The Compass Rose Pattern in Electricity Prices

Jonathan A. Batten and Mahmoud Hamada

Jonathan A. Batten
Department of Finance
Hong Kong University of Science & Technology
Clear Water Bay, Kowloon, Hong Kong
Tel: ++852-2358 8202
Fax: ++852-2358 1749
Email: jabatten@ust.hk

Graduate School of Management,
Macquarie University
CBD Campus Level 6, 51-57 Pitt St
Sydney, NSW 2000, Australia
Tel: ++61-2-8274-8344,
Fax: ++61-2-8274-8370
Email: jonathan.batten@mgsm.edu.au

Mahmoud Hamada
Energy Risk Management
Energy Australia
570 George St
Sydney, NSW 2000, Australia

31st October 2007

JEL: C22; C32; E31; F30; G15

Keywords: chaos; compass rose; electricity prices; fractal; high-frequency data; Hurst

Abstract

The compass rose pattern appears in high-frequency returns of spot electricity prices. Once these returns are filtered using autoregressive filtering, no pattern remains. This is unexpected and suggests that factors other than discreteness contribute to the compass rose pattern. Though the series are non-normal in terms of their distribution, statistical tests fail to identify significant chaos, or fractal structures, in price returns. The phase diagram of the filtered returns provides a useful visual check on independence, a property necessary for pricing and trading derivatives and portfolio construction, as well as providing useful insights into the market dynamics.

The Compass Rose Pattern in Electricity Prices

1. Introduction

The "Compass Rose" pattern was originally identified in stock returns by Crack and Ledoit (1996) and appears in phase portraits or scatter diagrams as rays emanating from a centre when returns are plotted against their first order lagged value (in a two dimensional space), or first and second order lagged values (in a three dimensional space). While Crack and Ledoit (1996) recognized that this pattern was an artefact of market microstructure – due to price clustering, discreteness (Kramer and Runde, 1997; Wang and Wang, 2002; Mitchell and McKenzie, 2006), tick size (Szpiro, 1998) and data frequency (Wang, Hudson and Keasey, 2000; Cai, Hudson and Keasey, 2003; Lee, Mathur and Gleason, 2005), technical traders, amongst others, have hoped that this technique could be used to identify deterministic structure in returns, which could then be used to improve predictability and provide abnormal trading returns.

We undertake an experiment from the alternate perspective. Rather than attempt to identify a nonlinear, or deterministic structure in a financial series, we investigate a return series well known for having such a structure: the spot electricity market. For example, Bessembinder and Lemmon (2002) and Longstaff and Wang (2004) identify daily, weekly and annual cycles in spot and forward electricity prices. It is not surprising that when phase portraits of electricity price returns are plotted, a compass rose is revealed. Second, we then model the spot returns using autoregressive techniques. These and other models have been shown to successfully capture serial correlation (autocorrelation) for use in forecasting future spot prices (Bessembinder and Lemmon, 2002). Finally, we investigate the residuals from this modelling for the presence of a compass rose. Having removed linear dependence through autoregressive filtering of returns we then compare the two series and find that the compass rose, evident in the unfiltered series, has now disappeared. Various statistical tests are conducted to explain this finding, including the test for independence described by Brock, Dechert, Scheinkman and Le Baron (1996), which highlight differences in randomness in the two series.

The paper is set out as follows. The next section (2) describes the data and modelling method used, section three provides the results and the final section allows for some concluding remarks.

2. Data and Method.

The financial series investigated are returns on half-hourly prices in the Australian spot electricity market, termed the National Electricity Market (NEM). Prices are traded in markets as an amount of Australian dollars (1A$ = 0.88 US$) per megawatt hour. The NEM started operation on the 13 December 1998 and now comprises the electricity industries in the major Australian states of Queensland, New South Whales, Victoria, South Australia and Tasmania, which together form one competitive wholesale pool. In 2006, the NEM oversaw the electricity production, network transmission and trading functions of 180,000 Giga Watt hours (GWh) per annum, an industry with annual turnover of over A$6 billion. This comprises private corporations in Victoria and South Australia, and state-owned corporations in NSW and Queensland. The NEM is governed by the National Electricity Code which, among other goals, ensures competition among power suppliers in the multi-state service area in an effort to reduce overall energy costs for consumers and end users. The National Electricity Market Managing Company (NEMMCO) is in charge of operating and administering the market according to the code (see AFMA, 2005).

Our sample data comprises 30-minute spot prices per megawatt hour (MWh) of electricity from the state of New South Wales, in Australian dollars (A$), from the 1st July 2003 to 4th July 2006, over the entire 24-hour day. There were a total of 52,752 observations. The mean price over the sample period was A$36.33 per MWh with the lowest price of A$4.59 and the highest price of A$9,909.03 per MWh. The economic implication of this range in price is significant and justifies significant resources directed towards hedging and risk management (Weron, 2000; Chang and Park, 2007). Figure 1 shows a time series plot of the mean spot electricity price (P_t) over the 24-hour day. The average price per MWh peaks at A$105.7 in the afternoon around 16:00, then rises again to A$71.2 around 18:30pm. The lowest prices occur during the early hours of the morning, with the average lowest price (A$15.2) occurring at 4:00am.

(Insert Figure 1)

We first estimate the intraday 30-minute returns (ΔP_t) for the spot electricity price (P_t) allowing

$$\Delta P_t = \log(P_t) - \log(P_{t-1}) \qquad (1)$$

which is common in financial time series analysis, where the interval $t\text{-}1 \rightarrow t$, is 30-minutes. Over the entire sample period the mean of ΔP_t was 0.00001 and indistinguishable from zero at the 95% confidence interval. The series returns had a standard deviation of 0.07723 with slight positive skewness (0.122) and significant kurtosis (177.51) and consequently failed standard tests for normality (Anderson-Darling test was 3.642, $p < 0.005$).

(Insert Figure 2)

As expected from the literature, the returns show a regular cycle across the 24-hour day with significant variation between the forty-eight, 30-minute intervals. An *F-test* for variation in these 48 averages was 192.11 (*p-value* 0.000) with a pooled standard deviation of 0.0714. A plot of the mean and standard deviation of 30-minute returns is provided in Figure 2. The minimum average return was -0.0741 (19:30pm) and the highest was 0.0932 (18:00pm), while the lowest standard deviation was 0.0244 (4:00am) and the highest was 0.2167 (18:00pm). This suggests the lowest variation in price occurs in the early morning and the highest in the late afternoon, early evening.

(Insert Figure 3a and 3b)

Time dependence in returns is clearly evident from a correlogram of returns, displayed in Figure 3a for 240 lags, a number corresponding to 5-days of 48 intervals. Note the significant declining positive autocorrelation of the five peaks at lags 48, 96, 144, 192 and 240 with values of 0.464 (*t-statistic* 94.25), 0.414 (*t-statistic* 68.57), 0.374 (*t-statistic* of 55.16), 0.352 (*t-statistic* 6.77) and 0.338 (*t-statistic* 43.04) respectively. This is consistent with the declining influence of previous day's returns on current day returns for a specific 30-minute return interval.

Figure 3b highlights the structure across 48 lags, corresponding to one day. Within each day the autocorrelation is initially positive at lag 1 (0.199, *t-statistic* = 45.74), negative at lag 3 (-0.091, *t-statistic* = -19.97) and thereafter increasing to -0.025 at lag 10 (*t-statistic* = -5.16). This autocorrelation pattern resembles an autoregressive process and is consistent with patterns identified by other researchers in other electricity markets (e.g. Rambharat et al. 2005). This feature of the time-series generally invalidates the usual assumption of independent identically distributed (iid) residuals. Also, the normality assumption for residuals does not hold because of the return spikes. However, when taking weekly averages, the spikes and the serial correlation are smoothed and the data can be modelled using conventional time series techniques. Weekly averaging also eliminates the presence of daily and weekly cycles.

Next we filter the short-term autocorrelated innovations in the return process via an autoregressive (AR) model. Of specific interest is the residual ψ_t after applying various filters (AR(1) →AR(10) to ΔP_t given the autoregressive behaviour evident in the correlogram (Figures 3a and 3b). This model has the following mathematical form, with the general autoregressive term $\beta_n \Delta P_{t-n}$ = AR(n):

$$\Delta P_t = \alpha_0 + \beta_1 \Delta P_{t-1} + \beta_2 \Delta P_{t-2} + \beta_3 \Delta P_{t-3} ......\beta_{10} \Delta P_{t-10} + \psi_t \qquad (2)$$

(Insert Table 1 about here)

Table 1 provides the coefficients for Equation 2 in the spot electricity returns. Note that the coefficient α_0 = 0.0000 (*t-statistic* = 0.016), which is statistically and economically insignificant. The two most significant terms are β_1 = 0.1956 (*t-statistic* = 44.965) and β_3 = -

0.102 (*t-statistic* = -22.960) being the prior 30-minute and 90 minute time periods. To model electricity prices it would be necessary to apply longer, lagged moving average terms (for example one–day lagged residuals, ψ_{t-48}) to remove the effects of longer time dependence (due to regularity in demand based upon the time of day shown in Figure 1). However, our interest is on the lower order price dependence as revealed in the short-term autocorrelation structure. It is this short-term structure that will influence the presence of the compass-rose pattern in a phase portrait of returns.

(Insert Figure 4a, 4b and 4c)

3. The Compass Rose in Electricity Prices

A three dimensional phase portrait, or scatter plot, of electricity returns (ΔP_t) plotted against first (ΔP_{t-1}) and second (ΔP_{t-2}) order lagged values is provided in Figure 4c. Note the rays emanating from the centre of the cube clearly form a star or compass rose pattern. This is clearer in Figures 4a and 4b which shows a two dimensional plot of electricity returns (ΔP_t) plotted against the first lag (ΔP_{t-1}). Figure 4a and 4b plots the range of returns from 0.00 to ± 0.50, while 4b plots a smaller range of returns from 0.00 to ± 0.15). Though clear in both figures, the compass rose pattern is very much a function of both the scale and the symbols used to record each observation: if both the scale and symbols are too large, then the compass rose pattern is lost. As Kramer and Runde (1997) note, the former is due to the relative relationship between the changes in price to the price. Plotting the change in electricity price to price level (not reproduced here) also highlights the clustering of prices in the area between zero and A$200. The importance of price clustering and its impact on asset pricing has long been a concern for market researchers (e.g. Osbourne, 1962; Niederhoffer, 1965; Antoniou and Vorlow, 2005).

The economic interpretation of the compass rose pattern provides an interesting insight into the electricity market dynamics. First, the most obvious of the compass rose lines occur when x and y = 0 (i.e. $\Delta P_t = 0$ and $\Delta P_{t-1} = 0$, termed no-change prices) and when x = y (i.e. $\Delta P_t = \Delta P_{t-1}$, when the change is the same). Further analysis using a two-way cross-tabulation of the frequency of no-change and same prices versus half-hourly time highlights statistically significant daily variation in no-change and same prices (for example, the chi-square test for the no-change association was 715.11 with 47 degrees of freedom and *p*-value = 0.000). More specifically the greatest frequency of both no-change and same prices (a total of 1599 and 1176 observations respectively) occurs between 1200-430pm, while the least number is from 600am to 900am and then 600pm to 1030pm. This is linked to the daily demand cycle clearly visible in Figure 1.

Attention should also be drawn to the line with a negative parabolic form which declines through the origin in the two dimensional plot in Figure 4b (and a diagonal line in 4c due to the smaller scale). The pairs in the top-left quadrant indicate negative returns followed by positive returns in the next period (half hour). Having them clustered in this manner suggests that when prices fall below a certain level there is an immediate systematic response from some of the generators in the state (NSW) that would withdraw capacity, say 100MW from the market to decrease supply so that prices then rise as a result of the law of supply and demand.

For the pairs in bottom-right quadrant, this is the case of positive returns followed by negative ones. So, as soon as prices rise to a certain level, there is an immediate systematic capacity offered by the generator (100MW) to take advantage of high prices. This action of over supply then pushes prices down as a consequence. The curve is not linear because the effect of withdrawing, or supplying, of the systematic capacity (100MW) is not linear on price. So this may suggest inefficiency in the electricity market, and an arbitrage opportunity for some generators.

Recall that standard statistical tests reject normality in the sampled electricity returns largely due to the presence of positive kurtosis. From Equation 2, a series of residuals ψ_t may be obtained. While now lacking the low order autocorrelation structure of the original series, the ψ_t series remains non-normal, with high kurtosis (180.14) and slight skewness (0.536) and a slightly smaller standard deviation (0.0749). The mean remains close to zero. A sum of the returns from the initial series totals 0.3288, indicative of a positive price drift in the return series, whereas the filtered series has a sum of close to zero at four decimal places, which is more consistent with a random series without drift.

(Insert Figure 5a, 5b and 5c)

The conclusion of various authors, including Szpiro (1998) and Wang and Wang (2002), is that price discreteness (for example, the electricity price quotes are to two decimal places) was the sole and necessary condition for the appearance of the compass rose. This conclusion is challenged by the plots of the residuals ψ_t against first and second order lags in two and three dimensional phase portraits. These are illustrated in Figures 5a, 5b and 5c. What is immediately apparent is the disappearance of the compass rose in all figures. Note that care has been exercised to ensure that both the symbols for the observations and the scale are the same as those used in Figures 4a, 4b and 4c. Reducing the scale of the axes does not change this result and despite the discreteness of the series the compass rose pattern has now been removed. This is also clear from repeating the two-way cross-tabulation of the frequency of no-change and same prices versus half-hourly time on the filtered returns. Now, the number of observations along the major compass rose lines when $\Delta P_t = 0$ and $\Delta P_{t-1} = 0$ is 80 and $\Delta P_t = \Delta P_{t-1}$ is 75 (instead of 1599 and 1176 respectively). Fundamentally, though non-normal in terms of its distribution, the ψ_t series appears now to be more random, and more likely iid compared with the unfiltered returns.

Non randomness in the returns may be verified by estimating the test statistic described by Brock, Dechert, Scheinkman and LeBaron (1996), henceforth termed the BDSL statistic. This test is commonly used in time-series analysis and may also be used for checking the null hypothesis of white noise (iid process) against the unspecified alternative, which may suggest deterministic chaos, or a stochastic nonlinear process (see Antoniou Vorlow, 2005).

Under the assumption of independence, the test statistic lies close to zero. To estimate the BDSL statistic, it is necessary to specify the embedding dimension, or number of consecutive data points used - generally a number from two to five for long time series data - and a threshold term ranging from 0.5 to 2. For large samples distribution, the BDSL statistic is approximated by an asymptotic normal distribution, with a value of greater than 1.96 signifying a rejection of the null that the data are iid at the 5% level of significance (see McKenzie, 2001 for a discussion.)

(Insert Table 2)

The results from this analysis for both series are provided in Table 2. To simplify presentation, various embedding dimension values from 2 to 6, with a single threshold value of 0.7 are shown. In all cases the *z*-statistic fails to reject the null hypothesis that the observations of the series are independent, although the filtered series has values closer to zero (and therefore more "random") than the raw return series. The test coefficient also increases as the embedding dimension increases. Note that Crack and Ledoit (1996) and Mitchell and McKenzie (2006) suggest the presence of the compass rose pattern may distort the null distribution of the BDSL test for nonlinearity. Specifically, Kramer and Runde (1997) empirically test this proposition and find that the null distribution of the BDSL test is distorted in the presence of discreteness in the time series.

An alternate approach is to determine the presence of long-term dependence in the time-series of electricity price returns. Long-term dependence is frequently linked to the presence of fractal structures, which may be measured using the rescaled adjusted range technique of Hurst (1951). This technique yields an exponent (H), which under the assumption that the series follows a Gaussian random walk equals 0.5. In turn the failure to identify long-term dependent effects also lends support to the proposition that the time series investigated conforms to normally distributed, standard Brownian motion. The statistical method employed for measuring long-term dependence in this study is based on the adjusted local measure of the rescaled range of Mandelbroot and Wallis (1969) and Lo (1991), as adopted by Batten and Ellis (2005).

We begin by taking the unfiltered return ΔP_t or ψ_t, from the autoregressive equation (Equation 2). For simplicity in the following exposition we simply specify ψ_t (ΔP_t can be substituted). Thus, for each ψ_t the classical rescaled adjusted range $(R/\sigma)_n$ is calculated as

$$(R/\sigma)_n = (1/\sigma_n)\left[\underset{1 \le k \le n}{\text{Max}} \sum_{j=1}^{k} (X_j - \overline{X_n}) - \underset{1 \le k \le n}{\text{Min}} \sum_{j=1}^{k} (X_j - \overline{X_n}) \right] \quad (3)$$

where $\overline{X_n}$ is the sample mean $(1/n)\Sigma_j X_j$ of a ψ_t nd σ_n is the standard deviation of ψ_t over a particular series n

$$\sigma_n = \left[1/n \sum_{j=1}^{n} (X_j - \overline{X_n})^2 \right]^{0.5} \quad (4)$$

In order to capture the time-varying nature of dependence in ψ_t (the effects of daily and weekly cycles) a local measure of the Hurst exponent (h) is employed. This exponent is calculated as

$$h_n = \frac{\log(R/\sigma)_n}{\log n} \tag{5}$$

Under the null hypothesis of no long-term dependence, the value of $h_n = 0.5$. For time-series exhibiting positive long-term dependence, the observed value of the exponent $h_n > 0.5$. Time-series containing negative dependence are alternatively characterised by $h_n < 0.5$. Importantly for equilibrium reverting processes the local Hurst exponent should be negative, since a movement back towards the equilibrium should follow a movement away from equilibrium. For positively dependent processes another movement further away from equilibrium will follow the earlier movement away from equilibrium.

Estimates of the local Hurst exponent are calculated for (N – n + 1) times overlapping subseries of length n, with n having a set value. In this case, n is arbitrarily set to either 24 (equivalent to 12 hours), 48 (equivalent to 1 day), and then 240 (equivalent to 5-days). The procedure in effect creates a time-series of exponent values.

(Insert Table 3)

The local Hurst statistics for the unfiltered electricity returns (ΔP_t) and those returns filtered using the autoregressive model of Equation 2 (ψ_t,) are presented in Table 3. Under the null hypothesis of no long-term dependence, $h_n = 0.5$. What is significant about these results is the importance of the sample length in estimating h. When n is less than one day (e.g. n = 24) $h_{24} > 0.5$ for the 95% confidence interval. In this instance the sample average h_{24}= 0.79926 for the unfiltered and h_{24}= 0.79853 for the filtered series. This result suggests a positively dependent return process where intraday positive price changes tend to follow positive and vice versa.

On the other hand $h_{48} < 0.5$ for the 95% confidence interval, is suggestive of a mean reverting process with the sample average h_{48}= 0.49057 for the unfiltered and h_{24}= 0.49534 for the filtered series. That is for the one day sample length, positive price changes follow negative and vice versa. As the sample length becomes longer than one day it becomes less distinguishable from a random process. Thus for one week samples (h_{240}) the process is only slightly positively dependent and barely indistinguishable at the 95% level from a random process (the sample average h_{240} = 0.51250 for the unfiltered and h_{240} = 0.51391 for the filtered series). Filtering the series using autoregressive techniques removes the effects of short term serial or autocorrelation and moves the series closer to the random average of 0.5. However for n = 240, the h-statistic is less random once filtered. Overall, at longer sample lengths, the Hurst local mean is very close to 0.5000 which means there is no significant fractal structure, although since $h \neq 0.5000$ at the 95% level of confidence, we cannot accept that the series is random. This appears to be a more sensitive test for randomness than the BDSL statistic (which was unable to distinguish non-randomness estimated over the entire sample length).

To verify the h-statistic finding of non-randomness a non-parametric runs test was also conducted. In the case of the filtered series the observed number of runs O = 33,796, while the number of positive observations N_{+}= 16,937 and the number of negative observations N_- = 35,805. The expected number of runs μ= 22,997, while the variance of runs σ^2 = 10,026. The key *Z*-statistic estimated as $Z = (O - \mu)/\sigma$ = 107.85 rejects the null hypothesis of randomness. This also was the case for the unfiltered series.

4. Conclusions

This study began by discussing the presence of deterministic structure in the returns of spot electricity prices, due to the cyclic effects of demand and supply. This is evident in both low and high order correlograms, and the presence of a compass rose in a phase diagram of various lagged electricity returns. The presence of clear autoregressive behaviour allows the filtering of returns using conventional autoregressive modeling. We then compare the residual series from this modeling with the original series. Although both series remain non-normal, the residual series no longer displays –irrespective of the range of returns - the compass rose pattern in either a two, or three dimensional, phase diagram. This result is unexpected given the conclusions of Wang and Wang (2002) that discreteness is a necessary condition for the appearance of the compass rose.

The lack of a compass rose pattern is consistent with the short-term deterministic component in returns being successfully removed through autoregressive filtering, although higher order weekly and seasonal cycles remain. However, further tests of series independence (the BDSL and Hurst tests) favour independence and the absence of any significant unspecified structure, although the series remain both non-normal and slightly non-random by various measures. While the limitations of the BDSL test in the presence of discreteness have been well documented what has not been pointed out in the literature is the usefulness of the phase diagram as a visual check on independence. The use of this procedure lies beyond the electricity price example presented here and has important implications for asset pricing more generally, which relies on the key assumption of independence.

Table 1
AR(10) Filtering of Spot Electricity Returns

This Table reports the coefficients, standard error, t-statistics and probability of the regression $\Delta P_t = \alpha_0 + \beta_1 \Delta P_{t-1} + \beta_2 \Delta P_{t-2} + \beta_3 \Delta P_{t-3} ......\beta_{10} \Delta P_{t-10} + \psi_t$ where $\beta_n \Delta P_{t-n}$ is the general autoregressive term AR(n) for AR(1) →AR(10). N= 52,752

Variable	Coefficient	Standard Error	t-statistic	Probability
α	0.000000	0.000287	0.027784	0.9778
AR(1)	0.195666	0.004351	44.96559	0.0000
AR(2)	-0.019001	0.004433	-4.285737	0.0000
AR(3)	-0.101787	0.004433	-22.96039	0.0000
AR(4)	0.019941	0.004451	4.480292	0.0000
AR(5)	-0.077062	0.004449	-17.32083	0.0000
AR(6)	-0.034891	0.004449	-7.842232	0.0000
AR(7)	-0.045824	0.004451	-10.29541	0.0000
AR(8)	-0.022846	0.004433	-5.153308	0.0000
AR(9)	-0.016591	0.004434	-3.742040	0.0002
AR(10)	-0.033832	0.004352	-7.774538	0.0000

Table 2
BDS Statistics for the Spot Electricity Returns

This Table reports the BDSL (1996) test for independence of the spot electricity price returns and those filtered using the AR(10) model from Equation 2. The embedding dimension has values from 2 to 6. The test differentiates independent and identically distributed (iid) processes from deterministic chaos or stochastic nonlinear models. In all cases the z-statistic fails to reject the null hypothesis of independence.

Spot Electricity Returns				
Embedding Dimension	BDS Statistic	Standard Error	z-Statistic	Probability
2	0.044105	0.000496	88.92691	0.0000
3	0.080964	0.000790	102.4892	0.0000
4	0.104381	0.000943	110.6607	0.0000
5	0.114892	0.000986	116.5209	0.0000
6	0.116726	0.000954	122.3749	0.0000
AR(10) Filtered Electricity Returns				
2	0.041014	0.000475	86.30076	0.0000
3	0.075037	0.000755	99.43190	0.0000
4	0.096181	0.000898	107.0740	0.0000
5	0.105850	0.000936	113.0841	0.0000
6	0.107085	0.000903	118.6414	0.0000

Table 3
Local Hurst Statistics for the Spot Electricity Returns
This Table reports the average local Hurst exponent test for the unfiltered electricity returns and those returns filtered using the autoregressive model of Equation 2. Under the null hypothesis of no long-term dependence $h_n = 0.5$. Positive long-term dependence exists when the observed value of the exponent $h_n > 0.5$, while negative dependence is alternatively characterised by $h_n < 0.5$. Note that when n is less than one day (n= 24) $h_{24} > 0.5$, which suggests a positively dependent processes (positive price changes tend to follow positive and vice versa), while for $h_{48} < 0.5$, the process is equilibrium reverting (positive price changes follow negative and vice versa). For longer n (h_{240}) the process is slightly positively dependent although at a level barely distinguishable from a random process.

	Mean estimation of local Hurst (h)		
Sample	n = 24	n = 48	n = 240
ΔP_t	0.79926	0.49057	0.51250
ψ_t	0.79853	0.495343	0.51391
95% confidence interval			
ΔP_t	0.79786-0.80065	0.49229-0.49457	0.51197-0.51355
ψ_t	0.79713-0.79993	0.48943-0.49171	0.51312-0.51471

Figure 1
Time-series plot of the average spot Australian electricity prices across the 24-hour day.

The time series plot shows the mean spot electricity price (P_t), every 30-minutes over the 24-hour day. The prices are per megawatt hour (MWh) of electricity in Australian dollars (A$). The sample period is from the 1st July 2003 to 4th July 2006. The average price peaks at A$105.7 in the afternoon around 16:00, then rises again to A$71.2 around 18:30pm. The lowest prices are during the early hours of the morning, with the average lowest price (A$15.2) occurring at 4:00am.

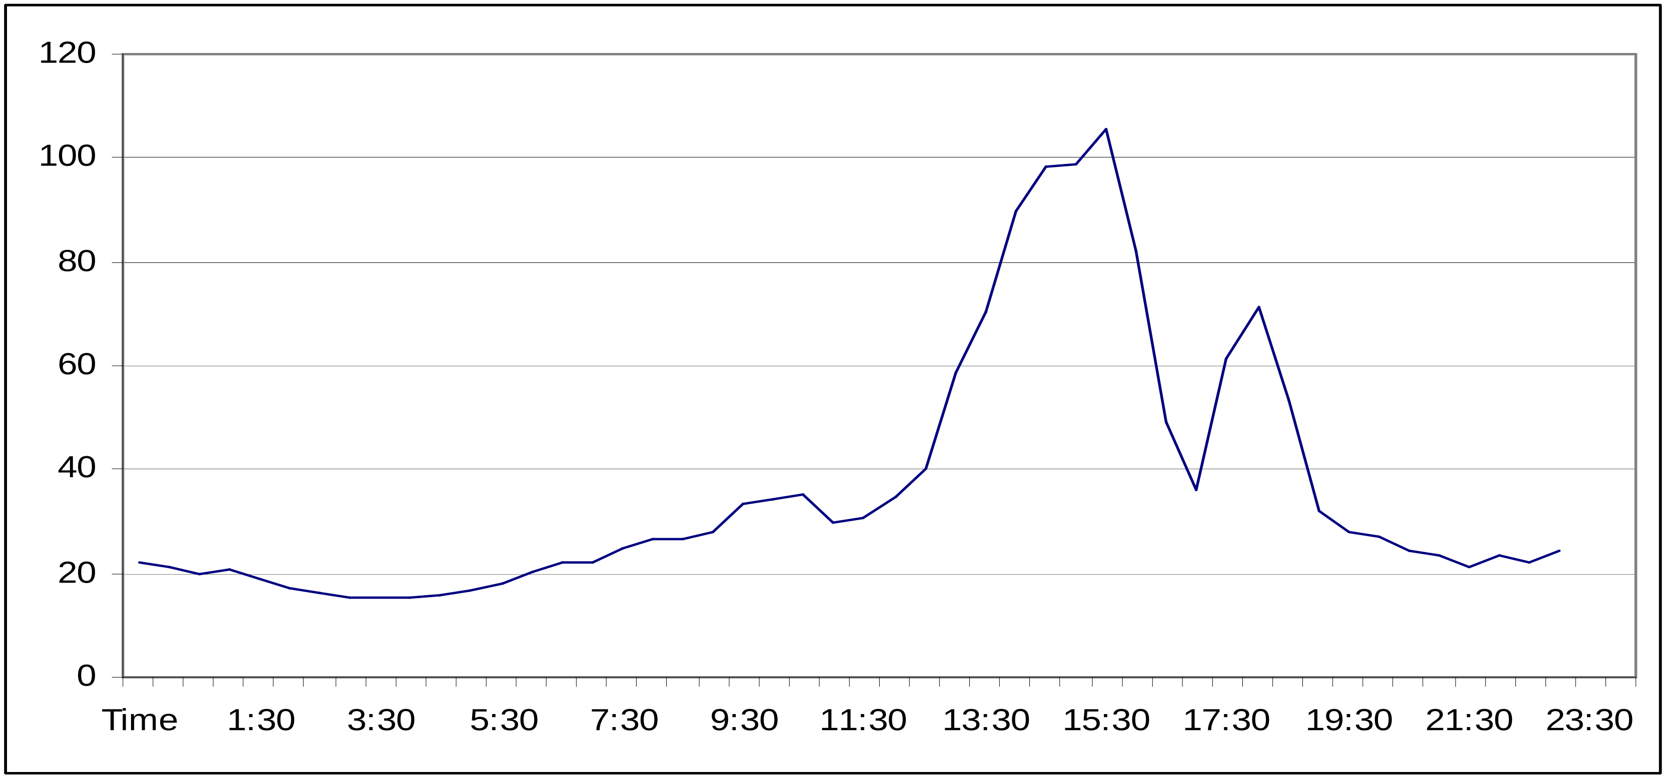

Figure 2

Time-series plot of the average 30-minute returns and standard deviation of spot Australian electricity prices across the 24-hour day.

The time series plot shows the mean and standard deviation of 30-minute spot price returns (ΔP_t) for the spot electricity price (P_t) where $\Delta P_t = log(P_t) - log(P_{t-1})$, where the interval $t\text{-}1 \rightarrow t$, is 30-minutes. The prices are per megawatt hour (MWh) of electricity over the 24-hour day. The sample period is from the 1st July 2003 to 4th July 2006. The mean is the bottom bold line and the standard deviation is the top dashed line. The average price return peaks at 17:30.

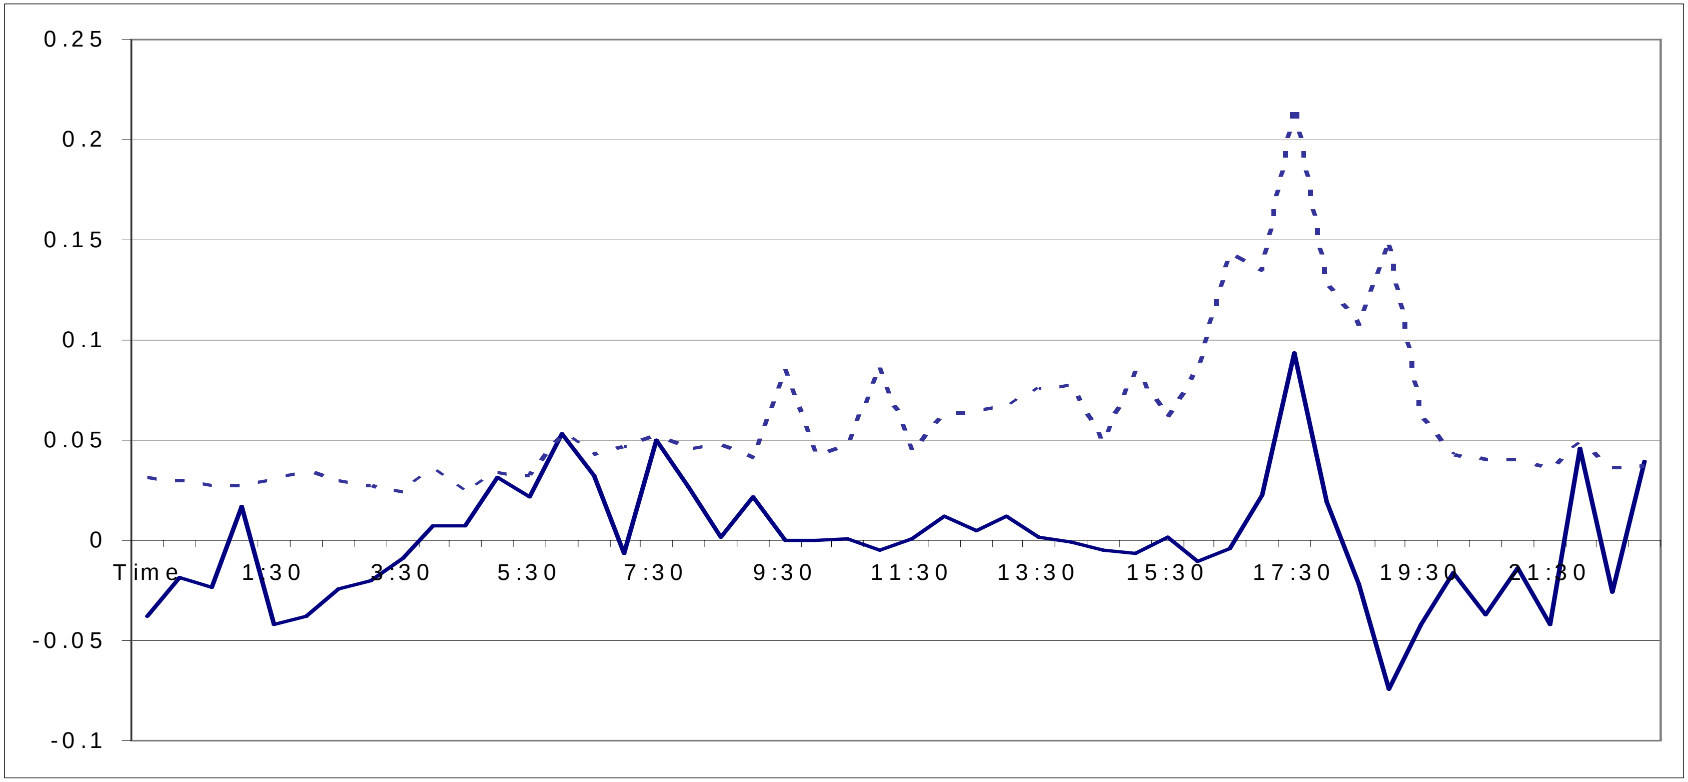

Figure 3a
Correlogram plot of the autocorrelation in electricity price returns over 240 lags

This plot shows the autocorrelation of 30-minute spot price returns (ΔP_t) for the spot electricity price (P_t) where $\Delta P_t = log(P_t) - log(P_{t-1})$ and the interval $t\text{-}1 \rightarrow t$, is 30-minutes, over 240 lags (corresponding to 5 days). The prices are per megawatt hour (MWh) of electricity over the 24-hour day. The sample period is from the 1st July 2003 to 4th July 2006. N= 52,752.

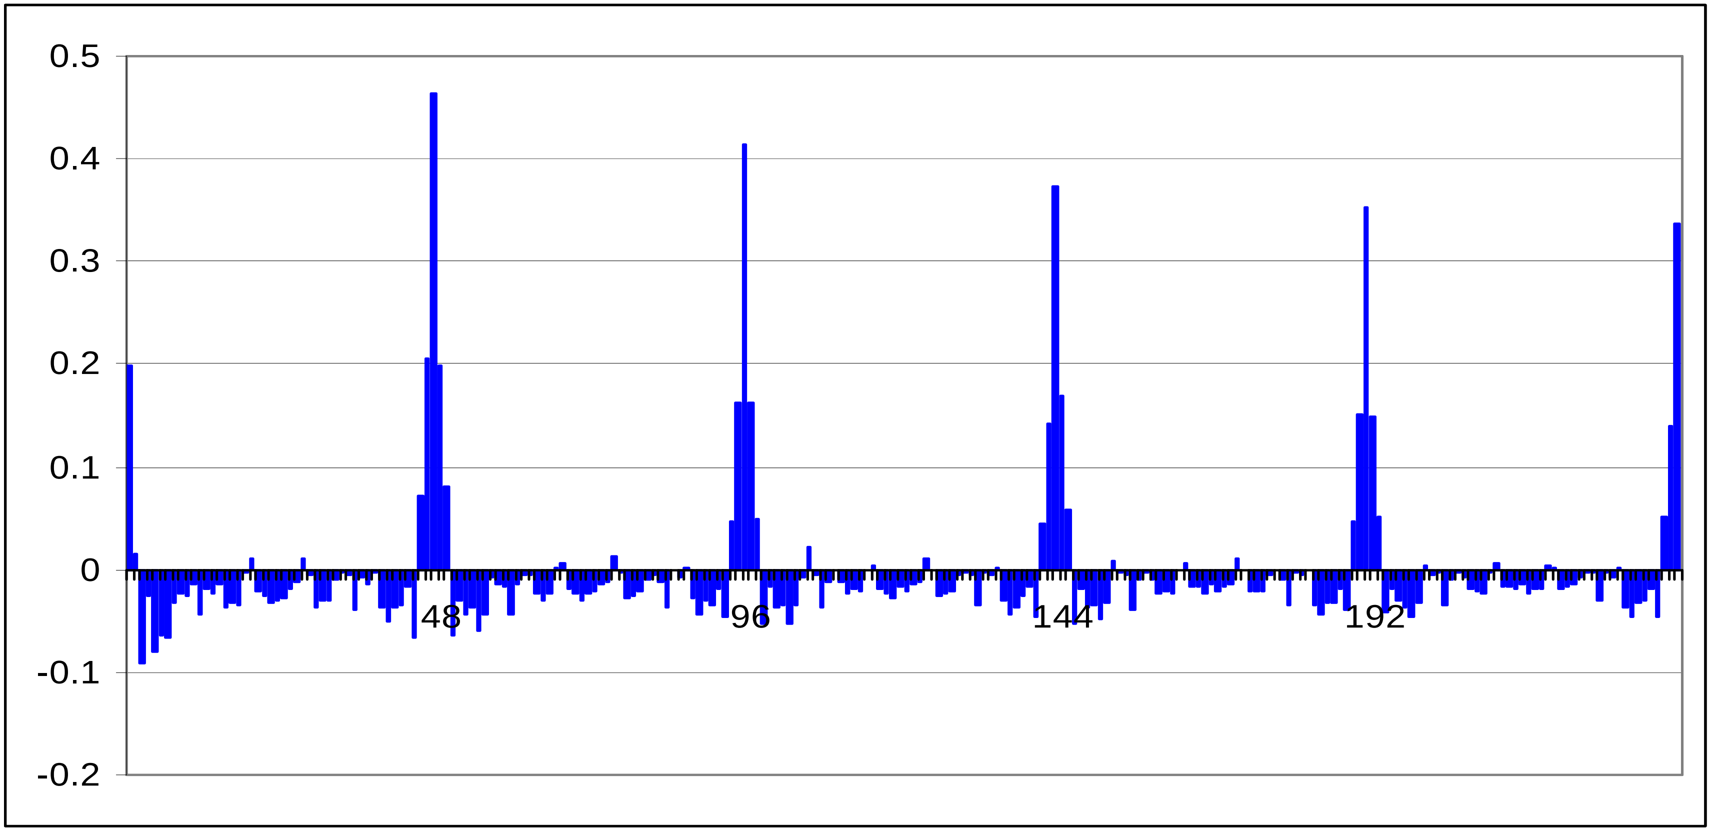

Figure 3b.
Correlogram plot of the autocorrelation in electricity price returns over 48 lags

This plot shows the autocorrelation of 30-minute spot price returns (ΔP_t) for the spot electricity price (P_t) where $\Delta P_t = log(P_t) - log(P_{t-1})$ and the interval $t\text{-}1 \rightarrow t$, is 30-minutes, over 48 lags (corresponding to 1 day). The prices are per megawatt hour (MWh) of electricity over the 24-hour day. The sample period is from the 1st July 2003 to 4th July 2006. N= 52,752.

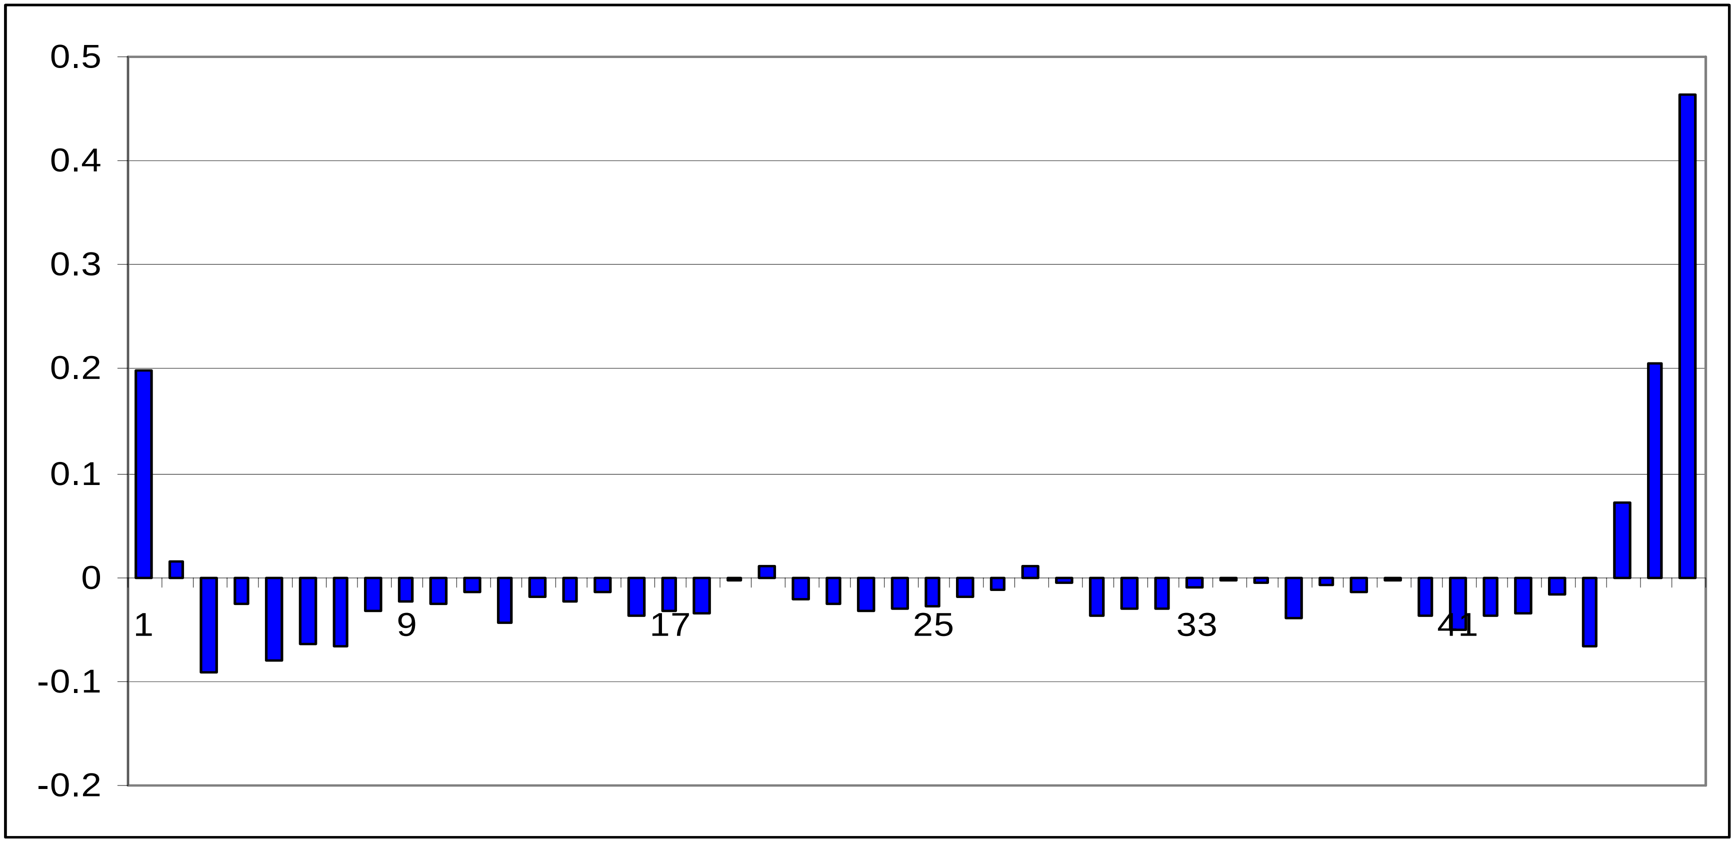

Figure 4a.
Three Dimensional Phase Portrait of Electricity Returns
Scatter plot of electricity returns (ΔP_t) plotted against first (ΔP_{t-1}) and second (ΔP_{t-2}) order lagged values. Note the rays emanating from the centre clearly form a star or compass rose pattern. N= 52,752.

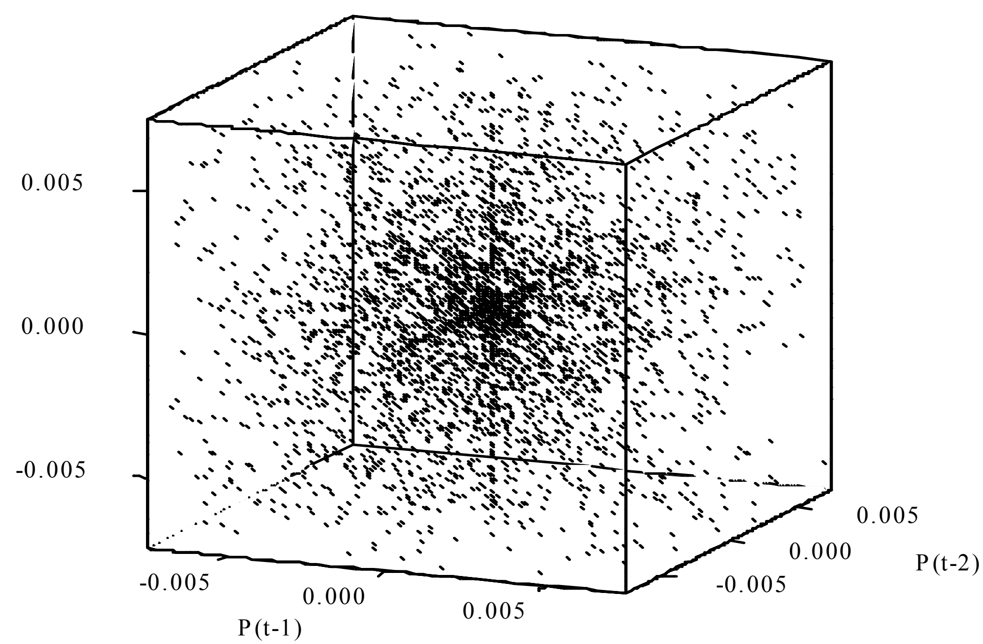

Figure 4b.

Two Dimensional Phase Portrait of Electricity Returns

Scatter plot of electricity returns (ΔP_t) plotted against first (ΔP_{t-1}) order lagged values with a return interval of 0.00 ± 0.50. Note the rays emanating from the centre clearly form a star or compass rose. N= 52,752.

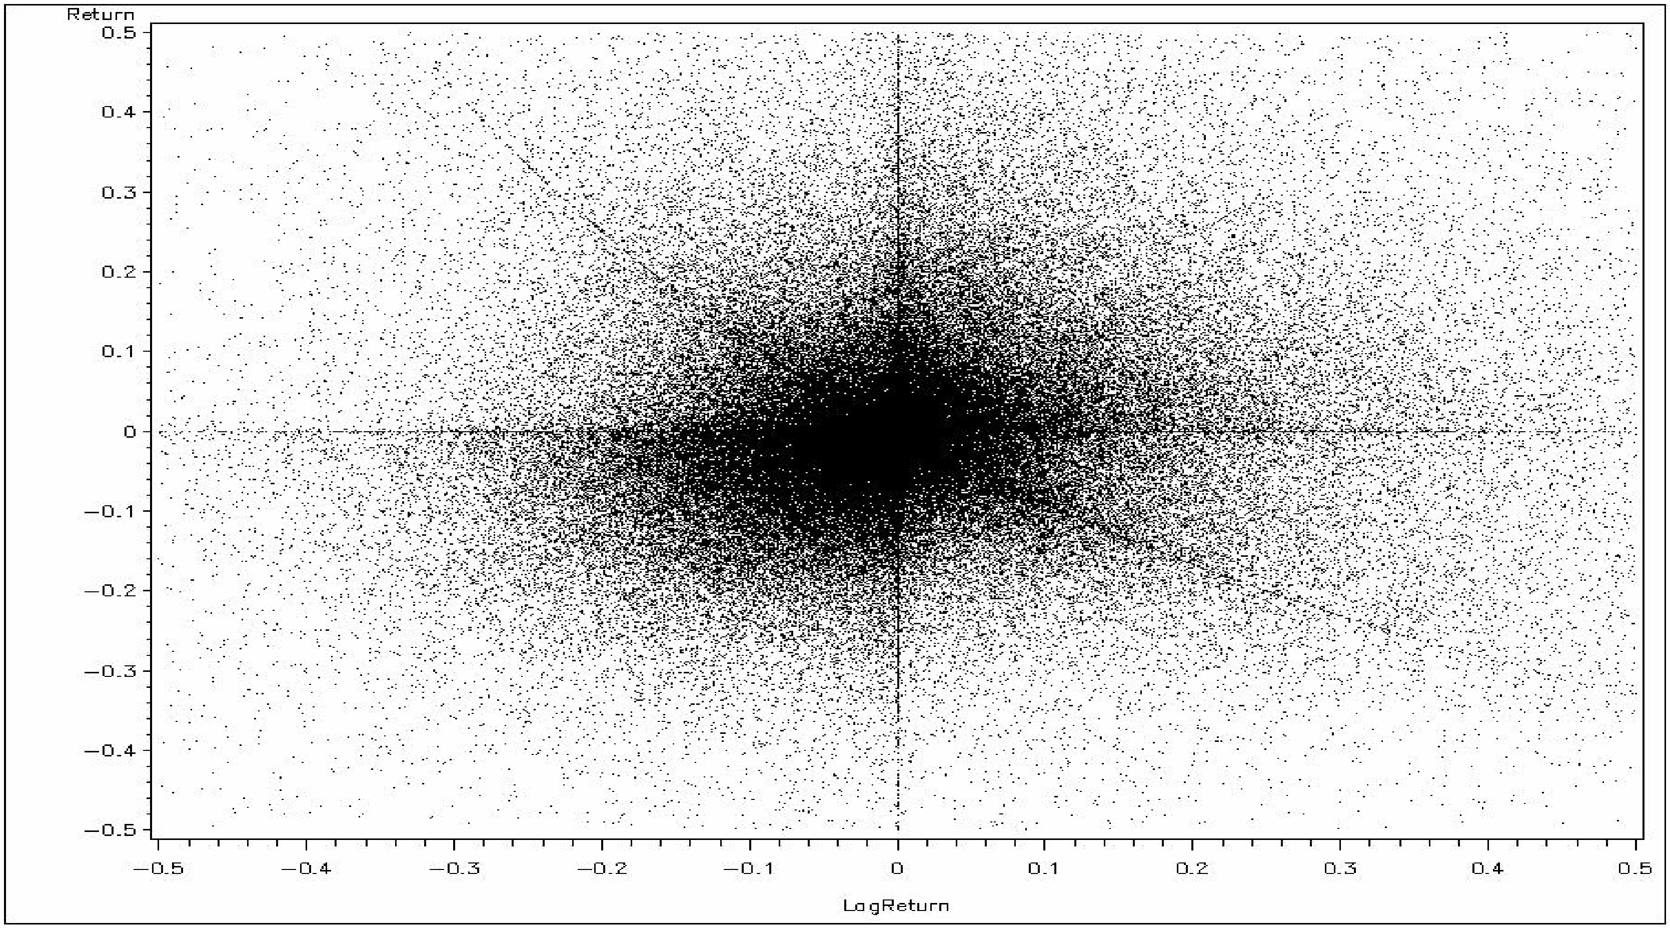

Figure 4c.
Two Dimensional Phase Portrait of Electricity Returns
Scatter plot of electricity returns (*ΔP$_t$*) plotted against first (*ΔP$_{t-1}$*) order lagged values, with a return interval from 0 ± 0.015 Note the rays emanating from the centre clearly form a star or compass rose. N= 52,772.

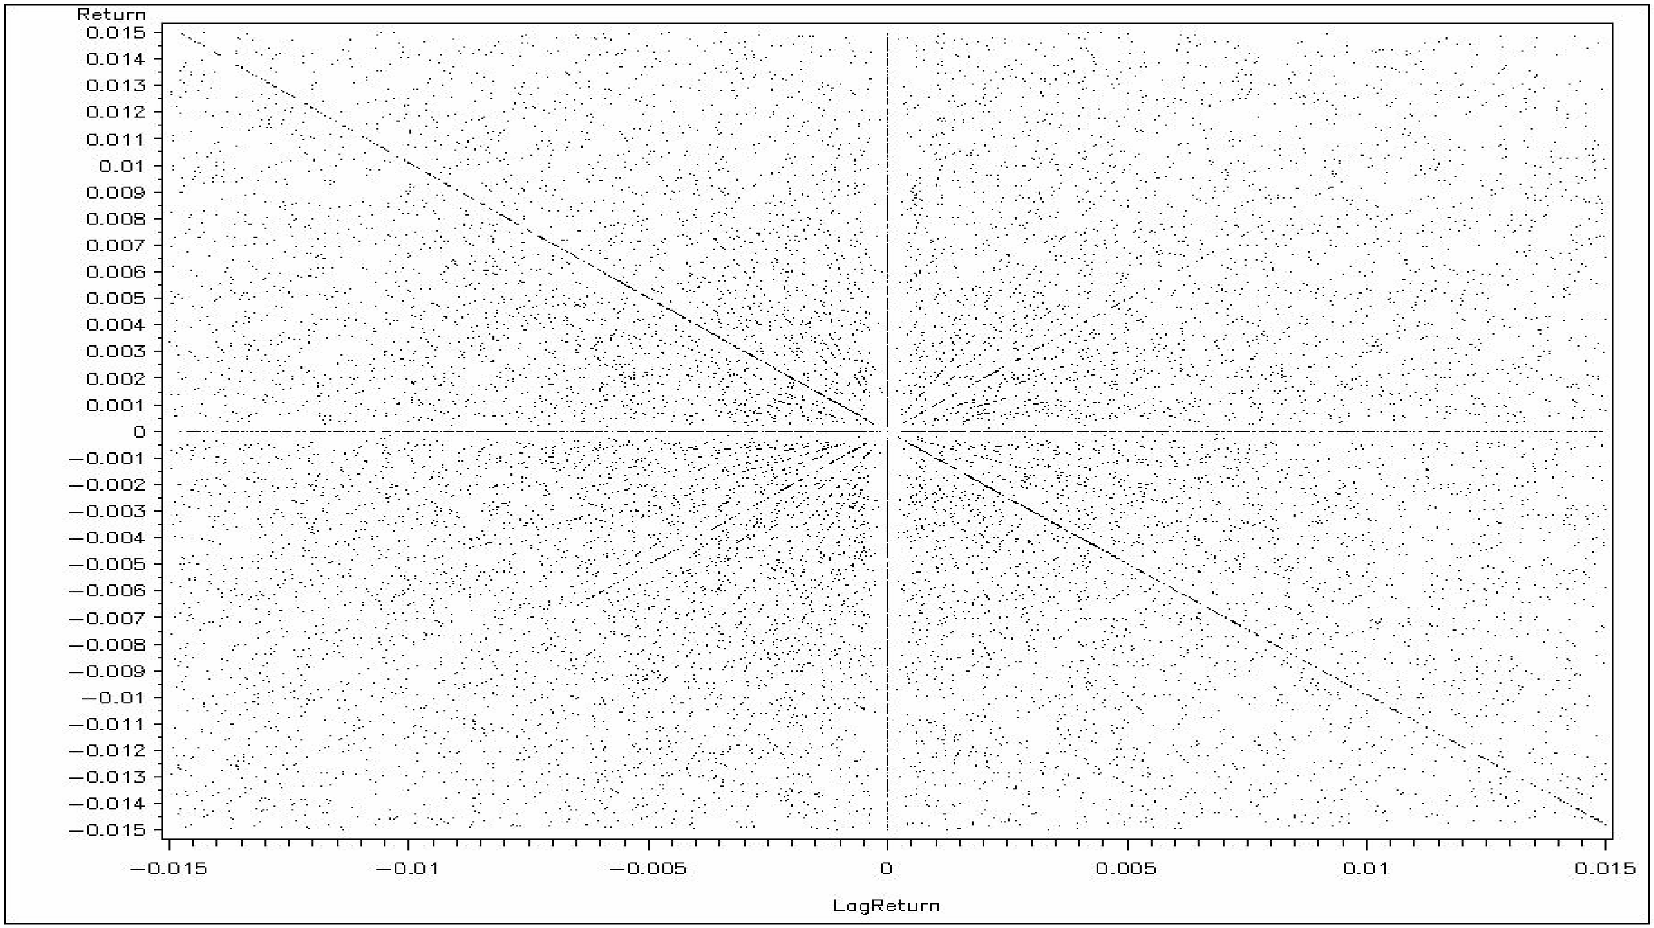

Figure 5a.

Three Dimensional Phase Portrait of Filtered Electricity Returns using an AR(10) Model

Scatter plot of filtered electricity returns (P_t AR10 = ψ_t from Equation 2) plotted against first (P_t AR10-1) and second (P_t AR10-2) order lagged values. N= 52,752. Note the rays emanating from the centre have disappeared.

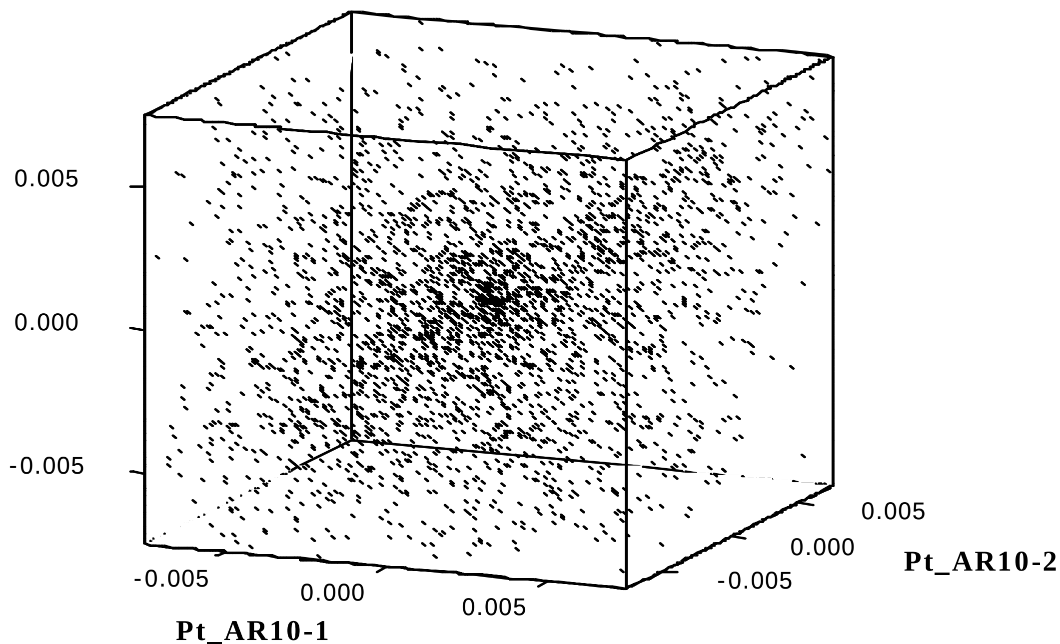

Figure 5b

Two Dimensional Phase Portrait of Filtered Electricity Returns using an AR(10) Model

Scatter plot of filtered electricity returns (P_t AR10 = ψ_t from Equation 2) plotted against first (P_t AR10-1) and second (P_t AR10-2) order lagged values. N= 52,752. The return interval is 0.00 ± 0.50. Note the rays emanating from the centre have disappeared.

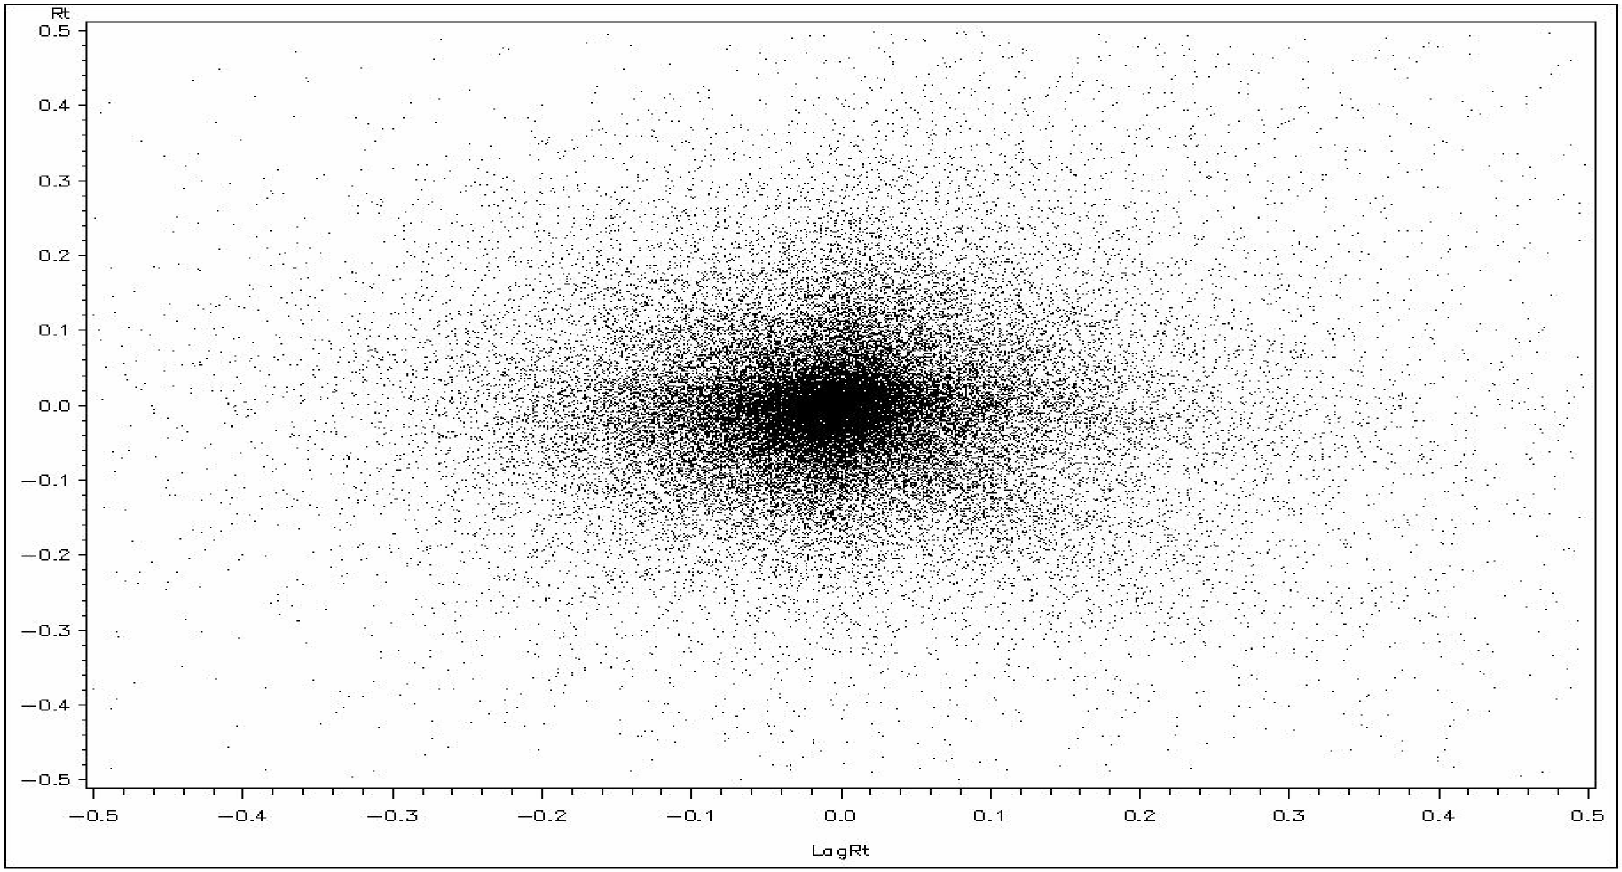

Figure 5b

Two Dimensional Phase Portrait of Filtered Electricity Returns using an AR(10) Model

Scatter plot of filtered electricity returns (P_t AR10 = ψ_t from Equation 2) plotted against first (P_t AR10-1) and second (P_t AR10-2) order lagged values. N= 52,752. The return interval is 0.00 ±0.150. Note the rays emanating from the centre have disappeared.

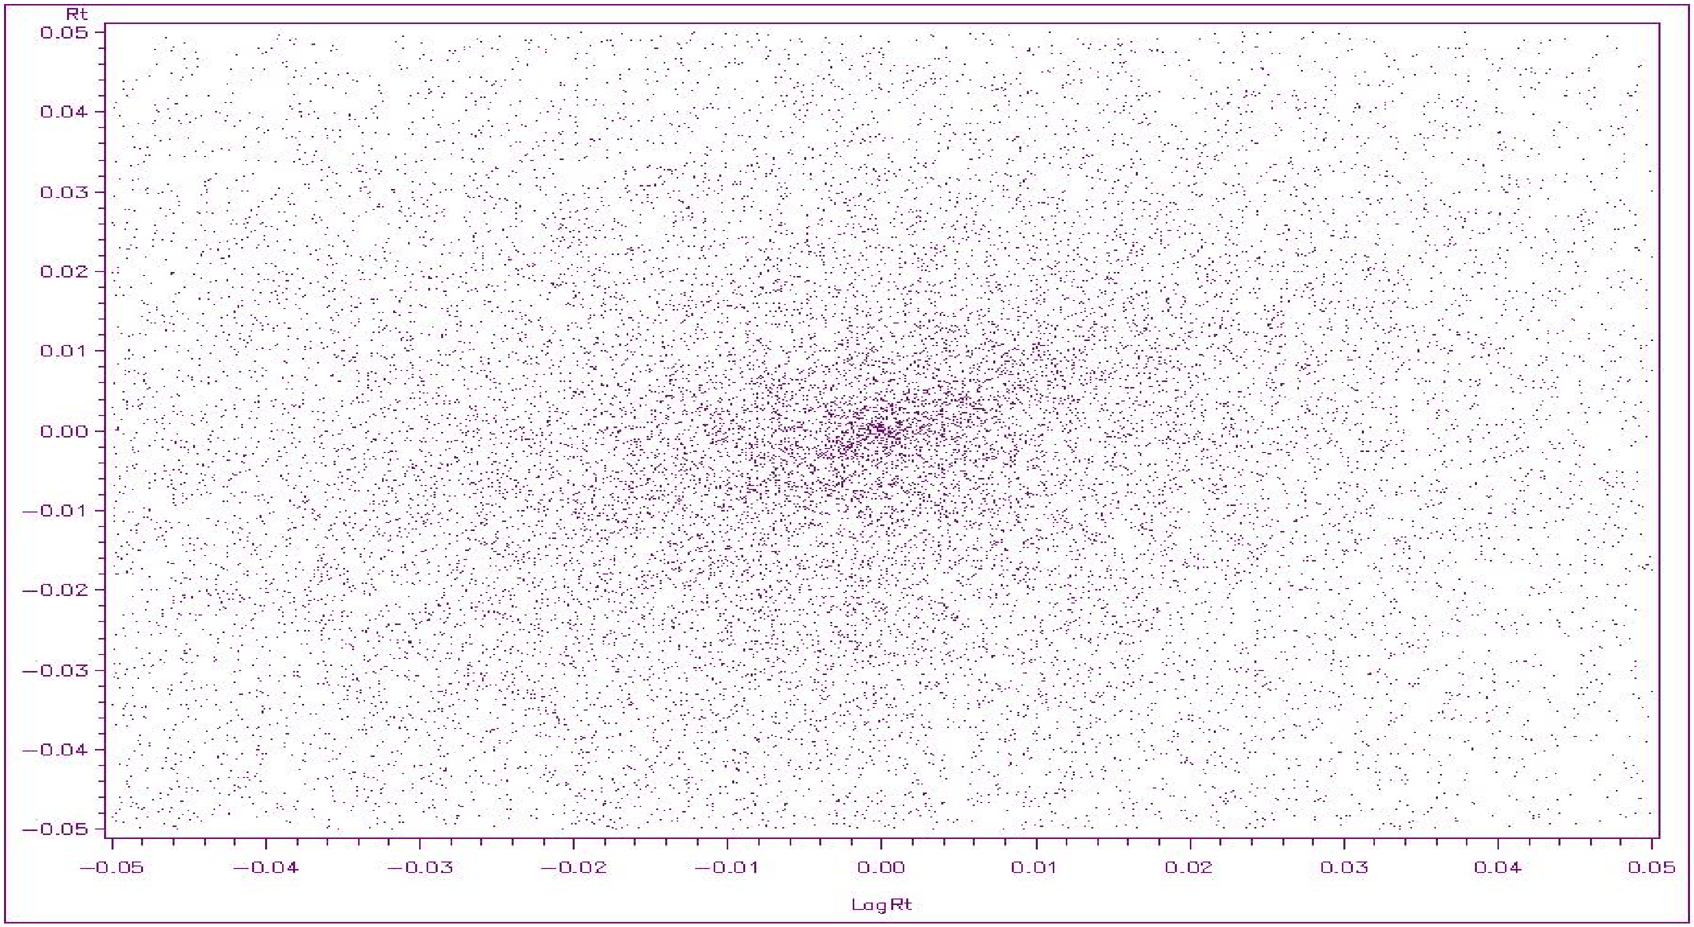

References:

AFMA Electricity Module - A module of the Financial Services Accreditation Program, 2005. www.afma.com.au

Antoniou, A., Vorlow, C.E., 2005. Price clustering and discreteness: Is there chaos behind the noise. Physica A 348, 389–403.

Batten, J. A., Ellis, C. A., Hogan, W. P., 2005. Decomposing intraday dependence in currency markets: evidence from the AUD/USD spot market. Physica A: Statistical and Theoretical Physics. 352(2-4), 558-572.

Bessembinder, H. and Lemmon, 2002. M., "Equilibrium Pricing and Optimal Hedging in Electricity Forward Markets", Journal of Finance 57 (June): 1347-82.

Brock W. A, Dechert D, Scheinkman J, LeBaron B. 1996 A test for independence based upon the correlation dimension. Econometric Reviews 15: 137–235.

Cai, C., R. Hudson and K. Keasey 2003 Trading frequency and the compass rose. Applied Economics Letters, 10: 511-527.

Crack, T. F. and Ledoit, O. 1996 Robust structure without predictability: The compass rose pattern of the stock market, Journal of Finance, 51: 751–762.

Chen, A.S., 1997. The square compass rose: The evidence from Taiwan. J. Multinational Financ. Manage. 7, 127–144.

Fang, Y., 2002. The compass rose and random walk tests. Computational Statistics and Data Analysis. 39, 299–310.

Gleason, K.C., Lee, C.I., Mathur, I., 2000. An explanation for the compass rose pattern. Economics Letters. 68, 127–133.

Gottleib, G., Kalay, A., 1985. Implications of the discreteness of observed stock prices. Journal of Finance. 40, 135–154.

Hurst, H.E., 1951. Long-term storage capacity of reservoirs. Transactions of the American Society of Civil Engineers 116, 770-799.

Koppl, R., Tuluca, S., 2002. Random walk hypothesis testing and the compass rose. Finance Letters 2, 14–17.

Kramer,W. and Runde, R. 1997. Chaos and the compass rose, Economics Letters, 54: 113–118.

Lee, C. I., Gleason, K. C. and Mathur, I. 1999. A comprehensive examination of the compass rose pattern in futures markets, Journal of Futures Markets, 19: 541–564.

Lo, A., 1991. Long-term memory in stock market prices, Econometrica. 59, 1279-1313.

Longstaff, F.A. and Wang, A.W. 2004. Electricity Forward Prices: a High-Frequency Empirical Analysis", Journal of Finance, 49: 1877-1900.

Mandelbrot, B. B., Wallis, J. R. 1969. Robustness of the rescaled range R/S in the measurement of noncyclic long-run statistical dependence, Water Resour. Res. 5, 967-988.

McKenzie, M. D. 2001. Chaotic behaviour in national stock market indices: new evidence from the close returns test. Global Finance Journal, 12: 35-53.

Mitchell H., and M. McKenzie. 2006. A note on the Wang and Wang measure of the quality of the compass rose. Journal of Banking & Finance 30: 3519-3524.

Niederhoffer. V. 1965. Clustering of Stock Prices. Operations Research, 13(2):258–265.

Osborne. M.F.M. 1962. Periodic Structure in the Brownian Motion of Stock Prices. Operations Research, 10(3):345–379.

Szpiro, G. G. 1998. Tick size, the compass rose and market nanostructure, Journal of Banking and Finance, 22: 1559–1569.

Wang, E., R. Hudson and K. Keasey. 2000. Tick size and the compass rose: further insights. Economics Letters 68: 119-125.

Wang, H. and Wang, C. 2002. Visibility of the compass rose in financial asset returns: a quantitative study, Journal of Banking and Finance, 26: 1099–1111.

Weron, R. 2000. Energy price risk management, Physica A 285: 127-134.