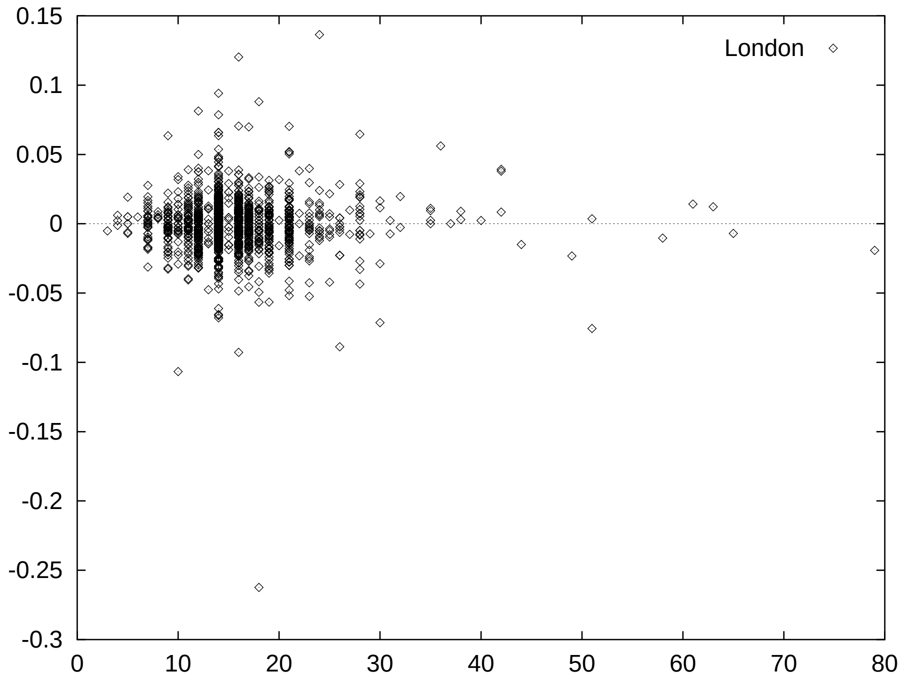

Figure 3: Logarithmic first differences of prices on the London exchange as a function of the number of days between successive data points.

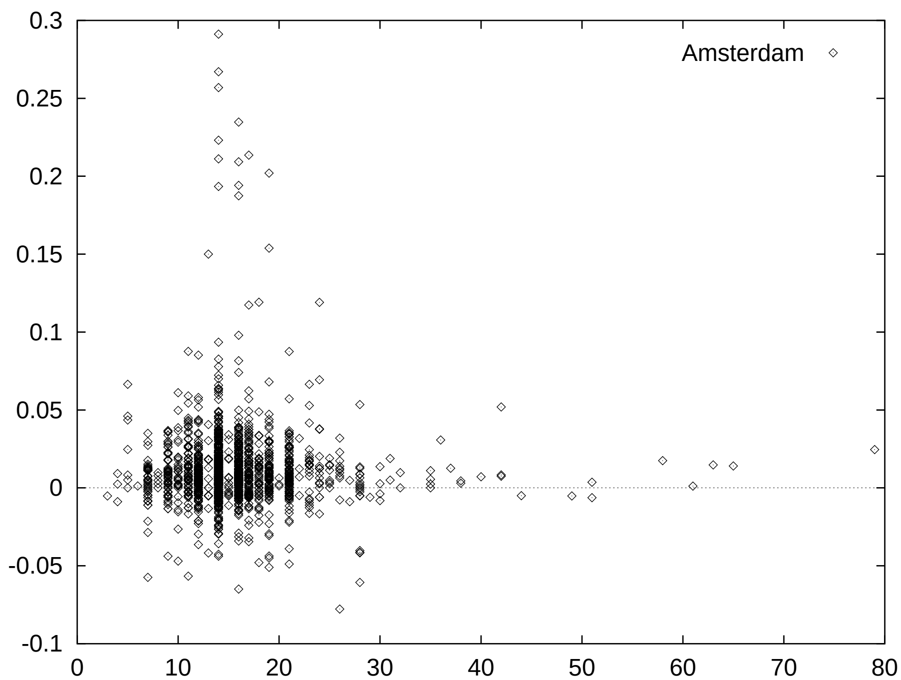

Figure 4: Logarithmic first differences of prices on the Amsterdam exchange as a function of the number of days between successive data points.

	London exchange				Amsterdam exchange		
	Min.	Median	Max.		Min.	Median	Max.
$Z(\phi_3)$	4.83	5.24	5.74		4.84	5.68	5.38
$Z(\phi_2)$	3.29	3.56	3.88		3.30	3.84	3.64
$Z(\phi_1)$	2.33	2.55	2.82		2.37	2.83	2.64
	Critical Values		10%	5%	2.5%	1%	
	$Z(\phi_3)$		5.34	6.25	7.16	8.27	
	$Z(\phi_2)$		4.03	4.68	5.31	6.09	
	$Z(\phi_1)$		3.78	4.59	5.38	6.43	

Table 4: Unit root tests for prices of the South Sea Trading Company listed on the London and Amsterdam exchanges.

The F statistic reported for the London data is significant at 5% confidence level. We also report results of the Tukey studentized range test [21, 10, 9]. The Tukey test is especially appropriate since there are 44 different categories of intervals resulting in $(44)(43)/2 = 946$ comparisons to be made. Because there are so many comparisons it is more likely to find a significant difference between a pair of means even when one does not exist. The Tukey test controls for this type of error. According to this test the means of only 2 out of the 946 comparisons are significantly different. Specifically, the mean of the time interval of 36 days differs from that of 51 days. However, the number of data points corresponding to the 36 and 51 day intervals is one and two, respectively. We conclude that any systematic biasing in the means is weak so that we do not adjust the data for biasing in the mean. We next test for systematic biasing in the variance. In Table 3 we report the F and studentized range statistics. Pairwise comparing the means of the dummy variables we find no significant differences between different time intervals. In summary we have found it unnecessary to adjust the data series because of any systematic biasing from a non-constant time interval between successive data points.

We first test each series for the presence of a unit root using the PP test. In Table 4 we present the results of the test.[5] In both price series we are unable to reject the hypotheses associated with $Z(\phi_3)$, $Z(\phi_2)$, or $Z(\phi_1)$. Thus, we conclude that each series possesses a unit root.

If we compare Figure 1 to Figure 2, we observe that the two price series, generally, follow each other very closely, suggesting that the two series may be cointegrated. To test this hypothesis we logarithmically transform each price series, obtaining $\tilde{s}_{it}$ where $i = 1, 2$ for prices on the London or Amsterdam exchanges, respectively. Using the PP

[5]In the PP test the various statistics depend on the correlation structure of residuals derived from a regression equation. To estimate the correlation structure Perron suggests that the test statistics be calculated for various values of a lag parameter, ℓ to insure stability of the various statistics. We have selected $\ell = 1, 2, \ldots 8$.

	Min.	Median	Max.	
$Z(\phi_3)$	142.2	182.5	210.3	
$Z(\tilde{\alpha})$	-45.8	-62.6	-73.9	
Critical Values	10%	5%	2.5%	1%
$Z(\phi_3)$	5.34	6.25	7.16	8.27
$Z(\tilde{\alpha})$	-18.3	-21.8	-25.1	-29.5

Table 5: Test for cointegration between stock prices on the London and Amsterdam exchanges (no adjustment for anomalous price behavior on the Amsterdam exchange).

Maxima			Minima		
Price	Time Step	Difference	Price	Time Step	Difference
123.25	52		74.00	259	
124.50	698	646	76.00	913	654
112.25	1062	364	62.00	1455	542
107.25	1625	563			

Table 6: Relative extrema in London stock prices. The length of the approximate cycle is obtained by averaging the numbers in the columns labeled difference, yielding a value of 554 periods or 21 years.

procedure we test whether the difference between the two series, $\tilde{v}_{2,t}$, i. e. $\tilde{s}_{1t} - \tilde{s}_{2t}$, is stationary. The results of the PP test are presented in Table 5. Based on these results we reject the hypotheses associated with $Z(\phi_3)$ and $Z(\tilde{\alpha})$ and therefore conclude that $\tilde{v}_{2,t}$ is stationary so that the two price series are cointegrated.

In Figure 1 the stock prices on the London exchange appear to exhibit long-term cyclical behavior. This is evidenced in the approximately constant elapsed time between relative extrema in prices. There are four relative maxima and three relative minima. The average elapsed time between consecutive maxima and consecutive minima is approximately 554 periods or 21 years. In Table 6 we present the stock price and the time step at which each relative extremum occurs. We also give the number of time steps between successive extrema. To prove the existence of the cycle rigorously we test for its presence by fourier decomposing the data from the London exchange. Since the stock price possesses a unit root and is nonstationary, we analyze the first differenced series $\tilde{s}_{1t}$. In Figure 5 we show the periodogram of the series $\tilde{s}_{1t}$. For ease of interpretation we have plotted the period rather than the frequency on the abscissa. From appearances there is very little evidence for the presence of any cycles. To confirm this we apply the Fisher-Kappa test to the largest periodogram ordinate, corresponding to a period of approximately 6 time steps [5]. The statistic is insignificant at the 5% level and therefore fails to confirm the presence of any cycle. The results of the test are summarized in Table 7. As an alternative to differencing

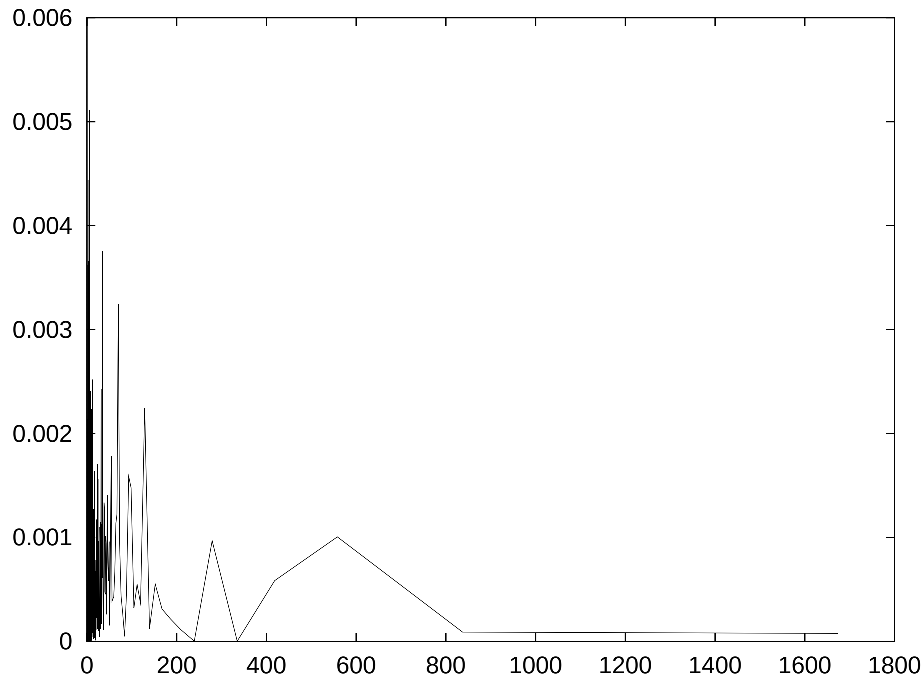

Figure 5: Periodogram of stock prices on the London exchange. The prices have been logarithmically transformed and then differenced. For ease of interpretation, period rather than frequency, has been plotted along the abscissa.

we detrend the logarithmically transformed data by regressing them against the year. In Figure 6 we show a plot of the residuals which result from removing the trend from the data. This Figure possesses characteristics similar to Figure 1 except the trend has been removed. In Figure 7 we show the periodogram of the detrended data as a function of period. Note the peak in the spectrum at a period of approximately 550 time steps. In Table 7 we report the Fisher–Kappa statistic. It indicates the presence of a highly significant long-term cycle of approximately 558 time steps. This supports our initial analysis based on differences between extrema. Nonetheless, we must emphasize that although this result is highly suggestive, the analysis is, strictly speaking, inappropriate since even after detrending, the series is nonstationary. This follows as a result of the PP test from which we have concluded that the data possess a unit root and not a deterministic trend.

In Section II we have described rescaled range analysis which provides an alternative methodology to spectral analysis for identifying long term cycles. We now apply this technique to the logarithmically transformed stock data on the London exchange. In Figure 8 we show a plot of the statistic $V_{n_j}(q)$ against $\log(n_j)$ (See Eq. 7.) for $q = 8$[6]. The dramatic drop in the value of the V statistic suggests the onset of anti-persistence in the data at about 558 time steps. For $n_j < 558$ the V statistic is consistent with a process whose variance expands linearly in time like brown noise.

[6] The value of the truncation lag is consistent with the SNP model to be described.

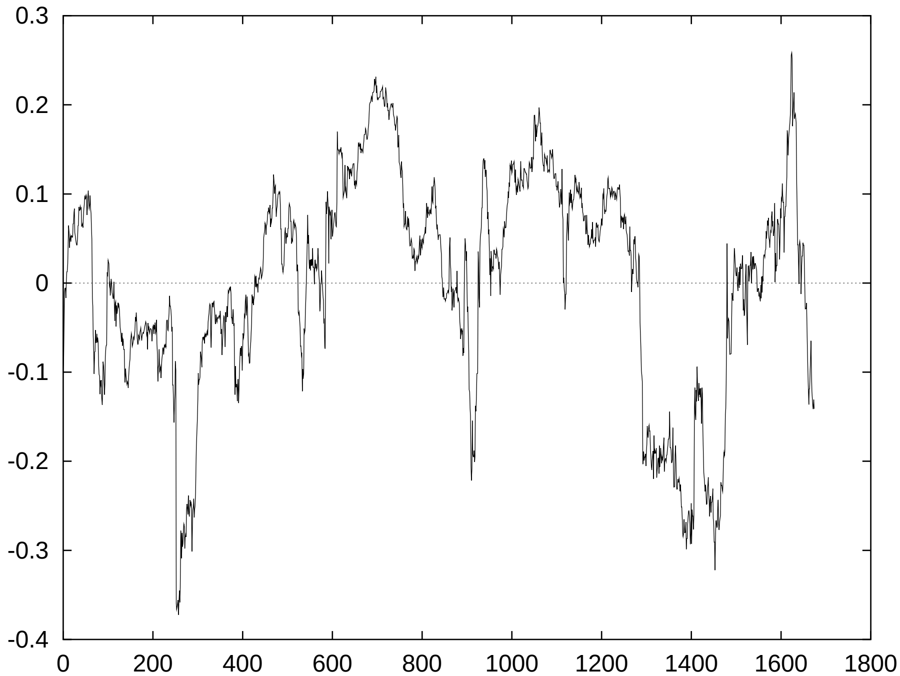

Figure 6: Residuals of detrended stock prices on the London exchange. The prices have been logarithmically transformed and then detrended.

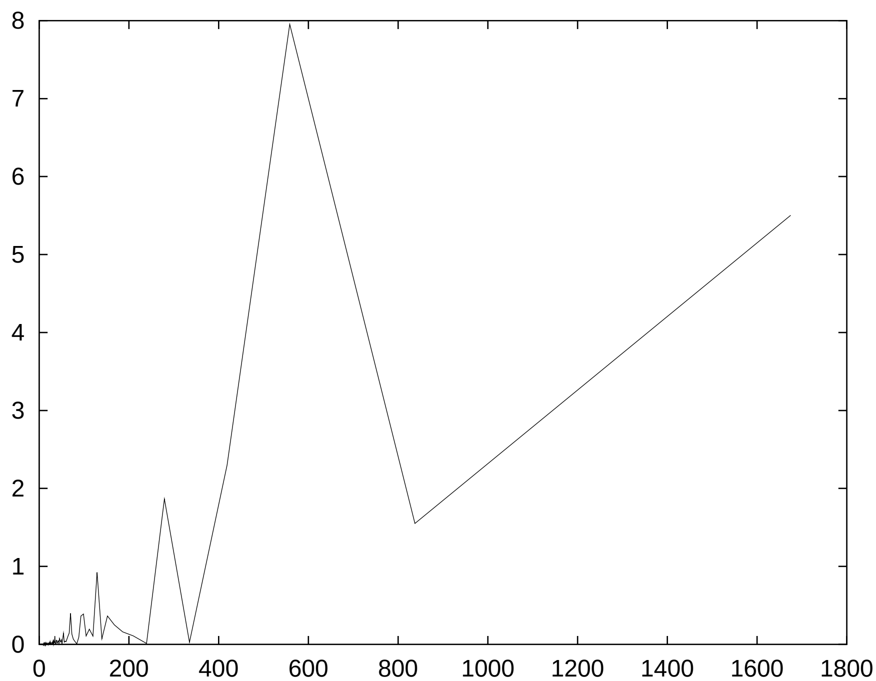

Figure 7: Periodogram of stock prices on the London exchange. The prices have been logarithmically transformed and then detrended. For ease of interpretation period, rather than frequency has been plotted along the abscissa.

	Differenced		Detrended
ξ	7.566		264.602
Period	6.022		558.330
Critical Values	10%	5%	1%
ξ	8.901	9.612	11.220

Table 7: The Fisher–Kappa statistics, ξ, are reported for the London prices in two cases: In each case the data have been logarithmically transformed and then either differenced or detrended. Also, given is the period corresponding to the value of the statistic. Critical values are a lower bound.

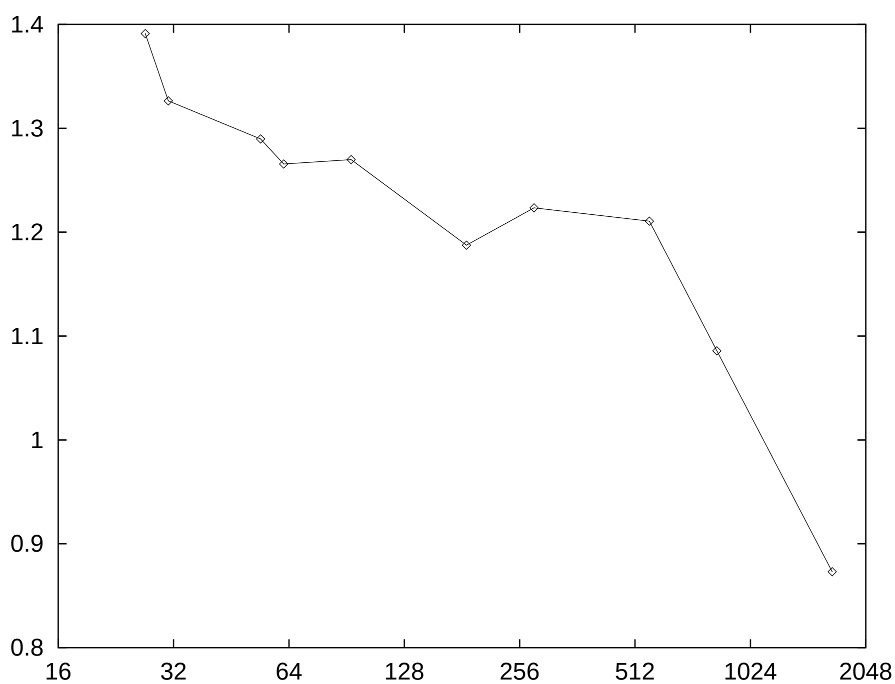

Figure 8: The V statistic versus the logarithm of the number of points for stock prices on the London exchange. The value of the V statistic decreases at faster rate when the number of points is greater than 558. This suggests the onset of anti-persistence in the data at about 558 time steps. The value of the V statistic for the entire set of stock prices ($n_j = 1774$) is .8782.

q	$V_{n_j}(q)$	q	$V_{n_j}(q)$
0	.88341230	9	.87311548
1	.89098346	11	.86693072
2	.89802933	12	.86245286
3	.90026760	13	.85899228
4	.89923990	14	.85546982
5	.89440978	15	.85591251
6	.88819647	16	.85912472
7	.88388050	17	.86016440
8	.87827218	20	.86613411

Table 8: The V statistic for various values of the truncation parameter, q . Listed are values of $V_{n_j}(q)$ where $n_j = 1774$, for the stock prices on the London exchange.

To determine whether anti-persistence is present, we test whether the value of the V statistic for the entire series, i.e. $V_{1764}(q)$, is significantly different from that associated with brown noise. In Table 8 we give the values of the statistic for various values of q. Comparing the value of the V statistic with $q = 8$ to the values in Table 1 we find this value to be significantly different from that of brown noise at slightly greater than the 5% confidence level. Consequently, we have fairly strong evidence for the presence of a long-term cycle, stronger than from analyzing the spectrum of first differenced logarithmically transformed data.[7] The corresponding V statistic for prices on the Amsterdam exchange is $V_{1764}(8) = .9235$, somewhat less significant than prices on the London exchange. However, after adjusting each data series for systematic effects and recalculating the V statistics, they are closer in value.[8] Specifically, for the London data $V_{1764}(8) = .8709$ and the for Amsterdam data $V_{1764}(8) = .8907$. It appears that the price data from the Amsterdam exchange are influenced more by systematic effects than the corresponding data on the London exchange. We restrict our subsequent analysis of long-term cycles to the London data.

Although the results of the rescaled range analysis provide strong support for anti-persistence, we may be able to strengthen the case for antipersistence. This can be accomplished by constructing confidence bands for the Lo test based on a specific model. The model should capture all characteristics of the data except the apparent long-term cyclical behavior. The power of the Lo test should increase when a specific alternative model is adopted as the null hypothesis.

The bivariate series to be modeled (See Eq. 12) consists of the London data,

[7] We have tested the stability of these results by calculating the V statistics for the London price series containing only $(s_{1,1}, s_{1,3}, s_{1,5} \cdots)$ or $(s_{1,1}, s_{1,6}, s_{1,11} \cdots)$. The value of the V statistic for the entire price series in each case is .85431 or .90735, respectively. This compares reasonably well to the value .87827 obtained when using all points in the series.

[8] Adjustments for systematic effects in the price data are discussed later in this section.

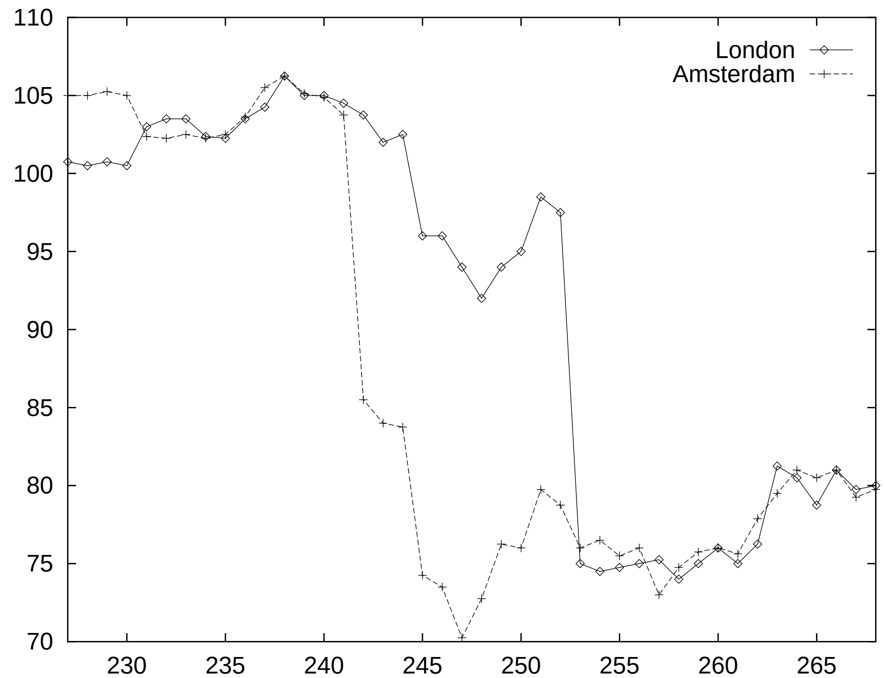

Figure 9: Stock prices of the South Sea Trading Company quoted on the London and Amsterdam Exchanges from Jan. 5, 1733 until Sept. 15, 1734. A divergence in prices occurred on Aug. 31, 1733 and persisted until Feb. 1, 1734. This precipitous drop in prices on the Amsterdam exchange may have been a harbinger of the Barnard Act of 1734.

first differenced, i. e. $v_{1,t} = s_{1,t} - s_{1,t-1}$, and the cointegrating relationship, i. e. $v_{2,t} = s_{1,t} - s_{2,t}$. Since the London and Amsterdam price series are cointegrated, they follow each other closely during the period under investigation. A notable exception, however, is during the period from Aug. 31, 1733 until Feb. 1, 1734. As shown in Figure 9, prices on the Amsterdam exchange have deviated considerably from prices on the London exchange. This anomalous price behavior may have been indicative of a structural shift in the data and therefore cannot be taken into account by a stochastic model. Consequently, we have divided the data into three periods:

Period–1. prior to Aug. 31, 1733,

Period–2. from Aug. 31, 1733 until Feb. 1, 1734, and

Period–3. after Feb. 1, 1734.

We have assigned dummy variables to the three periods and have tested for systematic differences in the means and variances, as described in Appendix A.1. In Tables 9 and 10 we summarize the results of the tests for $v_{1,t}$ and $v_{2,t}$ respectively. The Tables give the F statistic for each regression and the value of the Tukey test for each pairwise comparison. As is apparent from Figure 9, the mean value of the cointegrating

Mean				Variance			
F Stat.	Signif.	Pairwise Compar.	Range Test	F Stat.	Signif.	Pairwise Compar.	Range Test
0.63	.5346	Per-1,2	1.583	11.76	0.0001	Per-1,2	4.090
		Per-1,3	.2874			Per-1,3	4.612
		Per-2,3	1.546			Per-2,3	5.228

Table 9: Regression and Studentized Range tests for series $v_{1,t}$. The 5% confidence level of the range test is 3.317.

Mean				Variance			
F Stat.	Signif.	Pairwise Compar.	Range Test	F Stat.	Signif.	Pairwise Compar.	Range Test
758.38	0.0001	Per-1,2	55.08	40.50	0.0001	Per-1,2	6.398
		Per-1,3	10.35			Per-1,3	12.03
		Per-2,3	52.72			Per-2,3	3.748

Table 10: Regression and Studentized Range tests for series $v_{2,t}$. The 5% confidence level of the range test is 3.317.

relationship, $v_{2,t}$, during period 2 is significantly shifted from the other periods. In addition, we find that the mean of period 1 is significantly different from period 3. In the case of the differenced London price data, $v_{1,t}$, there is no difference in means between periods. In both series variances, and therefore the volatility, differ between periods. It seems unlikely that a stochastic model will be able to account for these differences. Consequently, we adjust the data to take account of these mean and variance shifts using Eqs. 17 and 18 in Appendix A.1. In Table 11 we give values of quantities used in the adjustment. In Figures 10 and 11 we show series $v_{1,t}$ before and after the adjustments. In Figures 12 and 13 we show the corresponding graphs for series $v_{2,t}$.

Comparing Figures 12 and 13 we observe that the adjustment process has homog-

	$v_{1,t}$		$v_{2,t}$	
	$\mathbf{g} \cdot \boldsymbol{\beta}$	$\mathbf{g} \cdot \gamma$	$\mathbf{g} \cdot \boldsymbol{\beta}$	$\mathbf{g} \cdot \gamma$
Per-1,2	0.00022291	-10.4872960	-0.00074386	-10.9764999
Per-1,3	0.00022291	-7.4589315	-0.22697933	-7.7974622
Per-2,3	0.00022291	-9.6257246	-0.01034503	-9.6257246
a	.010392486		-.00024556215	
b	.010298557		.0036976245	

Table 11: Adjustment coefficients for the series $v_{1,t}$ and $v_{2,t}$.

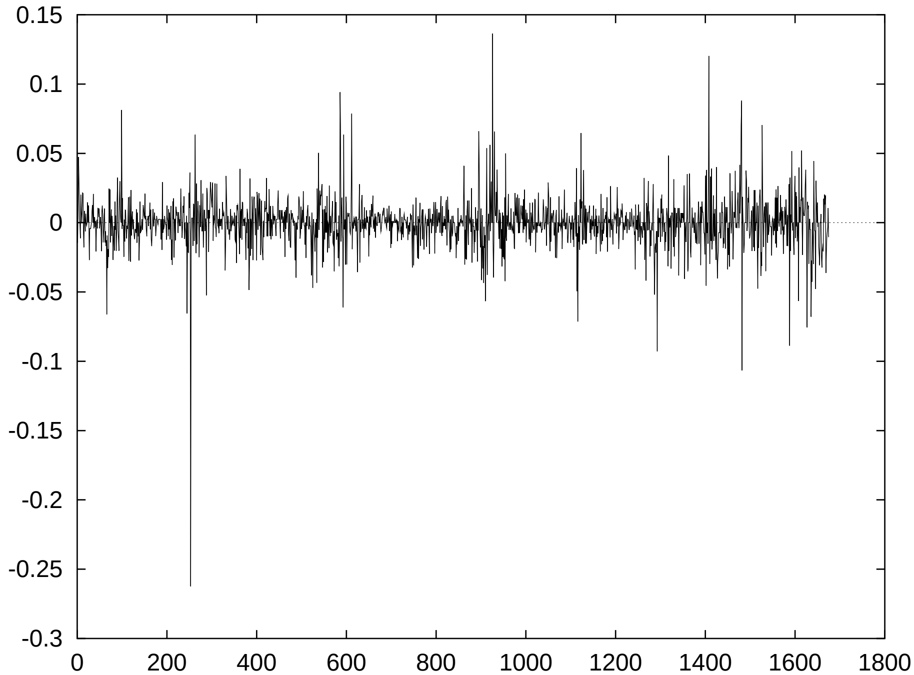

Figure 10: The series $v_{1,t}$ prior to mean and variance adjustments.

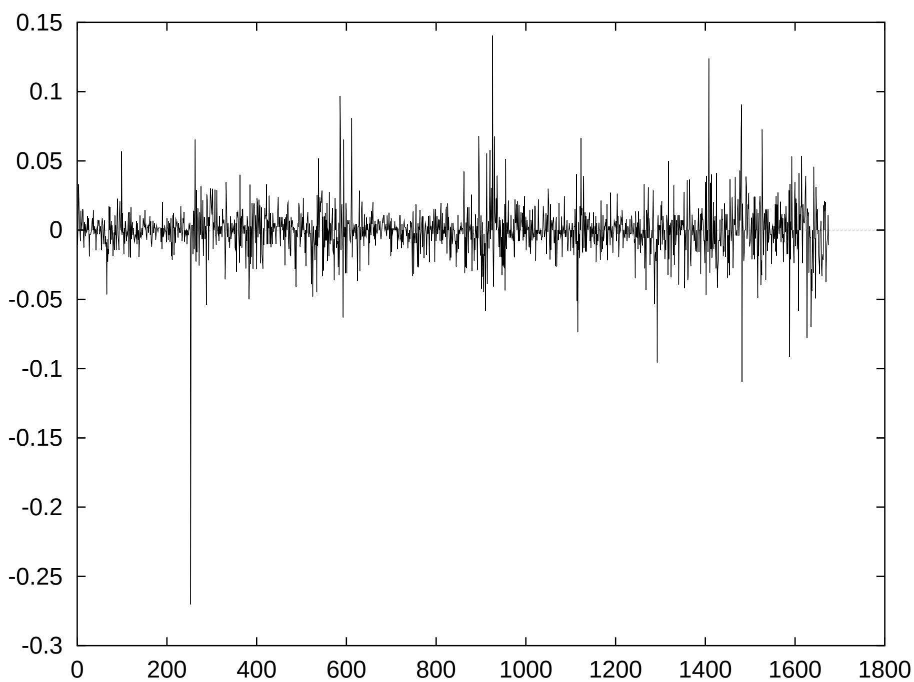

Figure 11: The series $v_{1,t}$ after mean and variance adjustments.

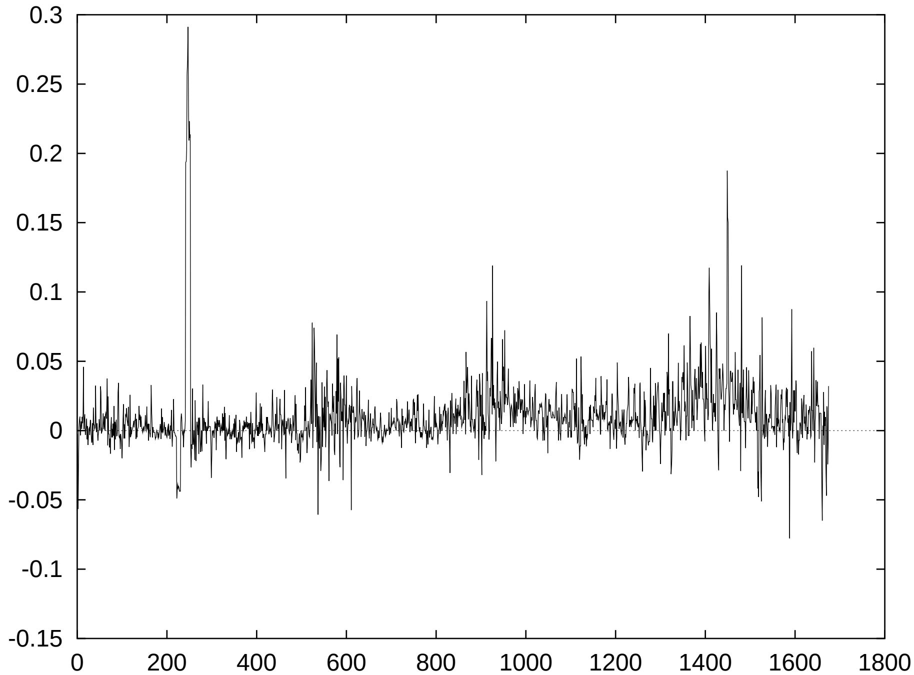

Figure 12: The series $v_{2,t}$ prior to mean and variance adjustments.

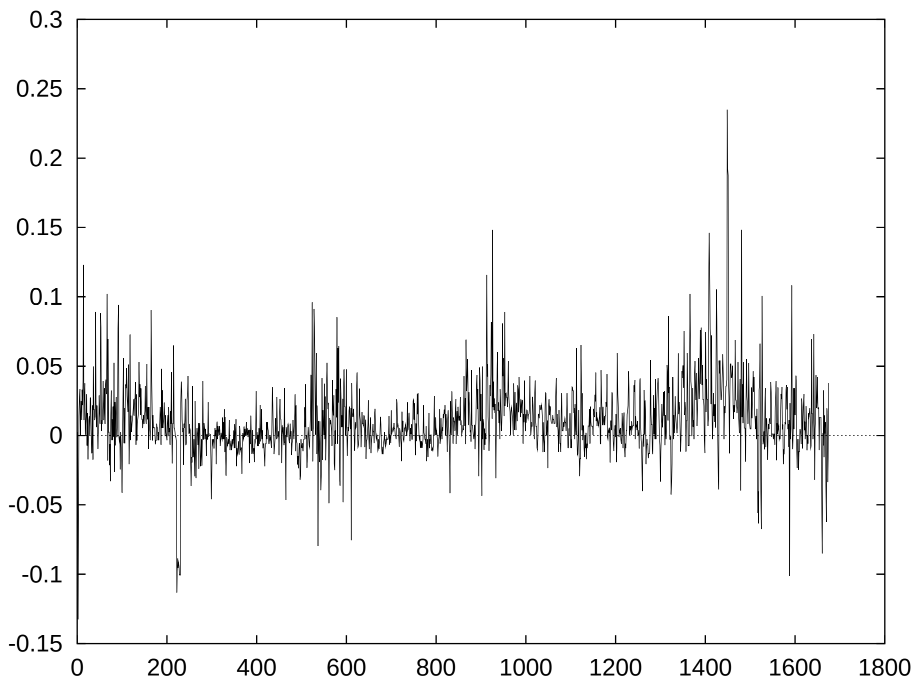

Figure 13: The series $v_{2,t}$ after mean and variance adjustments.

enized the series $v_{2,t}$ to some extent. This is not the case for series $v_{1,t}$. Figures 11 and 12 appear very nearly the same except for some subtle differences. The adjustment factors in Table 11 can be interpreted in a straightforward manner. Period 1 shows the least volatility, period 2 the most, and period 3 slightly more volatility than period 1. It appears that prior to the Barnard Act these markets were least volatile. Then the markets showed a marked increase in volatility in anticipation of the Barnard Act. Finally, the volatility subsided for the most part but does not diminish to the levels of period 1. This would imply that the implementation of the Barnard Act restricted the activity of traders, creating an increase in uncertainity in the markets. However, the sharp drop in the Amsterdam market as well as the persistent divergence in prices, well in advance of the price movement in London is somewhat of a puzzle. This may be the result of forward nature of the Amsterdam market, or it may hint at some communication inefficiency.

Having removed some systematic effects from the data we now are ready to construct a model. For this purpose we have adopted the SNP methodology of Gallant and Tauchen. [6, 7] The SNP methodology is a seminonparametric methodology for modeling an m–dimensional time series $\mathbf{v}_t$, $(1 \leq t \leq n)$. [9] The methodology is based on the fact that a probability density for m–dimensional innovations $\mathbf{z}_t$ can be approximated by an Hermite polynomial expansion. The lowest order polynomial corresponds to a multivariate normal distribution. As higher order polynomials are included in the expansion it can more closely approximate distributions which deviate from a normal distribution. The series $\mathbf{v}_t$ is assumed to be derivable from a VAR/ARCH-like process driven by the innovations $\mathbf{z}_t$. Furthermore, the coefficients of the Hermite polynomials can depend on lagged values of $\mathbf{v}_t$. There are a number of tuning parameters for controlling the nature of the distribution: L_r is the number of lags in ARCH component; L_p is the number of lags of $\mathbf{v}_t$ in the coefficients of the polynomial; the number of lags in the VAR component is the $\max(L_r, L_p)$; K_z is the highest degree in $\mathbf{z}_t$ of the Hermite polynomial; K_x is the highest degree in $\mathbf{v}_t$ of the coefficients of the Hermite polynomial and; I_z and I_x are used for reducing the number of higher order interaction terms, i.e. products between different components of $\mathbf{z}_t$ and $\mathbf{v}_t$, respectively, of the polynomial. The parameter K_z controls the shape of the homogeneous component of the distribution, and K_x the heterogeneous component. A suitable choice of the parameter L_r can greatly reduce the value of K_x in fitting a distribution which is conditionally heterogeneous. In Table 12 the types of processes characterized by these parameters are summarized.

The model building process is incremental in that one starts with a vector autoregressive model and gradually increases the number or value of parameters until a suitable fit to the data is achieved. The model takes account of cointegration between the two series, conditional heterogeneity, and deviations of the conditional probability density function from normality. The model does not take account of any long-term

[9] In our case $m = 2$ so that $\mathbf{v}_t = (v_{1,t}, v_{2,t})$.

Parameter setting	Characterization of $\mathbf{v}_t$
$L_r = 0,\ L_p = 0,\ K_z = 0,\ K_x = 0$	iid Gaussian
$L_r = 0,\ L_p > 0,\ K_z = 0,\ K_x = 0$	Gaussian VAR
$L_r = 0,\ L_p > 0,\ K_z > 0,\ K_x = 0$	non-Gaussian VAR, homogeneous innovations
$L_r > 0,\ L_p = 0,\ K_z = 0,\ K_x = 0$	Gaussian ARCH
$L_r > 0,\ L_p > 0,\ K_z > 0,\ K_x = 0$	non-Gaussian ARCH, homogeneous innovations
$L_r > 0,\ L_p > 0,\ K_z > 0,\ K_x > 0$	general nonlinear process, heterogeneous innovations

Table 12: Characterization of processes according to restrictions on the tuning parameters.

cyclical behavior. To evaluate the goodness-of-fit we use the Schwarz criterion,

$$BIC = s_n(\hat{\theta}) + (1/2)(p_\theta/n)\ln(n)\ , \tag{13}$$

where s_n is the objective function for the model, $\hat{\theta}$ are the values of the estimated parameters defining the model, s_n is the objective function for the fit to the data, p_θ is the number of parameters in the model. The Schwarz criterion is more conservative than other criteria like the Akaike information criterion in that it penalizes more for a model when a low value of the objective function is obtained at the expense of an excessive number of parameters. The model has been constructed using the FORTRAN code SNP 7.5 [7]. In Table 13 we summarize the results of constructing the model. The table gives model specifications and the value of the Schwarz criterion for each model. The optimum model according to the Schwarz criterion is a non-Gaussian ARCH with homogeneous innovations. The values of the tuning parameters are $(L_r, L_p, K_z, K_x, I_z, I_x)$ = (7,0,9,6,0,0).

Finally, because of the conservative nature of the Schwarz criterion the adequacy of the optimum model is checked for long-term and short–term misspecification as follows [7]. First, the one–day–ahead scaled residuals of the model are calculated

$$\hat{\mathbf{u}}_t = [\mathrm{Var}_c(\mathbf{v}_t)]^{-1/2}[\mathbf{v}_t - <\mathbf{v}_t>_c]\ , \tag{14}$$

where $\mathrm{Var}_c(\mathbf{v}_t)^{-1/2}$ is the Cholesky factor of the one–day–ahead conditional variance-covariance matrix, and $<\mathbf{v}_t>_c$ is the one–day–ahead conditional mean, both being estimated from the model. Then, to test for short-term misspecification the scaled residuals are regressed on the first, second, and third powers of all lagged components of $\mathbf{v}_t$ up to and including lag 10 (mean test). The squared components of the residuals are also regressed in a similar manner (variance test). If either multivariate regression accounts for a significant amount of variance according to the Wilks test, then the values of the tuning parameters are systematically increased, and the parameters θ re–estimated [17]. Testing for long-term misspecification proceeds in a similar manner.

L_r	L_p	K_z	I_z	K_x	I_x	p_θ	$s_n(\boldsymbol{\theta})$	Schwarz
1	0	0	0	0	0	10	2.596725	2.618885
2	0	0	0	0	0	20	2.518330	2.562649
3	0	0	0	0	0	30	2.491190	2.557670
4	0	0	0	0	0	40	2.477593	2.566233
5	0	0	0	0	0	50	2.461504	2.572303
6	0	0	0	0	0	60	2.421857	2.554816
7	0	0	0	0	0	70	2.399490	2.554609
8	0	0	0	0	0	80	2.388446	2.565725
9	0	0	0	0	0	90	2.373771	2.573210
10	0	0	0	0	0	100	2.345458	2.567057
11	0	0	0	0	0	110	2.330176	2.573935
7	0	1	0	0	0	72	2.395421	2.554972
7	0	2	1	0	0	74	2.218418	2.382401
7	0	3	2	0	0	76	2.204689	2.373105
7	0	4	3	0	0	78	2.149414	2.322261
7	0	5	4	0	0	80	2.145976	2.323255
7	0	6	5	0	0	82	2.119520	2.301231
7	0	7	6	0	0	84	2.100522	2.286665
7	0	8	7	0	0	86	2.098095	2.288670
7	0	9	8	0	0	88	2.087796	2.282803
7	0	10	9	0	0	90	2.076966	2.276406
7	0	9	7	0	0	89	2.076889	2.274112
7	0	9	6	0	0	91	2.072282	2.273937
7	0	9	5	0	0	94	2.070921	2.279224
7	0	9	4	0	0	98	2.067736	2.284903
7	0	10	6	0	0	96	2.062804	2.275539
8	0	9	6	0	0	101	2.062697	2.286512
9	0	9	6	0	0	111	2.064805	2.310780
7	1	9	6	1	0	113	2.062608	2.313015
7	10	9	6	0	0	103	2.069727	2.297974
7	8	9	6	0	0	95	2.069304	2.279823
7	9	9	6	0	0	99	2.070180	2.289563

Table 13: Diagnostics for the estimated model of the bivariate Series $\mathbf{v_t}$.

						Mean		Variance	
L_r	L_p	K_z	I_z	K_x	I_x	Wilks Λ	Prob.	Wilks Λ	Prob.
Raw Data						0.2557	.0000	0.0500	.0000
Adjusted data						0.5269	.0000	0.6576	.0000
7	0	9	6	0	0	0.9123	.0360	0.9571	.9999
7	8	9	6	0	0	0.9137	.0483	0.9572	.9999
7	9	9	6	0	0	0.9144	.0560	0.9569	.9999
7	10	9	6	0	0	0.9249	.3116	0.9567	.9998

Table 14: Short–term diagnostics for the bivariate series $\mathbf{v_t}$.

						Mean		Variance	
L_r	L_p	K_z	I_z	K_x	I_x	Wilks Λ	Prob.	Wilks Λ	Prob.
Raw data						0.74737833	.0001	0.69875647	.0001
Adjusted data						0.73885728	.0001	0.82521561	.0001
7	0	9	6	0	0	0.90955188	.2009	0.90496504	0.0984
7	8	9	6	0	0	0.90962413	.2030	0.90532303	0.1046
7	9	9	6	0	0	0.90927679	.1933	0.90521097	0.1026
7	10	9	6	0	0	0.91205138	.2807	0.90436815	0.0887

Table 15: Long–term diagnostics for the bivariate series $\mathbf{v_t}$.

The scaled residuals are regressed on dummy variables assigned to each year. Squared, scaled residuals are regressed on the annual dummies. In each case the Wilks test is again used to test for the significance of the regression. If either regression is significant, then the values of the tuning parameters are increased, and the model is re–estimated. This fine tuning process is continued until no regression accounts for a significant amount of variance at the 5% level. In order to satisfy the short- and long-term misspecification criteria we have had to augment the model. The diagnostics of the final model which is characterized by $(L_r, L_p, K_z, K_x, I_z, I_x)$ = (7,9,9,6,0,0) are presented in Tables 14 and 15.

In summary, the final model is adequate, exhibiting many features of modern financial markets. The derived conditional probability density function is extremely non-normal, as evidenced by the large value of K_z. The model is also markedly conditional heteroskedastic. The fact that $L_r = 7$ means that 7 lags are required in series $v_{2,t}$ and 6 lags in series $v_{1,t}$ to account for the conditional heteroskedasticity in the data. Finally, since $L_p = 9$, 9 lags are required in series $v_{2,t}$ and 8 lags in series $v_{1,t}$ to account for the autoregressive behavior in the data. The number of estimated parameters in the model is 99. Each data series contains 1675 points with 15 being used for initial lags in the model. Thus, there are 1660 points from each series used in the model construction. This results in a saturation ratio of approximately 34 observations per parameter. The fact that the model includes lagged values of $v_{1,t}$

up to lag 8 is the reason for setting the truncation lag, q, in $V_{n_j}(q)$ equal to 8 as discussed previously.

IV Analysis of Simulation Data

The primary motivation for modeling the bivariate price series has been to increase the power of the V statistic. Since the model does not take account of any cyclical behavior in the data, the model provides a null hypothesis against which to test for a long-term cycle. The testing procedure involves constructing confidence bands for the V statistic, using the SNP model as the null hypothesis. If the V statistic for the London data lies outside of the confidence bands, this would provide evidence for rejecting the SNP model, which does not take account of any long-term cycle in the data. Moreover, by using a model with specific short-range dependence we should be able to improve the power of the Lo test.

To implement the strategy we require simulations of stock prices on the London exchange. The simulation process involves two steps. First, we simulate data using the SNP model. The SNP code does this using the rejection method to generate random samples from the estimated probability distribution. [7] Then, we adjust the simulation data using the inverse transformation for removing systematic biases (See Table 11.). We have generated 1000 simulations of prices on the London exchange. Each simulation contains 1675 points, the first 15 of which are the first 15 points of the London data. The reason for this is that the SNP model is based on points 16 and beyond. The first 15 points were dropped in anticipation of retaining some data for initial values in the model. This means that the procedure is conservative. Specifically, we are less likely to reject our null hypothesis since each simulation contains prices from the London data series. In Figures 14 and 15 we present typical examples of simulations derived from the SNP model.

We have calculated the V statistic, i. e. $V_{1774}(8)$, for each simulation and obtained fractiles of its limiting distribution for the SNP model. The fractiles are given in Table 16. These values are compared to those of Table 1. It appears that fractiles presented in Table 1 may be positively biased with respect to those derived from the SNP model. Using Table 16 we find by interpolation that the probability of obtaining the value of $V_{1774}(8) = .8783$ for the London data to be .073. This value is slightly less significant than the value .063 obtained by interploating the data in Table 1. It seems that the power of the Lo test based on the SNP model is not increased appreciably from that of the standard Lo test. However, if there is any increase, it results in a slightly less significance level for the existence of a long–term cycle.

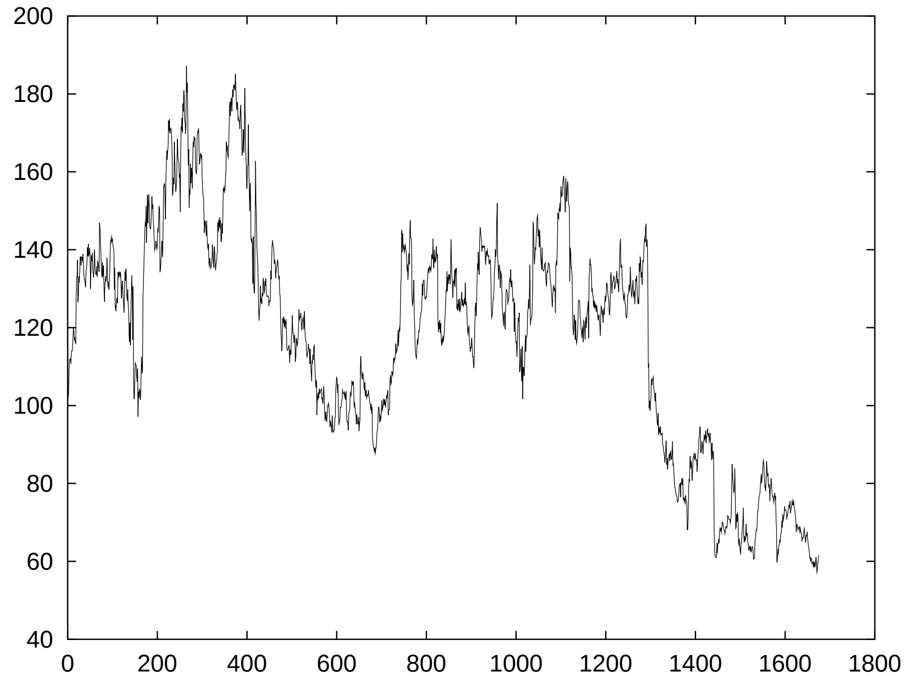

Figure 14: Simulation of London prices using the SNP model.

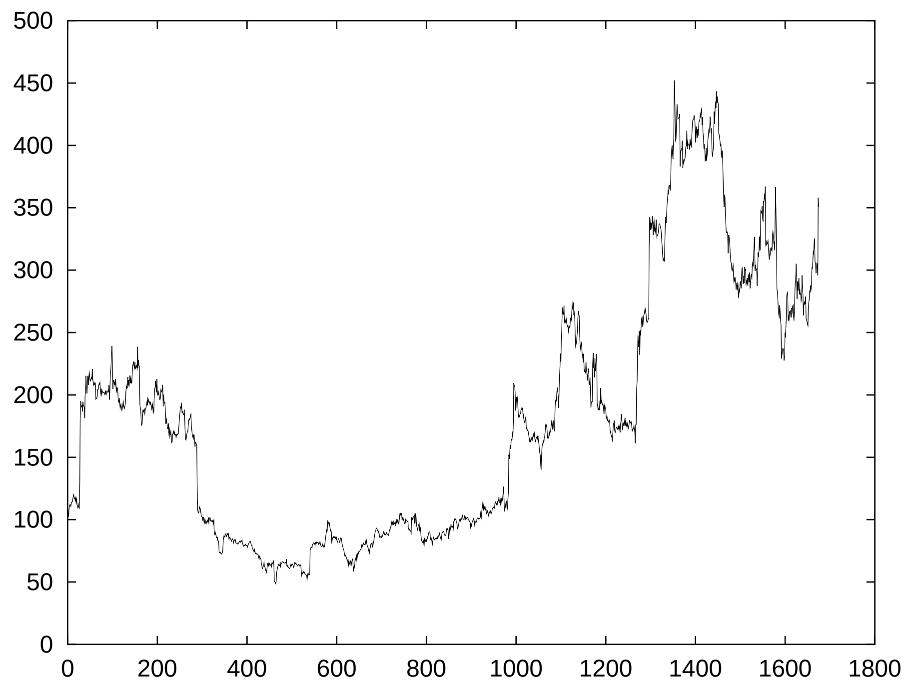

Figure 15: Simulation of London prices using the SNP model.

Prob($V < v$)	.005	.0250	.050	.100	.200	.300	.400	.500
v	0.732	0.795	0.845	0.916	1.006	1.069	1.137	1.197
Prob($V < v$)	.543	.600	.700	.800	.900	.950	.975	.995
v	1.231	1.268	1.345	1.442	1.576	1.715	1.807	1.992

Table 16: Fractiles of the limiting distribution of the V statistic under the assumption of the SNP model. The probability of the V statistic which has been obtained for the London data ($V_{1774}(8) = .878$) is .073.

V Conclusions

We have analyzed the bivariate time series of bi-monthly stock prices of the South Sea Trading Company as listed on the London and Amsterdam exchanges during the period 1723 through 1794. The prices quoted on the Amsterdam exchange are for a forward settlement of three weeks so that these prices are more similar to those of a futures contract than a stock. The bivariate series therefore has characteristics similar to a cash and futures price series. Both price series are non-stationary, each being integrated of order one. The two series are cointegrated, the cointegrating relationship being the difference between the two series. This is expected since the stock price on one exchange should be a proxy for the stock price on the other exchange. Any effects due to interest rate differences between the two countries should be negligible because of the fixed, short time period until settlement on the Amsterdam exchange. In addition, the data reveal that the introduction of financial legislation (Barnard's Act) that restricted forward trading activity in London increased the level of volatility in the market.

A graph of the data shows what appears to be a relatively long-term cycle. Using spectral analysis we have searched for the presence of cycles in the data listed on the London exchange. We have found none. As an alternative to spectral analysis we have applied a methodology based on the V statistic used in rescaled range analysis and found strong evidence for a long-term cycle of approximately 20 years. Based on the value of the calculated V statistic we can reject the hypothesis of no long-term cycle at the 6% confidence level, approximately. To increase the power of this technique we have constructed a model for the bivariate price series. The model provides a null hypothesis against which to apply the rescaled range analysis. The model is constructed using the seminonparametric methodology of Gallant and Tauchen, but modified to include the cointegration between the two series. The model does not take account, however, of the long-term periodicity in the data. Based on the model we conclude that the data are markedly conditionally heteroskedastic and that the conditional probability density fit to the data is extremely non-normal. Using the model to generate simulations of the London data we have constructed fractiles for the V statistic. The fractiles corresponding to the Lo test appear to be somewhat positively biased with respect to those derived from the model. According to fractiles derived from the model the value of the V statistic for the original price series occurs with a probability of approximately 7% or greater. This result is somewhat less significant than the result obained using the standard Lo test. The rescaled range analysis based on the simulation data, nonetheless, supports the existence of a long-term cycle in the London data.

The fact that our analysis provides strong evidence for the presence of a long-term cycle in these old data is motivation to apply these techniques to modern financial data. Indeed, Peters has reported evidence for a four year cycle in Dow Jones Industrial Average. Do long term cycles exist in futures prices? If uncovered, can profitable trading strategies be constructed that take advantage of this information? Such cycles could be related to the business cycle or even long-term weather cycles. Perhaps the futures markets have already discounted such cycles, perhaps not. Clearly these are questions which require further investigation.

A Appendix

A.1 Adjustment for Systematic Effects

Let $\hat{r}_t$ be a univariate time series. Assume that the time series has been divided into groups which are suspected of exhibiting systematic differences. We describe a procedure for analyzing and adjusting for systematic differences in the means or variances of the groups [6]. First, dummy variables are assigned to each group. Then $\hat{r}_t$ is regressed on the dummy variables denoted by the vector $\mathbf{g}$,

$$\hat{r}_t = \mathbf{g} \cdot \boldsymbol{\beta} + u_t \,. \tag{15}$$

When some regression coefficients, β^i, are found not to be significantly different from zero, the groups corresponding these coefficients are combined with other groups, as appropriate. The regression is then rerun. This process is continued until the number of groups has been reduced, and all regression coefficients are significantly different from zero.

Next, a variance adjustment $\mathbf{g} \cdot \boldsymbol{\gamma}$ is obtained by regressing $\log(u_t^2)$ on $\mathbf{g}$,

$$\log(u_t^2) = \mathbf{g} \cdot \boldsymbol{\gamma} + \epsilon_t \,. \tag{16}$$

If some regression coefficients, γ^i, are found not to be significant, some groups are combined and the regression is rerun as described previously. The series $\hat{r}_t$ is then adjusted so that

$$\tilde{r}_t = \frac{\hat{r}_t - \mathbf{g} \cdot \boldsymbol{\beta}}{\exp(\mathbf{g} \cdot \gamma/2)} \,. \tag{17}$$

Finally, each $\tilde{r}_t$ is transformed

$$r_t = a + b\tilde{r}_t \,, \tag{18}$$

where a and b are selected so that the mean and variance of r_t is the same as the unadjusted series $\hat{r}_t$.

References

[1] J. Beran. *Statistics for Long Memory Process.* Chapman and Hall, New York, 1994.

[2] T. Bollerslev, R. T. Chou, and K. F. Koner. Arch modeling in finance: A review of the theory and empirical evidence. Working Paper No. 97, Northwestern University, November 1990.

[3] R. Davies and D. Harte. Tests for the hurst effect. *Biometrika*, 74:95, 1987.

[4] E. F. Fama. The behavior of stock market prices. *Journal of Business*, 38:34–105, 1965.

[5] W. A. Fuller. *Introduction to Statistical Time Series.* John Wiley and Sons, Inc., New York, 1976. See page 283.

[6] A. R. Gallant, P. E. Rossi, and G. Tauchen. Stock prices and volume. Working Paper, Department of Statistics, North Carolina State University, January 1990.

[7] A. R. Gallant and G. Tauchen. Nonlinear time series analysis: Estimation and simulation. Working Paper, Department of Statistics, North Carolina State University, June 1990.

[8] H. E. Hurst. The long-term storage capacity of reservoirs. *Transactions of the American Society of Civil Engineers*, 116, 1951.

[9] SAS Institute Inc. *SAS/STAT Users Guide, Version 6, Fourth Edition, Volume 2.* Cary, NC: SAS Institute Inc., 1989.

[10] C. Y. Kramer. Extension of multiple range tests to groups means with unequal numbers of replications. *Biometrics*, 12:307–310, 1956.

[11] A. W. Lo. Long-term memory in stock market prices. NBER Working Paper # 2984, Sloan Shool of Management, Massachusetts Institute of Technology, May 1989.

[12] B. Mandelbrot. Statistical methodology for non-periodic cycles: From the covariance to R/S analysis. *Annals of Economic and Social Measurement*, 1:259, 1972.

[13] B. B. Mandelbrot and J. W. Van Ness. Fractional brownian motions, fractional noises and applications. *SIAM Review*, 10(4):422, October 1968.

[14] B. B. Mandelbrot and J. R. Wallis. Computer experiments with fractional gaussian noises. part 1, averages and variances. *Water Resources Research*, 5(1):229, February 1969.

[15] B. B. Mandelbrot and J. R. Wallis. Computer experiments with fractional gaussian noises. part 1, mathematical appendix. *Water Resources Research*, 5(1):260, February 1969.

[16] B. B. Mandelbrot and J. R. Wallis. Computer experiments with fractional gaussian noises. part 2, rescaled ranges and spectra. *Water Resources Research*, 5(1):242, February 1969.

[17] D. F. Morrison. *Multivariate Statistical Methods*, page 222. McGraw-Hill, New York, second edition, 1976.

[18] L. Neal. *The Rise of Financial Capitalism.* Cambrige University Press, New York, 1990.

[19] P. Perron. Trends and random walks in macroeconomic time series. *Journal of Economic Dynamics and Control*, 12:297–332, 1988.

[20] E. E. Peters. *Fractal Market Analysis.* John Wiley and Sons, Inc., New York, 1994.

[21] J. R. Tukey. The problem of multiple comparisons. Unpublished, 1953.