Realized Volatility and Price Spikes in Electricity Markets: The Importance of Observation Frequency

Carl J. Ullrich*

June 29, 2009

ABSTRACT

This paper uses high frequency wholesale electricity spot price data from Australia, Canada, and the United States to estimate realized volatility and the frequency of price spikes. I find similar levels of realized volatility in Australia and North America, with estimates ranging from 1,500% to 3,000%. I present evidence that nonparametric jump detection tests based on the difference between realized variance and bipower variation are not reliable for electricity prices. Because daily electricity prices are averages, fitting models to data sampled at the daily frequency can never lead to a "correct" specification for the underlying data generating mechanism.

JEL CLASSIFICATION CODES: C14, Q40, L94, C52, G10.

KEY WORDS: Realized volatility, bipower variation, observation frequency, electricity markets, price spikes.

*Ullrich is at the Department of Finance, Insurance, and Business Law, Pamplin College of Business, Virginia Polytechnic Institute and State University, Blacksburg, VA 24061. Tel: (540) 231-2343, and email: cullr@vt.edu. This paper has been improved by discussions with Doron Avramov, Gurdip Bakshi, Yong Chen, Mike Cliff, John Easterwood, Doug Patterson, and Nagpurnanand Prabhala. I thank Gokhan Sonaer for assistance in gathering the data. The author is responsible for any errors. I welcome comments, including references to related papers that I have inadvertently overlooked.

A thorough understanding of volatility is crucial for many topics of interest to financial economists including derivative pricing, corporate risk management, market efficiency, and many others. Recent developments in the financial econometrics literature, combined with the availability of intradaily data, have led to much better estimates of volatility. Volatility estimates based on high frequency data, so-called realized volatilities, improve our understanding of price formation. One of the most promising uses of high frequency data is to discern jumps in otherwise continuous price paths.

Electricity is not yet storable in economically meaningful quantities and as a result electricity prices are extremely volatile. Supply and demand shocks are transmitted into prices almost instantaneously - literally at the speed of light - resulting in the price spikes endemic to electricity markets, episodes during which the price can increase by a factor of 100 or more, followed by a relatively quick return to normal levels. Price spikes make wholesale electricity markets very risky. A front page article (Smith (2008)) in the 17 July 2008 edition of *The Wall Street Journal* reports that five retail electricity companies in Texas failed when wholesale electricity prices spiked up to $4000. Typical wholesale electricity prices are $40 to $100. In order for firms operating in the industry to effectively manage such risk, it is necessary to have a good understanding price spikes and volatility.

Given the importance of price spikes in electricity markets, combined with the advent of high frequency electricity price data, it seems natural to apply realized volatility techniques to electricity prices. One particularly attractive feature of realized volatility and the associated jump detection tests is that they are nonparametric. The correct specification for modelling electricity prices is still an open question. One goal of this paper is to shed light on the behavior of electricity prices without resorting to parametric specifications for the underlying data generating process.

I calculate realized volatility and estimate the frequency of price spikes for eight wholesale electricity markets - five in Australia, one in Canada, and two in the United States. The Australian data is observed at the half-hourly frequency; the North American data is observed at the hourly frequency.

I make several contributions. First, the estimates of daily electricity price volatility reported in this paper, expressed in annualized standard deviation form, range from 1,500% to 3,000%. These estimates are much larger than previous estimates reported in the literature, which range from approximately 300% to 900%. Further, previous results suggest that Australian electricity markets are more volatile than markets elsewhere. Using high frequency data, I show that electricity markets in North America are, on average, just as volatile as markets in Australia. The reason for these seemingly contradictory results lies in the observation frequency of the data used to make the estimates. Previous results in the literature are based on daily prices, while the results presented here are based on intradaily prices. Unlike, e.g., daily stock prices which are sampled once per day, available daily electricity prices are averages of intradaily prices. Averaging prices *across* the day necessarily attenuates price variations *within* the day.

In order to reconcile previous results with those presented here I calculate two estimates of monthly volatility. First, I create average daily prices from intradaily data and calculate monthly volatility as the standard deviation of daily (log) price changes, which I term low frequency volatility. The low frequency volatility estimates are similar to estimates reported in the literature. Second, I calculate monthly realized volatility, which I term high frequency volatility. Not surprisingly, the estimates of high frequency volatility are much greater than the corresponding low frequency estimates. Low frequency and high frequency estimates are based on exactly the same raw data.

Ranking the volatility of markets based upon low frequency volatility, I find that Australian markets are more volatile than North American markets. Based upon high frequency volatility, I find that markets in Australia and North America display similar levels of volatility. Because price spikes in Australian markets tend to be larger than price spikes in North American markets (owing to lower price caps in North America), the dilutive effect of averaging on volatility is less pronounced in Australian markets. This explains why previous authors, using daily data, have concluded that Australian electricity markets are more volatile.

Next I show that the jump detection techniques developed for use in financial markets and based on high frequency data are less effective in electricity markets. There are two primary reasons. The first reason is because jumps in high frequency electricity prices inevitably are reversed within a few hours. It is the upward jump and the reversal together that define the price *spike*.

In the case of high frequency stock price data, microstructure noise induces spurious negative serial correlation in returns. Serial correlation causes the jump detection test statistic to be biased downward, i.e., to incorrectly classify some days as non-jump days when in fact a jump occurred. Andersen, Bollerslev, and Diebold (2007) show that offsetting, or lagging, returns in the calculation of the jump detection statistic serves to break the negative autocorrelation. Reversals in electricity prices also induce negative serial correlation in returns. Though the cause is different than the microstructure noise in equity prices, the effect is the same - the jump detection statistic underrejects the null of no jumps. I show that, in the case of electricity prices, it is crucial to increase the lag length by several periods in order to account for reversals.

The second, and more troubling, reason that jump detection based on high frequency electricity price data does not work well is because electricity prices can jump more than once in a single day. The jump detection methodology is based on the assumption that price jumps are rare and thus there is at most one jump per day. I present evidence that multiple jumps can fool the jump detection test.

Many authors have pointed out that, because electricity is effectively not storable, it is crucial to develop a good model for spot electricity prices. A spot price model is required for pricing electricity derivatives, both real and financial. Much of the existing electricity literature is devoted to determining which stochastic model best fits electricity prices. Such efforts typically rely on mean-reverting, jump-diffusion, and/or regime-switching models common from the modelling of equities, foreign exchange, and interest rates. Almost all of these model fitting exercises use data sampled at daily (or lower) frequencies. Such models are important for valuing options that settle against the daily average price, such as baseload power plants and many power purchase agreements. However, daily models must necessarily undervalue options which settle against intradaily prices.[2]

[2] Specifically, models of daily electricity prices must undervalue flexible technologies such as gas-fired combustion turbines.

It is well known that the distribution of average prices differs from the distribution of the prices from which the average is taken.[3] Because daily electricity prices are averages, fitting models to data sampled at the daily frequency can never lead to a "correct" specification for the underlying data generating mechanism. This is an important point that has not been emphasized in the literature.

The layout of this paper is as follows. Section I provides a literature review. Section II describes and summarizes the data. Section III reviews the theory and calculation of realized volatility and jump detection based thereupon, and discusses its application to the case of electricity prices. Section IV presents estimates of realized electricity volatility and Section V presents estimates for jump frequencies. Section VI presents realized volatility and jump frequency estimates for two alternative datasets, hourly Australian data and five-minute New England data. Section VII considers the pitfalls associated with fitting models of spot electricity to daily average prices. Section VIII summarizes the results and gives directions for further research.

I. Literature Review

The nascent financial economics literature about electricity prices is growing rapidly owing to (i) deregulation of electricity markets worldwide, (ii) the unique behavior of electricity prices, and (iii) the increasing availability of price data. In this section I review the relevant literature on electricity prices.

Deng (2000) is one of the first widely-cited papers in the literature in which the author attempts to model electricity prices using stochastic models. Deng (2000) relies on affine jump-diffusions[4] to model spot electricity prices. He incorporates regime switching between "abnormal" and "normal" states to account for price spikes. In later work, Deng (2005) uses real options methods to value power plants in the presence of jumpy electricity prices.

[3] In a separate context Nielsen and Sandmann (2003) write "The main problem ... with the arithmetic average ... is that the law of the arithmetic average is unknown."

[4] See Duffie, Pan, and Singleton (2000).

Escribano, Pena, and Villaplana (2002) study daily electricity prices for six different markets. They document that electricity prices are highly variable, mean-reverting, right-skewed, and severely leptokurtic, and they also document price spikes.

Atkins and Chen (2002) study electricity prices in the Alberta, Canada market. They, too, document all the usual properties - right-skewed prices, leptokurtic price and return distributions, and price spikes. Atkins and Chen (2002) suggest that the realized volatility approach of Andersen, Bollerslev, Diebold, and Ebens (2001) might be of use in characterizing the stochastic process responsible for generating electricity prices.

Guthrie and Videbeck (2002) and Guthrie and Videbeck (2007) examine high frequency spot price data from New Zealand. They treat electricity delivered at different times of the day as separate commodities. They point out that an accurate model of electricity spot prices is necessary for the purposes of derivative valuation and real options analysis often applied to power plants. They emphasize the need for intradaily data.

Goto and Karolyi (2004) study electricity price volatility in markets in Australia, the Nordic Power Exchange (NORDPOOL), and the United States. They document that electricity price returns have near zero mean, high volatility, and large excess kurtosis. Even after including jumps and ARCH effects, the residuals from their model are still autocorrelated, leaving (p.22), "... important unexplained systematic components ...".

Hadsell, Marathe, and Shawky (2004) study electricity volatility by fitting ARCH-type models to electricity prices for five US markets. In their words (p.24), "Understanding the volatility dynamics of electricity markets is important in evaluating the deregulation experience, in forecasting future spot prices, and in pricing electricity futures and other energy derivatives." They find that the persistence of electricity volatility has decreased over time and speculate that this might be due to learning effects in relatively new markets.

Cartea and Figueroa (2005) develop a model for electricity spot prices incorporating mean reversion, seasonality and, discontinuous jumps. They calibrate the model using the daily average of intradaily spot prices from the UK and then develop closed-form expressions for the forward price.

Knittel and Roberts (2005) analyze spot electricity prices from California and subsequently fit several models common in the finance literature. They document high variability, positive skewness, and an inverse leverage effect, which they attribute to convex supply curves (marginal costs).[5] They point out that negative prices, while unique to electricity markets and quite novel, are of little importance in pricing of financial securities.

Higgs and Worthington (2005) use high frequency data from Austrailian markets to model volatility using ARCH-type models. Consistent with previous work, they find that electricity price return distributions have essentially zero mean, high volatility, and fat tails. Interestingly, they proxy for information arrivals using demand volume. To my knowledge, theirs is the only paper to attempt to capture such an effect in electricity markets.

Hlouskova, Kossmeier, Obersteiner, and Schnabl (2005) develop a real options model to value power plants in the German market. They point out that real options models can be used in the electricity industry to assess investment decisions and to analyze a firm's portfolio of physical assets (e.g., power plants) and financial assets (e.g., futures and options contracts). Most importantly, they point out that in the presence of a well-functioning spot market the decision to operate a power plant is independent from other assets in the portfolio, both physical and financial. From p.300, "...the spot market automatically splits a utility into two separate firms: a power producer and a power marketer. ...the optimization problems relating to the optimal exercise of contracts can be solved separately from the overall portfolio problem."

Geman and Roncoroni (2006) use data from three markets in the United States to fit a mean-reverting, jump-diffusion (in their words a "jump-reversion") to daily logarithmic electricity prices. They point out

[5] The leverage effect in equity markets refers to the asymmetrical response of volatility to price changes. Prices and volatility are inversely related, but the volatility increase that accompanies a price decrease is greater in magnitude than the volatility decrease that accompanies a price increase.

that nonstorability means that one cannot derive spot price fundamentals from forward prices because the standard spot-forward relationship familiar in financial markets does not hold. They also point out that the Pennsylvania-New Jersey-Maryland (PJM) market design has functioned well and that the PJM transmission system has been very efficient.

Mount, Ning, and Cai (2006) fit a stochastic regime-switching model to PJM electricity prices. They suggest that price spikes are more common in Australian electricity markets because regulators in Australia are less willing to "... modify the behavior of suppliers ...". Consistent with Ullrich (2008), they show that the difference between available supply and contemporaneous demand is crucial in price formation.

Hambly, Howison, and Kluge (2007) develop a model for spot electricity prices with three separate components - seasonality, an Ornstein-Uhlenbeck process for spot prices, and a mean-reverting jump process. By separately modelling the O-U process and the jump process, they can specify different speeds of mean reversion for continuous price movements and jumps.

Huisman, Huurman, and Mahieu (2007) point out that, because day-ahead electricity prices are determined simultaneuously for all hours of the next day, it is incorrect to model these prices as a time series.

Zareipour, Bhattacharya, and Canizares (2007) study electricity price volatility in the Ontario market. They conclude that the Ontario electricity market is one of the most volatile markets worldwide. One potential reason is that Ontario has no day-ahead forward market.

Perhaps the most sophisticated effort to model electricity prices is in Pirrong and Jermakyan (2008). These authors use hourly demand and fuel price as state variables for modelling the price of electricity. Using data from PJM, Pirrong and Jermakyan (2008) find a significant risk premium in forward prices, even in the very short-term.

An excellent recent paper by Karakatsani and Bunn (2008) uses half-hourly data from the UK market to study price formation subsequent to market restructuring in March 2001. Karakatsani and Bunn (2008) introduce a regime switching model for each half-hourly trading period and show that fundamentals matter more during off-peak periods, while they find evidence of strategic behaviour (market power) during periods of high prices. They emphasize the need to go beyond purely statistical models of electricity prices.

Benth and Koekebakker (2008) model forward contracts written on electricity. Benth and Koekebakker (2008) provides a cautionary tale about the applicability of standard models used for equities and even other commodities to the case of electricity.

Huisman (2008) fits several regime switching models to average day-ahead electricity prices from the Dutch APX market and finds that the inclusion of temperature as a forecasting variable improves the fit of the model.

The work most closely related to this paper is Chan, Gray, and van Campen (2008). These authors also compute estimates of realized electricity price volatility and jump frequencies based on high frequency data. Chan, Gray, and van Campen (2008) recognize that unlike, e.g., equity prices, intradaily electricity price changes are not mean zero, owing to distinct seasonalities across the day and across the year. Hence, they model the drift of the price process as mean-reverting, incorporating dummy variables to account for seasonalities, and use the resulting estimates to demean returns.

This paper is different from Chan, Gray, and van Campen (2008) in important ways. First, I emphasize the different information content of intradaily data versus daily data. Second, while Chan, Gray, and van Campen (2008) define return to be the first difference in price levels, I use the first difference in log prices. Each approach has its advantages, but I choose to use log prices in order to make my results comparable to the finance literature. The conclusions of this paper are qualitatively similar in either case. Third, I compare the performance of four separate schemes to account for intradaily patterns. Fourth, Chan, Gray, and van Campen (2008) deal exclusively with Australian electricity markets. I analyze data from markets in Australia and North America.[6] Fifth, I show that jump detection tests based on the difference between realized variance and bipower variation require modification in the case of electricity prices.

[6]Previous work based on low frequency data concludes that Australian electricity markets are more volatile than other markets. I find that volatility is similar across markets.

II. Data

The data used in this study are publicly available[7] high frequency wholesale electricity spot prices from eight markets - five in Australia and three in North America. The Australian markets are New South Wales (NSW), Queensland (QLD), South Australia (SA), the Snowy Mountains (SNOWY), and Victoria (VIC). The Australian data cover the period January 1999 through April 2008. For the Australian markets, the raw data are half-hourly.

The North American markets are New England (NEISO), the eastern hub of Pennsylvania-New Jersey-Maryland (PJM), and Ontario, Canada (ONT). The NEISO data cover the period May 1999 through February 2008. The PJM data cover the period June 2000 through May 2007. The ONT data cover the period May 2002 through April 2008. For the North American markets, the raw data are hourly.

Electricity is produced and consumed all day every day, thus the data include weekends and holidays, and cover all hours (half-hours) of the day. I drop days that have missing return observations.

Table 1 lists overall summary statistics for prices and returns, i.e., log prices differences. Price data have units of dollars per megawatthour ($/MWh). The return data are in percent. While it is well known that electricity prices can be less than or equal to zero, negative prices render the concept of return poorly defined.[8] Low positive prices can also result in very large returns which, while likely to be classified as jumps, are not economically important. I drop all hours for which the spot price is less than or equal $5.00. I experimented with various values for the cutoff ranging from $0.00 to $10.00 with little effect on the results.[9]

[7] For NEISO, www.iso-ne.com. For PJM, www.pjm.com and see Longstaff and Wang (2004) for details. For ONT, www.ieso.com and see Zareipour, Bhattacharya, and Canizares (2007) and Zareipour, Canizares, and Bhattacharya (2007) for details. For the Australian markets, www.nemmco.com.au and see Higgs and Worthington (2005) and Higgs and Worthington (2008) for details.

[8] One of the advantages of defining return as the difference in price levels as in Chan, Gray, and van Campen (2008), rather than the log difference used here is that negative prices cause no difficulty.

[9] I also drop 8-May-2000 in NEISO. This particular day saw hourly spot prices reach $6,000, and thus induces large increases in spot price mean, standard deviation, skewness, and kurtosis. Arguably, any effort to determine jumps in electricity prices should include such an extreme outlier, however, removal of this one day has little effect on the estimates of realized volatility presented in Table 2 and the estimates of jump frequencies given in Table 4.

Table 1 confirms stylized facts reported in Escribano, Pena, and Villaplana (2002), Atkins and Chen (2002), Goto and Karolyi (2004), Knittel and Roberts (2005), Geman and Roncoroni (2006), Karakatsani and Bunn (2008), Higgs and Worthington (2008), and many others. Properties of electricity prices include large standard deviations (relative to the mean), positive skewness, severe leptokurtosis, and large positive outliers. Returns are near zero on average, have large standard deviations (relative to the mean), are approximately symmetrical, and display severe leptokurtosis.

Prices in Australian markets have higher standard deviations, are more positively skewed, and have fatter tails than prices in North American markets. Return distributions are more similar across markets, though return kurtosis is higher for the Australian markets.

III. Realized Volatility and Jumps

Recent work in financial econometrics shows that realized volatilty calculated from high frequency data provides a superior estimate of true latent volatility than estimates based on daily (or lower) frequency data.[10] Realized volatility techniques also show great promise in separating continuous price movements from discontinuous jumps.[11] Given the focus on price spikes in the electricity literature and the advent of high frequency electricity price data, it seems natural to apply these techniques to electricity prices.

In this section I present a brief review of the theory and calculation of realized volatility and the jump detection techniques developed by Barndorff-Nielsen and Shephard (2004b) and further refined by Huang and Tauchen (2005). I focus on the implementation of the calculations in the case of electricity prices.

[10] See, e.g., Andersen, Bollerslev, Diebold, and Ebens (2001), Andersen, Bollerslev, Diebold, and Labys (2001), Barndorff-Nielsen and Shephard (2004a), and the references therein.

[11] See, e.g., Andersen, Bollerslev, and Diebold (2007), Barndorff-Nielsen and Shephard (2004b), Barndorff-Nielsen and Shephard (2006), and the references therein.

A. Realized Volatility

Let P_t be the price of an asset at time-t. Let the natural logarithm of the price be denoted by $y_t = \log(P_t)$. Consider the following continuous-time jump-diffusion process for y_t,

$$dy_t = \mu_t dt + \sigma_t dZ_t + \kappa_t dK_t, \tag{1}$$

where μ_t is the drift, σ_t is the local price volatility, dZ_t is a Wiener process, κ_t is the jump size, and K_t is a Poisson counting process with (possibly) time-varying intensity λ_t.

Suppose that the price is observed at discrete times $j = 1,\ldots,M$ within each day $t = 1,\ldots,T$ and let $r_{t,j} \equiv y_{t,j} - y_{t,j-1}$ be the $\Delta \equiv \frac{1}{M}$ period return. Realized variance is equal to the sum of squared intradaily returns,

$$RV_t = \sum_{j=2}^{M} r_{t,j}^2. \tag{2}$$

Andersen, Bollerslev, Diebold, and Labys (2001) justify theoretically the use of realized variance as a proxy for unobserved quadratic variation. They show that, as the time interval between observations goes to zero, $\Delta \to 0$ (or $M \to \infty$) RV_t converges in probability to quadratic variation,

$$\lim_{M\to\infty} RV_t = \int_{t-1}^{t} \sigma(u)^2 du + \sum_{j=1}^{N_t} (\kappa_{t,j}^2). \tag{3}$$

The first term on the right hand side of equation (3) is the (continuous) integrated variance. The second term is the portion of overall variation that is due to the presence of jumps.

Barndorff-Nielsen and Shephard (2004b) and Barndorff-Nielsen and Shephard (2006) define realized bipower variation as

$$BV_t = \mu_1^{-2}\left(\frac{M}{M-(1+i)}\right)\sum_{j=2+(i+1)}^{M}|r_{t,j}||r_{t,j-(i+1)}|, \tag{4}$$

where $\mu_1 \equiv \sqrt{\frac{2}{\pi}}$ and i is the lag length in the multiplication of returns. In the limit $\Delta \rightarrow 0$ (or $M \rightarrow \infty$) BV_t converges to integrated variance,

$$\lim_{M\rightarrow\infty} BV_t = \int_{t-1}^{t} \sigma^2(s)ds. \tag{5}$$

Barndorff-Nielsen and Shephard (2004b) recognize that the contributions to overall variation can be separated into continuous and discontinuous parts. In particular, the portion of variation due to jumps (J_t) is just the difference between RV_t and BV_t,

$$J_t = RV_t - BV_t. \tag{6}$$

B. Jump Detection

In practice, J_t as defined in equation (6) can be negative and may result in many small positive jumps. Many of these jumps may be attributable to measurement error. In order to alleviate this problem, I rely on the simulation evidence in Huang and Tauchen (2005) and use the ratio statistic $Z_t(\Delta)$, defined as

$$Z_t(\Delta) \quad = \quad \frac{1}{\sqrt{\Delta}}\frac{[RV_t(\Delta)-BV_t(\Delta)]/RV_t(\Delta)}{\sqrt{(\mu_1^{-4}+2\mu_1^{-2}-5)\max(1,TQ_t(\Delta)/BV_t(\Delta)^2)}}, \tag{7}$$

where $TQ_t(\Delta)$ is the standardized realized tri-power quarticity

$$TQ_t(\Delta) \quad = \quad \left(\frac{M^2}{M-2(1+i)}\right)\left(\frac{1}{\mu_{4/3}^3}\right)\sum_{j=2+2(1+i)}^{M}|r_{t+j\Delta,\Delta}|^{4/3}|r_{t+(j-(1+i))\Delta,\Delta}|^{4/3}|r_{t+(j-2(1+i))\Delta,\Delta}|^{4/3}, \tag{8}$$

$\mu_{4/3} \equiv 2^{2/3}\Gamma(7/6)\Gamma(1/2)^{-1} \approx 0.8308609$, and i is as before the lag length in returns. The ratio statistic $Z_t(\Delta)$ defined in equation (7) is asymptotically normal under the null of no jumps. Andersen, Bollerslev, and Diebold (2007) define significant jumps as those for which $Z_t(\Delta)$ a critical value, i.e.,

$$J_{t,\alpha}(\Delta) \quad = \quad I_{Z_t(\Delta)>\Phi_\alpha}\Big(RV_t(\Delta) - BV_t(\Delta)\Big), \tag{9}$$

where $I_{x>y}$ is the indicator function which takes a value of one if $x > y$, and Φ_α is the value of the inverse cumulative standard normal distribution evaluated at α. In the empirical work that follows I set $\alpha = 0.99$.

The intuition behind the jump detection statistic Z_t in equation (7) is relatively simple. Realized variance (RV_t) and bipower variation (BV_t) both estimate variance. Realized variance is simply the sum of squared high frequency returns. Any abnormally large (positive or negative) return will be squared and thus have a large impact on RV_t. Bipower variation is the sum of the product of lagged (absolute) returns. Given that price jumps are rare events, the lag in returns in bipower variation means that a large return in one period will be multiplied by a small return from a nearby period and hence not have a large impact on BV_t. An abnormally large return therefore results in a large difference $RV_t - BV_t$, a correspondingly large value of Z_t, and the day is classified as a jump day.

Microstructure noise in high frequency stock prices induces negative serial correlation in returns. Andersen, Bollerslev, and Diebold (2007) demonstrate that this autocorrelation causes the jump detection statistic in equation (7) to be biased downward. They suggest increasing the lag in returns, i.e., increasing i in equations (4) and (8) above, in order to break the autocorrelation and improve the performance of the test.

C. The Case of Electricity Prices

Electricity prices behave differently than stock prices. A large upward movement in electricity prices inevitably is followed shortly thereafter by a reversal. Also, electricity prices have intraday patterns that vary by day of week and time of year. These peculiarities must be accounted for in the calculation of realized volatility and subsequent jump detection.

C.1. Accounting for Reversals

In the case of electricity prices, a large positive return, or jump, is followed soon thereafter by a large negative return, or reversal. That is, when prices jump, the subsequent returns effectively display negative serial correlation. The source of the autocorrelation is different than the microstructure noise in equity returns, but the effect is the same. The autocorrelation means that bipower variation is larger than it would be in the absence of the reversal, the difference $RV_t - BV_t$ and hence the Z_t statistic in equation (7) are reduced, and the test underrejects the null of no jumps. An example serves to illustrate the problem.

Figure 1 plots half-hourly spot prices and log returns for the SA market on 19 March 2003. In half-hours 16, 16.5, and 17, the spot price goes from \$34.27 to \$3,870, and then back to \$30.65. Thus, in half-hour 16.5 the return is large and positive, and in half-hour 17 the return is large and negative. This series of prices and returns surely qualifies as a price spike by any reasonable definition.

The realized variance for the day is $RV_t = 47.19$.[12] With the lag set to $i = 0$, the large positive return in half-hour 16.5 is multiplied by (the absolute value of) the large negative return in half-hour 17 in the calculation of BV_t, resulting in $BV_t = 42.44$. The Z_t statistic takes the value $Z_t = 0.841$ and the day is classified as having no jump.

Andersen, Bollerslev, and Diebold (2007) find that setting $i = 1$ (what they call "skip-one" returns) is sufficient to break the serial correlation due to microstructure noise. Accordingly, I recalculate BV_t and Z_t with the lag set to $i = 1$. Increasing the lag does not affect RV_t. However, lagging the returns means that the large returns in hours 16.5 and 17 are not multiplied together in the calculation of BV_t. In this case, $BV_t = 8.04$, $Z_t = 7.37$, and the day is classified as a jump day.

[12]These calculations use raw returns, with no drift adjustment. See the next subsection.

C.2. Intradaily Patterns

Chan, Gray, and van Campen (2008) (hereafter CGC) recognize that, because electricity prices vary throughout the day in predictable ways, in order to apply realized volatility techniques to electricity prices one must first demean returns. CGC specify and estimate a mean-reverting drift function to account for known seasonalities in electricity prices. They specify the drift $\mu_{t,j}$[13] in equation (1) as

$$\mu_{t,j} = \gamma(\theta_{t,j} - y_{t,j}), \tag{10}$$

where γ is the speed of mean reversion, the conditional mean $\theta_{t,j}$ is given by

$$\theta_{t,j} = \beta_0 + \beta_1 I_{offpeak} + \beta_2 I_{weekend} + \beta_3 I_{fall} + \beta_4 I_{winter} + \beta_5 I_{spring}, \tag{11}$$

and I is the indicator function. For example, $I_{offpeak} = 1$ if hour (half-hour) $j < 6$ or $j > 22$ and $I_{offpeak} = 0$ otherwise.[14] They estimate the coefficient vector $\Theta \equiv (\gamma, \beta_0, \beta_1, \beta_2, \beta_3, \beta_4, \beta_5)$ for each market via nonlinear regression and therefrom form an estimate of the drift $\widehat{\mu_{t,j}}$. They demean returns by replacing $r_{t,j}$ in equations (2), (4), and (8) by

$$r^*_{t,j} = r_{t,j} - \widehat{\mu_{t,j}}. \tag{12}$$

CGC point out (see their footnote 11) that the drift in equations (10) and (11) may be misspecified, and that any such misspecification affects their estimates of RV_t and BV_t. I repeat the calculations using four different demeaning schemes, i.e., different proxies for $\widehat{\mu_{t,j}}$.

[13] The additional subscript j emphasizes that the observation is in hour (half-hour) j on day t.

[14] Peak periods are defined to be from 6 am until 10 pm. For the half-hourly Australian data, $j = 0, 0.5, 1, 1.5, \ldots, 23, 23.5$.

A. Raw Returns: While it is true that intradaily returns are not mean zero, in every case they have very large standard deviations. As a baseline I perform the analysis with raw, unadjusted returns.

$$\widehat{\mu_{t,j}} = 0. \tag{13}$$

B. Demeaned Returns: In this case I demean returns by month of year, day of week, and hour (half-hour) of day. Return distributions for individual hours (half-hours) have non-zero skew and large excess kurtosis. One very large positive (or negative) return observation can exert undue influence on the mean return, hence I 'demean' high frequency returns using the hourly (half-hourly) median return.

$$\widehat{\mu_{t,j}} = \overline{r}_{mn,dy,hr}, \tag{14}$$

where $\overline{r}_{mn,dy,hr}$ is the median return for day-t in month mn, on day of the week dy, and in hour (half-hour) $j = hr$.

C. CGC Drift Specification: The specification of the drift is given by equation (10) and the conditional mean $\theta_{t,j}$ is given by equation (11).

D. CGC Drift Specification with Hourly Dummies: In this specification of the drift, I include hourly dummies rather than an offpeak dummy. The specification accounts for the fact the intradaily returns vary across hours within peak and off-peak periods. The specification of the drift is given by equation (10) and the conditional mean is

$$\theta_{t,j} = \beta_0 + \sum_{j=1}^{23} \beta_{1j} I_j + \beta_2 I_{weekend} + \beta_3 I_{fall} + \beta_4 I_{winter} + \beta_5 I_{spring}. \tag{15}$$

IV. Realized Volatility Results

Table 2 presents summary statistics for daily realized volatility in annualized[15] standard deviation form ($\sqrt{RV_t}$), for each of the four demeaning schemes (Panels A-D, corresponding to schemes A-D) from the previous section. For example, using the original CGC specification for the drift, from Panel C the mean daily (annualized) realized volatility for SNOWY is 1,986%. Figure 2 plots mean daily realized volatlity for each market by demeaning scheme. Several interesting facts emerge from Table 2 and Figure 2.

1. Electricity prices are extremely volatile. The lowest mean daily volatility reported in Table 2 is 1,516%, in NEISO using demeaned returns (Panel B). The highest is 3,066%, in PJM using raw returns (Panel A). These volatility estimates are roughly two orders of magnitude greater than similar estimates for equities, foreign exchange, and interest rates.[16]. Extreme volatility in electricity markets is caused primarily by the fact that electricity is not storable. Demand and supply shocks are transmitted into prices almost instantaneously, with no inventory to cushion the blow.

2. Electricity markets in Australia and North America display similar levels of realized volatility. For any demeaning scheme, PJM is the most volatile market[17] and NEISO is the least volatile market, as measured by $\sqrt{RV_t}$.

3. While the average level of realized volatility is similar in Australia and North America, the standard deviation of realized volatility is much higher in Australia. That is, intradaily volatility is itself more volatile in Australia than in North America. Also, the maximum observed daily realized volatility is higher in Australian markets than in North American markets. The likely reason is the level of the price cap in each market. In each of the North American markets, prices are capped at $1,000, while in Australia the cap is set at $10,000.

[15]Because the sample includes weekends and holidays, I assume a 365 day year.

[16]For example, Andersen, Bollerslev, and Diebold (2007) report daily realized volatilties of 17.7% for the S&P500, 12.8% for the DM$ exchange rate, and 9.7% for United States T-bonds. These numbers are based on the column labeled $RV_t^{1/2}$ in their Table 1A. I assume a 365 day year to ensure consistency with the results reported in herein.

[17]These results stand in contrast to the results of Zareipour, Bhattacharya, and Canizares (2007) who find that ONT is more volatile than PJM. However, these authors use day-ahead (forward) prices for PJM and real-time prices for ONT. I use real-time prices for both.

4. On average, each of the demeaning schemes B, C, or D reduces daily realized volatility relative to using raw returns, scheme A. Except for ONT, simply demeaning returns by month of year, day of week, and hour of day (Panel B) results in the lowest estimates for daily realized volatility. The inclusion of hourly dummies in the CGC drift specification (Panel D) decreases daily realized volatility relative to the original CGC drift specification (Panel C).

5. Examining the minimum values of realized volatility reported in Table 2 points to one problem introduced by demeaning raw returns. For every market under demeaning schemes C and D, and five of the eight markets under scheme B, the minimum daily realized volatility is higher than the minimum observed using raw returns. The reason is that prices do not fluctuate much on these very mild days, so they do not display the intraday patterns that the demeaning schemes are designed to eliminate. In this case demeaning the returns serves to *increase* volatility.

A. Comparison with Previous Results

Previous results reported in the literature suggest that Australian electricity markets are more volatile than North American electricity markets. Higgs and Worthington (2008) write that (p.3173) “In fact, the Australian electricity market is regarded as significantly more volatile and spike-prone than many comparable systems.” As documented in Table 2 and Figure 2, realized volatility calculated from intradaily data is similar across markets in North America and Australia.

Also, the volatility estimates reported in Table 2 are much greater than previous estimates reported in the literature. The Federal Energy Regulatory Committee (2004) estimates volatility at approximately 300% in United States markets. Booth (2004) estimates volatility in Australian markets at 900%. Figure 3.10 (p.86) in Eydeland and Wolyniec (2003) plots time series of annualized monthly electricity volatility for several United States markets, with values ranging from roughly 100% to 1,000%. The key to reconciling these seemingly contradictory results lies in the observation frequency of the data. The results reported in the literature are based upon *daily average* prices, while the results presented in this work are based upon intradaily data.

Daily electricity prices are different than, e.g., daily stock prices. Daily stock prices are sampled once per day, usually at closing. The realized volatility literature shows that sampling stock prices at intradaily frequencies provides better estimates of the true underlying volatility than does sampling at the daily frequency. Daily electricity prices are averages of intradaily prices. Volatility estimates based upon intradaily electricity prices measure a different quantity than volatility estimates based upon daily electricity prices.

B. Low Frequency vs. High Frequency Monthly Volatility

Table 3 reports summary statistics for two estimates of monthly volatility. For both estimates, I demean using scheme B. Low frequency volatility (σ_L) is the standard deviation of logarithmic daily price changes, where daily prices are the simple average of intradaily prices. Thus σ_L is directly comparable to previous volatility estimates reported in the literature such as those in Federal Energy Regulatory Committee (2004), Booth (2004), and Eydeland and Wolyniec (2003). High frequency volatility (σ_H) is the square root of monthly realized variance. For example, in NEISO the mean monthly low frequency volatility (σ_L) is 347%. The mean monthly high frequency volatility (σ_H) is 1,669%.[18] σ_L and σ_H are based upon exactly the same raw data.

The estimates of σ_L reported in Table 3 range from 347% (NEISO) to 779% (QLD). These estimates are similar in magnitude to those reported elsewhere in the literature and are consistent with previous evidence indicating that Australian electricity markets typically are more volatile than North American markets.

[18]For each market, monthly volatility (σ_H) reported in Table 3 exceeds the corresponding daily volatility ($\sqrt{RV_t}$) reported in Panel B of Table 2. The daily volatility in Table 2 is the square root of daily realized variance. The monthly volatility in Table 3 is the square root of monthly realized variance. Effectively, the mean in Table 2 is the average of the square root, while the mean in Table 3 is the square root of the average. Because the square root function is concave, the latter must always exceed the former.

Monthly volatility estimates based on high frequency data are much larger, with estimates of σ_H ranging from 1,669% (NEISO) to 2,988% (SA). It is not surprising that σ_H exceeds σ_L. Averaging prices *across* the day necessarily must reduce the impact of price variations *within* the day. Ye (2005) succinctly writes "... the average price of a stock over a period is less volatile than the stock price at a particular time."

Table 3 also reports correlation coefficients (ρ) and Spearman rank-order correlations between the time series of σ_L and σ_H. In all cases, $0 < \rho < 1$. Correlations between σ_L and σ_H are larger in Australian markets than in North American markets. The Spearman correlations uniformly are closer to one, and are statistically significant at the 1% level for all markets except PJM. In other words, high volatility months as measured by σ_L also tend to have high volatility as measured by σ_H, particularly in Australian markets. But there are some months that have relatively high volatility as measured by σ_L, but relatively low volatility when measured by σ_H, and vice versa.

Figure 3 plots the time series of σ_L and σ_H for each of the eight markets. The figure makes clear that, while the level of high frequency volatility is similar across markets in Australia and North America, σ_H is itself much more volatile in Australia.

C. Why Does Low Frequency Data Rank Australia as more Volatile?

Owing to higher price caps, the largest spot price observations in Australia are much larger than those in North America. (See Panel A of Table 1.) Averaging prices across the day has a greater effect in North America than in Australia. Hence, σ_L, which is based on average daily prices, is higher in Australia. This is why previous authors conclude that Australian electricity markets are more volatile than North American markets. However, using high frequency data retains within day variations. Based on these measures, North American markets are just as volatile, on average, as Australian markets.

V. Jump Results

In this section, I examine jump detection based on high frequency electricity price data. Previous results reported in the literature suggest that the frequency of jumps in electricity prices range from roughly 5% to 20%. I show that increasing the length of the lag i in the calculation of bipower variation improves the detection of jumps, but that the test still performs poorly when applied to electricity prices.

A. The Effect of Lag Length

Because the price increase on 19 March 2003 in SA is immediately reversed (see Figure 1 and the discussion in Section III.C), increasing the lag from $i = 0$ to $i = 1$ is sufficient to ensure that the day is classified as a jump day. However, not all price jumps reverse within one half-hour. Price spikes can persist for more than one half-hour, resulting in higher order serial correlation. Andersen, Bollerslev, and Diebold (2007) write that (p.711) "... higher-order serial dependence could be broken in an analogous fashion by further increasing the lag length." I therefore expect that increasing the lag length in the calculation of BV_t will improve the performance of the jump detection test. In the empirical work that follows, I examine the performance of the jump detection statistic Z_t for lag lengths ranging from $i = 0$ to $i = 5$. In total, I examine 24 separate cases, four demeaning strategies and six lag lengths.[19]

B. Jump Frequency Estimates

Table 4 presents jump frequencies for each of the four demeaning schemes (Panels A-D, corresponding to schemes A-D) at lag lengths ranging from $i = 0$ to $i = 5$. For example, based on the CGC demeaning scheme (Panel C) and lag length $i = 2$, jumps occur on 15.5% of the sample days in NEISO. The total sample size is 3,000 days (from Table 2) and 465 of these days are classifed as jump days. The jump

[19] Huang and Tauchen (2005) analyze the performance of the jump detection statistic for various values of lag length i using Monte Carlo simluations.

frequency is calculated as $\frac{465}{3,000} = 0.155$. Figure 4 presents the same jump frequency data in graphical form. All jump frequencies are calculated at the $\alpha = 0.99$ level of significance.

The frequency of jumps increases nearly monotonically with the lag length i in every market. Longer lag lengths improve the performance of the jump detection test by breaking the negative serial correlation in returns induced by price spikes. An example is 23 July 2001 in NEISO, the first plot in Figure 5, which is not classifed as a jump day with lag length $i = 1$, but is classified as a jump day with lag $i = 2$.

C. Comparison with Previous Results

Using daily Australian data and different methodology, Higgs and Worthington (2008) report jump frequencies ranging from 5.16% (NSW) to 9.44% (VIC). Goto and Karolyi (2004) use daily data from the United States, NORDPOOL, and Australia and find jump frequencies ranging from 1.34% (Mid Columbia, Washington) to 10.05% (PJM) in the United States, 4.39% (Helsinki, Finland) to 18.28% (Copenhagen, Denmark) in NORDPOOL, and 4.77% (VIC) to 12.46% (QLD) in Australia. CGC use a shorter sample of the Australian data together with the same jump detection techniques used here and find jump frequencies ranging from 7.5% (NSW) to 14.6% (SA).

To focus the discussion, and to facilitate comparison with the CGC results, consider the original CGC setup, demeaning scheme C with lag length $i = 1$. From Panel C of Table 4, the jump frequency estimates are 11.5% (NEISO), 4.1% (PJM)[20], and 10.3% (ONT) in North America, and 11.4% (NSW), 11.8% (QLD),

[20]PJM behaves differently than other markets. PJM is the the most volatile market as measured by daily realized volatility. The coefficient of variation (the coefficient of variation is the ratio of standard deviation to mean) of $\sqrt{RV_t}$ (based on demeaning scheme B) is only 0.33, the lowest any market considered in this paper. ONT is the only other market for which the coefficient of variation of $\sqrt{RV_t}$ is less than 0.50. The level of realized volatility in PJM is relatively high, but the volatility of realized volatility is relatively low. While PJM is highly volatile, it displays the lowest frequency of price spikes. Several papers, including Bessembinder and Lemmon (2002), Longstaff and Wang (2004), Mount, Ning, and Cai (2006), and Ullrich (2008), rely exclusively on PJM data.

15.6% (SA), 11.2% (SNOWY), and 11.1% (VIC) in Australia.[21] These estimates are similar in magnitude to those reported in the literature. However, there are several reasons to question these results.

First, the results in Table 4 and those reported in CGC are based upon intradaily data, while the results of Higgs and Worthington (2008) and Goto and Karolyi (2004) are based upon daily data. Averaging prices across the day should wash out some price jumps, and thus one would expect that jumps would be more frequent in intradaily data.

Second, using exactly the same methodology used here, Andersen, Bollerslev, and Diebold (2007) report jump frequencies (at the $\alpha = 0.99$ significance level; see their Table 3A) of 14.1% for the S&P500, 25.4% for the $\frac{DM}{\$}$ exchange rate, and 25.4% for United States T-bonds. Hence the jump frequency estimates for financial markets exceed the estimates for electricity markets. Given the unstorable nature of electricity, it seems unlikely that jumps occur less frequently in electricity prices than in financial prices.

Third, and most telling, given the CGC demeaning scheme and $i = 1$ lag length, 29 of the 50 most volatile days in the Australian markets are classfied as not having a jump. The spot price reaches at least $800 on each of those 29 high volatility, non-jump days. Clearly these are days that should be classified as jump days. The situation is similar for the North American markets, where 30 of the 50 most volatile days are classifed as non-jump days. Figure 5 plots the spot price and return for several days which are not classified as jump days.

Increasing the lag length dramatically improves the performance of the jump detection test, but there are still many days which the test misclassifies as non-jump days. Again using the CGC demeaning scheme B, and increasing the lag length to $i = 5$, the jump detection test still classifies 21 of the 50 most volatile days in Australian markets as non-jump days. The spot price reaches at least $1,500 on each of these days.

[21] CGC report (see their Table 4) jump frequency estimates of 7.5% (NSW), 12.0% (QLD), 14.6% (SA), 8.2% (SNOWY), and 10.0% (VIC). There are two reasons for the differences. First, I use a longer sample than CGC. Their sample ends in December 2006, mine continues through April 2008. Second, CGC define returns to be the first difference in prices, whereas I define return as the first difference in log prices. I have also performed the analysis using the CGC definition of return and I find estimates of jump frequencies which are for every market less than the CGC estimates. Hence, jumps were less frequent in the January 2007 - April 2008 time period than in the CGC sample period. Defining returns as log price differences increases the jump frequency estimates relative to defining returns to be price differences. However, the conclusions of this section are unchanged by the use of log prices.

The reason is because electricity prices can jump more than once in a single day. The jump detection statistic Z_t in equation (7) is based on the assumption that jumps are rare events and that there can be at most one jump in any one day. This assumption clearly does not hold in the case of electricity prices. Electricity prices can jump more than once in any one day, and multiple jumps can fool the test by inflating the bipower variation (BV_t) estimate. An example from North America is 25 February 2003 in ONT (the second plot in the first column in Figure 5) which is not classified as a jump day even with the lag length set to $i = 5$.

VI. Robustness

In this section, I repeat the calculations of realized volatility and the jump detection tests for (i) hourly Austrailian data, and (ii) five-minute NEISO data.

A. Hourly Australian Data

I average half-hourly Australian data across each hour, thereby producing data at the hourly frequency. I then compute realized volatility and jump frequency estimates from this hourly Australian data. The results (based on demeaning scheme B) are presented in Table 5.

The realized volatilities for Australian markets presented in Panel A of Table 5 are less than the corresponding results from Panel B of Table 2. Aggregating the data up to the hourly frequency reduces realized volatility by approximately 10% to 15%. Based upon hourly data, both PJM and ONT are more volatile than any of the Australian markets. Based upon hourly data, both NSW and SNOWY are less volatile than any of the North American markets.

From Panel B of Table 5, using hourly data also reduces the jump frequency estimates for the Australian markets, by approximately 30% compared to the corresponding results in Panel B of Table 4. Based on these results, it is tempting to conclude that North American markets (except for PJM) are just as prone to price spikes as Australian markets. However, such a conclusion is premature given the poor performance of the jump detection test. Still, these results dramatically illustrate the effects of using average price data.

B. Five Minute NEISO Data

As a further illustration of the importance of averaging, I recompute the realized volatility estimates and jump frequencies for one year (2003) in NEISO using data sampled at the five minute frequency.[22] In order to save space, I do not report the results in a table. Focusing again on demeaning scheme B, the realized volatility estimate is 2,783%. Based upon hourly data, the realized volatility for 2003 is 1,508%. I also estimate the jump frequency at lag lengths ranging from $i = 0$ to $i = 5$. The estimates of the jump frequency exceed 89% in every case.

VII. Implications for Modelling Electricity Prices

The fact that daily electricity prices are averages has implications for the model fitting exercises common in the literature. Many authors, for example Guthrie and Videbeck (2002), Hadsell, Marathe, and Shawky (2004), and Hlouskova, Kossmeier, Obersteiner, and Schnabl (2005), point out the need to develop good models of electricity spot prices. Much of the existing electricity literature attempts to fit stochastic models to electricity prices. In almost every case, the data used to estimate the model parameters and to compare competing models are daily average prices. Examples include Deng (2000), Goto and Karolyi (2004), Hadsell, Marathe, and Shawky (2004), Knittel and Roberts (2005), Cartea and Figueroa (2005), Mount, Ning, and Cai (2006), and Geman and Roncoroni (2006).[23]

[22] The five minute price data are also available on the NEISO website. I average the five minute data across hours to confirm that the hourly data are indeed built from the five minute data.

[23] Two exceptions are Karakatsani and Bunn (2008) who use half-hourly prices from the UK, and Pirrong and Jermakyan (2008) who use hourly prices from PJM. Higgs and Worthington (2008) recognize the loss of information caused by the use of daily average prices, but choose to use daily data anyway due to the " ... unwieldiness of intraday information."

Power plants are call options written on the "spark spread," the spread between (i) the cost of fuel required to generate electricity and (ii) the sales price of electricity. If the spot price of electricity exceeds the cost to generate the electricity, then the option is in-the-money and will be exercised. That is, the owner of the power plant will purchase fuel, convert the fuel into electricity, and sell the resulting electricity. If the cost to generate electricity exceeds the spot price of electricity, then the option is out-of-the-money and the power plant will not operate.

Flexible power generation technologies, such as natural gas-fired combustion turbines, can respond to price changes at the intradaily level. The owner of such a flexible plant effectively owns a collection of spark spread call options. Foe example, the owner of a gas turbine in PJM has, for each day, 24 hourly spark spread call options.[24]

Baseload technologies, such as coal-fired and nuclear power plants, are not designed to respond to intradaily prices. These technologies are effectively Asian spark spread options, i.e., they settle against the average daily (or weekly) price. Similarly, many power purchase agreements (financial contracts) are settled based on average prices. It is well known that Asian options are worth less than collections of otherwise similar individual options.[25] The reason is because the average price is less volatile than the prices from which the average is calculated (i.e., $\sigma_L < \sigma_H$) and option value increases in volatility.

In practical applications, the model should be fitted to data observed at the frequency that is relevant to the problem at hand. Fitting a model to daily average prices is a reasonable and practical exercise if the goal is to value baseload power plants and PPAs which settle against daily average prices. However, such a model must necessarily undervalue technologies which respond to intradaily prices.

Consider the simple example an option with a strike price is \$50, and a day for which the price is \$40 for hours 1-12 and \$60 for hours 13-24. An hourly option would be exercised in hours 13-24 and earn \$10

[24]Operational constraints (frictions) such as minimum up- and down-times, ramp rates, startup and shutdown costs, etc., reduce the frequency with which these options may be exercised.

[25]See, for example, the textbook treatment of McDonald (2006) on pp.444-449.

per hour, or $120. An Asian option which settles against the daily average price would be at-the-money and earn $0.

Further, because available electricity prices are averages, the search for the “correct” stochastic model for electricity prices depends crucially on the frequency at which the data is observed. Even if a particular model cannot be rejected when fit to daily average prices, it cannot be said to be correctly specified relative to the underlying data generating process. This is an important point that has not been emphasized in the literature.

VIII. Conclusions

This paper reports estimates of realized volatility and the frequency of price spikes for eight wholesale electricity markets - five in Australia, one in Canada, and two in the United States. The estimates of daily realized volatility, expressed in annualized standard deviation form, range from 1,500% to 3,000%. These estimates are much larger than previous estimates reported in the literature, which range from approximately 300% to 900%. Further, previous results suggest that Australian electricity markets are more volatile than markets elsewhere. Using high frequency data, I show that electricity markets in North America are, on average, just as volatile as markets in Australia. The reason for the differences is the observation frequency of the data. Previous results are based on daily data, while the results reported here are based on intradaily data.

I present evidence that jump detection techniques based on the difference between realized volatility and bipower variation are not reliable when applied to electricity prices because (i) reversals in electricity prices induce negative serial correlation in returns, and (ii) electricity prices can jump more than once in a single day. Adjusting the lag length in the calculation of bipower variation can overcome negative serial correlation. One potential avenue for future research is to develop modified jump detection statistics to account for the unique properties of electricity prices.

Daily electricity prices are averages of intradaily prices. Model fitting exercises that use data sampled at daily (or lower) frequencies, while useful in certain applications, can never lead to the correct specification for the underlying data generating mechanism. Similarly, hourly (half-hourly) electricity prices are averages of intrahourly prices, so fitting models to hourly data suffers from the same problem. The lesson is that the modelling exercise should be based upon data observed at a frequency that corresponds to the particular problem at hand.

References

Andersen, T. G., T. Bollerslev, and F. X. Diebold, 2007, "Roughing it up: Including jump components in the measurement, modeling, and forecasting of return volatility," *Review of Economics and Statistics*, 89, 701–720.

Andersen, T. G., T. Bollerslev, F. X. Diebold, and H. Ebens, 2001, "The distribution of realized stock return volatility," *Journal of Financial Economics*, 61, 43–76.

Andersen, T. G., T. Bollerslev, F. X. Diebold, and P. Labys, 2001, "The distribution of realized exchange rate volatility," *Journal of the American Statistical Association*, 96, 42–55.

Atkins, F. J., and J. Chen, 2002, "Some statistical properties of deregulated electricity prices in Alberta," working paper, University of Calgary.

Barndorff-Nielsen, O., and N. Shephard, 2004a, "Econometric analysis of realized covariation: high frequency based covariance, regression, and correlation in financial economics," *Econometrica*, 72, 885–925.

——— , 2004b, "Power and bipower variation with stochastic volatility and jumps," *Journal of Financial Econometrics*, 2(1), 1–48.

——— , 2006, "Econometrics of testing for jumps in financial economics using bipower variation," *Journal of Econometrics*, forthcoming.

Benth, F. E., and S. Koekebakker, 2008, "Stochastic modeling of financial electricity contracts," *Energy Economics*, 30, 1116–1157.

Bessembinder, H., and M. L. Lemmon, 2002, "Equilibrium pricing and optimal hedging in electricity forward markets," *Journal of Finance*, 57, 1347–1382.

Booth, R., 2004, "Too much volatility is bad for you," working paper, Bardak Management Services.

Cartea, A., and M. G. Figueroa, 2005, "Pricing in electricity markets: A mean reverting jump diffusion model with seasonality," *Applied Mathematical Finance*, 12, 313–335.

Chan, K. F., P. Gray, and B. van Campen, 2008, "A new approach to characterizing and forecasting electricity volatility," *International Journal of Forecasting*, 24, 728–743.

Deng, S.-J., 2000, "Stochastic models of energy commodity prices and their applications: Mean-reversion with jumps and spikes," working paper, Georgia Institute of Technology.

———, 2005, "Valuation of investment and opportunity-to-invest in power generation assets with spikes in electricity price," *Managerial Finance*, 31, 94–114.

Duffie, D., J. Pan, and K. Singleton, 2000, "Transform analysis and asset pricing for jump-diffusions," *Econometrica*, 68, 1343–1376.

Escribano, A., J. I. Pena, and P. Villaplana, 2002, "Modeling electricity prices: International evidence," working paper, Universidad Carlos III de Madrid.

Eydeland, A., and K. Wolyniec, 2003, *Energy and Power Risk Management: New Developments in Modeling, Pricing, and Hedging*. John Wiley and Sons, Hoboken, New Jersey.

Federal Energy Regulatory Committee, 2004, "State of the Markets Report," http://www.ferc.gov/market-oversight/st-mkt-ovr/som-rpt-2004.pdf.

Geman, H., and A. Roncoroni, 2006, "Understanding the fine structure of electricity prices," *Journal of Business*, 79, 1225–1269.

Goto, M., and G. A. Karolyi, 2004, "Understanding electricity price volatility within and across markets," working paper, Ohio State University.

Guthrie, G., and S. Videbeck, 2002, "High frequency electricity spot price dynamics: An intra-day markets approach," working paper, New Zealand Institute for the Study of Competition and Regulation.

———, 2007, "Electricity spot price dynamics: Beyond financial models," *Energy Policy*, 35, 5614–5621.

Hadsell, L., A. Marathe, and H. A. Shawky, 2004, "Estimating the volatlity of wholesale electricity spot prices in the US," *The Energy Journal*, 25, 23–40.

Hambly, B., S. Howison, and T. Kluge, 2007, “Modelling spikes and pricing swing options in electricity markets,” working paper.

Higgs, H., and A. C. Worthington, 2005, “Systematic features of high-frequency volatility in Australian electricity markets: Intraday patterns, information arrivals and calendar effects,” *The Energy Journal*, 26, 23–41.

——— , 2008, “Stochastic price modeling of high volatility, mean-reverting, spike-prone commodities: The Australian wholesale spot electricity market,” *Energy Economics*, 30, 3172–3185.

Hlouskova, J., S. Kossmeier, M. Obersteiner, and A. Schnabl, 2005, “Real options and the value of generation capacity in the German Electricity market,” *Review of Financial Economics*, 14, 297–310.

Huang, X., and G. Tauchen, 2005, “The relative contribution of jumps to total price variance,” *Journal of Financial Econometrics*, 3, 456–499.

Huisman, R., 2008, “The influence of temperature on spike probability in day-ahead power prices,” *Energy Economics*, 30, 2697–2704.

Huisman, R., C. Huurman, and R. Mahieu, 2007, “Hourly electricity prices in day-ahead markets,” *Energy Economics*, 29, 240–248.

Karakatsani, N. V., and D. W. Bunn, 2008, “Intra-day and regime-switching dynamics in electricity price formation,” *Energy Economics*, (forthcoming).

Knittel, C. R., and M. R. Roberts, 2005, “An empirical examination of restructured electricity prices,” *Energy Economics*, 27, 791–817.

Longstaff, F. A., and A. W. Wang, 2004, “Electricity forward prices: A high-frequency empirical analysis,” *Journal of Finance*, 59, 1877–1900.

McDonald, R. L., 2006, *Derivatives Markets*. Addison-Wesley, 2nd edn.

Mount, T. D., Y. Ning, and X. Cai, 2006, “Predicting price spikes in electricity markets using a regime-switching model with time-varying parameters,” *Energy Economics*, 28, 62–80.

Nielsen, J. A., and K. Sandmann, 2003, "Pricing bounds on Asian options," *Journal of Financial and Quantitative Analysis*, 38, 449–473.

Pirrong, C., and M. Jermakyan, 2008, "The price of power: The valuation of power and weather derivatives," *Journal of Banking and Finance*, 32, 2520–2529.

Smith, R., 2008, "Deregulation jolts Texas electric bills," *The Wall Street Journal*, July 17, 2008, A1.

Ullrich, C. J., 2008, "Equilibrium forward premia in electricity markets," working paper, Virginia Polytechnic Institute and State University.

Ye, G. L., 2005, "Asian options can be more valuable than plain vanilla counterparts," *Journal of Derivatives*, pp. 56–60.

Zareipour, H., K. Bhattacharya, and C. A. Canizares, 2007, "Electricity market price volatility: The case of Ontario," *Energy Policy*, 35, 4739–4748.

Zareipour, H., C. A. Canizares, and K. Bhattacharya, 2007, "The operation of Ontario's competitive electricity market: Overview, experiences, and lessons," *IEEE Transactions on Power Systems*, 22, 1782–1793.

TABLE 1: Hourly Price and Return Summary Statistics

The table summarizes wholesale electricity spot prices and returns. Price data have units of dollars per megawatthour ($/MWh). The return data are in percent. NEISO is New England, PJM is the eastern hub of Pennsylvania-New Jersey-Maryland, and ONT is Ontario, Canada. The NEISO data cover the period May 1999 through February 2008. The PJM data cover the period June 2000 through May 2007. The ONT data cover the period May 2002 through April 2008. The raw data from the North American markets is observed at the hourly frequency. NSW is New South Wales, QLD is Queensland, SA is South Australia, SNOWY is the Snowy Mountains, and VIC is Victoria. The Australian data cover the period January 1999 through April 2008. The raw data from the Australian markets is observed at the half-hourly frequency.

Panel A: Prices

	North America			Australia				
	NEISO	PJM	ONT	NSW	QLD	SA	SNOWY	VIC
NOBS	76,784	60,643	52,377	163,488	163,505	163,351	163,472	162,819
Mean	$52.11	$45.37	$53.20	$37.26	$39.62	$45.98	$35.12	$34.55
Stdev	$33.62	$38.37	$33.07	$188.10	$187.55	$222.84	$134.17	$126.92
Skew	10.40	7.70	5.11	35.19	29.45	32.56	36.81	41.02
Kurt	251.5	148.7	77.6	1,435	1,068	1,265	1,606	2,230
Min	$5.01	$5.00	$5.00	$5.00	$5.09	$5.00	$5.10	$5.00
Max	$1,003	$1,020	$1,028	$9,936	$9,921	$10,000	$7,716	$10,000

Panel B: Returns

	North America			Australia				
	NEISO	PJM	ONT	NSW	QLD	SA	SNOWY	VIC
NOBS	76,543	60,308	52,247	163,452	163,458	163,268	163,424	163,590
Mean	-0.101%	-0.025%	-0.005%	0.001%	0.007%	-0.003%	0.002%	-0.005%
Stdev	20.87%	34.24%	27.21%	19.02%	25.98%	26.32%	18.44%	20.08%
Skew	0.193	0.068	0.047	0.479	0.335	-0.525	0.635	0.350
Kurt	17.48	5.17	9.43	109.5	76.93	80.95	92.38	78.84
Min	-281.5%	-300.8%	-265.8%	-572.3%	-530.9%	-610.0%	-481.3%	-487.9%
Max	331.8%	202.5%	290.4%	544.7%	591.4%	597.1%	496.5%	496.6%

TABLE 2: Daily Realized Volatility Summary Statistics

The table summarizes the distributions of daily realized volatility ($\sqrt{RV_t}$), expressed in annualized percentage terms, assuming a 365 day year, for each of the four demeaning schemes detailed in Section III. NEISO is New England, PJM is the eastern hub of Pennsylvania-New Jersey-Maryland, and ONT is Ontario, Canada. The NEISO data cover the period May 1999 through February 2008. The PJM data cover the period June 2000 through May 2007. The ONT data cover the period May 2002 through April 2008. NSW is New South Wales, QLD is Queensland, SA is South Australia, SNOWY is the Snowy Mountains, and VIC is Victoria. The Australian data cover the period January 1999 through April 2008.

Panel A: Raw Returns

	North America			Australia				
	NEISO	PJM	ONT	NSW	QLD	SA	SNOWY	VIC
NOBS	3,000	2,270	2,100	3,367	3,356	3,324	3,359	3,210
Mean	1,709%	3,066%	2,294%	2,054%	2,519%	2,721%	2,026%	2,238%
Median	1,522%	2,988%	2,142%	1,727%	1,733%	2,008%	1,731%	1,921%
Stdev	849%	941%	977%	1,463%	2,311%	2,156%	1,329%	1,415%
Skew	2.18	0.51	1.48	3.94	2.90	2.78	3.67	3.56
Kurt	10.53	3.53	8.25	25.68	13.80	12.38	21.46	20.08
Min	356%	497%	227%	326%	337%	576%	362%	470%
Max	7,891%	6,801%	10,290%	20,344%	21,850%	18,836%	13,290%	14,306%

Panel B: Demeaned Returns

	North America			Australia				
	NEISO	PJM	ONT	NSW	QLD	SA	SNOWY	VIC
NOBS	3,000	2,270	2,100	3,367	3,356	3,324	3,359	3,210
Mean	1,516%	2,712%	2,087%	1,675%	2,215%	2,394%	1,638%	1,785%
Median	1,305%	2,617%	1,928%	1,321%	1,381%	1,616%	1,305%	1,417%
Stdev	846%	901%	968%	1,403%	2,305%	2,215%	1,283%	1,397%
Skew	2.32	0.67	1.65	4.42	3.02	2.79	4.05	3.90
Kurt	11.25	3.61	8.88	30.47	14.48	12.33	24.52	22.64
Min	321%	709%	411%	392%	351%	520%	397%	412%
Max	7,976%	6,571%	10,141%	20,342%	21,698%	18,816%	13,237%	14,248%

TABLE 2: Daily Realized Volatility Summary Statistics - continued

The table summarizes the distributions of daily realized volatility ($\sqrt{RV_t}$), expressed in annualized percentage terms, assuming a 365 day year, for each of the four demeaning schemes detailed in Section III. NEISO is New England, PJM is the eastern hub of Pennsylvania-New Jersey-Maryland, and ONT is Ontario, Canada. The NEISO data cover the period May 1999 through February 2008. The PJM data cover the period June 2000 through May 2007. The ONT data cover the period May 2002 through April 2008. NSW is New South Wales, QLD is Queensland, SA is South Australia, SNOWY is the Snowy Mountains, and VIC is Victoria. The Australian data cover the period January 1999 through April 2008.

Panel C: CGC Drift Specification

	North America			Australia				
	NEISO	PJM	ONT	NSW	QLD	SA	SNOWY	VIC
NOBS	3,000	2,270	2,100	3,367	3,356	3,324	3,359	3,210
Mean	1,632%	2,852%	2,105%	2,015%	2,480%	2,643%	1,986%	2,183%
Median	1,441%	2,761%	1,942%	1,691%	1,713%	1,959%	1,686%	1,869%
Stdev	790%	827%	877%	1,417%	2,191%	2,034%	1,287%	1,366%
Skew	2.26	0.68	1.68	3.94	2.94	2.80	3.66	3.57
Kurt	10.96	3.49	8.92	25.48	14.04	12.48	21.28	20.02
Min	487%	981%	521%	435%	489%	729%	442%	513%
Max	7,488%	6,251%	9,439%	19,655%	21,115%	17,701%	12,828%	13,785%

Panel D: CGC Drift Specification with Hourly Dummies

	North America			Australia				
	NEISO	PJM	ONT	NSW	QLD	SA	SNOWY	VIC
NOBS	3,000	2,270	2,100	3,367	3,356	3,324	3,359	3,210
Mean	1,564%	2,746%	2,062%	1,911%	2,391%	2,514%	1,871%	2,040%
Median	1,363%	2,636%	1,889%	1,570%	1,596%	1,789%	1,554%	1,701%
Stdev	796%	810%	881%	1,383%	2,175%	2,053%	1,263%	1,356%
Skew	2.34	0.77	1.74	4.20	3.04	2.83	3.87	3.73
Kurt	11.31	3.66	9.22	27.90	14.68	12.62	22.87	21.25
Min	466%	919%	518%	660%	705%	697%	656%	719%
Max	7,486%	6,381%	9,523%	19,683%	21,109%	17,690%	12,913%	13,870%

TABLE 3: Monthly Realized Volatility Summary Statistics

The table summarizes the distributions of monthly volatility. Low frequency volatility (σ_L) is the standard deviation of logarithmic daily price changes. High frequency volatility (σ_H) is the square root of monthly realized variance. Both estimates of monthly volatility are expressed in annualized percentage terms, assuming a 365 day year. AC_m is the autocorrelation at lag m months. Panel A presents data for North American wholesale electricity spot markets. NEISO is New England, PJM is the eastern hub of Pennsylvania-New Jersey-Maryland, and ONT is Ontario, Canada. The NEISO data cover the period May 1999 through February 2008. The PJM data cover the period June 2000 through May 2007. The ONT data cover the period May 2002 through April 2008. I demean raw returns by month of year and day of week (scheme B).

Panel A: North America

	NEISO		PJM		ONT	
	σ_L	σ_H	σ_L	σ_H	σ_L	σ_H
NOBS	106	106	84	84	72	72
Mean	347%	1,669%	561%	2,785%	428%	2,267
Stdev	184%	469%	161%	497%	123%	625.9
Skew	3.22	0.97	1.73	0.84	1.69	1.37
Kurt	15.91	3.96	7.32	3.52	9.88	5.74
Min	156%	745%	331%	1,976%	186%	1,331
Max	1,390%	3,131%	1,289%	4,346%	1,037%	4,638
AC_1	0.363	0.474	0.440	0.668	0.383	0.628
AC_2	-0.054	0.237	0.297	0.384	0.201	0.349
AC_3	-0.159	0.198	0.020	0.212	0.205	0.243
AC_6	-0.046	0.037	0.069	-0.097	0.133	-0.010
AC_{12}	0.114	-0.045	0.201	-0.073	-0.053	-0.111
ρ		0.414		0.189		0.585
Spearman		0.542***		0.253**		0.573***

TABLE 3: Monthly Realized Volatility Summary Statistics - Continued

The table summarizes the distributions of monthly volatility. Low frequency volatility (σ_L) is the standard deviation of logarithmic daily price changes. High frequency volatility (σ_H) is the square root of monthly realized variance. Both estimates of monthly volatility are expressed in annualized percentage terms, assuming a 365 day year. AC_m is the autocorrelation at lag m months. Panel B presents data for Australian wholesale electricity spot markets. NSW is New South Wales, QLD is Queensland, SA is South Australia, SNOWY is the Snowy Mountains, and VIC is Victoria. The Australian data cover the period January 1999 through April 2008. I demean raw returns by month of year and day of week (scheme B).

Panel B: Australia

	NSW		QLD		SA		SNOWY		VIC	
	σ_L	σ_H	σ_L	σ_H	σ_L	σ_H	σ_L	σ_H	σ_L	σ_H
NOBS	112	112	112	112	112	112	112	112	112	112
Mean	656%	1,994%	779%	2,828%	776%	2,988%	559%	1,903%	607%	2,097%
Stdev	473%	894%	478%	1,490%	496%	1,302%	374%	838%	394%	844%
Skew	1.16	1.26	0.80	0.82	1.29	0.57	1.10	1.15	1.48	1.14
Kurt	3.49	4.47	2.96	2.82	4.13	2.39	3.28	3.90	4.80	3.78
Min	135%	850%	151%	862%	167%	1,048%	144%	869%	177%	933%
Max	2,082%	5,137%	2,203%	6,827%	2,413%	6,195%	1,747%	4,574%	1,889%	4,872%
AC_1	0.346	0.354	0.219	0.345	0.458	0.534	0.274	0.393	0.198	0.270
AC_2	0.107	0.155	-0.007	0.301	0.280	0.366	0.006	0.140	0.170	0.114
AC_3	-0.139	-0.050	-0.109	0.204	0.047	0.297	-0.268	-0.078	0.067	0.076
AC_6	0.006	0.082	-0.102	0.163	-0.129	0.172	0.085	0.176	0.013	-0.026
AC_{12}	0.032	-0.096	0.048	0.089	0.268	0.126	-0.052	-0.014	0.014	-0.005
ρ		0.698		0.746		0.750		0.771		0.754
Spearman		0.785***		0.814***		0.830***		0.827***		0.757***

TABLE 4: Jump Frequencies

The table summarizes jump frequencies, by market, for each of the four demeaning schemes detailed in Section III as a function of the lag length i in the calculation of BV_t and TQ_t, equations (4) and (8) in the text. A day is classified as a jump day if the ratio statistic $Z_t(\Delta)$, equation (7) in the text, exceeds the value of the inverse cumulative standard normal distribution evaluated at 0.99. That is, the day is classified as a jump day if $I_{Z_t(\Delta)>\Phi_{0.99}}$ takes a value of one, where $I_{x>y}$ is the indicator function which takes a value of one if $x > y$, and $\Phi_{0.99}$ is the value of the inverse cumulative standard normal distribution evaluated at 0.99. NEISO is New England, PJM is the eastern hub of Pennsylvania-New Jersey-Maryland, and ONT is Ontario, Canada. The NEISO data cover the period May 1999 through February 2008. The PJM data cover the period June 2000 through May 2007. The ONT data cover the period May 2002 through April 2008. NSW is New South Wales, QLD is Queensland, SA is South Australia, SNOWY is the Snowy Mountains, and VIC is Victoria. The Australian data cover the period January 1999 through April 2008.

Panel A: Raw Returns

	North America			Australia				
Lag (i)	NEISO	PJM	ONT	NSW	QLD	SA	SNOWY	VIC
0	0.065	0.036	0.088	0.044	0.054	0.100	0.046	0.039
1	0.175	0.096	0.207	0.151	0.183	0.224	0.153	0.141
2	0.257	0.142	0.285	0.234	0.270	0.342	0.242	0.251
3	0.286	0.157	0.323	0.405	0.429	0.483	0.425	0.430
4	0.321	0.153	0.337	0.514	0.527	0.557	0.519	0.528
5	0.358	0.178	0.349	0.537	0.572	0.566	0.532	0.547

Panel B: Demeaned Returns

	North America			Australia				
Lag (i)	NEISO	PJM	ONT	NSW	QLD	SA	SNOWY	VIC
0	0.075	0.031	0.097	0.079	0.075	0.134	0.081	0.071
1	0.165	0.078	0.186	0.158	0.188	0.259	0.160	0.177
2	0.241	0.128	0.271	0.247	0.298	0.385	0.257	0.290
3	0.289	0.147	0.297	0.357	0.389	0.472	0.364	0.388
4	0.313	0.153	0.324	0.423	0.471	0.527	0.431	0.446
5	0.357	0.185	0.333	0.440	0.500	0.545	0.447	0.470

TABLE 4: Jump Frequencies - continued

The table summarizes jump frequencies, by market, for each of the four demeaning schemes detailed in Section III as a function of the lag length i in the calculation of BV_t and TQ_t, equations (4) and (8) in the text. A day is classified as a jump day if the ratio statistic $Z_t(\Delta)$, equation (7) in the text, exceeds the value of the inverse cumulative standard normal distribution evaluated at 0.99. That is, the day is classified as a jump day if $I_{Z_t(\Delta)>\Phi_{0.99}}$ takes a value of one, where $I_{x>y}$ is the indicator function which takes a value of one if $x > y$, and $\Phi_{0.99}$ is the value of the inverse cumulative standard normal distribution evaluated at 0.99. NEISO is New England, PJM is the eastern hub of Pennsylvania-New Jersey-Maryland, and ONT is Ontario, Canada. The NEISO data cover the period May 1999 through February 2008. The PJM data cover the period June 2000 through May 2007. The ONT data cover the period May 2002 through April 2008. NSW is New South Wales, QLD is Queensland, SA is South Australia, SNOWY is the Snowy Mountains, and VIC is Victoria. The Australian data cover the period January 1999 through April 2008.

Panel C: CGC Drift Specification

	North America			Australia				
Lag (i)	NEISO	PJM	ONT	NSW	QLD	SA	SNOWY	VIC
0	0.049	0.023	0.046	0.037	0.044	0.063	0.035	0.034
1	0.115	0.041	0.103	0.114	0.118	0.156	0.112	0.111
2	0.155	0.056	0.140	0.200	0.199	0.265	0.213	0.239
3	0.180	0.061	0.157	0.323	0.287	0.334	0.337	0.350
4	0.208	0.065	0.166	0.387	0.330	0.382	0.392	0.415
5	0.259	0.090	0.185	0.383	0.349	0.378	0.392	0.406

Panel D: CGC Drift Specification with Hourly Dummies

	North America			Australia				
Lag (i)	NEISO	PJM	ONT	NSW	QLD	SA	SNOWY	VIC
0	0.052	0.022	0.051	0.040	0.044	0.069	0.040	0.043
1	0.115	0.039	0.107	0.124	0.132	0.157	0.124	0.137
2	0.166	0.059	0.153	0.235	0.231	0.263	0.240	0.259
3	0.194	0.070	0.170	0.347	0.299	0.337	0.362	0.382
4	0.219	0.070	0.172	0.375	0.343	0.381	0.376	0.408
5	0.244	0.093	0.187	0.415	0.375	0.382	0.419	0.440

TABLE 5: Daily Realized Volatility Summary Statistics using Hourly Data

The table summarizes (Panel A) the distributions of daily realized volatility ($\sqrt{RV_t}$), expressed in annualized percentage terms, assuming a 365 day year, and (Panel B) jump frequency estimates. The half-hourly Australian data have been aggregated up to hourly data, i.e., averaged within each hour. NEISO is New England, PJM is the eastern hub of Pennsylvania-New Jersey-Maryland, and ONT is Ontario, Canada. The NEISO data cover the period May 1999 through February 2008. The PJM data cover the period June 2000 through May 2007. The ONT data cover the period May 2002 through April 2008. NSW is New South Wales, QLD is Queensland, SA is South Australia, SNOWY is the Snowy Mountains, and VIC is Victoria. The Australian data cover the period January 1999 through April 2008. I demean raw returns by month of year, day of week, and hour of day (scheme B).

Panel A: Realized Volatility $\sqrt{RV_t}$

	North America			Australia				
	NEISO	PJM	ONT	NSW	QLD	SA	SNOWY	VIC
NOBS	3,000	2,270	2,100	3,379	3,371	3,349	3,378	3,252
Mean	1,516%	2,712%	2,087%	1,511%	1,893%	2,010%	1,456%	1,558%
Median	1,305%	2,617%	1,928%	1,171%	1,204%	1,376%	1,142%	1,229%
Stdev	,846%	,901%	,968%	1,348%	1,886%	1,849%	1,219%	1,259%
Skew	2.32	0.67	1.65	4.44	2.92	2.90	4.23	4.12
Kurt	11.25	3.61	8.88	28.59	13.45	12.96	25.24	24.09
Min	321%	709%	411%	357%	314%	367%	337%	340%
Max	7,976%	6,571%	10,141%	17,753%	17,440%	14,349%	11,283%	11,730%

Panel B: Jump Frequencies

	North America			Australia				
Lag (i)	NEISO	PJM	ONT	NSW	QLD	SA	SNOWY	VIC
0	0.075	0.031	0.097	0.046	0.060	0.049	0.046	0.048
1	0.165	0.078	0.186	0.086	0.124	0.148	0.090	0.105
2	0.241	0.128	0.271	0.180	0.236	0.228	0.169	0.190
3	0.289	0.147	0.297	0.210	0.300	0.261	0.205	0.212
4	0.313	0.153	0.324	0.249	0.342	0.298	0.245	0.256
5	0.357	0.185	0.333	0.301	0.389	0.338	0.297	0.301

FIGURE 1: Half-Hourly Prices and Returns for 19 March 2003 in South Australia (SA)

The figure plots half-hourly spot prices (solid line, left axis) and returns (dashed line, right axis) for a single day, 19 March 2003, in the South Australia electricity market. The figure illustrates the importance of increasing the lag length in the calculation of bipower variation BV_t (see equation (4) in the test) for the case of electricity prices. Because an upward price jump in electricity markets inevitably is followed by a reversal, or downward jump, setting the lag length equal to $i = 0$ inflates the value of BV_t and causes the jump detection statistic Z_t (see equation (7) in the test) to underreject the null of no jumps. With $i = 1$ this day is not classified as a jump day.

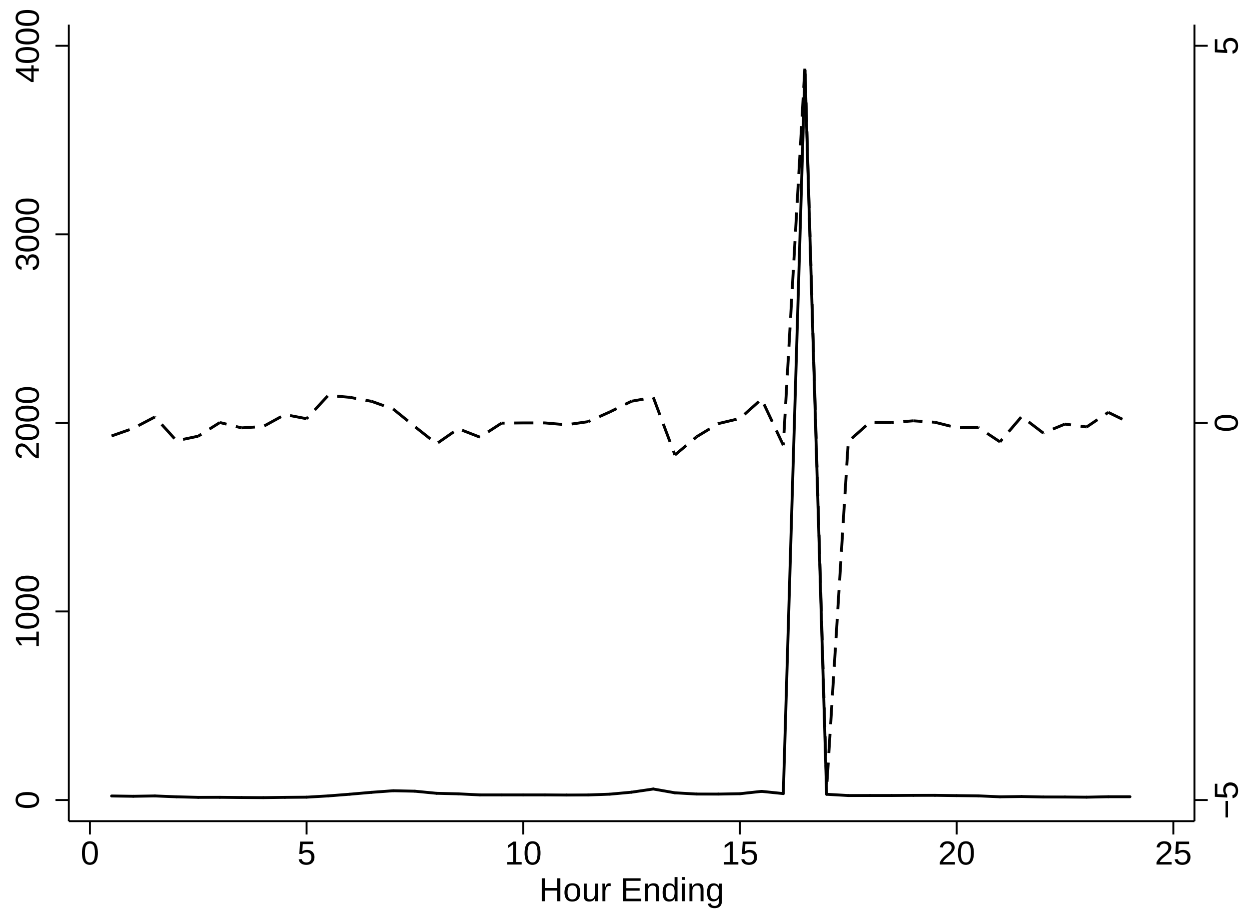

FIGURE 2: Mean Daily Realized Volatility ($\sqrt{RV_t}$)

The figure plots mean daily realized volatility $\sqrt{RV_t}$, expressed as an annualized standard deviation, for each of the demeaning schemes defined in Section III in the text. NEISO is New England, PJM is the eastern hub of Pennsylvania-New Jersey-Maryland, and ONT is Ontario, Canada. The NEISO data cover the period May 1999 through February 2008. The PJM data cover the period June 2000 through May 2007. The ONT data cover the period May 2002 through April 2008. NSW is New South Wales, QLD is Queensland, SA is South Australia, SNOWY is the Snowy Mountains, and VIC is Victoria. The Australian data cover the period January 1999 through April 2008.

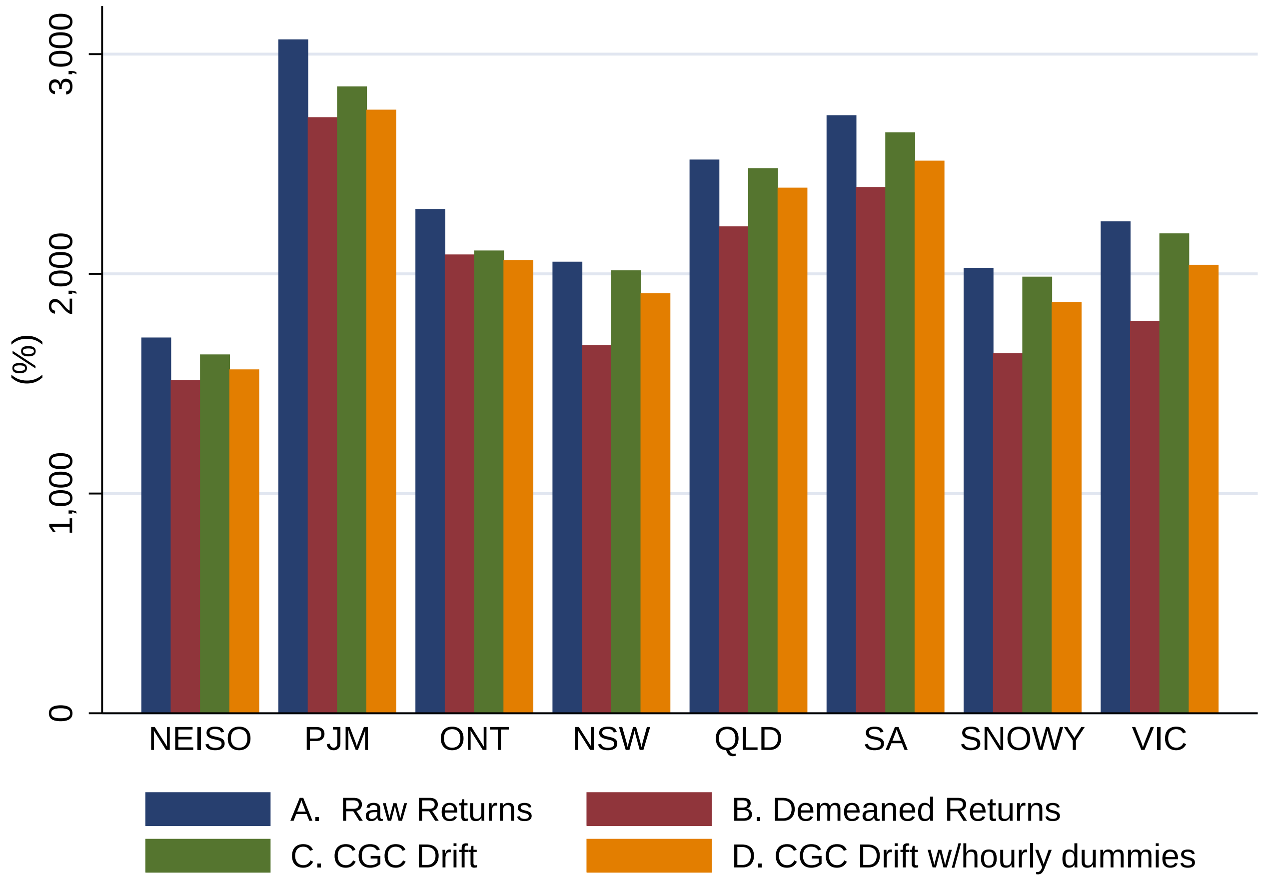

FIGURE 3: Monthly Electricity Volatility

The figure plots the time series of two estimates of monthly volatility. Low frequency volatility (σ_L, solid line) is the standard deviation of logarithmic daily price changes. high frequency volatility (σ_H, dashed line) is the square root of monthly realized variance. Both estimates of monthly volatility are expressed in annualized percentage terms, assuming a 365 day year. I demean raw returns by month of year and day of week (scheme B).

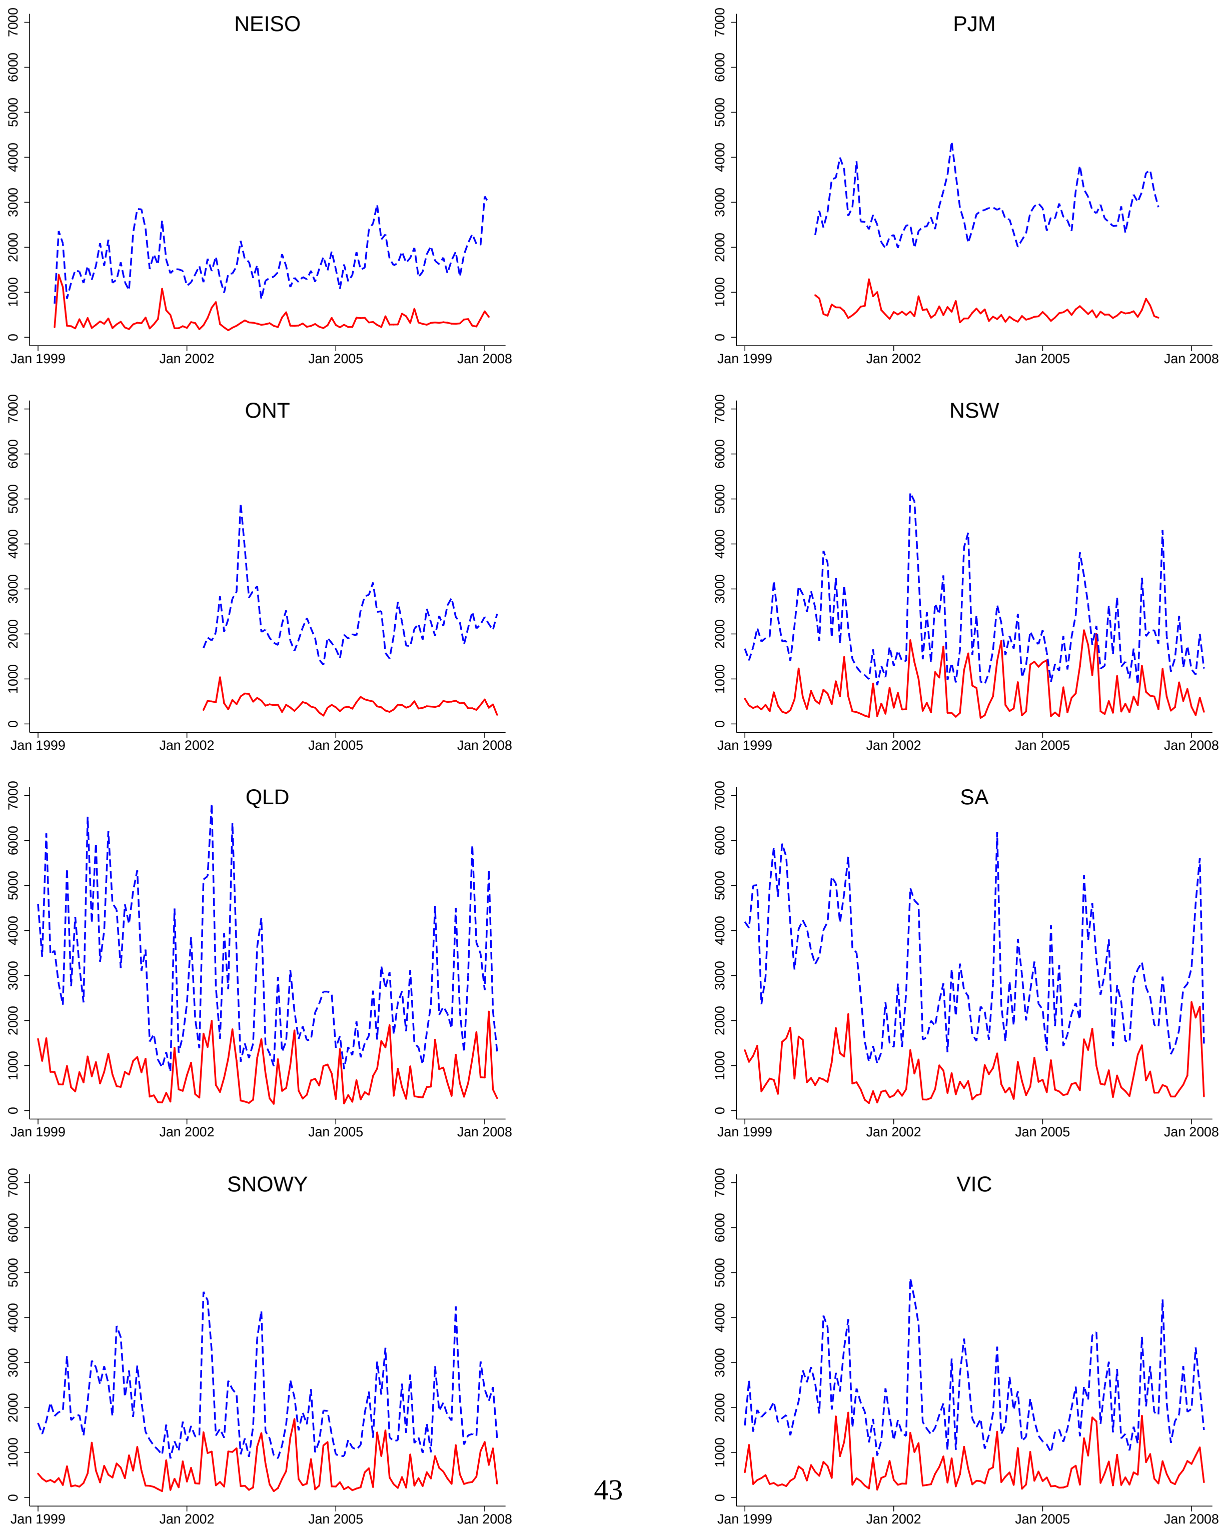

FIGURE 4: Jump Frequencies

The figure plots estimated jump frequencies for each of the demeaning schemes defined in Section III in the text and for lag lengths $i = 0$ through $i = 5$. NEISO is New England, PJM is the eastern hub of Pennsylvania-New Jersey-Maryland, and ONT is Ontario, Canada. The NEISO data cover the period May 1999 through February 2008. The PJM data cover the period June 2000 through May 2007. The ONT data cover the period May 2002 through April 2008. NSW is New South Wales, QLD is Queensland, SA is South Australia, SNOWY is the Snowy Mountains, and VIC is Victoria. The Australian data cover the period January 1999 through April 2008.

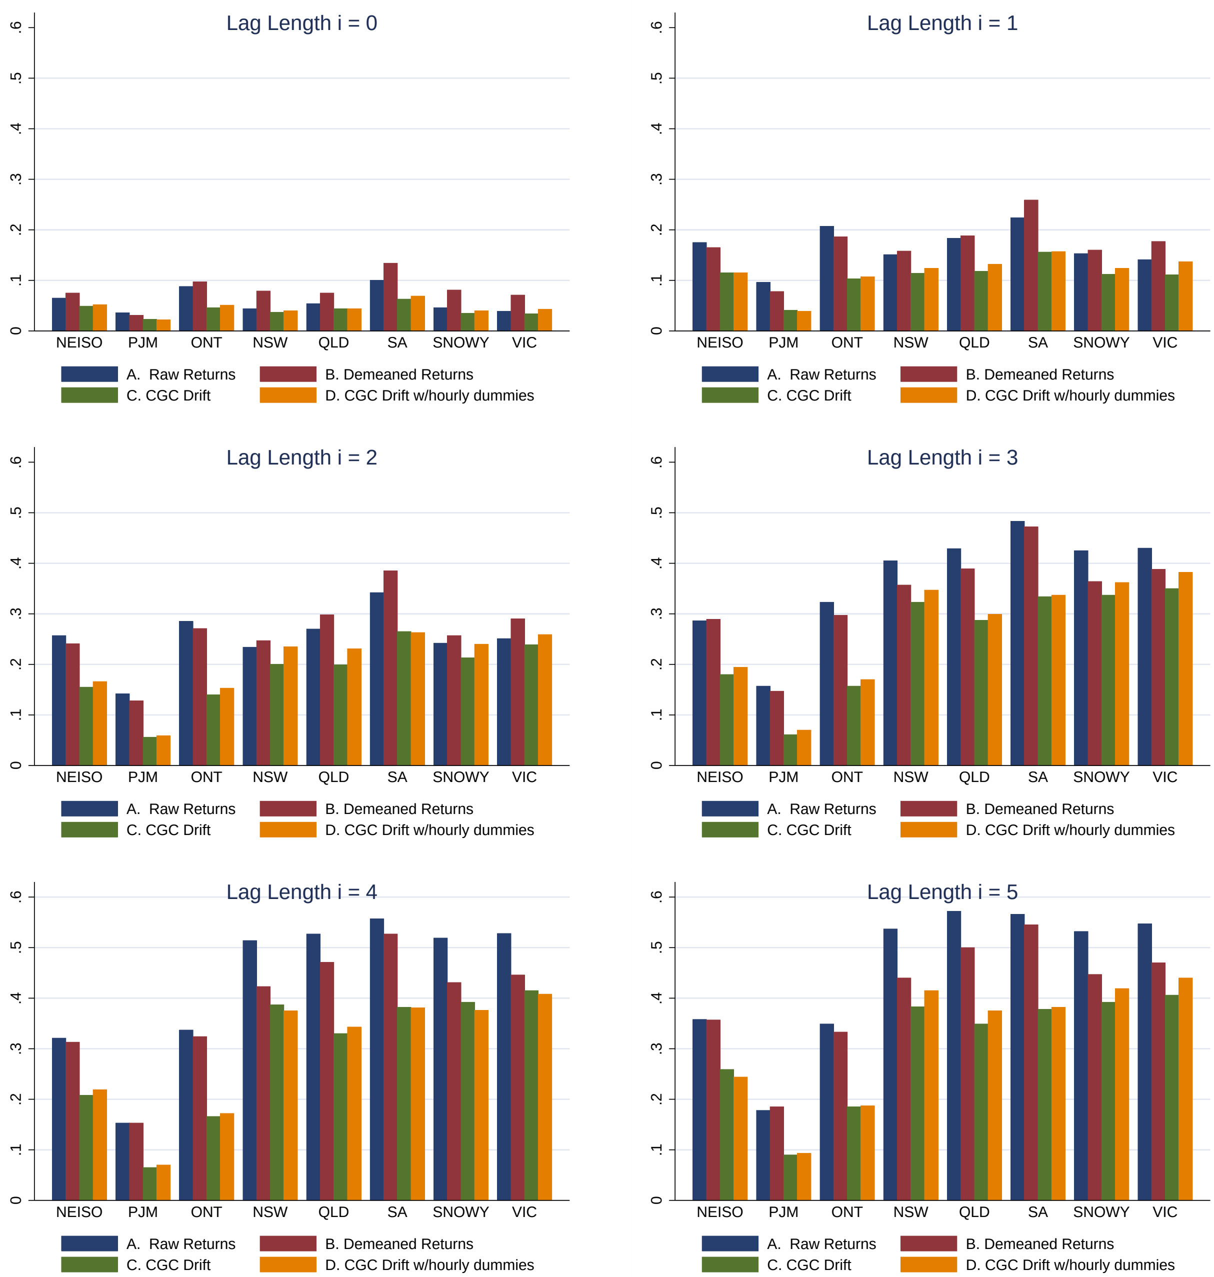

FIGURE 5: Non-Jump Days

The figure plots spot prices and (raw) returns for days classified as not having a jump. Spot prices (solid line) are on the left axis and returns (dashed line) are on the right axis. The jump detection is based upon demeaning scheme C and lag length $i = 1$.