Studies in Nonlinear Dynamics & Econometrics

Volume 10, *Issue* 3 2006 *Article* 4

NONLINEAR ANALYSIS OF ELECTRICITY PRICES

Risk Management and the Role of Spot Price Predictions in the Australian Retail Electricity Market

Maxwell J. Stevenson* Luiz Felipe Moreira do Amaral† Maurice Peat‡

*The University of Sydney, m.stevenson@econ.usyd.edu.au
†Pontificia Universidade Catolica Do Rio De Janeiro, lfelipe@ele.puc-rio.br
‡The University of Sydney, m.peat@econ.usyd.edu.au

Risk Management and the Role of Spot Price Predictions in the Australian Retail Electricity Market*

Maxwell J. Stevenson, Luiz Felipe Moreira do Amaral, and Maurice Peat

Abstract

This study investigates the extent to which predicted electricity spot prices from a statistical model, along with consensus forecasts issued by the Australian Financial Market Association (AFMA), provide unbiased price estimates of a forward contract price over a specified time to expiration. The statistical model is a regime switching time series model which is based on the dynamics of the market mechanism. To evaluate a price estimate, two criteria are utilized in order to conclude appropriateness for use in the marking-to-market process. First is the requirement that the predicted prices converge to the spot price at expiration of a hedging contract. The second criterion refers to the mis-pricing due to the price estimates over the days leading up to the contract expiration. Over the data period under consideration, the ranking of alternatives for generating price predictions is clear. On both criteria the Stevenson (2001) model is preferred. Of significance is the lack of support for the consensus (market) prices. They do not converge to the spot price at equilibrium and, further, they generate a considerable overvaluation of the risk management portfolio.

*Acknowledgements: We would like to thank Dr Adam Kucera, Integral Energy, for his encouragement and many worthwhile discussions. We acknowledge the School of Business at the University of Sydney for providing a research grant, and acknowledge the support of the Discipline of Finance at the School of Business, The University of Sydney, where one of the co-authors completed this paper as a Visiting Scholar. We also thank the CAPES foundation from Brazil for the financial support that assisted his travel and stay in Australia. Finally, but importantly, this paper has been enhanced by the insightful and helpful comments and suggestions of two anonymous referees.

1. Introduction

In Australia's increasingly deregulated National Electricity Market (NEM), electricity is traded with prices set every half-hour. Because electricity is a non-storable commodity and inventories cannot be used to smooth demand and supply shocks, market-clearing prices are characteristically volatile over time. One characteristic of this volatility is a marked variability in prices during a day, across days in the week and weeks in the year. Further, large transient price increases appear to occur in a random fashion and, on occurrence, exhibit rapid reversion to mean price levels.

Capacity constraints within electricity markets contribute to the characteristically extreme price volatility. Due to its non-storability, electricity supply can be inelastic at times. Intermittent changes to supply can arise as a result of network constraints caused by generator outages, transmission failures or generators exercising market power by implementing gaming strategies. Further, electricity retailers in the Australian market face fixed prices from their customers while being exposed to a floating wholesale price. Apart from the usual seasonal and diurnal variations, changes in demand are often precipitated by transient extreme weather events. As a consequence, short term demand can be extremely inelastic. The often dramatic effect that these changes have on market-clearing prices and volatility should be recognized in the modelling and forecasting of the electricity pricing process.

Derivatives are a means used by generators and retailers of electricity to hedge their exposure to these often abnormally high and rapidly mean reverting spot prices. With no transmission rights currently available in the NEM, derivative contracts are financial contracts and not contracts for physical delivery. These contracts include swaps, caps, floors, collars and swaptions. The ability to mark-to-market the entire "book" of derivative positions held by an electricity generator or retailer on a regular basis, is a requirement under the recently declared international accounting standard, IAS 39[1]. What is critical for this task is the availability of a price prediction mechanism that can be depended upon.

A forward price in the electricity industry indicates the price (or rate) at which market participants will transact today for the exchange of delivered electricity and its cash value at some specific time in the future. If the forward price is not equal to the spot price at expiration, arbitrage profits can be made using the standard no-arbitrage or cost-of-carry forward price model. This is only possible if inventories of the underlying asset can be held in order to arbitrage prices over time. As electricity is a non-storable commodity, an arbitrageur is unable to take a long position in the underlying asset and hold it until the contract

[1] IAS 39 requires companies to differentiate between either their hedging or speculative positions.

expiration date. It follows that the cost-of-carry model that links forward to spot prices cannot be used to price electricity forward contracts, except for the one-day-ahead forward market where generators can either deliver or purchase power from the pool on the following day.[2] For securities that have longer maturation times than a day, an accurate prediction of future realizations of electricity prices is needed.

The present study investigates the extent to which predicted electricity spot prices from a statistical model, along with consensus forecasts issued by the Australian Financial Market Association (AFMA), provide unbiased estimates of the realised price over a specified time to expiration. As previously noted, an important function of predictions of electricity prices is their role in marking-to-market hedging derivatives held for risk management. If the predicted price is a biased estimator of the subsequent forward price at maturity, then the effectiveness of a risk management system based on these prices will be severely diminished.

The paper is organized as follows. In the following section we describe the need for accurate predictions of future electricity spot prices for risk management purposes. In section 3, a statistical price prediction model and an Australian consensus Forward Price Curve are discussed, along with an explanation of the procedure followed for evaluating the different price predictions. A description of the data used in the study is found in section 4. Model estimation and price prediction evaluation results are presented in section 5, while section 6 contains our conclusions.

2. The Forward Price Market For Electricity In Australia

The structure of the Australian National Electricity Market (NEM) has generators offering energy to buyers (retailers) via a wholesale spot market. They are required to submit generation supply curves to a system operator twenty-four hours ahead of dispatch. A supply curve is created for each of the forty-eight half-hours in a day. This indicates how much energy each generator is willing to generate from each generating unit as the price varies. The system operator decides which generators should be dispatched when the spot price exceeds their offer price and determines the price every five minutes. Expected load (demand) is matched with the industry supply curve for a particular half-hour, with the price offered by the last generator dispatched forming the five-minute price. A time-weighted average of the five-minute prices sets the half-hour pool price. Generators and retailers use derivative contracts to reduce their exposure to pool price risk. In Australia, they are usually over-the-counter (OTC) contracts of

[2] This situation will be taken up in the following section.

varying maturity ranging from months to years, with the cash value negotiated between the parties without the involvement of intermediaries.

The Australian electricity market is continuing to undergo restructuring with increased importance being placed on the modeling and forecasting of prices. Under the old regulated market, prices were largely determined by average cost with minimal variation due to the strict control of regulators. Market entry was barred and investment in new generation facilities by existing generators linked to demand forecasts. Given the lack of volatility in prices under this regulated market, there was little need for hedging electricity price risk. Restructuring has brought about an abandonment of price controls, along with a removal of the regulatory barriers of entry for new generator and retail electricity companies. As a result, price volatility has increased dramatically. Hence, there is a need for derivative contracts to enable purchasers of electricity to hedge their price risk, and to enable existing and new entrants to facilitate the capitalization of strategic generating opportunities on offer in the restructured market.

A simple derivative that is often used to hedge price risk is a forward contract. The cash value of a forward contract is fixed at the time of contract signing. If a forward contract is traded on an exchange then it is known as a futures contract. At present, the futures market for electricity contracts traded on the Sydney Futures Exchange[3] is in an embryonic stage of development. Hedging and speculative opportunities are inhibited by a lack of liquidity. As reported by Bessembinder and Lemmon (2002), this is also the case for the larger US market. Even though the New York Mercantile Exchange lists futures prices for power delivery at the Pennsylvania, New Jersey, Maryland (PJM) market, activity levels are extremely low. One important explanation for this lack of liquidity is that trading by outside speculators is constrained to their taking positions in cash-settled contracts. The prices of these contracts are linked to physical-delivery contracts with liquidity in the latter restricted to the few trades by generators who can accomplish delivery.

The standard development of forward prices for commodities (Hull 2003) yields;

$$F_{t,T} = E_t\left[S_T\right]e^{(c-y)(T-t)} \qquad \ldots (1)$$

c is the cost-of-carry, which is defined to be the cost of storage plus the interest incurred in financing a long position in the commodity. The variable y is the convenience yield, which reflects the markets expectation of the future availability of the commodity. Electricity retailers are unable to store electricity against future demand, so the cost-of-carry component drops out of the calculation. In regulated electricity markets retailers are required to provide a continuous service, so under normal operations the market expects electricity to

[3] www.sfe.com.au

be available and convenience yield will be zero[4]. Under these conditions the variable which drives forward prices is the expected spot price at T. For electricity retailers, it follows that the expected spot price is a key input for marking-to-market derivative positions. As such, it is a critical component in their risk management processes.

Alternatively, there is a well-established literature that focuses on the implications for the relation between forward and expected spot prices based on equilibrium considerations (Bessembinder and Lemmon (2002) and Longstaff and Wang (2004)). The forward premium, often defined as the difference between a forward price and the corresponding expected spot price, is linked to economic risks and is represented as the equilibrium compensation for market participants accepting either one or both of demand and price risk. In adopting the equilibrium approach, Bessembinder and Lemmon (2002) explicitly model the economic determinants of market clearing forward electricity prices. The implication they draw from their study is that electricity forward prices will be generally biased forecasts of the future spot electricity price. Further, they claim that an anticipated increase in the volatility of wholesale spot prices will lead to a decrease in the forward premia, while there is a positive relation between forward premia and an anticipated increase in the skewness of wholesale spot prices.

Longstaff and Wang (2004) examine the pricing of electricity forward contracts in the day-ahead PJM electricity market and their relation to corresponding spot prices. In an empirical study, they confirm the implications from Bessembinder and Lemmon (2002) concerning the correlations between the forward premia and both the volatility and skewness of wholesale electricity prices. Further, they establish that risk measures such as unexpected changes in demand, spot prices and total revenues of market participants play a significant role in explaining the forward risk premium, as well as finding support for the existence of time-varying forward premia that are predominantly positive.[5]

In the PJM day-ahead forward market that forms the basis for the methodological and empirical implications drawn from the Bessembinder and Lemmon (2002) and Longstaff and Wang (2004) studies, clearing prices for next-day delivery are announced at 4 pm, along with production schedules for dispatching generation and trades between buyers and sellers. This market functions as a forward market run in parallel with the spot market. On a day-

[4] In cases of abnormal market conditions, when supply is limited due to equipment failure or when demand is high due to weather events, convenience yield can be high for short periods. As these events are random in nature they cannot be predicted and are therefore ignored in the following analysis.

[5] While the classical literature (see Keynes (1930) and Hicks (1939) for example) suggests that systematic hedging pressure effects imply a negative premium, Hirshliefer (1990) has provided examples that indicate the premium can be of either sign.

ahead basis, market participants can hedge against price risk by selling or buying electricity contracts. This market is modeled as a closed system with the only participants being generators and retailers. The demand and supply functions that determine equilibrium prices are relevant only for a short-term (day-ahead) forward price market.

This study is primarily concerned with pricing long-term contracts with maturity dates of up to five years for future delivery at a specified half-hour during the day. These contracts are written by generators and purchased by electricity retailers who have no generation capacity. The contracts are predominantly fixed-to-float swaps but do include other derivative products such as options, caps, swaptions and sculptured swaps. The market for these forward contracts is predominantly an over-the-counter (OTC) market that provides a very important hedging capacity for electricity retailers.

The AFMA Forward Price Curve is a resource available to participants in the Australian electricity market for pricing forward contracts. It is produced by the Australian Financial Management Association (AFMA), an association consisting of major participants in the electricity industry whose charter is to act as the registry of OTC derivative positions transacted, as well as to offer market enhancing resources. The AFMA Forward Price Curve is, essentially, a consensus forecast. The quorum of contributors required before a forward price estimate can be made must total 14 and is made up mainly from generators, as well as retailers and traders. On a daily basis, AFMA seeks forward estimates[6] for the next two calendar months, four quarters and for each year for the next four years. The contributor assesses the bid-offer in the market for the specified days at 3.30pm Australian Eastern Standard Time (AEST) for a 10 Mw fixed electricity swap.[7]

Generators who make up a large number of the contributors to the consensus-based AFMA forward price curve are also responsible for writing the hedging derivatives.[8] With a large proportional representation among contributors, generators are in a position to influence the AFMA forward curve estimates. Given the volatile nature of electricity prices, any upward bias in the forward price over different maturity dates will bring with it an increase in forecast spot price uncertainty. It follows that there should be a corresponding increase in the value of the derivative pool and, therefore, an increase in price of most derivative products that make up the pool. Conflict could arise if contributors offered inflated forward prices that translated into increased

[6] Estimates are sought for the peak load period (Peak) and all other periods (Flat).

[7] There are restrictions on the spread with no restriction for the next calendar month, two dollars for all periods up to and including calendar year two and five dollars for calendar years three and four.

[8] As generators can deliver electricity at a future date, it is not unreasonable to view the AFMA curve as a forward price curve.

derivative prices. This is a potential source of observed bias in the AFMA forward price estimates, and a possible reason for many market participants regarding these prices as misleading.[9]

From the perspective of an electricity retailer with no generation capacity, an appropriate measure of the price of a forward contract is the expected spot price on a future expiration date. One approach to predicting future spot price realizations is to model and forecast the underlying data generating process. Knittel and Roberts (2005) model the pricing process using various parametric models from the asset-pricing literature. One class of models is based on stochastic differential equations that incorporate various jump-diffusion processes designed to capture the volatile and mean reverting characteristics of the data. Another class uses time series models to model the persistence in the data and the effect of weather on the evolution of the pricing process. They conclude that the superior forecasting performance is realized from models that control for the characteristics of electricity prices that distinguish them from other asset prices.

In this study we use an existing time series model from the threshold autoregressive class to price forward contracts by estimating future spot price realizations at half-hour granularity. The model was first introduced by Stevenson (2001) to predict electricity prices. In the following section we describe the model and discuss the economic motivation for it in the context of electricity prices. Further, we discuss our approach for evaluating the forward price estimates from the time series model at 4.00pm each day over a defined contract period, with the corresponding consensus-based AFMA Forward Prices that are released daily to the market at the same time.

3. Pricing Forward Contracts in the Australian Electricity Market

The challenge in modeling and predicting electricity prices is to find a class of models able to take account of the distinct characteristics of this data, rather than be overwhelmed by them. The first step in this process is to filter the data so as to mitigate the effect of the extreme volatility while retaining the fundamental signal. We achieve this by applying wavelet analysis.[10]

We examined both the electricity price and demand (load) series at different time locations and levels of resolution to differentiate between what was signal and what was noise. Firstly, we cleansed the data of leakage from the high frequency mean reverting price spikes into the lower and more fundamental levels of frequency resolution. We then selected the level of resolution where the

[9] For further details see Anderson, Hu and Winchester (2006).

[10] We recognize there are several ways of filtering the spikes in the data other than using wavelet analysis. Spline fitting, simple ARMA modelling and state-space models are among the alternatives.

extreme volatility in the data was removed without compromising the integrity of the signal.[11] It is the reconstructed data at this level of resolution that we refer to as the filtered data. Using the filtered data, we used a time series model from the threshold autoregressive (TAR) class to price a forward contract. This model chosen was the Stevenson (2001) model.

To evaluate its performance, we compared the predicted prices from the model with the benchmark Australian market prices that were constructed by consensus and published by AFMA.[12] What is of relevance in the determination of the effectiveness of pricing forward contracts is to be able to accurately predict spot prices at the maturity of the contract, and for estimates of the forward prices over the life of the contract to gradually converge to the spot price at expiration. Additionally, an important consideration is how mis-pricing impacts on the cost of the hedging component of the risk management process. Furthermore, if estimates of forward contract prices are used to mark-to-market derivative positions within a portfolio, then the volatility in the bias of prior day-to-day forward price estimates from the actual price at expiration (that is, forecast error) needs to be minimal.

To determine the extent to which these desirable properties are inherent in a model used to forecast forward contract prices is essentially an evaluation of model risk. Under the Basel II convention, model risk has to be evaluated and documented before approval can be sought from the Bank of International Settlements. In Australia, this requirement is encompassed in the recently announced accounting standard, AASB7 (Financial Instruments: Disclosures).[13]

We complete this section by describing the economic motivation for, and the structure of the Stevenson (2001) model, as well as discussing our evaluation approach from a risk management perspective.

3.1 The Stevenson Model

The Stevenson (2001) model is a regime switching autoregressive model for electricity prices which is based on the properties of the supply and demand for electricity. Bessembinder and Lemmon (2002) note that the underlying state variable in electricity markets is power demand (load). The Stevenson model uses forecast change in demand for electricity to switch between normal and more extreme price regimes.

---

[11] A more detailed description of the steps in this procedure is contained in the Appendix.

[12] Given the lack of liquidity in electricity derivative contracts traded on the Sydney Futures Exchange (SFE), the accepted market forward price is the AFMA consensus price rather than a market-traded price.

[13] For further information concerning AASB7, see www.aasb.com.au

The normal state of the market is characterised by an inelastic demand curve and a relatively elastic supply curve. The inelasticity in demand is due to consumers' preference for constant availability of electricity. Whenever producers are working within their capacity constraints the supply curve remains relatively flat. From this base position there can be a shift in demand, which will in general be related to a weather event and will increase the total demand for electricity. Any problems in infrastructure (plant failure and transmission network failure, for example) which leads to capacity constraints being binding will result in increased inelasticity of supply.

Normal price variation is characterised by the occurrence of supply or demand factors in isolation. When factors which affect demand and supply occur together, extreme price movements of the kind documented in Knittel and Roberts (2005) will occur. The outward movement of the demand curve combined with increasing inelasticity of supply over short time horizons leads to rapid increases in the spot price. In the case of regulated markets these rapid increases in price can lead to caps being applied by the system operator.

The modeling strategy adopted by Stevenson (2001) was to fit a model from the threshold autoregressive (TAR) class to filtered changes in electricity prices. The model fitted from the TAR class is a threshold autoregressive switching (TARSW) model. This is a piecewise-linear autoregressive model which is implemented with two regimes. What determines whether a price change belongs to one regime or the other is whether the forecast change in demand for electricity, which represents the expected change in the state of the market, has a positive or negative value. If a previous forecast change in demand was positive (negative), then the price change was assigned to the regime where previous price changes are positive (negative). These assignments are based on the notion that increases in demand indicate movement away from market equilibrium, while decreases in demand indicate movement towards equilibrium. The switching model has intuitive appeal as a model capable of capturing the high number and different degrees of price increases and decreases that are consistent with observed asymmetric price behaviour.

Let DL_t be the first-difference of the electricity prices and DD_t be the corresponding forecast first difference of the demand for electricity. The specification of the threshold autoregressive switching (TARSW) model for estimating forward price changes is given by equation 2 below.

$$DL_t = \begin{cases} \alpha_0 + \alpha_1 DL_{t-1} + \alpha_2 DL_{t-2} + \cdots + \alpha_p DL_{t-p} + \\ +\gamma_1 DL_{t-p-48} + \gamma_2 DL_{t-p-2(48)} + \cdots + \gamma_q DL_{t-p-q(48)} + \varepsilon_1 \\ \text{if } DD_t \le 0 \\ \beta_0 + \beta_1 DL_{t-1} + \beta_2 DL_{t-2} + \cdots + \beta_s DL_{t-s} + \\ +\lambda_1 DL_{t-s-48} + \lambda_2 DL_{t-s-2(48)} + \cdots + \lambda_m DL_{t-s-m(48)} + \varepsilon_2 \\ \text{if } DD_t > 0 \end{cases} \quad \ldots \; (2)$$

The data for the prices and the demand for electricity is high frequency (half-hourly) and characterized by seasonal, day-of-week and diurnal effects. The combination of the switching process and the autoregressive lag structure in each of the regimes is used to account for the seasonal effects. The assignment of observations to regimes, partitions the data into periods of increasing demand which are common in winter and summer, and decreasing demand in the remaining seasons. The number of autoregressive lags present in each regime, p and s, will be large and a multiple of 48. With 48 half-hours in a day, it is easily recognized how many days are represented by the values of p and s. The number of lags, s, was set equal to p and different models were estimated with lag structures for p that varied from 96 to 672. In order to capture any daily persistence in price changes, the lag structure in both regimes was extended to include further lags spaced at multiples of 48 out to q (equal to m) multiples of 48. The best-fitting model was chosen on the basis of the Minimum Akaike Information Criterion. It had a lag structure with s and p equal to 576 and q and m equal to 88 in equation 2. This resulted in an autoregressive structure extending over one hundred days of daily price changes.

To forecast filtered price differences using the Stevenson (2001) switching model, an accurate forecast of lagged demand difference is needed. An autoregressive model was chosen for this task. The lag structure for the AR model of demand changes was chosen to be 672, or two weeks of half-hour data. Figure 1 depicts demand differences and their forecasts for the last five days of the contract forecast period. While Stevenson (2001) noted the advantages of parsimonious models, he also recognized that given the high frequency nature of the data and the need to estimate forward prices over long time horizons, lag structures incorporating half-hour prices at multiples of 48 for up to one hundred days were not uncommon for electricity data sets.[14]

[14] In any event, implementations of multiplicative and seasonal threshold models that specify a parsimonious parameter space are part of ongoing research.

The combination of a nonlinear time series model based on the characteristics of the electricity market and the use of wavelet-filtered data to ensure against an abnormal and transient price effect at any half-hour tick, provides a sound basis for the generation of the price forecasts needed in risk management applications.

3.2 Evaluation of Forward Price Estimates – A Risk Management View

The situation faced by electricity retailers involves purchasing electricity at market prices and selling at prices that are fixed by regulators. This arrangement exposes retailers to considerable price risk. Retailers seek to manage this risk by establishing a hedging portfolio of derivative positions, typically swaps and various options. Standard approaches to risk management require the hedging portfolio to be valued (marked-to-market) at regular intervals. The estimated price at each of the event dates of the contracts that make up the portfolio is typically used in this valuation process.

Figure 1 Demand Difference Forecast For the Last Five Days Of Contract Period

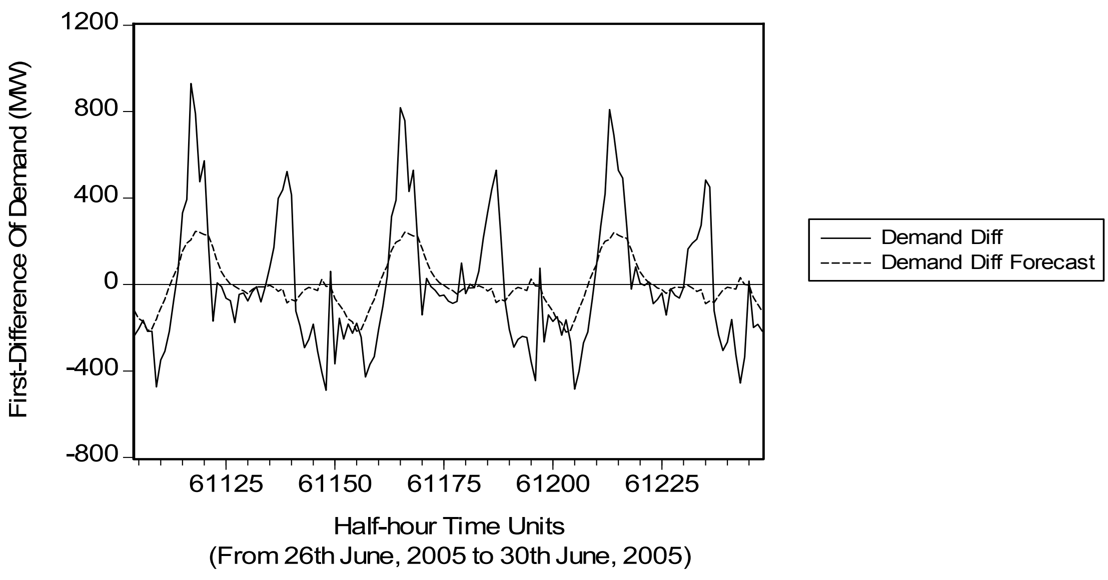

The portfolio valuation process will be sensitive to the nature of the price data, particularly transitory price spikes. A price spike can occur around the half-hour for which contracts are settled each day. These price spikes are usually caused by a transitory disruption to supply. As the price mean-reverts quickly, marking-to-market derivatives using an abnormally high price can lead to an over-valuation of the risk management portfolio. Given that the price schedules that determine what retailers can sell electricity for are fixed, as is their gross margin, overvaluation at any time over the contract horizon can result in an unnecessary unwinding of positions to reduce apparent hedging costs. As a consequence, the hedging portfolio may be rebalanced to the extent that a breach of the Value-at-Risk (VaR) or Earnings-at-Risk (EaR) targets are potentially triggered at some time before expiration. Such a breach would be based on a small price block that is unrepresentative of the fundamental price, both throughout a particular day and across days for the specified half-hour designated for marking-to-market. A triggering of the VaR or EaR limits due to overvaluation from mispricing, is likely to lead to management undertaking remedial action that can entail unnecessary and costly adjustment to the risk management portfolio.

To accurately value the hedging portfolio, the price forecasting mechanism used should generate predicted values which converge to the actual forward contract (spot) price at maturity as the expiry date of the contract approaches. Any difference between these two values is the bias at maturity. In this study we evaluate price estimates from the approaches discussed previously, over a time period of six months. Not only are we interested in the bias at maturity in six months time, but also the bias as measured on a day-to-day basis over the days from the start of the contract period and leading up to maturity. Minimizing the total bias of price estimates in the days running up to the maturity date is important for accuracy in the marking-to-market process.

4. Original and Reconstructed Data

The data used in this study includes half-hour electricity spot prices for the Australian state of New South Wales, as well as the corresponding load or demand for electricity. Our sample consisted of observations of system marginal prices and the quantity of electricity demanded between the 1st January, 2002 and 30th June, 2005.[15] This is a total of 61,248 observations.[16] For each series, we broke this sample into two; an estimation sample from 1st January, 2002 to 31st December, 2004 (52,560 observations) and a forward price estimation sample from 1st January, 2005 to 30th June, 2005 (8,688 observations).

[15] The electricity market in NSW is being gradually deregulated. As a result, the structure of the market exhibits fluidity over time. There is no reason to think that if another estimation and forecast period were to be chosen, the structure would be exactly the same.
[16] In this study we made no attempt to bucket the series into either time-of-day or day-of-the-week categories. This task we left for further research.

Figure 2 NSW State Electricity Price Estimation Sample

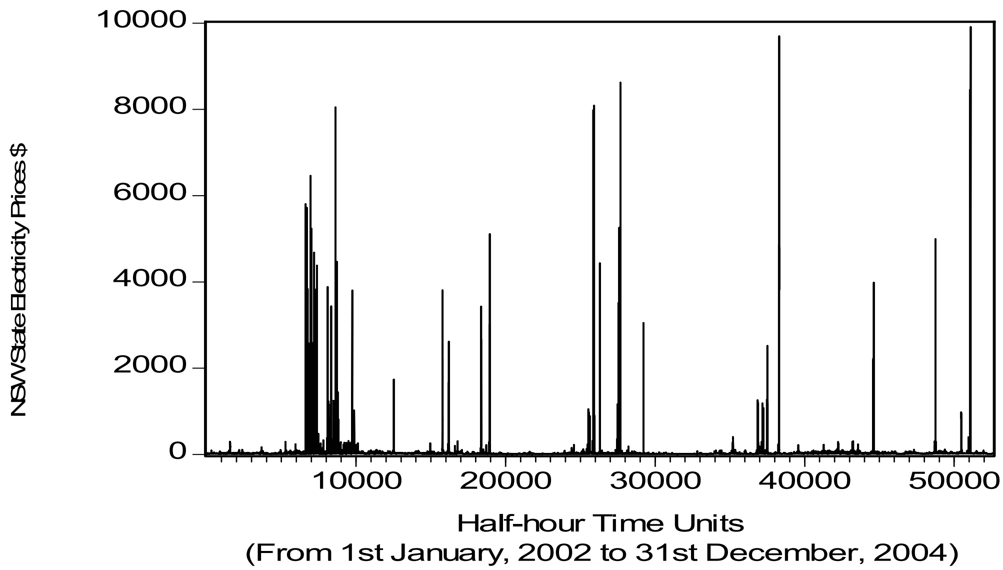

As previously noted in the introduction to the previous section, the data was de-noised, decomposed to lower levels of resolution and reconstructed into half-hour times at different levels of resolution.

Figure 3 NSW State Electricity Load

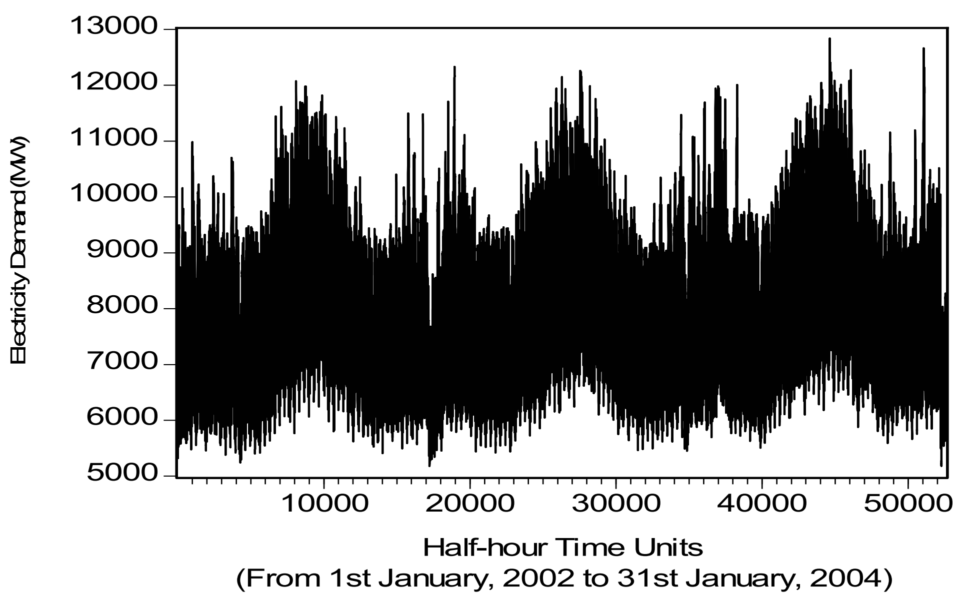

Given that wavelet analysis requires the length of the data set to be of the order 2^n, an extended dataset from 4:00pm on the 3rd October, 2001, through to 30th June, 2005 (65,536 observations) was used to perform the wavelet analysis. This allowed for the removal of data contaminated by edge effects that occur when using a wavelet transform to decompose and reconstruct each series. Figure 2 graphically depicts the original estimation price series, while the corresponding demand series is depicted in Figure 3.

The effect of the smoother-cleaner wavelet transform on reducing the effect of outlier patches in the price data by preventing leakage from higher to lower levels of resolution, can be seen in Figure 4. There, the original price series and the reconstructed series at the second level of resolution are compared over a shorter sub-segment of the time horizon.

The important message to take from Figure 4, as with Figure A.1 in the Appendix, is the effect that the smoother-cleaner wavelet transform has on reducing the effect of outlier patches in the data by preventing leakage from higher to lower levels of resolution.

Figure 4 A Sub-sample Of Original And Reconstructed (Level 2) Prices

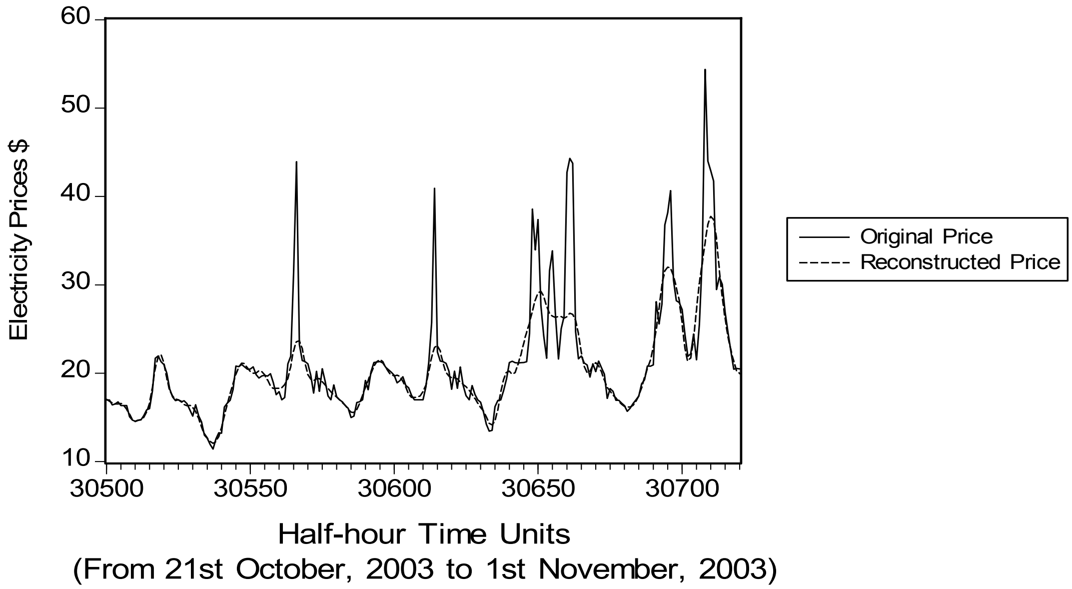

5. Model Estimation and Forward Price Evaluation

The evaluation of a risk management portfolio involves a regular "mark-to-market" process for the derivatives in the portfolio. In this section the time-series and consensus forecasts are evaluated in the context of a "mark-to-market" exercise.

The Stevenson (2001) model is estimated and used to forecast the price at 4.00pm on the last day of the forecast horizon in a rolling fashion. From one day to the next, the updated estimation sample adds the data for the next day (48 half-hours) ahead and drops the previous day's data, the model is then re-estimated and a forecast for the contract expiration date is generated. This entails estimating the model and forecasting the price on the expiration date 180 times. As a result, through the bias we can track the properties of the forecast generating process, starting 180 days out and moving closer to the expiration date day-by-day. These forecasts are compared with the equivalent published consensus forecasts.

As detailed in section 3.2, predicted prices examined in this study were evaluated according to two basic criteria. The first was whether they converged to the settlement price of a derivative contract. Convergence implies a zero bias and increasing accuracy in contract valuation as an event date approaches. The remaining criterion was the minimization of the total bias of the price predictions that are used for marking-to-market in the days leading up to a derivative's expiration date.

In addressing the total bias criterion, we measured the level of mis-pricing associated with the use of a price prediction mechanism. Further, we assumed that the retailer is valuing a floating-to-fixed swap for one megawatt hour of electricity with the expiration day price equal to the realized price on that day. In this case, the prediction error on any day is equal to the dollar value of the mispricing of the swap contract that would result from using the predicted price to value the swap contract. To evaluate the total mis-pricing due to a sequence of price predictions for the expiration day price, the sum of the absolute values of the prediction errors should be used. This will ensure that large positive and negative errors do not "net" out. The total of undervaluation and overvaluation should be minimized.

Table 1 presents the summary statistics for the absolute values of the bias sequences from the Stevenson and AFMA forecasting procedures.

Table 1: Statistical Properties of Absolute Step-Ahead Biases.

This table contains descriptive statistics for the sequence of the absolute values of the bias (mis-pricing) generated by the AFMA and Stevenson price forecasting mechanisms.
The sample has been truncated for days where the AFMA Consensus Forecasts were unavailable.

	Stevenson	**AFMA**
Mean	4.38	20.76
Median	3.79	20.99
Maximum	20.03	25.92
Minimum	0.03	17.62
Std. Dev.	3.38	1.28
Skewness	1.39	-0.07
Kurtosis	6.20	4.03
Sum	521.34	2470.58
Sum Sq. Dev.	1350.53	191.97
Expiration Day Bias	-1.16	20.56
Observations	119	119

These statistics show that, according to the criterion of minimizing total mispricing, the Stevenson model with a total of $521.34 per megawatt hour clearly outperforms the AFMA consensus forecast which would generate mis-pricing of $2470.58 per megawatt hour over the six month evaluation period. The average daily mis-pricing from the Stevenson model is one fifth of the level of mis-pricing that would result from using the AFMA price predictions. The expiration day bias is also significantly smaller for the Stevenson model. In terms of minimizing mis-pricing over time, the Stevenson model is clearly superior. These findings are consistent with the perception by market participants that the AFMA prices are misleading (Anderson, Hu and Winchester, 2005).

The evolution of the bias is also presented graphically. Figure 5 depicts the daily step-ahead biases from the Stevenson and AFMA models. The first observation of both graphs is the bias corresponding to an estimate of the price 181 days from the expiration date. Subsequent step-ahead observations record the biases of daily price estimates in the run-up to the end of the six month contract period.

Figure 5 A Comparison of Step-ahead Biases For The AFMA and Stevenson Models

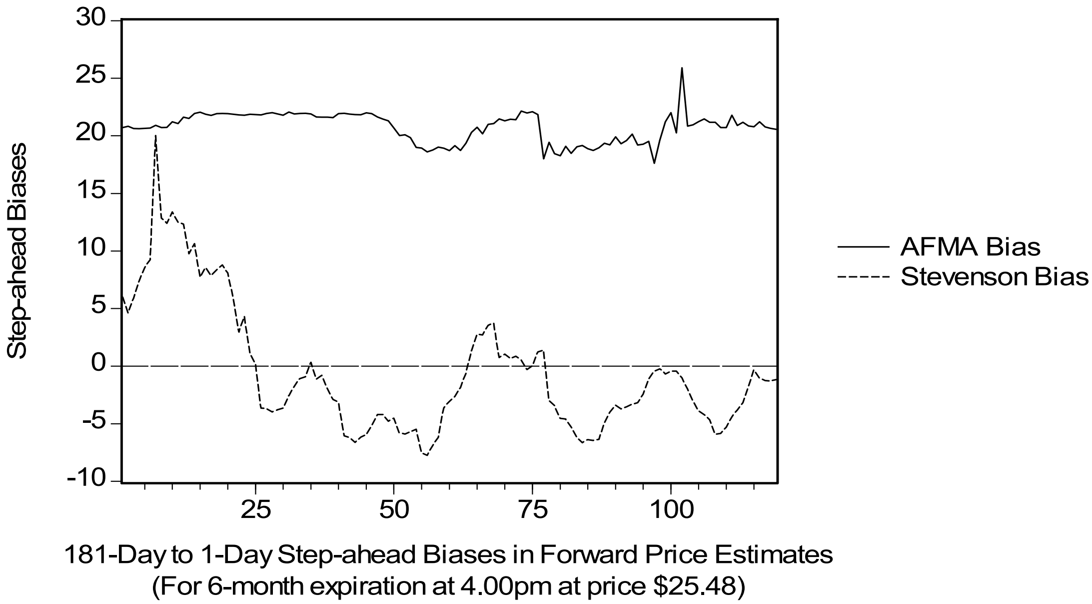

The second criterion of interest is the convergence of the bias series towards zero as the expiration date approaches. This criterion will be evaluated in two parts. First, a simple AR(1) model is examined to evaluate the stability of the bias sequence. A value of the AR(1) parameter of less than one indicates that the sequence is stable. If a bias series is stable, then a simple time trend regression is used to measure its rate of convergence. The trend regression is given by,

$$Bias_t = a + b \times time \,.$$

If the estimated slope coefficient, b, is significant and has a negative sign then the bias sequence is convergent. The correction factor is given by,

$$correction\ factor = \frac{1}{|b|},$$

where b is the slope coefficient of the time trend regression. The correction factor indicates the number of time units that would be required by the forecasting mechanism to correct a one unit bias. Table 2 contains the results for the convergence criterion.

Table 2 – Convergence Results

This table contains the parameter estimates from simple the autoregressive and time trend regressions used to gauge stability, and rate of convergence of the realized biases of the competing forecasting systems. Estimated parameter values and their t statistic are presented for the bias sequence from the AFMA and Stevenson price forecasting mechanisms.

Variable	AFMA Coefficient (t)	Stevenson Coefficient (t)
C	21.44	4.89
	(96.63)	(5.89)
Time Trend	-0.011	-0.0915
	(-3.53)	(-7.53)
Adjusted R-squared	0.08858	0.320498
AR(1) Coefficient	0.99	0.95
	(240.38)	(34.58)

Both forecasting procedures generate a bias sequence that is stable, each having a significant AR(1) parameter which is less than one. The slope coefficients of the time trend regressions on the bias sequences are significant and negative for both forecasting approaches. This indicates that the bias sequences are converging to zero over the contract window. For the Stevenson model, the value of $1/|b|$ (or the time required to correct for a unit of bias) is 10.93 days. This is substantially less than the value of 87.15 days for the AFMA forecasting procedure. Forecasting procedures with small values for the correction factor will provide forecast sequences which converge quickly, with a smaller value preferred. On the basis of this criterion, the Stevenson approach is again preferred for valuation purposes.

The proposed evaluation criteria that are applied to a sequence of forecasts over a decreasing forecast horizon are designed to consider the extent of mis-pricing, along with assessing the convergence properties of a predicted price sequence. The purpose of the evaluation is to compare price prediction mechanisms. Those with the smallest level of mispricing combined with small values of the correction factor are preferable for the marking-to-market exercise. The evaluation of the Stevenson and AFMA price prediction mechanisms clearly shows that the former approach, which is based on the properties of the market mechanism, outperforms the AFMA consensus forecasts on both criteria.

6. Conclusions

This study has been concerned with examining the role of estimating the prices of forward contracts in the Australian electricity market. To evaluate forward price estimates, two criteria were utilized in order to conclude appropriateness for use in the marking-to-market process for the derivative products that are typically used in the Australian electricity industry to form risk management portfolios. First, was the requirement that the forward price converged to the spot price at expiration of a hedging contract. The second criterion referred to the mis-pricing of the forward price estimates over the days leading up to the contract expiration.

Forward price estimates were generated from a time series model and the consensus forecasts of electricity market operatives. The model was the Stevenson (2001) time series model that is a threshold autoregressive model. The consensus estimates were the AFMA Forward Prices that serve as market forward prices due to the extent of their adoption by risk managers throughout the Australian electricity industry.

The ranking of the two alternatives for generating forward price curves was clear. On both criteria the Stevenson (2001) model was preferred. It converged close to the spot price at expiration. Further, on a step-ahead basis, the Stevenson (2001) model was characterised by small positive biases relative to that of the AFMA prices. Therefore, using the Stevenson (2001) model obviates the tendency for costly rebalancing of the risk portfolio that is likely to occur due of the mis-pricing at the long-end of the AFMA curve.

The results of this study suggest that the preferred forward price estimates are model-driven. The length of the hold-out period over which the derivatives that constituted the risk management portfolio were marked-to-market was six months. This period was determined by data availability. In time, in order to verify these results it will be important to repeat the analysis with a longer period of days before expiration.

7. References

1. Anderson, E., Hu, X. and D. Winchester, (2005): “Forward contracts in electricity markets: the Australian experience,” *Working Paper*, Centre for Energy and Environmental Markets, University of New South Wales.

2. Bessembinder, H. and M. L. Lemmon, (2002): “Equilibrium pricing and optimal hedging in electricity forward markets,” *Journal of Finance*, 57, 3, 1347-82.

3. Bruce, A. and H. Gao, 1994: *S+Wavelets User Manual*. StatSci Division: MathSoft Inc., Seattle, Washington, U.S.A.

4. Bruce, A., Donoho, D., Gao, H., and R. Douglas Martin, (1994): "Denoising and robust nonlinear wavelet analysis," *SPIE Proceedings*, Orlando, Fl, U.S.A.

5. Daubechies, I., (1992): "Ten Lectures on Wavelets", *Society for Industrial and Applied Mathematics*, Philadelphia, U.S.A.

6. Hicks, J. R., (1939): *Value and Capital*. Oxford University Press, Cambridge.

7. Hirshleifer, D, (1990): “Hedging pressure and future price movements in a general equilibrium model”, *Econometrica*, 58, 441–28.

8. Hull, John C., (2003): *Options, Futures, and Other Derivatives*, 5th edition. Prentice Hall.

9. Keynes, J. M., (1930): *Treatise on Money.* Macmillan, London.

10. Knittel, C. R. and M. R. Roberts, (2005): “An empirical examination of restructured electricity prices,” *Energy Economics*, 27, 5, 791-817.

11. Lin, S. and M. Stevenson, (2001): "Wavelet analysis of the cost-of-carry model”, *Studies in Nonlinear Dynamics and Econometrics*, 5, 1, 87-102.

12. Longstaff, F. A. and A. W. Wang, (2004): “Electricity forward prices: a high-frequency empirical analysis,” *Journal of Finance*, 59, 4, 1877-1900.

13. Ramsey, J. and C. Lampart, (1998): "The decomposition of economic relationships by time scale using wavelets: expenditure and income," *Studies in Nonlinear Dynamics and Econometrics*, 31, 23-42.

14. Stevenson, M., (2001): "Filtering and forecasting spot electricity prices in the increasingly deregulated Australian electricity market", *Quantitative Finance Research Papers*, 63, University of Technology Sydney.

Appendix Wavelet Analysis

Briefly, using wavelet transforms[17], a signal can be decomposed into a parsimoniously countable set of basis functions at different time locations and resolution levels. Unlike Fourier analysis, which assumes the same frequencies hold at the same amplitudes for any sub-segment of an observed time series, wavelet analysis captures the more localized behaviour in a signal. Trigonometric functions (with infinite support or waves) serve as functions on which a Fourier decomposition of time series data is based in the frequency domain. In contrast, wavelet analysis is characterised by basis functions that are not trigonometric and that have their energy concentrated within a short interval of time. These 'small waves', or wavelets, are defined over the square functional space, $L^2(\mathcal{R})$, and they have compact support. It is the property of compact support that enables wavelet analysis to capture the short-lived, often transient components of data that occur in short time intervals. Further, they are not necessarily homogenous over time in that the same frequencies will not hold at the same amplitudes over all subsets of the observed time series.

Wavelets belong to families and it is these families that provide the building blocks for wavelet analysis. Just as sine and cosine functions are functional bases onto which we project data to extract information belonging to the frequency domain, wavelet functions are functional bases that allow for extraction of information available in both the time and frequency domains. A wavelet family comes in pairs; a father and mother wavelet. The father wavelet, $\phi(t)$, represents the smooth, low-frequency part of the signal, while the mother wavelet, $\psi(t)$, captures the detail or high-frequency component.

[17] For an introduction to wavelets and wavelet analysis in an economic and financial market context, see Ramsey and Lampart (1998) and Lin and Stevenson (2001).

A continuous function, *f(t)*, can be approximated by the orthogonal wavelet function given by:

$$f(t) \approx \sum_k s_{J,k}\phi_{J,k}(t) + \sum_k d_{J,k}\psi_{J,k}(t) + \sum_k d_{J-1,k}\psi_{J-1,k}(t) + \ldots + \sum_k d_{1,k}\psi_{1,k}(t), \qquad ......(2)$$

where J is the number of multi-resolution components (or scales), and k ranges from one to the number of coefficients in a multi-resolution component. The coefficients, $s_{J,k}$, $d_{J,k}$, $d_{J-1,k}$,...,$d_{1,k}$ are the wavelet transform coefficients, while $\phi_{J,k}(t)$ and $\psi_{j,k}(t)$ are the approximating father and mother wavelet functions, respectively. The wavelet approximation to *f(t)*, given by equation (2), is orthogonal since the basis functions, ϕ and ψ, are orthogonal by construction.[18] Wavelet functions usually do not have a closed functional form. After firstly imposing desired mathematical properties and characteristics, they are generated through dilation and translation according to the following normalised [19] functions.

$$\phi_{j,k}(t) = 2^{-j/2}\phi\left(\frac{t-2^j k}{2^j}\right) \qquad ......(3)$$

$$\psi_{j,k}(t) = 2^{-j/2}\psi\left(\frac{t-2^j k}{2^j}\right) \qquad ....(4)$$

The wavelet transform coefficients measure the contribution of the corresponding wavelet function to the approximating sum. At a particular level of time resolution, j, the impact of the information subset on the signal is reflected in the number and magnitude of the wavelet coefficients, and is roughly equal to the sampling interval at that resolution level. Information corresponding to finer detail in the signal than that at resolution level, j, can only be incorporated into the signal by considering shorter sampling intervals which are associated with higher levels of resolution than j. Such information will not contribute to approximating the signal at lower levels.

To prevent outliers from leaking into the wavelet coefficients at levels of higher resolution the robust smoother-cleaner transform developed by Bruce, Donoho, Gao and Martin (1994) was used. This involved decomposing our series using a biorthogonal wavelet that is robust against leakage of outlier patches in the data into the smooth coefficients.[20] The biorthogonal wavelet used comes

---

[18] A detailed mathematical exposition of how the basis functions are constructed can be found in Daubechies (1992).

[19] The factor, $2^{-j/2}$, in equations (3) and (4) serves to normalise the functions.

[20] A biorthogonal wavelet transform utilises both low-pass and high-pass filters. The low-pass filters are short and avoid outlier leakage to the smooth coefficients. The high-pass filters are long and ensure sufficient smoothness of the underlying basis functions. While biorthogonal wavelets are not orthogonal, for the most part we can use them as we would an orthogonal wavelet.

from the "b-spline" family and is coded as bs3.5 in the S+ Wavelets package produced by the StatSci Division of MathSoft and written by Bruce and Gao (1994).[21] Implementation of this wavelet decomposition started with a set of smooth wavelet coefficients, s_j. After calculating a robust set of coefficients, $\hat{s}_j$, using running medians of length 5, we derived a robust set of residuals, r_j, where

$$r_j = \delta(s_j - \hat{s}_j)$$

and δ is a shrinkage function which shrinks the coefficients such that

$$\delta(x) = \begin{cases} 0 & \text{if} \quad |x| \le a, \\ sign(x)\dfrac{b(|x|-a)}{b-a} & \text{if} \quad a < |x| < b, \\ x & \text{if} \quad |x| \ge b \end{cases}$$

Thresholds levels a and b were chosen to ensure that most of the robust residuals were zero. The next level of smooth wavelet coefficients, s_{j-1}, were obtained after applying the usual low-pass wavelet filter to the cleaned smooth coefficients,

$$u_j = s_j - r_j$$

while the detail wavelet coefficients, d_{j-1}, were the realization from the application of the high-pass wavelet filter. This procedure was repeated with the smooth coefficients at the next highest level of resolution. By using the robust smoother-cleaner wavelet transform we removed outlier patches from the decomposition.[22] An example of how a more fundamental signal was derived using a smoother-cleaner wavelet transform, without leakage of outlier patches into the signal at lower levels of resolution, is depicted in Figure A.1 below. There, a subsection of the original estimation price series is graphed against the more fundamental signal reconstructed at the lower fourth level of resolution.

---

[21] This is the computer package used to decompose and reconstruct the electricity price and demand series.

[22] A key property of the above procedure makes it extremely useful for filtering electricity price series. Outlier patches of length $(2^n + 2)$ are isolated to the wavelet coefficients in lower resolution levels than n. It has the effect of removing the high and rapidly mean-reverting prices from the lower-frequency levels.

Figure A.1 A Sub-sample Of Original And Reconstructed (Level 4) prices