GARCH Modelling of High-Frequency Volatility in Australia's National Electricity Market

Stuart Thomas* and Heather Mitchell*

Abstract

This paper considers the underlying volatility process in Australian electricity prices and examines the applicability of a range of GARCH specifications to modelling volatility in 5 regional pool markets in the NEM. The GARCH variants considered include the basic GARCH, TARCH, EGARCH and PARCH specifications. The approach used in this study differs from the previous Australian ARCH-based studies in that discrete half-hourly returns are used over a six-year sample period, across each of five regional pools in the NEM. Seasonal effects and outliers (price spikes) are filtered prior to fitting the various GARCH models in order to investigate the underlying volatility process without the noise contributed by these effects. Results show that the PARCH specification is favoured in the NSW, QLD and SNOWY regions but in QLD and SA, the EGARCH specification is preferred as it more reliably describes the volatility processes in those two regions.

Keywords: Electricity, Power, Price, Returns, Volatility, GARCH, Outliers.

JEL Classification Numbers: C13, C22, L94, Q41

** School of Economics, Finance and Marketing, RMIT University*

The authors wish to acknowledge the financial support of The Melbourne Centre for Financial Studies, via grant 18/2005. We also wish to thank Richard Heaney and Michael McKenzie at RMIT University for helpful comments. Binesh Seetanah provided invaluable research assistance. Any remaining errors are the responsibility of the authors.

Address for Correspondence:
Stuart Thomas
School of Economics, Finance and Marketing
RMIT University
Level 12, 239 Bourke Street
Melbourne, Australia, 3000.
Tel: +61 3 9925 5828
Fax: +61 3 9925 5986
Email: stuart.thomas@rmit.edu.au

1. Introduction

The deregulation and restructuring of the Australian electricity industry and the development of the National Electricity Market (NEM) have introduced greater competition in Australia's electricity supply sector but at the cost of higher price volatility. The non-storable nature of electricity, difficulties in forecasting and managing demand, poorly maintained transmission and distribution infrastructure, and the potential for market power and information asymmetry, among other factors, can provide a load-matching problem for market operators[1]. These conditions create clear seasonal patterns in wholesale electricity prices by time of day, weekday, month and to a lesser extent yearly patterns (see *inter alia* Lucia and Schwartz, 2002). Intra-day patterns are the predominant form of seasonality and they are frequently accompanied by short-run spikes that are a significant and challenging feature of price behaviour [see Wolak (1997), Goto and Karolyi (2004), Higgs and Worthington (2005)].

High volatility has been a feature of price behaviour in the NEM since its establishment in 1998. When prices are highly volatile, it creates uncertainty about generators' revenues[2] and retailers' and distributors' costs. Further, high price volatility can make capacity planning and investment decisions difficult for generators and for owners and operators of the distribution grid. System operators and industry regulators also need to understand volatility to ensure that markets are designed and

---

[1] In the Australian summers of 2005/6 and 2006/7, persistent high temperatures have resulted in "load shedding" becoming a recurrent feature of electricity supply, particularly in southern Victoria and southeast Queensland. The system operator, NEMMCO, may from time to time direct electricity providers to disrupt supply in order to maintain system balance (see section 3.1).On January 16, 2007, VIC1 prices reached the $10,000 market cap for two hours after bushfires disrupted the interconnector between NSW and VIC. On the afternoon of that day, Victoria's demand peaked at 9100 megawatts, 4.2 per cent above the previous high observed on February 2006. (R Myer, "Power Meltdown Could Savage Snowy Hydro" http://www.theage.com.au/articles/2007/01/17/1168709832134.html, accessed 18/1/07).

[2] Taxpayer-owned Snowy Hydro faced possible large losses from the failure of the NSW-to-Victoria interconnector on Tuesday 16/1/07. VIC1 prices soared dramatically, sitting on $9000 a megawatt-hour for most of the afternoon and then hit the $10,000 mark for two hours after bushfires disrupted the cross-border transmission line. The average power price in normal times is about $30 a MWh. Snowy, a peak-power provider, has hedge contracts with Victorian power retailers and can usually deliver up to 1900 MW to Victoria via the interconnector. Being unable to deliver power via the interconnector, it had to meet its hedge contract obligations by buying wildly expensive power on the spot market and selling it to the retailers at their contract price. (ibid)

operated in a way that limits market power and promotes confidence and safety for market participants. Information about volatility informs measures of risk that are critically important to managers of energy commodity portfolios. Valuing derivatives and hedge contracts meaningfully and accurately requires meaningful and accurate measures and forecasts of price volatility over the life of an instrument.

In addition to generally high levels of volatility (see Bunn and Karakatsani, 2003), volatility clustering has been identified as a characteristic of electricity markets (Lucia and Schwartz, 2002). Autoregressive conditional heteroskedasticity (ARCH) models allow volatility shocks to cluster in time and may offer some insight into the volatility observed in electricity markets. To date, a relatively small number of ARCH-based studies have been undertaken in electricity markets. Lucia and Schwartz (2002) report volatility clustering in the form of GARCH effects and seasonality in both the deterministic component of prices and in spike intensity. Escribano *et al.* (2002) show volatility to be time-varying with evidence of heteroskedasticity in conditional variance for daily spot prices in Argentina, New Zealand, Nordpool (Norway and Sweden) and Spain. Knittel and Roberts (2001) apply a range of models of asset prices to hourly prices in the California market, including an EGARCH specification.

Of interest in the Australian context is the ARCH-based study of Worthington, Kay-Spratley and Higgs (2005). This paper examines electricity prices and price volatility among the five Australian electricity markets in the NEM by applying a multivariate generalised autoregressive conditional heteroskedasticity (MGARCH) model to identify the source and magnitude of spillovers, in a sample of half-hourly spot prices for the period December 1998 to June 1991. The authors find a large number of significant own volatility and cross-volatility effects in all five markets, indicating the presence of strong ARCH and GARCH effects. It should be noted that for the purposes of their analysis a series of daily arithmetic means is drawn from the trading interval data (following Lucia and Schwartz, 2002). The authors recognise that this treatment will entail the loss of at least some ‘news’ impounded in more frequent trading interval data, but correctly note that "…daily averages play an important role in electricity markets, particularly in the case of financial contracts…"[3].

Higgs and Worthington (2005) presents an investigation of the intra-day price volatility process in Australian electricity markets by applying five different ARCH processes: GARCH (generalised ARCH), Risk Metrics (normal integrated GARCH), normal APARCH (asymmetric power ARCH), Student-APARCH and skewed Student-APARCH (following Ding, Granger, and Engle, 1993; and Giot and Laurent, 2003a, 2003b). The authors include the documented systematic features – intra-day and monthly patterns (calendar effects), intra-day innovation and volatility spillovers (ARCH and GARCH effects) and market activity (demand and information asymmetry effects), with a view to providing a characterisation of the volatility process. The data employed consists of half-hourly electricity price relatives and demand volumes from 1 January 2002 to 1 June 2003 for the New South Wales, Queensland, South Australia and Victoria pools.[4]. The natural log of the price for each half-hourly interval is used to produce a time series of price relatives for analysis. In their analysis, the inclusion of news arrival is proxied by the contemporaneous volume of demand, time-of-day, weekday and monthly effects as exogenous explanatory variables. The authors find that on the basis of the log-likelihood, Akaike Information (AIC) and Schwartz Criteria (SC), the skewed Student-APARCH form is the best model for all four markets under consideration. Their results also indicate significant innovation (ARCH effects) and volatility (GARCH effects) in the conditional standard deviation equation, even with market and calendar effects included. They further observe significant asymmetric news responses in intra-day price volatility.

The previous Australian research typically confines its analysis to one regional market in the NEM over a relatively short time horizon (less than two years). This study considers a much larger data sample that is broader in scope than the previous papers – covering a six-year sample of higher-frequency half-hourly data, across five regions

---

[3] For example, the electricity futures contracts traded via the Sydney Futures Exchange (SFE) are settled against the arithmetic mean of half hourly spot prices in a given month.

[4] The SNOWY region is not included in the Higgs and Worthington (2005) study.

in the NEM, as compared to 1-2 year samples using daily average data from an abridged set of regions. We believe the use of a very-much-larger data set better characterises the volatility process by examining the market over a wider range of conditions and a broader market base. The treatment of seasonal effects and outliers is precise and specific and differs markedly from the generalised functional forms applied in the earlier studies; and this study finds that preferred ARCH model specifications and conditional error distributions differ when using large sample, high-frequency data.

Half-hourly trading-interval prices for the period from the commencement of the NEM in December 1998 to March 2005 are used and five NEM regions (NSW1, QLD1, SA1, SNOWY1 and VIC1[5]) are included. The GARCH variants considered include the "basic" GARCH specification (Bollerslev, 1986), the Threshold GARCH (TARCH) model of Glosten, Jaganathan and Runkle (1993), Nelson's (1991) Exponential GARCH (EGARCH) and the Power ARCH (PARCH) model proposed by Ding *et al.* (1993). The approach used in this study differs further from the previous Australian ARCH-based studies in that discrete half-hourly returns are used rather than log-based price relatives, to allow for the presence of negative prices, which were identified in Thomas et al. (2006) as a significant feature of the data. This study is further distinguished from previous work in that seasonal effects and individual spikes are treated by pre-whitening the data by removing seasonalities and outlier effects in an OLS framework *before* fitting the various GARCH models to the data. The reasons for doing so are twofold: firstly, after accounting for spikes and seasonalities, significant residual ARCH effects are observed in the whitened returns and we are interested in developing better understanding of underlying volatility process in the returns series *without* the noise contributed by seasonalities and outliers; and secondly, a model specified with a conditional mean and variance process that includes a very large number of variables (up to 260 to account for seasonal and outlier effects and serial correlation) *and* over a very large sample size (>110,000 returns observations in each region), is too unwieldy for available computing capabilities and as such a two-stage procedure is called for. The rest of the

[5] "NSW1", "QLD1", "SA1", "SNOWY1" and "VIC1" are the regional pool designations used by NEMMCO and will be used from this point forward.

paper is organised as follows: Data and preliminary statistical analysis is provided in section 2. Models and econometric methodology and main estimation results are presented in sections 3 and 4 respectively. Section 5 summarises findings and suggests further related research.

2. Data

This study employs half-hour discrete returns for each of the five NEM regions for the period from the commencement of the NEM at 2:00am on December 7th, 1998 to 12:00 am on April 1st, 2005. In the context of commodity futures contracts, Black (1976) notes that because futures contracts require no initial investment, futures positions cannot be said to yield rates of return as they are generally understood, i.e. as a result of change in value of the holder's initial investment over time. Because there is no ability to hold a unit of electricity and there is no "initial investment" in the commodity as such, spot electricity also does not yield a rate of return to an investor in the traditional sense. In light of this characteristic the term "returns" is used to denote proportionate price change over a trading interval. In general, attempts to model or forecast prices in financial markets should be based on successive variations in price and not on the prices themselves (see*, inter alia*, de Bodt, Rynkiewicz & Cottrell, 2001). The half-hourly rate of return series (hereafter referred to as "returns") is of interest because there are a growing number of over-the-counter and exchange-traded derivative products available for hedgers and speculators in the Australian and overseas electricity markets and pricing models for derivatives are informed by the behaviour of returns.

The base returns data are discrete half-hourly returns as defined by equation 1. A discrete returns specification[6] is preferred over log returns because the spot market in the NEM trades at discrete half-hourly intervals – it is not a continuous market in the way of most conventional financial markets. Further, a log returns specification may dampen the extreme spike effects we are attempting to capture, and is not defined in the presence of negative prices. With these market characteristics in mind, the returns series used in this study were generated as half-hourly discrete returns, ie:

[6] The analysis on returns reported in this chapter was also performed on first-difference (change in level) in the price series and the results were not found to be materially different from those found for the returns series.

$$RP_t = \frac{(P_t - P_{t-1})}{|P_{t-1}|}. \qquad (1)$$

Where RP_t represents the half-hourly discrete proportionate change in price ("return") at time t, P_t is half-hourly price at time t and $|P_{t-1}|$ is the absolute value of the previous half-hourly price, i.e. at time $t-1$. The denominator is specified as the absolute value to allow for the presence of negative prices.

Descriptive statistics for the half-hourly returns series are shown in table 1. The mean, standard deviation, minimum, maximum, range, skewness, kurtosis and Augmented Dickey-Fuller (ADF) statistics are reported for each region's returns series.

Insert Table 1 about here.

Mean half-hourly returns vary widely between regions, from 2.72% for VIC1 to 9.55% for SNOWY1. The standard deviation of returns is generally high, is widely dispersed across the regions and is consistent with the pattern of means, ranging from 111% for VIC1 to an extremely high 1700% for SNOWY1. The highest maximum return of 454,250% is observed in SNOWY1 and lowest in VIC1 of 14,243%. SNOWY1 also exhibits a markedly wider range of returns than the other regions. The extreme character of returns evident for SNOWY1 is most likely attributable to the nature of generation technology employed. All generation plant in SNOWY1 is hydroelectric, whereas more than 80% of generation capacity in NSW1, VIC1 and QLD1 is provided by coal-fired plant. Coal-fired generation is described as a "slow-start" technology, with orderly shutdown and start-up procedures taking up to 48 hours. By contrast, hydroelectric plant is a "fast-start" technology that can be called into production and shut down within a few minutes, with the result that hydroelectric generators are able to behave more opportunistically than coal-fired generators, with the ability to opt out of supply when pool prices are low and respond rapidly when prices are high.

The distributions of returns for QLD1, SA1 and SNOWY1 demonstrate positive skewness with NSW1 and VIC1 demonstrating a relatively low degree of negative skewness. Distributions of returns in all regions demonstrate extremely high positive kurtosis. Jarque-Bera (JB) statistics reject the null hypothesis of normal distribution at the 1% level of significance for all five regions. This fat-tailed character is consistent with earlier studies (see Huisman and Huurman (2002), Higgs and Worthington (2005) and Wolak (2000) and like price, is driven by the prevalence of extremely large spikes in returns. Augmented Dickey-Fuller (ADF) statistics clearly reject the hypothesis of a Unit Root at the 1% level of significance for all five regions, also consistent with the findings of the earlier studies.

2.1 Spike Behaviour

The presence of extreme spikes[7] in prices is a widely recognised characteristic of electricity markets [see Clewlow and Strickland, 2000a; Higgs and Worthington (2003, 2005); Bunn (2004); Alvaro, Peña, and Villaplana (2002); Hadsell, Marathe and Shawky (2004); and Goto and Karolyi (2004)]. Figure 1 illustrates the occurrence of extreme spikes in the price series in the VIC1 region over the sample period.

Insert Figure 1 about here.

For the purposes of this study a spike in returns is defined as any observed return more than four standard deviations larger than the mean. While the conventional practice is to apply a filter for outliers at three standard deviations from the mean, an initial survey of the data indicated that there is sufficient incidence of high prices and returns around and above the threshold at three standard deviations to justify applying a filter for outliers at four standard deviations. Table 2 collates the occurrences of spikes as defined, in aggregate for weekday, month and year. There are 566 extreme returns spikes observed across all regions during the sample period. QLD1 has the

[7] Note that electricity prices 'spike' rather than 'jump'. A 'jump' process in financial markets usually suggests that prices move rapidly to a new level and remain there, however electricity prices tend to move abruptly to an extremely high level and revert to mean levels just as abruptly (Blanco and Soronow, 2001).

highest incidence of extreme price spikes by state with 190 occurrences (34% of the total sample of spikes), followed by SA1 with 162 (29%), both have a markedly higher incidence than VIC1 with 98 (17%), NSW1 with 90 (16%) and SNOWY1 with 26 occurrences (5%). By day of the week, Monday shows the highest incidence with 121 observations (21%) and Friday the lowest with 49 occurrences (9%).

June shows the highest incidence by month with 87 (15%). The highest incidences by year occur in 2002 with 158 spikes (28%) and 2000 with 130 spikes (23%), both markedly higher than any other full year in the study period. It should be noted that the incidence of extreme price spikes appears to be declining from 2003 onwards. At the time of writing it is believed that this "settling down" is a feature of a maturing market combined with the development and wider use of bilateral hedge contracts between generators and distributors/retailers.

2.2 Negative Prices

Figure 1 also indicates that negative prices are a feature of the data. An observed characteristic of the price series for all 5 regions (NSW1, QLD1, SA1, SNOWY1, and VIC1) is the occurrence from time to time of negative prices in the pool. Figure 2 illustrates an extreme occurrence of a negative price in the pool, where the pool price fell to -$161.48 at 12:30 a.m, 15/4/2000.

Insert Figure 2 about here.

Negative prices are not usually encountered in financial time series but do occur from time to time in the NEM and are found to be significant (see Thomas et al., 2006). Occurrences of negative price are rare and typically short-lived, usually persisting for ½ to 1 hour, and are attributable to the market practices of generators. It is usual for generators to provide offers to supply electricity to the pool one day ahead of actual supply. The supply bid specifies a minimum level of generator output known as the "self-dispatch" level. A generator may lower its self-dispatch price to ensure that it is called in to generate. A generator may bid a negative price into the pool for its self-dispatch quantity (in effect, an offer to pay to generate) as a tactical move to ensure that they are among the first to be called in to generate. Demand usually outstrips the self-dispatch level of supply, so it is rare that generators actually pay to generate but on occasion a generator may not be called in and is “caught short”, effectively paying to generate for a short period. Trades are settled in the NEM daily on a net basis, so “paying to generate” does not usually require a cash outlay on the part of the generator. The occurrence of a negative price may also be a function of the nature of technology used to generate electricity. In Victoria and New South Wales, “base load” generation capacity employs what is generally referred to as “slow-start” generation technology, typically brown coal (VIC1) or black coal (NSW1) fired generation plant, usually located adjacent to a coalmine. The process of extracting coal from the earth, conveying it to a furnace, starting a furnace to boil water to make steam to ultimately turn a generator turbine may require up to 48 hours for an orderly start-up and a similar time for an orderly shut-down, at great cost to the generator. Given that electricity is not storable, if demand does not meet the minimum base-load capacity of these generators for a short period, a generator may be prepared to risk having to “pay” the pool to take their excess capacity for a short time rather than incur the greater cost of shut-down and start-up and the opportunity cost of lost production.

2.3 Filtered Returns:

The filtered data used in this study are derived by capturing the residuals (ε_t) from the model defined by equation 2:

$$RP_{R,t} = \alpha_0 + \sum_{j=1}^{6} \beta_{2,j} DAY_j + \sum_{k=1,\neq 9}^{12} \beta_{3,k} MTH_k + \sum_{l=1999,\neq 2001}^{2006} \beta_{4,l} YR_l + \sum_{m=1,\neq 23}^{48} \beta_{5,m} HH_m + \sum_{o=1}^{N_{R,S}} \beta_{6,o} SPIKE_{R,o} + \sum_{p=1}^{N_{R,N}} \beta_{7,p} NEG_{R,p} + \varepsilon_t \quad (2)$$

Where:

$RP_{R,t}$ represents the discrete return for region R at time t;

DAY_j represents the dummy variable for each day of the week (j=1 for Monday, 2 for Tuesday,…,6 for Saturday). MTH_k represents the dummy variable for each month (k=1 for January, 2 for February,…,12 for December).
YR_l represents the dummy variable for each year included in the sample period (l=1999,…,2006).
HH_m represents the dummy variable for each half-hourly trading interval (m=1 for 00:00hrs, 2 for 00:30hrs…,48 for 23:30hrs)
$SPIKE_{R,S}$ represents a set of $N_{R,S}$ dummy variables, one for each extreme spike as previously defined, with $N_{R,S}$ representing the number of extreme returns observed in region R for the period of the study (see table 2);
$NEG_{R,N}$ represents the dummy variable for the return associated with an occurrence of a negative price (p=1,…,$N_{R,N}$), with $N_{R,N}$ representing the number of occurrences of a negative price for region R during for the period of the study.

This model represents a relatively simple but highly effective method for controlling for seasonalities and the effects of individual spikes. Descriptive statistics for the filtered returns series are shown in table 3. Augmented Dickey-Fuller (ADF) statistics and results of ARCH-LM tests for each region's filtered returns series are also included in table 3.

The standard deviation is generally high relative to the mean and takes on a range of values across the regions, indicating a high degree of variability in the filtered returns and considerable variation between the 5 regions. The highest mean filtered returns are found in QLD1, SA1 and SNOWY1 and these regions exhibit the largest standard deviations. Consistent with the findings of the earlier studies, ADF statistics reject the hypothesis of a Unit Root at the 1% level of significance for all five regions.

Insert Table 3 about here.

The distributions of filtered returns for all 5 regions demonstrate positive skewness and extremely high kurtosis. This fat-tailed character persists despite the removal of extreme spikes from the raw returns data. The *p*-values for the Jarque-Bera (JB) statistics reject the null hypothesis of a normal distribution at the 1% level of significance for all five regions. It follows that these filtered returns are not well approximated by the normal distribution and that it may be appropriate to fit ARCH-type volatility models. The *p*-values for *F*-statistics in the ARCH-LM test results confirm the presence of ARCH effects in all five regions. Figure 3 suggests that volatility clustering is a feature of the data and the high positive skewness values suggest that there is a significant asymmetric response to positive shocks. It appears that VIC1 exhibits a different pattern of returns to the other regions, evidenced by smaller kurtosis and skewness.

Insert Figure 3 about here.

3. Methodology

The descriptive statistics presented in the previous section indicate that it is appropriate to use ARCH models to describe the volatility process in returns to electricity prices in the NEM. This is consistent with the findings of the earlier Australian studies and studies of foreign electricity markets that find temporal variation in electricity price volatility, with evidence of heteroskedasticity in conditional variance [see Bunn and Karakatsani, (2003) and Escribano *et al.,* (2002)].

Since Bollerslev (1986) proposed the Generalised ARCH (GARCH) model, there have been numerous developments in the ARCH literature to refine the mean and variance equations, in order to better capture temporal variations in financial market volatility[8]. An important innovation has been development of ARCH model specifications to describe the asymmetry present in financial data, where the current conditional volatility estimate for an asset is often dependent on the size and sign of past observations. For stock markets, this phenomenon was initially attributed to

[8] Developments in ARCH modelling and its application are surveyed in Bollerslev, Chou and Kroner (1992), Bera and Higgins (1993), Ding, Granger, and Engle (1993), Diebold and Lopez (1995), Pagan (1996), Giot and Laurent (2003a, 2003b), and Mitchell and McKenzie (2003, 2006).

leverage effects (see *inter alia* Black 1976, Christie 1982 and Nelson 1991). The presence of asymmetry in other financial markets such as foreign exchange markets required a different explanation. Bekaert and Wu (2000) suggest that volatility feedback mechanisms are a more likely explanation. Several ARCH models capture this characteristic, including Nelson's (1991) Exponential GARCH (EGARCH) and the Threshold ARCH and Threshold GARCH specifications that were introduced independently by Zakoïan (1994) and Glosten, Jaganathan, and Runkle (1993). In their examination of electricity price relatives, Higgs and Worthington (2005) argue that there is evidence of significant asymmetric effects in the volatility process.

A fat-tailed or leptokurtic distribution of prices and returns is a well-documented characteristic of electricity markets (see inter alia. Huismann and Huurman, 2003, Worthington, Kay-Spratley and Higgs, 2005 and Higgs and Worthington, 2005). Further innovations have been made to ARCH models to accommodate this characteristic in asset prices in conventional financial markets. Typically, this modification to the standard class of model involves replacing the standard normal density with some other assumed distribution such as a *t*-density (see Engle, 1982; and Bollerslev, 1986), the GED density (see Nelson 1991) and the autoregressive conditional density of Hansen (1994).

A further modification to the standard class of the ARCH model focuses on the specification of the power term that is used to transform the data to emphasise periods of volatility and relative tranquility. The standard class of ARCH model uses a squared power term, which may stem from the normality assumption traditionally invoked when describing the data. The presence of leptokurtosis suggests that this assumption may be invalid in which case the potential superiority of a squared transformation is lost and other power terms may be more appropriate (Mitchell & McKenzie, 2006). For example, Taylor (1986) and Schwert (1990) argued in favour of the standard deviation GARCH model, where a power term of unity is specified. Ding *et al.* (1993) introduced a new class of power-ARCH (PARCH) model in which the power parameter is estimated rather than imposed, thereby allowing an infinite number of transformations of the data. Higgs and Worthington (2005) find support for the Student-APARCH model in favour of other ARCH specifications. The Asymmetric Power-ARCH (APARCH) model proposed by Ding, Granger and Engle

(1993) extends the PARCH model to capture the asymmetric volatility response to negative and positive shocks. The Student-APARCH model is a further extension designed to account for the acute leptokurtosis in Australian electricity markets (see Bauwens and Giot, 2001, and Giot and Laurent, 2003a and 2003b).

ARCH models by and large are purely adaptive models and provide no clear theoretical basis for favouring one model specification over another[9]. As such they are regarded more as descriptive tools than representative of the actual data generating process. There are a wide variety of features found in the financial markets data and a very large number of models have been proposed to describe these features (Mitchell and McKenzie, 2005, 2006). Further, there are a growing number of univariate and multivariate extensions to the basic ARCH model, as such it would not be feasible to include all of the different specifications in this analysis. With this in mind and given that the purpose of this study is to examine the relative efficacy of univariate ARCH models in describing the filtered data, we have chosen to limit the scope of this study to the more basic forms of ARCH model that may reflect established characteristics of electricity prices, these being the TARCH, EGARCH and PARCH specifications as discussed above. The basic GARCH specification is included for comparison, to determine if a relatively simple model describes the data adequately, despite the documented characteristics of the data.

3.1 Model Specification

The Mean Equation

The basic GARCH (1,1) specification of Bollerslev (1986), gives the mean equation as follows:

$$Y_t = X_t^{'}\theta + \varepsilon_t \tag{3}$$

[9] For a discussion of the atheoretic nature of ARCH models see Diebold and Lopez (1995), Goodhart and O'Hara (1997) and Ackert and Racine (1997).

As such the mean equation (3) is a function of exogenous variables with an error term. For the purposes of this study the mean equation is modified to include appropriate AR and MA terms to control for autocorrelation in the data:

$$Y_t = X_t^{'}\theta + \sum_{i=1}^{p} \alpha_i Y_{t-i} + \sum_{i=1}^{q} \beta_i \varepsilon_{t-i} \quad (4)$$

Where p and q are chosen to capture significant spikes in the autocorrelation function.

GARCH Model Specification

The GARCH (1,1) conditional variance equation is given by equation 5:

$$\sigma_t^2 = \omega + \alpha\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2 \quad (5)$$

In which ω is a constant term, the ARCH term, ε_{t-1}^2, is given as the first lag of the squared residual from the mean equation and represents news about the volatility from the previous period, and the GARCH term, σ_{t-1}^2, represents last period's forecast variance. The specification of this model is consistent with the volatility clustering often seen in financial returns data, where large changes in returns are likely to be followed by further large changes.

Threshold ARCH (TARCH)

It is often observed in financial markets research that a downward price movement in the market will generate a higher volatility response than an equivalent upward movement. This is described as asymmetric news impact. The TARCH specification proposed by Glosten, Jaganathan, and Runkle (1993) and Zakoian (1994) is used to test for this asymmetric news impact. The occurrence of extremely short-lived spikes followed by periods of relative calm is a well-established feature of electricity price behaviour[10]. Evidence in support of the existence of *volatility* spikes is found by Wolak (1997) and Goto and Karolyi (2004). Higgs and Worthington (2005) find that price spikes, early-morning, late-afternoon and early evening hours are associated with high volatility and that negative price spikes, and other times of the day, week and year are associated with relatively lower volatility.

The GJR TARCH specification for the conditional variance is:

[10] Interestingly, Wolak (1997) and Goto and Karolyi (2004) also find evidence in support of the existence of *volatility* spikes.

$$\sigma_t^2 = \omega + \alpha\varepsilon_{t-1}^2 + \gamma\varepsilon_{t-1}^2 d_{t-1} + \beta\sigma_{t-1}^2 \quad (6)$$

The basic GARCH model of equation 7.4 is extended to include a threshold term $\gamma\varepsilon_{t-1}^2 d_{t-1}$. In this model, d_t =1 if ε_t < 0, and 0 otherwise. In this model, an upward spike (ε_t<0) has an impact of α and a downward or negative spike (ε_t>0) has an impact of $\alpha+\gamma$. If γ >0, a negative spike increases volatility and a leverage effect is present. If $\gamma\neq0$, the impact of news on the series' returns is asymmetric.

The asymmetric volatility response identified by Higgs and Worthington (2005) indicates that volatility tends rise in response to 'good news' for traders (proxied by positive price spikes) and fall in response to 'bad news' (negative spikes)', which is a perverse asymmetry that runs counter to the effects generally observed in conventional financial markets.

Exponential GARCH

Nelson's (1991) Exponential GARCH (EGARCH) model is formulated to capture news in the form of leverage effects. In the EGARCH model the specification for the conditional covariance is given by:

$$\log(\sigma_t^2) = \omega + \sum_{j=1}^{q}\beta_j \log(\sigma_{t-j}^2) + \sum_{i=1}^{p}\alpha_i \left|\frac{\varepsilon_{t-i}}{\sigma_{t-i}}\right| + \sum_{k=1}^{r}\gamma_k \frac{\varepsilon_{t-k}}{\sigma_{t-k}} \quad (7)$$

The left-hand side is the *log* of the conditional variance, implying that any leverage effects are exponential and that forecasts of conditional variance are guaranteed to be non-negative. In interpreting the model, the presence of leverage effects is indicated by γ_k<0, and the impact is asymmetric if $\gamma_k \neq 0$. While the basic GARCH model requires the restrictions in estimation that σ_t^2 >0, for *t* = 1…T , the EGARCH model allows unrestricted estimation of the variance, i.e. -∞ < log(σ_i^2) < ∞, *implying* that σ_t^2 >0, and so the assumption that σ_t^2 >0, for *t* = 1…T is automatically satisfied.

Power ARCH

The power-ARCH (PARCH) specification proposed by Ding *et al.* (1993) generalises the transformation of the error term in the models. The PARCH specification is given by equation 8:

$$\sigma_t^{\delta} = \omega + \sum_{j=1}^{q} \beta_j \sigma_{t-j}^{\delta} + \sum_{i=1}^{p} \alpha \left(\left| \varepsilon_{t-1} \right| - \gamma_i \varepsilon_{t-i} \right)^{\delta} \qquad (8)$$

The power parameter, δ, is estimated rather than imposed, and an optional threshold parameter, γ, may be included to capture asymmetry. The Bollerslev (1986) model sets δ=2, γ=0, and the Taylor (1986) model sets δ=1and γ,=0. Empirical estimates indicate the power term is sample dependent and values of near unity are common in the case of stock data (see Ding *et al.* 1993), while for foreign exchange data the power term varies between unity and two (see McKenzie and Mitchell, 2002). When fitting a PARCH model to electricity price data, the choice of power parameter is not obvious. Higgs and Worthington (2005) find some variation between regions in the estimated power term of a model for electricity prices.

3.2 Model Estimation Procedure

Preliminary attempts to estimate the four GARCH specifications to order (1,1) with a mean equation inclusive of the very large number of seasonal and outlier control variables identified in Thomas et al. (2006) created convergence problems that meant the model could not be reliably estimated[11]. In view of this constraint a decision was taken to undertake a two-stage estimation process by firstly pre-filtering the data by controlling for seasonalities and outliers in the returns series using OLS estimation in equation 2, capturing the residuals (ε_t) from the model (referred to herein as the "filtered returns"), then simultaneously estimate the mean and conditional variance

[11] Initial attempts using EVIEWS and SPSS statistical software packages to fit GARCH(1,1), TARCH(1,1), EGARCH(1,1) and PARCH(1,1) models to each region's discrete returns series, while controlling carefully for seasonal an outlier effects and serial correlation required that the mean equation be specified to include as many as 260 dummy variables (47 intra-day, six weekday, 11 monthly, six for year and up to 190 spikes, depending on region – see Thomas et al, 2006), over a sample size of 110,718 observations for each region. Almost all attempts to fit basic GARCH(1,1) processes to such large models failed to converge after 1000 iterations.

equation over the filtered returns, incorporating appropriate AR and MA terms in the mean equation to control for serial correlation. In view of the very large sample size, and according to Engle (1982), this two-stage approach should not result in loss of asymptotic efficiency in model estimation.

ARCH specification also requires that an assumption be made about the conditional distribution of the error term. There are three assumptions commonly employed when working with ARCH models: normal (Gaussian) distribution, *t*-distribution, and Nelson's (1991) Generalized Error Distribution (GED). Preliminary analysis of the full data set found that the Generalised Error Distribution (GED) was the most appropriate for model estimation[12].

Given a distributional assumption, ARCH models are typically estimated by the method of maximum likelihood. For the GED, the contribution to the log-likelihood for observation *t* is:

$$l_t = -\frac{1}{2}\log\left(\frac{\Gamma(1/r)^3}{\Gamma(3/r)(r/2)^2}\right) - \frac{1}{2}\log\sigma_t^2 - \left(\frac{\Gamma(3/r)(y_t - X_t'\beta)^2}{\sigma_t^2\Gamma(1/r)}\right)^{r/2} \qquad (9)$$

Where the tail parameter $r > 0$. The GED is a normal distribution if $r = 2$ and leptokurtic if $r < 2$.

4. Empirical Results

For clarity, the empirical results are presented in five tables, numbered 4 to 8, presenting the results for the five regions in alphabetical order. Each table presents the estimated coefficients, standard errors and *p*-values for the conditional mean and variance equations for the four different GARCH processes under examination, along with the estimated GED parameter (*r*), Akaike Information (AIC) and Schwartz Bayes

[12] In preliminary analysis, GARCH(1,1), TARCH(1,1), EGARCH(1,1) and PARCH(1,1) models were estimated over full data samples for all five regions, assuming normal (Gaussian), Student-*t*, constrained Student-t (with degrees of freedom set at 2 and 4) and GED. In all five regions, the distributions of standard errors were found to be significantly non-normal and significantly asymmetric and as such neither the assumption of a normal distribution nor a Student-*t* distribution for the standard errors was supported. Note also that Nelson's (1991) EGARCH specification as represented by equation 7 assumes that the standard errors ε_t follow a Generalised Error Distribution.

Criteria (SBC), Durbin-Watson (DW) statistic and *F*-Statistic and *p*-value showing the results of ARCH-LM tests.

The uppermost section of each table describes the ARMA structure required to account for serial correlation for each region. There is some variation between regions, with NSW1 demonstrating significant AR effects the 1% level at lags 1,5,7 and 8; QLD1 demonstrating significant AR effects at lags 1, 47 and 48 and significant MA effect at lag 48; SA1 showing significant AR effects at lags 1, 47 and 48. Interestingly SNOWY1 exhibits the same type of AR structure as QLD1 with significant AR effects at lags 1, 47 and 48 and similar magnitude of coefficients across all four GARCH specifications, while VIC1 demonstrates significant AR effects at lags 1, 5, 7, 47 and 48. Durbin-Watson statistics for all model specifications in all regions are close to 2, indicating a lack of significant residual serial correlation after model estimation.

Insert Table 4 about here.

The results from fitting the various GARCH specifications vary somewhat from region to region. For NSW1 (see table 4), the ARCH parameter (α), GARCH parameter (β), and parameter for the asymmetric volatility response (γ) (where applicable) in the GARCH, TARCH and PARCH specifications sum to less than one, and EGARCH imposes no constraints on the parameter estimates, indicating in NSW1 all four models result in stable ARCH processes and are viable models. The ARCH parameter (α), and GARCH parameter (β) are positive and significant in all four models, indicating the presence of ARCH and GARCH effects in the filtered returns. The estimated GED parameter (r) is in the order of 0.88, falling between 0 and 2, indicating that the distribution of standard errors is leptokurtic[13], which is consistent the with the fat-tailed character observed in Australian electricity prices and returns.

[13] See Nelson (1991).

The parameter for the asymmetric volatility response (γ) is negative and significant in the TARCH and PARCH models indicating an asymmetric response for positive returns in the conditional variance equation. This result is broadly consistent with the skewness values shown in table 3 and suggested by Figure 3, reflecting the condition that volatility tends rise in response to positive spikes and fall in response to negative spikes. This lies counter to the usual expectation in stock markets where downward movements (falling returns) are followed by higher volatility than upward movements (increasing returns)[14]. Ranking by Akaike Information (AIC) and Schwartz Bayesian Criteria (SBC) favours the PARCH model over the other three specifications. *F*-Statistics resulting from the ARCH-LM test are significant for all four models, indicating that heteroskedasticity has largely been accounted for by all four model specifications, with the p-value for the PARCH model's *F*-Statistic indicating that this model may be accounting for a greater proportion of heteroskedasticity than the other models. Finally, the power coefficient (δ) of the standard deviation process in the PARCH model is significantly different from one, indicating it is more appropriate to model the conditional standard deviation of electricity markets in a non-linear form.

Results for QLD1 differ markedly from those found for NSW1 (see table 5).

Insert Table 5 about here

The sum of the α, β and γ values is markedly greater than 1 in the GARCH and TARCH models (summing to 1.29 and 1.19, respectively) suggesting an explosive ARCH process. The α and β values for the PARCH specifications sum to 1.09, less marked than the other two specifications but still indicating a potentially unstable model, leaving the EGARCH which imposes no constraints on the parameter estimates as the only viable model, despite the AIC and SBC estimates slightly favouring the PARCH specification. In the EGARCH model, the α and β estimates are significant and positive, indicating the presence of ARCH and GARCH effects in the filtered returns.

[14] See *inter alia* Thomas and Brooks, 2001.

While the *F*-Statistics resulting from the ARCH-LM test are not significant for all four models, indicating that heteroskedasticity has been accounted for by all four model specifications, the *p*-values for the test statistic suggest that the PARCH model may be taking better account of the ARCH effects. Like NSW1, the estimated GED parameter (r) is in the order of 0.88, falling between 0 and 2, indicating that the distribution of standard errors is leptokurtic; and the parameter for the asymmetric volatility response (γ) is negative and significant, indicating an asymmetric response for positive returns in the conditional variance equation.

SA1 yields results that are more consistent with the QLD1 than NSW1 (see table 6). Again, F-Statistics resulting from the ARCH-LM test are all significant, indicating that heteroskedasticity has largely been accounted for by all four model specifications. The sum of the ARCH parameter (α), and GARCH parameter (β), in the GARCH model is greater than 1 (1.09). For the TARCH model, the α and β values sum to 0.93 but the addition of the γ parameter results in a value of 1.28.

Insert Table 6 about here.

The estimated GED parameter (r) is in the order of 0.7 for both models, indicating that the distribution of standard errors is leptokurtic. The parameter for the asymmetric volatility response (γ) is negative and significant in both the EGARCH and PARCH models, indicating an asymmetric response for positive returns in the conditional variance equation, consistent with NSW1 and SA1. Finally, ranking by Akaike Information (AIC) and Schwartz Bayesian Criteria (SBC) favours the PARCH model over the EGARCH specification. The *p*-values for the ARCH-LM test statistic suggest that the PARCH model may be taking better account of the ARCH effects than the EGARCH model. The power coefficient (δ) of the standard deviation process in the PARCH model is significantly different from one, indicating that the conditional standard deviation of electricity markets should not be modelled in a linear framework.

Results for the conditional variance equation for SNOWY1 are shown in Table 7.

Insert Table 7 about here.

Like QLD1 and SA1, results for SNOWY1 eliminate the GARCH and TARCH specifications on the basis of the *α, β* and *γ* values adding to values greater than 1 (≈1.06 for both models). This notwithstanding, *F*-Statistics resulting from the ARCH-LM test are all significant, indicating that heteroskedasticity has largely been accounted for by all four model specifications.

Of the remaining EGARCH and PARCH models, the AIC and SBC estimates favour the PARCH model. Interestingly, the *p*-values for the ARCH-LM test statistic are markedly larger for SNOWY1 than the other regions and although the AIC/SBC ranking favours the PARCH model, the *p*-value suggests that the EGARCH model might have better overcome the problem of heteroskedasticity in the returns series.

The ARCH parameter (*α*), and GARCH parameter (*β*) in the viable EGARCH and PARCH models are positive and significant, indicating that the filtered returns exhibit ARCH and GARCH effects. Consistent with the previous three regions, the estimated GED parameter (*r*) is between 0 and 2 for both models, indicating that the distribution of standard errors is leptokurtic, the parameter for the asymmetric volatility response (*γ*) is negative and significant in both the EGARCH and PARCH models, and the power coefficient (*δ*) of the standard deviation process in the PARCH model is significantly different from unity.

In VIC1 (Table 8), the estimated GARCH, EGARCH and PARCH models are stable, and the TARCH model is the only one rejected on the basis of instability (*α, β* and *γ* sum to 1.05).

Insert Table 8 about here.

As in the other 4 regions, ranking by Akaike Information and Schwarz-Bayes Criteria favours the PARCH model over the remaining viable models and like NSW1, QLD1 and SA1, the p-values for the ARCH-LM test statistic suggest that the PARCH model might better address heteroskedasticity in the data. The ARCH parameter (α), and GARCH parameter (β) in the EGARCH and PARCH models are positive and significant, consistent with the other regions. The leptokurtic character of the distribution of the standard is supported by the GED parameter (r) estimate of 0.84. The parameter for the asymmetric volatility response (γ) is negative and significant in both the EGARCH and PARCH models, and like the other regions the power coefficient (δ) of the standard deviation process in the PARCH model indicates that a non-linear conditional standard deviation equation is appropriate.

In summary, ranking by AIC and SBC favours the Power-ARCH (PARCH) specification in all five regions, although it should be noted that in QLD1 the sum of the ARCH, GARCH and Threshold (asymmetric volatility response) parameters sum to a value slightly greater than unity which signals an unstable model, in which case the choice of model defaults to the EGARCH specification. Model instability rejects the GARCH model in QLD1, SA1 and SNOWY1; and the TARCH model is rejected on the basis of instability in all but NSW1. In the generally-favoured PARCH model, strong ARCH effects, and strong lagged volatility or GARCH effects are evident.

In all five regions the PARCH models indicate that the estimated asymmetric coefficients (γ_1) are significant and negative for all four regional markets indicating that positive shocks are associated with higher volatility than negative shocks. This outcome is consistent with the findings of Higgs and Worthington (2005), but is contrary to what is generally observed in equity markets. Interestingly, in an examination of the Nordpool spot price, Solibakke (2002) found "…insignificant asymmetric volatility coefficient for all specifications… suggesting equal reaction patterns to positive and negative shocks" while in their application of a Threshold ARCH model to North American regional markets, Hadsell, Marathe and Shawky (2004) estimated that the asymmetric effect was also significant and *negative* thus capturing a strong market response to 'negative' news in US electricity prices.

In all five regions, the estimated GED parameter (r) associated with all models lies between 0 and 2, consistent with the generally observed leptokurtic character of distributions relating to electricity prices. Finally, in all five regions, the estimated power coefficients (δ) of the standard deviation process in the PARCH model were positive and significantly different from one and two, thus indicating it is more relevant to model the conditional standard deviation in a non-linear form.

5. Conclusion:

The study presented in this paper investigates the efficacy of four different GARCH model specifications in describing the underlying intra-day volatility processes in returns on electricity prices, in five regional pools (designated NSW1, QLD1, SA1, SNOWY1 and VIC1) in Australia's National Electricity Market (NEM). Four GARCH specifications, Generalised ARCH (GARCH), Threshold ARCH (TARCH), Exponential GARCH (EGARCH) and Power-ARCH (PARCH) models are applied to half-hourly returns on electricity prices for the period 7 December 1998 (commencement of the NEM) to 31 March 2005. Unlike previous GARCH-based studies on electricity prices, which seek to incorporate seasonal factors and outlier (price spike) effects in their models of the conditional mean equation, the very large data set used and the desire to investigate the underlying volatility process in the absence of these structural effects required that the returns data be deseasonalised and stripped of extreme spike effects prior to estimating the conditional mean and conditional variance equation in the GARCH estimation process.

The results show that significant ARCH and GARCH effects are present in the data and that Power ARCH specification with a Generalised Error Distribution applied to the standard errors generally describes the volatility process better than the other three GARCH models. The asymmetric volatility response captured by the PARCH model generally indicates that volatility tends rise in response to 'good news' for traders (positive price spikes) and fall in response to 'bad news' (negative spikes), which is counter to the effects generally observed in conventional financial markets but consistent with the findings of previous Australian GARCH-based studies. Finally, the estimated GED parameter (*r*) for each region confirms the fat-tailed properties that are generally observed in electricity market data in Australia and overseas.

A possible extension to this work might be an investigation of the application of other ARCH specifications to the data, possibly including higher-order GARCH specifications. Remembering that most ARCH/GARCH specifications and estimation procedures have been developed for more "conventional" financial markets and it may be that a further extension to the GARCH family is warranted, or that a different distributional assumption about the standard errors in the model is required. A further extension is possibly suggested by the recent work of Mitchell and McKenzie (2005, 2006) in the field of GARCH model selection criteria. This study and the previous Australian studies use the scheme of ranking models by Akaike Information (AIC) and Schwartz Bayesian Criteria (SBC). This approached may be well-established in practice but it has its detractors (see *inter alia* Pagan and Schwert, 1990) and it may be that there are more appropriate model selection criteria for electricity markets.

References:

Ackert L.F., and Racine, M.D., 1997, “The Economics of Conditional Heteroskedasticity: Evidence From Canadian and US Stock and Futures Markets”. *Atlantic Economic Journal;* vol. 25, no. 4, pp. 371–85.

Bauwens, L. and P. Giot, 2001, “Econometric Modelling of Stock Market Intra-day Activity (Kluwer).

Bekaert, G., and Wu, G., 2000, “Asymmetric Volatility and Risk in Equity Markets”, *Review of Financial Studies,* vol. 13, no.1, pp. 1–42.

Bera, A.K. and Higgins M.L., 1993, “ARCH Models: Properties, Estimation and Testing.” *Journal of Economic Surveys*, vol. 7, no. 4, pp. 305-366.

Black, F. (1976), “Studies Of Stock Price Volatility Changes” *Proceedings of the 1976 Meetings of the Business and Economics Statistics Section, American Statistical Association,* pp. 177-181.

Bollerslev, T., 1986, “Generalised Autoregressive Conditional Heteroskedasticity”, *Journal of Econometrics*, vol. 31, pp. 307-327.

Bollerslev, T., 1987, “A Conditionally Heteroskedastic Time Series Model For Speculative Prices And Rates Of Return”. *Review of Economics and Statistics*, vol. 69, no. 3, pp. 542–547.

Bollerslev, T., Chou, R.Y., and Kroner. K.F., 1992, “ARCH Modelling In Finance: A Review Of The Theory And Empirical Evidence”, *Journal of Econometrics,* vol. 52, no. 1, pp. 5-59.

Bunn, D. and Karakatsani, N. 2003, “Forecasting Electricity Prices”, v1, Working Paper, London Business School.

Christie, A., 1982, “The Stochastic Behaviour of Common Stock Variances: Value, Leverage and Interest Rate Effects” *Journal of Financial Economics,* vol. 10, no. 4, pp. 407–32.

Clewlow, L. and Strickland, C., 2000a, “Valuing Energy Options in a One-Factor Model Fitted To Forward Prices”. *Working paper*, University of Sydney.

Clewlow, L. and Strickland, C., 2000b, *Energy Derivatives – Pricing and Risk Management*, Lacima Publications, London.

de Bodt, E., Rynckiewicz, J. and Cottrell, M. “Some Known Facts About Financial Data, C*onference Proceedings*, European Symposium on Artificial Neural Networks, Bruges, 25-27 April 2001, D-Facto public., ISBN 2-930907-01, pp. 223-236

Diebold, F.X., and Lopez, J.A., 1995, "Modelling Volatility Dynamics", *Macroeconometrics: Developments, Tensions and Prospects* ed K D Hoover (Boston, MA: Kluwer).

Ding, Z., Granger, C.W.J., and Engle, R.F., 1993, "A Long Memory Property Of Stock Market Returns And A New Model." *Journal of Empirical Finance,* vol. 1, no. 1, pp. 83-106.

Engle, R.F., 1982, "Autoregressive Conditional Heteroscedasticity With Estimates Of The Variance Of UK Inflation." *Econometrica*, vol. 50, no. 4, pp. 987-1008.

Escribano A.J., Peaea J. and P. Villaplana, P., 2002, "Modelling Electricity Prices: International Evidence". Working paper, Universidad Carlos III de Madrid.

Giot, P. and Laurent, S., (2003a). "Market Risk In Commodity Markets: A VaR approach." *Energy Economics,* vol. 25, no. 5, pp. 435-457.

Giot, P. and S. Laurent (2003b). "Value-At-Risk for Long and Short Trading Positions." *Journal of Applied Econometrics*, vol. 18, no. 6, pp. 641-664.

Glosten, L.R., Jagannathan, R. and Runkle, D.E., 1993, "On the Relation between the Expected Value and the Volatility of the Nominal Excess Return on Stocks. *Journal of Finance*, vol. 48, no. 5, pp. 1779–1801.

Goodhart C.A.E., and O'Hara, M., 1997, "High Frequency Data In Financial Markets: Issues And Applications", *Journal of Empirical Finance,* vol. 4, no. 2, pp. 73–114.

Goto, M. and Karolyi, G.A., 2004 , "Understanding Electricity Price Volatility within and Across Markets." Dice Centre Working Paper, 2004-12, 1-41.

Hadsell, L., Marathe, A., and Shawky, H.A., 2004, "Estimating The Volatility Of Wholesale Electricity Spot Prices In The U.S", *The Energy Journal*, vol. 25, no. 4, pp. 23-40.

Hansen, B.E., 1994, "Autoregressive Conditional Density Estimation", International Economic Review, vol. 35, no. 3, pp. 705-730.

Higgs, H. and Worthington, A.C, 2003, "Evaluating the Informational Efficiency of Australian Electricity Spot Markets: Multiple Variance Ratio Tests of Random Walks" *Pacific and Asian Journal of Energy,* vol. 13, no. 1, pp. 1-16.

Higgs, H., and Worthington, A.C., 2005, "Systematic Features of High-Frequency Volatility in Australian Electricity Markets: Intraday Patterns, Information arrival and Calendar Effects", *The Energy Journal* vol. 26, no. 4, pp. 1-20.

Huisman, R., and Huurman, C., 2003, "Fat Tails in Power Prices", ERIM Report Series, Research in Management, Erasmus University.

Huisman R. and Mahieu, R., 2003, "Regime Jumps In Power Prices", *Energy Economics*, vol. 25, no. 5, pp. 425-434.

Knittel C.R. and Roberts, M., 2001, "An Empirical Examination of Deregulated Electricity Prices." *Working Paper, PWP-087*, University of California Energy Institute.

Lucia, J. and Schwartz, E.S., 2002, "Electricity Prices and Power Derivatives: Evidence from the Nordic Power Exchange". *Review of Derivatives Research*, vol. 5, no. 1, pp. 5-50.

McKenzie M.D., and Mitchell, H., 2002, "Generalised Asymmetric Power ARCH Modelling Of Exchange Rate Volatility", *Applied Financial Economics,* vol. 12, no. 8, pp. 555–564.

McKenzie, M.D., and Mitchell, H., 2004, "GARCH Model Selection Criteria". *Quantitative Finance*, vol. 3, no. 4, pp. 262–284.

Mitchell, H, and McKenzie, M.D., 2006, "A Comparison Of Alternative Techniques For Selecting An Optimum ARCH Model". *Journal of Statistical Computation and Simulation*, (*Forthcoming)*.

Nelson, D.B., 1991, "Conditional Heteroskedasticity in Asset Returns: A New Approach" *Econometrica,* vol. 59, no. 2, pp. 347–370.

Pagan, A., 1996, "The Econometrics of Financial Markets." *Journal of Empirical Finance,* vol. 3, no. 1, pp. 15-102.

Pagan, A.G. and Schwert, W., 1990, "Alternative Models For Conditional Stock Volatility". *Journal of Econometrics*, vol. 45, July 1990, pp. 267–290.

Schwert, W., 1990, "Stock Volatility and The Crash Of '87", *Review of Financial Studies,* vol. 3, no. 1, pp. 77-102.

Solibakke, P., 2002, "Efficient Estimated Mean and Volatility Characteristics for the Nordic Spot Electricity Power Market" *International Journal of Business,* vol. 7, no. 2, pp. 17-35.

Taylor, S., 1986, "Modelling Financial Time Series" (Wiley).

Thomas, S., and Brooks, R., 2001, "GARCH-Based Hedge Ratios for Australian Share Price Index Futures: Does Asymmetry Matter?" *Accounting, Accountability and Performance*, vol. 7, no. 2, pp. 61-76.

Thomas, S., Ramiah, V., Mitchell, H., and Heaney, R., 2006, "Seasonal Factors and Outlier Effects in Returns on Electricity Spot Prices in Australia's National Electricity Market", *Working Paper*, Melbourne Centre for Financial Studies.

Wolak, F.A., 2000, “Market Design and Price Behavior in Restructured Electricity Markets: An International Comparison”. *NBER-East Asia Seminar on Economics*, vol. 8, Chicago and London: University of Chicago Press

Wolak, F., and Patrick, R., 1997, “The Impact of Market Rules and Market Structure on the Price Determination Process in the England and Wales Electricity Market”, *Working Paper*, University of California Energy Institute.

Worthington, A.C. Kay-Spratley, A. and Higgs, H. (2005) “Transmission of Prices and Price Volatility in Australian Electricity Spot Markets: A Multivariate GARCH Analysis”, *Energy Economics,* vol. 27, no. 2, pp. 337-350.

Zakoian J. M., 1994, “Threshold Heteroskedastic Models”, *Journal of Economic Dynamics and Control,* vol. 18, no. 5, pp. 931–952.

Table 1: Descriptive Statistics for Half-Hourly Returns (by Region, December 1998 to March 2005).

Returns	NSW1	QLD1	SA1	SNOWY1	VIC1
Mean	0.0284	0.0643	0.0651	0.0955	0.0272
S.D.	1.22	1.72	1.77	17.00	1.11
Maximum	151.93	369.24	390.80	4542.50	142.43
Minimum	-222.00	-11.25	-32.94	-220.00	-209.50
Range	152.37	256.10	273.24	4300.99	157.29
Skewness	-0.44	113.14	117.56	241.51	-14.86
Kurtosis	15213.09	20540.27	22339.69	59676.37	16338.64
JB Stat P-Value	0.000	0.000	0.000	0.000	0.000
ADF Stat*	-38.16	-40.66	-341.20	-332.65	-39.84
N	110718	110718	110718	110718	110718

*All Augmented Dickey-Fuller test statistics reject the hypothesis of a Unit Root at the 1% level of confidence.

Table 2: Summary of Occurrences of Extreme Spikes in Returns by Region, by Weekday, Month and Year.

	NSW	QLD	SA	Snowy	VIC	Total
Sun	16	24	10	8	19	77
Mon	23	30	36	5	27	121
Tue	14	40	25	6	18	103
Wed	18	24	27	5	10	84
Thu	8	27	20	2	14	71
Fri	2	19	23	0	5	49
Sat	9	26	21	0	5	61
Total	***90***	***190***	***162***	***26***	***98***	***566***
Jan	5	23	13	0	3	44
Feb	2	9	20	2	7	40
Mar	0	13	18	0	2	33
Apr	0	4	11	0	2	17
May	14	19	11	4	16	64
Jun	23	30	11	3	20	87
Jul	13	29	12	5	12	71
Aug	6	14	7	3	6	36
Sep	7	2	6	1	6	22
Oct	9	21	13	7	12	62
Nov	7	11	14	0	6	38
Dec	4	15	26	1	6	52
Total	***90***	***190***	***162***	***26***	***98***	***566***
1998	2	2	16	0	2	22
1999	10	34	32	8	10	94
2000	17	59	31	3	20	130
2001	4	15	19	0	13	51
2002	37	55	29	5	32	158
2003	15	15	11	6	16	63
2004	4	10	21	2	5	42
2005	1	0	3	2	0	6
Total	***90***	***190***	***162***	***26***	***98***	***566***

Table 3: Descriptive Statistics for Filtered Half-Hourly Returns ε_t.
(by Region, December 1998 to March 2005).

ε_t	NSW1	QLD1	SA1	SNOWY1	VIC1
Mean	$-2.44\text{x}10^{-17}$	$1.29\text{x}10^{-16}$	$1.57\text{x}10^{-17}$	$1.48\text{x}10^{-17}$	$2.72\text{x}10^{-17}$
Median	-0.006	-0.014	-0.019	-0.008	-0.007
Maximum	4.82	6.86	7.00	19.82	4.39
Minimum	-1.13	-1.81	-1.46	-1.21	-1.30
Std. Dev.	0.18	0.30	0.29	0.28	0.19
Skewness	6.52	8.92	7.97	29.87	4.81
Kurtosis	121.53	139.03	122.96	1459.27	72.50
JB stat	$6.56\text{x}10^{7}$	$8.68\text{x}10^{7}$	$6.76\text{x}10^{7}$	$9.80\text{x}10^{9}$	$2.27\text{x}10^{7}$
(p-value)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
ADF Stat	-38.34	-32.22	-34.27	-36.10	-40.02
ARCH-LM Test					
F-Stat	95.06	57.06	49.53	44.89	62.12
(p-value)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
N	110718	110718	110718	110718	110718

The critical values of significance for skewness and kurtosis at the 0.05 level are 0.0305 and 0.0610, respectively. JB is the Jarque-Bera statistic. The critical value for the ADF statistic at the 0.01 level is -3.43

Table 4: Estimated Coefficients for Conditional Mean Returns and Variance Equations (NSW1)

	GARCH(1,1)		TARCH(1,1)		EGARCH(1,1)		PARCH(1,1)	
	Coeff.	**P-val**	**Coeff.**	**P-val**	**Coeff.**	**P-val**	**Coeff.**	**P-val**
Mean Equation								
ω	-0.006	0.000	-0.006	0.000	-0.006	0.000	-0.005	0.000
AR(1)	0.048	0.000	0.048	0.000	0.048	0.000	0.047	0.000
AR(5)	-0.025	0.000	-0.025	0.000	-0.024	0.000	-0.023	0.000
AR(7)	-0.030	0.000	-0.030	0.000	-0.029	0.000	-0.029	0.000
AR(48)	0.402	0.000	0.402	0.000	0.402	0.000	0.405	0.000
Variance Equation								
ω	0.003	0.000	0.003	0.000	-1.265	0.000	0.019	0.000
α	0.474	0.000	0.497	0.000	0.553	0.000	0.379	0.000
β	0.482	0.000	0.481	0.000	0.793	0.000	0.575	0.000
γ			-0.046	0.001	-0.002	0.616	-0.086	0.000
δ							1.065	0.000
r	0.875	0.000	0.875	0.000	0.873	0.000	0.876	0.000
AIC	-1.665		-1.665		-1.663		-1.671	
SBC	-1.664		-1.664		-1.662		-1.670	
DW-Stat	2.054		2.055		2.054		2.054	
ARCH-LM Test								
F-Stat	0.119	0.731	0.144	0.704	0.055	0.815	0.055	0.815
#Obs	110670		110670		110670		110670	

This table provides the estimated coefficients and *p*-values for the mean and conditional standard deviation equations for the NSW1 regional electricity pool in the NEM. *ω* is the constant in the conditional mean equation, αis the ARCH coefficient, γ is the leverage effect, δis the power of the conditional standard deviation process, AIC and SBC are Akaike Information and Schwartz-Bayes Criteria respectively. DW stat is the Durbin-Watson Statistic. ARCH-LM tests were specified with 48 lags, representing one full trading day.

Table 5: Estimated Coefficients for Conditional Mean Returns and Variance Equations (QLD1)

	GARCH(1,1)		TARCH(1,1)		EGARCH(1,1)		PARCH(1,1)	
	Coeff.	**P-val**	**Coeff.**	**P-val**	**Coeff.**	**P-val**	**Coeff.**	**P-val**
Mean Equation								
ω	-0.020	0.000	-0.022	0.000	-0.025	0.000	-0.021	0.000
AR(1)	0.017	0.000	0.016	0.000	0.015	0.000	0.016	0.000
AR(47)	0.014	0.000	0.014	0.000	0.014	0.000	0.014	0.000
AR(48)	0.934	0.000	0.937	0.000	0.941	0.000	0.937	0.000
MA(48)	-0.841	0.000	-0.846	0.000	-0.853	0.000	-0.845	0.000
Variance Equation								
ω	0.003	0.000	0.003	0.000	-1.209	0.000	0.013	0.000
α	0.899	0.000	0.798	0.000	0.696	0.000	0.658	0.000
β	0.390	0.000	0.391	0.000	0.806	0.000	0.479	0.000
γ			0.196	0.000	-0.183	0.000	0.043	0.000
δ							1.248	0.000
r	0.741	0.000	0.741	0.000	0.723	0.000	0.741	0.000
AIC	-1.510		-1.510		-1.481		-1.514	
SBC	-1.509		-1.509		-1.480		-1.513	
DW-Stat	2.183		2.182		2.178		2.181	
ARCH-LM Test								
F-Stat	0.182	0.669	0.172	0.678	0.129	0.720	0.094	0.760
#Obs	110670		110670		110670		110670	

This table provides the estimated coefficients and *p*-values for the mean and conditional standard deviation equations for the QLD1 regional electricity pool in the NEM. ω is the constant in the conditional mean equation, α is the ARCH coefficient, γ is the leverage effect, δ is the power of the conditional standard deviation process, AIC and SBC are Akaike Information and Schwartz-Bayes Criteria respectively. DW stat is the Durbin-Watson Statistic. ARCH-LM tests were specified with 48 lags, representing one full trading day.

Table 6: Estimated Coefficients for Conditional Mean Returns and Variance Equations (SA1)

	GARCH(1,1)		TARCH(1,1)		EGARCH(1,1)		PARCH(1,1)	
	Coeff.	**P-val**	**Coeff.**	**P-val**	**Coeff.**	**P-val**	**Coeff.**	**P-val**
Mean Equation								
ω	-0.017	0.000	-0.017	0.000	-0.018	0.000	-0.017	0.000
AR(1)	0.058	0.000	0.055	0.000	0.053	0.000	0.051	0.000
AR(47)	0.067	0.000	0.069	0.000	0.071	0.000	0.067	0.000
AR(48)	0.239	0.000	0.242	0.000	0.242	0.000	0.243	0.000
Variance Equation								
ω	0.005	0.000	0.005	0.000	-0.919	0.000	0.028	0.000
α	0.625	0.000	0.451	0.000	0.544	0.000	0.423	0.000
β	0.464	0.000	0.480	0.000	0.836	0.000	0.602	0.000
γ			0.347	0.000	-0.160	0.000	0.075	0.000
δ							0.964	0.000
r	0.698	0.000	0.702	0.000	0.698	0.000	0.704	0.000
AIC	-1.098		-1.101		-1.095		-1.108	
SBC	-1.097		-1.100		-1.094		-1.107	
DW-Stat	2.186		2.180		2.177		2.174	
ARCH-LM Test								
F-Stat	0.416	0.661	0.540	0.462	0.278	0.598	0.115	0.735
#Obs	110670		110670		110670		110670	

This table provides the estimated coefficients and *p*-values for the mean and conditional standard deviation equations for the SA1 regional electricity pool in the NEM. ω is the constant in the conditional mean equation, α is the ARCH coefficient, γ is the leverage effect, δ is the power of the conditional standard deviation process, AIC and SBC are Akaike Information and Schwartz-Bayes Criteria respectively. DW stat is the Durbin-Watson Statistic. ARCH-LM tests were specified with 48 lags, representing one full trading day.

Table 7: Estimated Coefficients for Conditional Mean Returns and Variance Equations (SNOWY1)

	GARCH(1,1)		TARCH(1,1)		EGARCH(1,1)		PARCH(1,1)	
	Coeff.	**P-val**	**Coeff.**	**P-val**	**Coeff.**	**P-val**	**Coeff.**	**P-val**
Mean Equation								
C	-0.014	0.000	-0.015	0.000	-0.017	0.000	-0.013	0.000
AR(1)	0.007	0.000	0.007	0.000	0.007	0.000	0.006	0.000
AR(47)	0.008	0.000	0.009	0.000	0.012	0.000	0.008	0.000
AR(48)	0.949	0.000	0.950	0.000	0.948	0.000	0.952	0.000
MA(48)	-0.828	0.000	-0.831	0.000	-0.826	0.000	-0.833	0.000
Variance Equation								
ω	0.003	0.000	0.003	0.000	-1.504	0.000	0.020	0.000
α	0.575	0.000	0.527	0.000	0.581	0.000	0.428	0.000
β	0.431	0.000	0.431	0.000	0.738	0.000	0.546	0.002
γ			0.099	0.000	-0.111	0.000	0.022	0.000
δ							1.061	0.000
r	0.806	0.000	0.807	0.000	0.804	0.000	0.805	0.000
AIC	-1.717		-1.717		-1.704		-1.723	
SBC	-1.716		-1.716		-1.703		-1.722	
DW-Stat	1.975		1.975		1.975		1.974	
ARCH-LM Test								
F-Stat	0.002	0.989	0.001	0.986	0.001	0.989	0.018	0.894
#Obs	110670		110670		110670		110670	

This table provides the estimated coefficients and *p*-values for the mean and conditional standard deviation equations for the SNOWY1 regional electricity pool in the NEM. ω is the constant in the conditional mean equation, α is the ARCH coefficient, γ is the leverage effect, δ is the power of the conditional standard deviation process, AIC and SBC are Akaike Information and Schwartz-Bayes Criteria respectively. DW stat is the Durbin-Watson Statistic. ARCH-LM tests were specified with 48 lags, representing one full trading day.

Table 8: Estimated Coefficients for Conditional Mean Returns and Variance Equations (VIC1)

	GARCH(1,1)		TARCH(1,1)		EGARCH(1,1)		PARCH(1,1)	
	Coeff.	**P-val**	**Coeff.**	**P-val**	**Coeff.**	**P-val**	**Coeff.**	**P-val**
Mean Equation								
ω	-0.005	0.000	-0.006	0.000	-0.007	0.000	-0.006	0.000
AR(1)	0.053	0.000	0.051	0.000	0.050	0.000	0.051	0.000
AR(5)	-0.026	0.000	-0.026	0.000	-0.024	0.000	-0.023	0.000
AR(7)	-0.022	0.000	-0.023	0.000	-0.021	0.000	-0.021	0.000
AR(47)	0.073	0.000	0.075	0.000	0.076	0.000	0.074	0.000
AR(48)	0.330	0.000	0.330	0.000	0.330	0.000	0.330	0.000
Variance Equation								
ω	0.004	0.000	0.004	0.000	-1.320	0.000	0.029	0.000
α	0.513	0.000	0.417	0.000	0.591	0.000	0.391	0.000
β	0.448	0.000	0.457	0.000	0.772	0.000	0.554	0.000
γ			0.179	0.000	-0.081	0.000	0.031	0.002
δ							0.962	0.000
r	0.841	0.000	0.807	0.000	0.843	0.000	0.844	0.000
AIC	-1.456		-1.457		-1.458		-1.464	
SBC	-1.455		-1.456		-1.457		-1.463	
DW-Stat	2.109		2.106		2.104		2.106	
ARCH-LM Test								
F-Stat	0.921	0.337	0.553	0.457	0.248	0.618	0.029	0.865
#Obs	110670		110670		110670		110670	

This table provides the estimated coefficients and *p*-values for the mean and conditional standard deviation equations for the VIC1 regional electricity pool in the NEM. ω is the constant in the conditional mean equation, α is the ARCH coefficient, γ is the leverage effect, δ is the power of the conditional standard deviation process, AIC and SBC are Akaike Information and Schwartz-Bayes Criteria respectively. DW stat is the Durbin-Watson Statistic. ARCH-LM tests were specified with 48 lags, representing one full trading day.

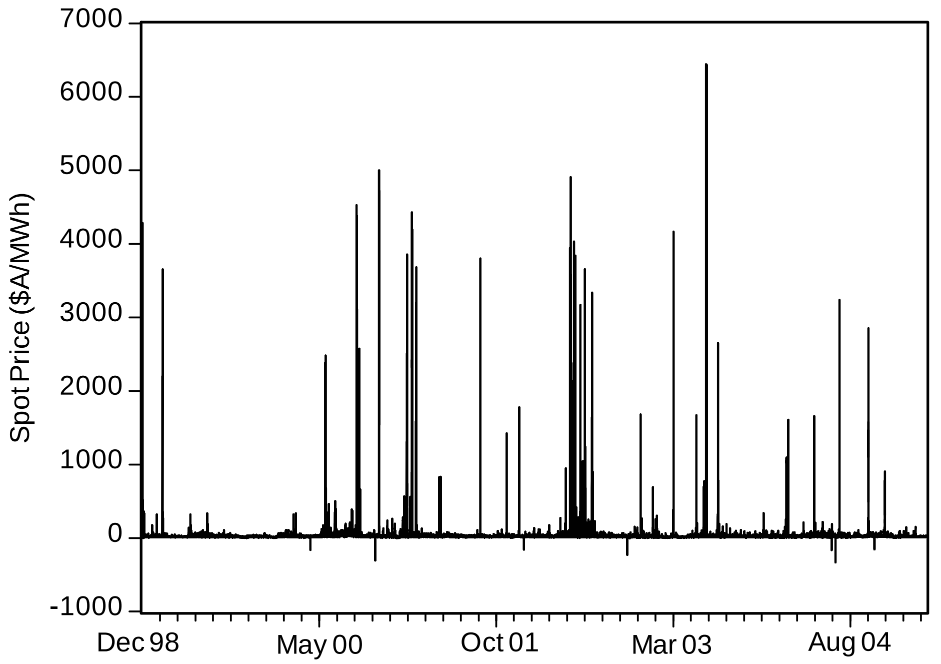

Figure 1: VIC1 Spot Price for the Period December 1998 to End March 2005 Illustrating the Occurrence of Extreme Spikes and Negative Values in the Price series

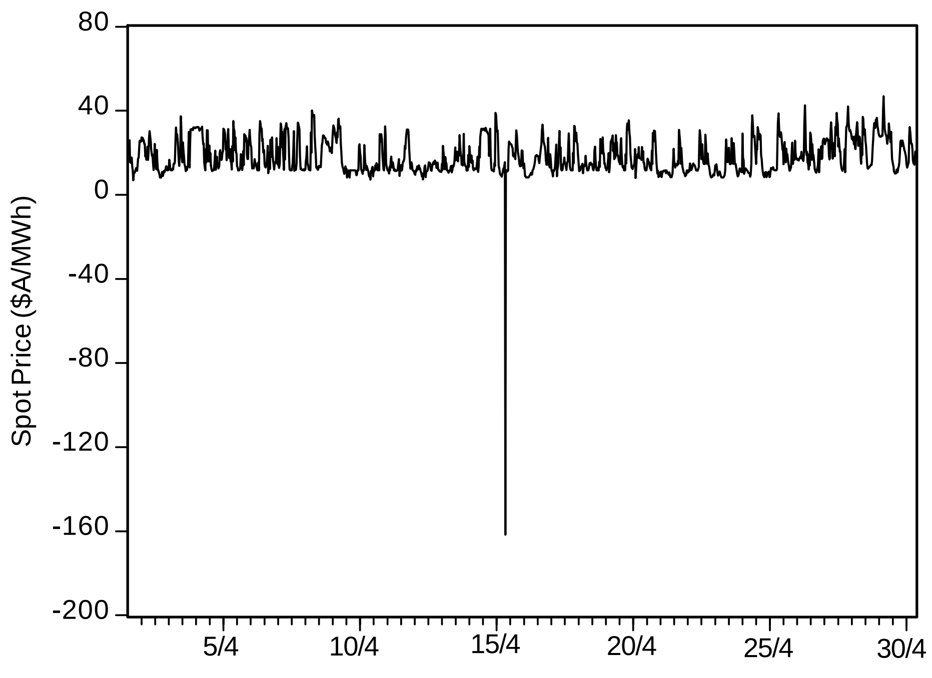

Figure 2: VIC1 Half-Hourly Spot Price for the Month of April, 2000

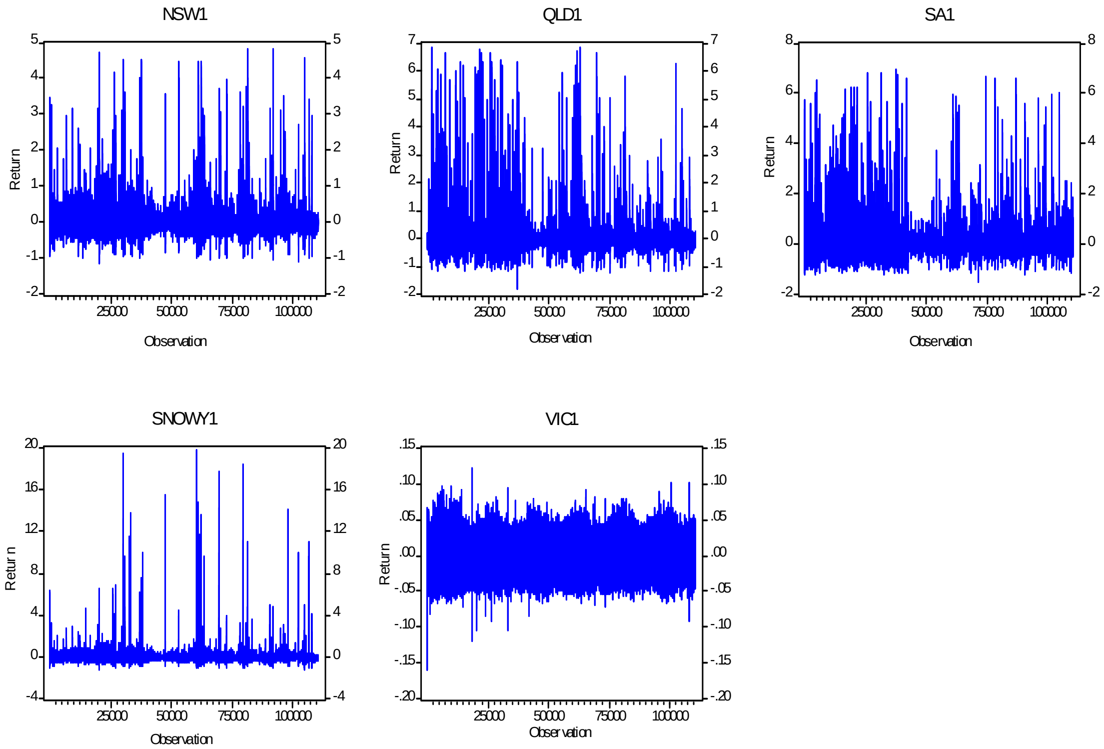

Figure 3: Filtered Discrete Returns by Region, for the period 7/12/1998 to 1/4/2005