Essays on Australian Wholesale Electricity Price Spikes and the Australian Pre-Dispatch Process

A thesis submitted to The University of Manchester for the degree of

Doctor of Philosophy

in the Faculty of Humanities

2013

Wan Nur Rahini Aznie Binti Zainudin

Manchester Business School

Table of Content

This thesis contains 54, 967 words including title page, tables, and footnotes

Abstract

The University of Manchester
Wan Nur Rahini Aznie Binti Zainudin
Doctor of Philosophy (PhD)
Essays on Australian Wholesale Electricity Price Spikes and the Australian Pre-Dispatch Process
September 2013

In the first essay I examine whether the occurrences of the extreme price events display any regularities that can be captured using an econometric model. Here I treat these price events as point processes and apply Hawkes and Poisson autoregressive models to model the dynamics in the intensity of this process. I use load and meteorological information to model the time variation in the intensity of the process. The models are applied to data from the Australian wholesale electricity market, and a forecasting exercise illustrates both the usefulness of these models and their limitations when attempting to forecast the occurrence of extreme price events.

In the second essays I explain that in the past doubts have been raised as to whether the pre-dispatch process in Australia Electricity Market is able to give market participants and market operator good and timely quantity and price information. It is the purpose of the second essay to introduce a framework to analyse whether the pre-dispatch process is delivering biased predictions of the actual wholesale spot price outcomes. Here I investigate the bias by comparing the actual wholesale market spot price outcome to pre-dispatch sensitivity prices established the day before dispatch and on the day of dispatch. I observe a significant bias (mainly indicating that the pre-dispatch process tends to underestimate spot price outcomes) and I further establish the seasonality features of the bias across seasons and/or trading periods. I also establish changes in bias across the years in our sample period (1999 to 2007). In the formal setting of an ordered probit model I establish that there are some exogenous variables that are able to explain increased probabilities of over- or under-predictions of the spot price. It transpires that meteorological data, expected pre-dispatch prices and information on past over- and under-predictions contribute significantly to explaining variation in the probabilities for over- and under-predictions. The results allow me to conjecture that some of the bids and re-bids provided by electricity generators are not made in good faith.

Finally, the third essay investigates whether information from this pre-dispatch process can be useful when predicting next-day price spikes. In a preliminary analysis I establish the effect of pre-dispatch prices on the quantiles of the spot price distribution. A Quantile regression approach reveals that higher pre-dispatch prices signal only to a certain extent an increased probability of higher spot price outcomes. They also signal a higher uncertainty about the resulting spot price outcomes. I further establish whether the inclusion of information from the pre-dispatch process can significantly improve the predictability of price spikes when these are modelled as a point process (as in the first essay). The models used here are Hawkes and Poisson autoregressive models which allow for time variation (correlated to exogenous information) in the intensity process that governs the occurrence of price spikes. It transpires that the pre-dispatch process of the Australian Electricity Market does not provide any information that can be used in a systematic manner to help predicting on what days price spikes are more likely to occur.

Declaration

I, Wan Nur Rahini Aznie Binti Zainudin, declare that no portion of the work referred to in the thesis has been submitted in support of an application for another degree or qualification of this or any other university or other institute of learning.

Copyright Statement

i. The author of this thesis (including any appendices and/or schedules to this thesis) owns certain copyright or related rights in it (the "Copyright") and s/he has given The University of Manchester certain rights to use such Copyright, including for administrative purposes.

ii. Copies of this thesis, either in full or in extracts and whether in hard or electronic copy, may be made only in accordance with the Copyright, Designs and Patents Act 1988 (as amended) and regulations issued under it or, where appropriate, in accordance with licensing agreements which the University has from time to time. This page must form part of any such copies made.

iii. The ownership of certain Copyright, patents, designs, trade marks and other intellectual property (the "Intellectual Property") and any reproductions of copyright works in the thesis, for example graphs and tables ("Reproductions"), which may be described in this thesis, may not be owned by the author and may be owned by third parties. Such Intellectual Property and Reproductions cannot and must not be made available for use without the prior written permission of the owner(s) of the relevant Intellectual Property and/or Reproductions.

iv. Further information on the conditions under which disclosure, publication and commercialisation of this thesis, the Copyright and any Intellectual Property and/or Reproductions described in it may take place is available in the University IP Policy (see http://documents.manchester.ac.uk/DocuInfo.aspx?DocID=487), in any relevant Thesis restriction declarations deposited in the University Library, The University Library's regulations (see http://www.manchester.ac.uk/library/aboutus/regulations) and in The University's policy on Presentation of Theses.

Acknowledgements

I would like to extend my gratitude to both of my supervisors, Dr. Ralf Becker, Prof. Sydney Howell and Prof. Adam Clements for the commitment and patience shown to me throughout my PhD. I also would like to thank my PhD committee Prof. Stuart Hyde and Dr. Alex Taylor for their insightful comments and suggestion during my six-monthly PhD reviews.

I am deeply indebted to my family especially both of my parents, Sharifah Mizan Binti Syed Omar and Zainudin Bin Wan Omar for the unconditional love and supports and my mentor, Prof. Jalani Sukaimi for continuous assistance and motivation. I want to thank my fellow PhD colleagues, Nida Abdioglu and Zhenmei Zhu for moments that we shared. To my housemates, Dr. Asma Zamanuri, Noor Qamariah Zainal, Syikin Saiful Bahri and Dr. Jannaf Mohd Radzif, thanks for sharing the tears and laughters.

Finally, I would like to acknowledge all the staffs in Islamic Science University of Malaysia who manage my study leave extension for their efficiency and the Government of Malaysia for their financial support.

Chapter 1

1. Introduction

1.1. Motivation

Until the 1990s, the electricity sector in most industrial countries was vertically integrated and often in state ownership. This means the electricity had been a regulated natural monopoly with a guaranteed rate of return in exchange for an obligation to supply electricity to the public. In this setting, the regulators fixed prices as a function of generation, transmission and distribution costs. Due to little uncertainty in prices, generators could make decisions by only applying standard deterministic valuation tools such as discounted cash flow analysis (Heydari & Siddiqui 2010). In recent years, many countries have deregulated their electricity sectors with the aim of introducing competition in generation and retail activities. One of them is Australia. During the 1990s, the Australian electricity market started a gradual deregulation by opening up their electricity sectors to competition. By December 1998 a National Electricity Market (NEM) had been formed between the regions of New South Wales, Queensland, South Australia and Victoria. The electricity grids of these four states were fully connected by February 2001. The fifth region, Tasmania joined the NEM by 2005 and operations today are based in five interconnected regions (Australian Energy Market Operator 2010).

The transformation from a regulated monopoly to private ownership of generation and market liberalisation may result in lower prices and better use of resources (Heydari & Siddiqui 2010). For example, in terms of the composition of the supply side, low marginal cost of production such as coal-fired generators and hydroelectric production supply 81% and 5% of the NEM's capacity respectively while higher marginal cost of production, gas turbines and oil-fired plants supply 12.2% and around 0.2% of the market capacity (Australian Energy Market Operator 2010). When deregulation came into place, electricity prices in the wholesale market were subject to interaction of supply and demand, which now have become highly volatile resulting in unexpected price spikes. These sudden spike events may be caused by unexpected increases in temperature, supply or transmission shocks. Given that electricity is a necessity commodity for the economy, the general public are still insulated from these spikes and pay at a regulated price. As for electricity retailers, however, ignoring these spikes event would be costly and results to them incurring huge

losses, as they have to purchase electricity on the deregulated wholesale market but sell on the regulated retail market.

Wholesale markets for electricity are complex, with electricity demand highly influenced by the time of day and weather. The volatility in price is exacerbated by the way consumers are usually charged at a fixed retail price, independent of the spot market price and hence reducing the effective price elasticity of demand, and the inability to store electric power at any significant scale (Anderson et al. 2007). Altogether these factors resulted in electricity spot prices that have the following features; seasonality, skewness, high and clustering volatility as well as the presence of jumps. In this situation, generators reserves play an important role to ensure the stability of the supply and the system operator plays a role in balancing the market by dispatching the generators reserves as cost-efficient as possible to satisfy the escalating demand.

In the NEM setting, by 12.30pm on the day before dispatch, electricity generators are obliged to submit supply schedules for all 48 half hours of the following trading day. These supply schedules are then aggregated and matched with demand forecasts to produce expected wholesale electricity prices for the next day. Therefore the predicted price is the cost of the marginal (most expensive) unit of electricity uses in this matching process.

Within this framework, market participants need to engage in bids and offers in advance of real time simply to enable the Australian Energy Market Operator (AEMO) to look ahead and ensure physical feasibility of the proposed schedule (Anderson et al. 2007). In order to inform AEMO with good and timely quantity and price information (Australian Energy Market Operator 2010), they perform a pre-dispatch process. In order to evaluate how sensitive the predicted price is with respect to changes in demand, AEMO will repeat the matching exercise for different demand quantities (demand forecast +/- 200 MWh) and reports the resulting sensitivities of the predicted price outcomes with respect to different demand forecast. Generators are allowed to change their bids (re-bidding) after 12.30pm on the day before dispatch up until shortly before dispatch. However, any bid or rebid made, should reflect the generators "genuine intention to honour" (Australian Energy Market Commission, 2013, paragraph 3.8.22A.b) their bid.

In the past doubts have been raised as to whether the pre-dispatch process in NEM is able to give market participants and AEMO good and timely quantity and price information. In

the Chapter 3 I introduce a framework to analyse whether the pre-dispatch process is delivering biased predictions of the actual wholesale market spot price outcome. As I can observe in the results of Chapter 3, there are some significant biases in the pre-dispatch process, mainly indicating that the predicted prices tend to be lower than the actual spot price outcomes. The results from this investigation allow us to conjecture that some of the bids and re-bids provided by electricity generators are not made in good faith and it is clear that the pre-dispatch prices have clear deficiencies as a price forecast.

The 'good faith' requirement was mentioned by the Australia National Electricity Code Administrator's (NECA) in their report published in 2001. In the report, NECA drew attention to a number of a very short-term price spikes either directly as a result of inefficiencies in the market rules or as a result of individual generators taking advantage of those inefficiencies to drive up prices. When a generator has the opportunity to bid higher than marginal cost to drive up prices by taking advantage of a poor market design, this behaviour is called strategic bidding. NECA concluded some very short-term price spikes do not represent a genuine price signal to either the supply side, in terms of the need for new investment, or the demand side of the market. In the report, NECA also stated these very short-term spikes would be the effects of inappropriate bidding and rebidding behaviour. In order to rectify this problem, NECA and the Code Change Panel require generators' bids and rebids to be made in good faith.

Hence, in Chapter 3 I discuss this requirement in full length and how it relates to our conjectures that the failure of the generators to comply with the good faith requirement means they are practicing inappropriate bidding and rebidding behaviour which causes significant biases in the pre-dispatch process.

Since wholesale electricity markets have been deregulated, modelling and forecasting wholesale electricity prices has attracted a considerable amount of attention. Most of the studies that proposed approaches to model electricity prices focused on modelling the trajectory of the spot price or its return across time (Weron 2007). However, studies that focus on the price spikes or extraordinary price events are still scarce. Therefore in this thesis I focus on one particular aspect of the electricity price distribution, the so called price spikes.

There is no particular definition of what constitutes a price spike, but for the purpose of this thesis I follow the lead of Christensen et al. (2009) who define price spikes as a binary variable defined by whether the wholesale price exceeds some price threshold. They found evidence of significant persistence in the occurrence of such price spikes. This feature is commonly neglected by most traditional price series models since they, if they allow for a spike feature, treat these as a memory-less jump process with an intensity which is independent of its own history.

More recently, the issue of price spikes has been tackled in a more direct manner by defining a binary series that identifies the instances in which the electricity price exceeds a certain price threshold and subsequently modelling the dynamics of this series. To the best of our knowledge the first study using this approach was Christensen et al. (2009) who used a modified Poisson autoregressive (PAR) framework to forecast next day's price spike occurrence. They model and forecast the probability of price spikes (as defined through the binary series). In Chapter 2 I tackle the same task using Hawkes models. Both approaches essentially allow for the probability of a price spike occurring to vary with (weakly) exogenous covariates and to display persistence.

Although I know many of the price spikes are the results of mainly unpredictable event (e.g. unexpected generator breakdown), in Chapter 2 I build an econometric models using PAR and Hawkes models with the aid of load and meteorological information, in hope to find some element of predictability in price spikes. The forecasting exercise illustrates the usefulness of these models but perhaps more importantly their limitations when attempting to forecast the occurrence of extreme price events.

The pre-dispatch process described above is one element of the NEM architecture that aims to increase market transparency and therefore to facilitate the prediction of potential price spikes. In our quest to uncover the potential of pre-dispatch information for the spike forecasting exercise, in Chapter 4, I initially conduct a preliminary analysis. This consists of a careful analysis of correlations and quantile regressions. From the results of this analysis I find evidence that pre-dispatch prices are a weak predictor of price outcomes. Despite this finding, using the methodology framework set up in Chapter 2, I investigate whether information from this pre-dispatch process can be useful when predicting next-day price spikes. In the light of our findings in Chapter 3, it is it is possibly not surprising to find that the pre-dispatch process of the Australian Electricity Market does not provide any

information that can be used in a systematic manner to help predicting on what days price spikes are more likely to occur.

1.2. Thesis Structure

The thesis structure follows the required format by the Manchester Accounting and Finance Group, Manchester Business School, at The University of Manchester. The structure allows each chapter in the thesis to follow a format suitable for submission and publication in peer-reviewed academic journals. Therefore, this thesis is structured based on three essays containing original research in chapters 2, 3, and 4. Work presented in Chapter 2 had been published in Journal of Energy Market, hence there is great similarity of the chapter with the published paper. However, I add further explanations on the figures and tables in the chapter where appropriate.

The chapters are self-contained with separate literature reviews, using different dataset and have different research objectives. Chapter 2 aims to examine whether the occurrence of the spike events display any regularities that can be captured using an econometric model. In order to explore the potential lies in the information of pre-dispatch process in predicting spike events, in Chapter 3 I introduce a framework to analyse whether the pre-dispatch process is delivering a good prediction of the actual wholesale spot price outcomes. Finally in Chapter 4 I investigate whether information from this pre-dispatch process can be useful when predicting spike events. Chapter 5 concludes.

1.3. Summary and Conclusions

This thesis contributes to the emerging literature in three ways:

1. Forecasting price events is a difficult task and neither PAR nor Hawkes models are able to predict price events that occurred in isolation.

2. In general the pre-dispatch prices tend to underestimate the actual price outcomes of the wholesale spot prices. This allows us to conjecture that some of the bids and re-bids provided by generators are not made in good faith as required by the National Electricity Rules (Australian Energy Market Commission, 2013, paragraph 3.8.22A).

3. The pre-dispatch process of the Australian National Electricity Market does not provide any information that can be used in a systematic manner to help predicting price spikes on the next day.

In the empirical chapters I use the term "we" rather than "I", reflecting that each empirical chapter is associated with either a published paper or a working paper co-authored with my supervisors: Ralf Becker and Sydney Howell.

References

Anderson, E.J., Hu, X. & Winchester, D., 2007. Forward contracts in electricity markets: The Australian experience. *Energy Policy*, 35(5), pp.3089–3103. Available at: http://linkinghub.elsevier.com/retrieve/pii/S0301421506004289 [Accessed August 14, 2013].

Australian Energy Market Commission, 2013. National Electricity Rules Version 56 Chapter 3: Market Rules. , p.146. Available at: http://www.aemc.gov.au/Electricity/National-Electricity-Rules/Current-Rules.html.

Australian Energy Market Operator, 2010. An introduction to australia's national electricity market. *Available at AEMO website http://http://www.aemo.com.au/About-the-Industry/Energy-Markets/National-Electricity-Market*.

Christensen, T., Hurn, S. & Lindsay, K., 2009. It never rains but it pours: modeling the persistence of spikes in electricity prices. *Energy Journal*, 30(1), p.25.

Heydari, S. & Siddiqui, A., 2010. Valuing a gas-fired power plant: A comparison of ordinary linear models, regime-switching approaches, and models with stochastic volatility. *Energy Economics*, 32(3), pp.709–725.

Weron, R., 2007. *Modeling and forecasting electricity loads and prices: A statistical approach*, Wiley. com.

Chapter 2

2. Modelling electricity price events as point processes

ABSTRACT

Energy prices are highly volatile and often feature unexpected spikes. It is the aim of this chapter to examine whether the occurrence of these extreme price events display any regularities that can be captured using an econometric model. Here we treat these price events as point processes and apply Hawkes and Poisson autoregressive models to model the dynamics in the intensity of this process. We use load and meteorological information to model the time variation in the intensity of the process. The models are applied to data from the Australian wholesale electricity market, and a forecasting exercise illustrates both the usefulness of these models and their limitations when attempting to forecast the occurrence of extreme price events.

Published: joined with Ralf Becker and Adam Clements in: Journal of Energy Markets Volume 6/Number 2, Summer 2013 (99–140)

2.1. Introduction

Wholesale electricity prices display a number of interesting features. These properties include daily, weekly and seasonal cycles, high volatility, mean-reversion and, most interestingly, frequent price spikes. The price spikes are the result of a number of factors. In the short run demand is largely price-inelastic (Anderson et al. 2007). The electricity supply curve, however, is based on increasing marginal production cost. The larger the demand for electricity, the more expensive is the marginal unit of electricity. Unlike most other commodities, electricity is largely nonstorable (other than through pump-storage hydro schemes). These factors, in combination with the requirement that supply and demand are matched instantaneously and at all times, explain why significant price variations are required to clear the market, resulting in occasional price spikes.

Significant efforts have been made to model the variation in wholesale electricity prices. Most of these literatures model the wholesale electricity (spot) price, or its logarithm. However, the features of the electricity price series listed above make this a very difficult exercise. In this chapter we focus on one particular aspect of the electricity price process: the price spikes or extraordinary price events. In the Australian wholesale electricity market, for instance, prices for one Megawatt hour (1 MWh) can move very quickly from a normal price level of around A$30 to a price of A$10 000. It is the aim of this chapter to establish whether the occurrence of price events can be predicted through the time series features of the event series itself and/or variation in exogenous variables. To establish an appropriate modeling framework, we treat the time-series of (suitably defined) price events as a binary series or a point process.

There are only two papers in the literature that adopt this approach: Christensen et al. (2012) and Christensen et al. (2009). In the first, Christensen et al. (2009) propose a modified Poisson autoregressive (PAR) model to incorporate the historical components of the spiking process together with exogenous factors that drive the occurrence and intensity of price spikes. The Hawkes (1971) model used in this chapter follows a similar approach of modeling price spikes in a binary fashion. The Hawkes model has been applied successfully in the context of high-frequency financial econometrics. Both models present different technical difficulties and allow for different dynamics in the occurrence of price spikes.

The objectives of this chapter are

1. To establish which exogenous variables are relevant to explain variation in the intensity of price spikes,
2. To evaluate the Hawkes model and its merits as compared to the PAR model (Christensen et al. 2009).
3. To establish a framework in which to compare the forecast performance of these models.

This chapter uses data from the Australian National Electricity Market (NEM). This market has been operating since mid-1990s and can be considered a mature deregulated market for wholesale electricity. The regional markets used for the purpose of this study are New South Wales (NSW), Queensland (QLD), Victoria (VIC), and South Australia (SA).

The remainder of the chapter is organised as follows. We present a literature review in Section 2.2. In Section 2.3, we present the methodology of the univariate Hawkes model as it is applied in this chapter, as well as a brief summary of the Poisson autoregressive model. The data sets used are presented, together with their descriptive statistics, in Section 2.4, which also includes a discussion of the covariates that are considered when modeling the price event intensities. Section 2.5 presents the estimation results of our chosen models and discusses the role of the covariates, while in Section 2.6 the model's forecast performance is discussed. Section 2.7 concludes Chapter 2.

2.2. The Market for Spot Electricity and Literature Review

In the Australian NEM, supply and demand of electricity are matched through a centrally coordinated dispatch process. All generators have to bid their supply curves, consisting of a maximum of ten price-quantity pairs with a price floor of -A$1000 and a price ceiling of A$12 500/MWh (A$10 000/MWh prior to July 1, 2010) before 12.30pm on the day before delivery of electricity. The bids placed by generators to sell electricity are then matched with the electricity demand. In general the matching occurs such that electricity offered at the cheapest prices receives priority in the dispatch process. The dispatch spot price represents the cost to supply the last megawatt of electricity to meet demand (Australian Energy Market Operator 2010) and is calculated for every half hour. This price is then used

to settle all electricity supplied in the half hour and henceforth is called the wholesale spot price.

Before discussing the relevant literature, it is important to establish a few stylized facts for the Australian NEM wholesale spot price (many of which are replicated in wholesale electricity markets in other countries). Figure 2.2-1 displays the weekly profile of median spot prices for spring and summer months (September 21 to March 20) and for fall and winter months (March 21 to September 20). A clear weekly and intra-daily pattern can be identified, where the latter changes significantly for the different seasons of the year. During spring and summer, the highest prices tend to occur during the middle of the day, reflecting the increased use of air-conditioners. In fall and winter months, high prices tend to occur in the early morning and toward the end of the day as a result of increased domestic electricity demand in colder periods.

(INSERT Figure 2.2-1 HERE)

In Figure 2.2-2 we show the time series of (log) daily average spot prices from 2001 to 2010 in NSW (the time series of wholesale spot prices for the other regions share the same main features). This series displays all the common features of deregulated electricity prices, especially the extreme temporary price jumps. It also appears as if there is significant temporal dependence in the occurrence of the extreme price events.

The extreme price fluctuations are commonly associated with unplanned generator outages or unexpected demand jumps often caused by extreme weather conditions. These abrupt, short-lived and generally unanticipated extreme price changes are known as price spikes or jumps and have attracted significant attention in the empirical electricity price modelling literature (Park et al. 2006; Weron 2009; Weron 2008; Simonsen 2005; Christensen et al. 2012; Christensen et al. 2009; M. T. Barlow 2002; Jong & Huisman 2002; Lucia & Schwartz 2002; Byström 2005; Cartea & Figueroa 2005; Burger, Klar, Müller, et al. 2004).

(INSERT Figure 2.2-2 HERE)

As we are ultimately interested in the occurrences of these extreme price events we define a binary series of price spikes to indicate in which periods the wholesale spot price exceeds a certain threshold price (see formal definition in Section 2.3.1). In Figure 2.2-3 we display the weekly pattern for the frequency of spot prices exceeding A$100/MWh. It is apparent that the general results from Figure 2.2-1 are mirrored in the occurrence of price spikes. In

spring and summer months these are more likely to occur in the middle of the day, whereas in fall and winter months price spikes are most likely to occur in the early hours of the evening. No extreme price events occur during the night.

(INSERT Figure 2.2-3 HERE)

2.2.1. Literature Review

Most papers modeling electricity prices have applied fairly traditional econometric techniques, including traditional autoregressive time series models and nonlinear Markov switching models, in order to capture the unique features of electricity prices (e.g (Ethier & Mount 1998; Huisman & Mahieu 2003; Huisman & De Jong 2003)). A different category of approaches has been adapted from the continuous-time diffusion or jump-diffusion models often used in financial econometrics (e.g (Deng 2000; Bhanot 2000; Knittel & Roberts 2005; Pirrong & Jermakyan 2008; Barone-Adesi & Gigli 2002; Lucia & Schwartz 2002; M. T. Barlow 2002; Escribano Sáez et al. 2002; and Burger et al., 2008, for a good overview of the different approaches)).

However, all of these models operate on the price or log of the price series. When doing so, the presence of price spikes constitutes a significant complicating factor and has to be dealt with by adding a jump component to the price process (see (Barz & Johnson 1998; Deng 2000; Deng & Jiang 2005; Huisman & Mahieu 2003; Ethier & Mount 1998)). In doing so they regard price jumps as a memory-less process and its intensity as independent of its own history.

However, it was established by Christensen et al. (2012) and Christensen et al. (2009) that the price spike process contains some significant persistence. In this strand of research they concentrate on a binary series of price spikes. In Christensen et al. (2009) a modified PAR framework to model price spikes is proposed. Therein the occurrence of a spike is a result of the latent arrival and survival of system stresses. Christensen et al. (2012) used an autoregressive conditional hazard (ACH) model to treat the conditional intensity of the price process as a function of the log duration between previous price spike events. Here the approach is to model the durations as an ARMA (p, q) process.

In the area of high-frequency finance a related literature has been developing. Financial market events such as the occurrence of a transaction or the update of a price quote have been successfully modeled as a binary series (or point process) via a time-varying intensity-based Hawkes model. Some of the works based on this type of model include the autoregressive conditional intensity (ACI) model proposed by Russell (1999) and the bivariate ACI model to estimate the arrival of buy and sell trades on a limit order market by Hall & Hautsch (2007). Extensions of the ACI model include the addition of a latent, Gaussian autoregressive component to the log intensity by Bauwens & Hautsch (2006). Bowsher (2007) proposes an extension of the ACI model by using vector-conditional intensity to model the two-way interactions of trades and quote changes in continuous time.

2.3. Methodology

In this section we discuss how price spikes, as a typical feature of deregulated electricity prices, can be modelled using a univariate Hawkes point process. The model is also used to form predictive probabilities for the potential occurrence of price spikes based on the past occurrence of the price events and the values of other exogenous variables. In Section 2.3.2 we will provide a short introduction to the PAR model, completing the set of models used in our forecast evaluation.

2.3.1. Hawkes point process model

In this chapter, we focus on modelling one particular aspect of the wholesale electricity price process: the occurrence of price spikes. Accordingly, let $s_{t,j}$ be the electricity spot price for the jth half hour ($j = 1, \dots, 48$) on the tth day. Here we define the daily series of price events as follows:

$$y_t = \begin{cases} 1 & \text{if any } s_{t,j} > \tau, \text{for any } j = 1, \dots, 48 \\ 0 & \text{if any } s_{t,j} \leq \tau, \text{for all } j = 1, \dots, 48 \end{cases} \qquad (2.3\text{-}1)$$

where τ is a threshold value. Clearly the choice of the threshold is crucial. To enable comparison with the work of Christensen et al. (2009), a threshold of A\$100/MWh[1] is used. Under "normal" conditions the spot price fluctuates between A\$20 to A\$40 per MWh, and less than 2% of half-hourly spot prices exceed the threshold of A\$100/MWh. In what follows, the daily time series of $\{y_t\}$, as defined in (2.3-1), is treated as the dependent variable in our analysis. It will also be useful to identify the times at which a price event

[1] For robustness test, we also investigate the impact on the model when different thresholds are used. The result is given in Section 2.5.4

occurs, with the additional time index $\{t_i\}, i = 1, \dots, m$, where we assume that m price events occurred and t_i is the time at which the ith price event occurred. Further, we define the information set I_{t-1} to consist of all price events up to time $t-1$ and values for all other available exogenous variables x available at time $t-1$.

The econometric model presented in this section will model the probability of a price event occurring at time t, $P(y_t = 1|I_{t-1}) = \lambda_t$, based on information available at time $t-1$. The exogenous variables that will later be considered in the modeling of λ_t are the unanticipated increase/decrease in electricity load ($load_t$), unusually hot and cold temperatures ($Tmax_t$ and $Tmin_t$), and the number of half-hourly spot prices that exceed the threshold ($count_t$). They are further explained in Section 2.4. We will also allow for deterministic seasonal variation.

The modelling approach adopted follows that of a self-exciting process, allowing for the probability λ_t to be dependent on past event occurrences and on realization of exogenous covariates in a way that may cause clusters of events. In particular, we allow the time path of λ_t to follow a univariate Hawkes process

$$\lambda_t = \mu_t + \sum_{t_i<t} w(t - t_i) \qquad (2.3\text{-}2)$$

where μ_t is a positive ($\mu_t > 0$) deterministic function and $w(t - t_i) > 0$ to ensure nonnegativity. In the notation of Hawkes processes, λ_t is also called intensity, and for the purpose of this chapter we will use the terms "event probability" and "intensity" interchangeably[2].

The intensity λ_t is modeled as the sum of two parts: a predictable or deterministic (seasonal) component,μ_t, and a stochastic component, $\sum_{t_i<t} w(t - t_i)$, which depends on the occurrence of price events (at times $t_i < t$) that occurred prior to the current time t. We will now discuss the specification of the deterministic and stochastic components in turn.

From the initial observation of weekly and annual seasonality in Section 2.2, it appears appropriate to allow for some deterministic, seasonal variation in μ_t. Following Lucia & Schwartz (2002); Knittel & Roberts (2005); Heydari & Siddiqui (2010) and Becker et al.

[2] Formally, the intensity is defined as the expected number of events in a time interval. It therefore can exceed 1. However, by definition of our series y_t we have either 0 or 1 event per day. Therefore, $p(y_t = 1|I_{t-1})$ = min(λ_t,1). It will transpire in the empirical application in Section 2.5 that only on very rare occasions does $\lambda_t > 1$. We shall interpret the intensity as the probability of this event occurring.

(2004), we specify μ_t to be the sum of trigonometric functions with frequencies designed to allow for a weekly and annual patterns. The trigonometric function is estimated by fitting these functions to the binary events series $\{y_i\}$ using ordinary least squares. Only statistically significant trigonometric functions are retained and incorporated into the deterministic components, μ_t, of the Hawkes model[3]. We can express this component as

$$\mu_t = trig_t \zeta \tag{2.3-3}$$

where $trig_t$ is a $(1 \times q_{tr})$ vector with the remaining trigonometric terms and ζ is a $(q_{tr} \times 1)$-parameter vector. This parameter vector will be estimated along with the remaining parameters of the Hawkes models described below.

The stochastic component of the intensity, λ_t (in equation (2.3-2)), is defined as follows:

$$\begin{aligned} \lambda_t &= \mu_t + \sum_{t_i<t} w(t - t_i) \\ w(t - t_i) &= \alpha_{t_i} e^{-\beta(t-t_i)} \end{aligned} \tag{2.3-4}$$

The intensity process is described by a left-continuous sample path, jumping up by an amount α_{t_i} in response to the occurrence of a price event at time t_i. The intensity λ_t then decays according to $e^{-\beta(t-t_i)}$, until we see another jump in the intensity by $\alpha_{t_{i+1}}$ as a response to the price spike at time t_{i+1}. At any period t, the sum of contributions from past price events can introduce persistence into the price events process, as more recent price events increase the probability of an event occurring in the next period.

In its basic form, with $\alpha_{t_i} = \alpha$, the intensity process resembles an autoregressive process in which the intensity merely depends on deterministic seasonality and the occurrence of previous price spikes. It is, however, plausible to conjecture that the probability of a price event will also depend on other (stochastic) covariates. Therefore, Bowsher (2007) suggested making the scalar parameter α_{t_i} dependent on a set of covariates at time t_i. We consider the $(1 \times q_x)$ vector x_{t_i}, which collates the information of q_x covariates at time, t_i, and specify

$$\alpha_{t_i} = \Phi(\omega_{t_i})\beta, \tag{2.3-5}$$

$$\omega_{t_i} = x_{t_i}\gamma, \tag{2.3-6}$$

[3] In principle, the selection of the relevant trigonometric terms could be done in the context of the nonlinear optimisation of the entire Hawkes model. The latter, however, is a fairly complicated model to optimise, and the problem is greatly facilitated by the prior selection of relevant trigonometric terms.

where γ is a $(q_x \times 1)$ vector of linear parameters and $\Phi(\cdot)$ is the standard normal cumulative distribution function. Since $\omega_{t_i} \in \mathbb{R}$, $\Phi(\omega_{t_i})\beta$ ensures that $\alpha_{t_i} \leq \beta$, which is a sufficient condition for the point process to be stationary (see (Ogata 1978)).

In the above specification the influence of covariates is restricted to the times, t_i, at which price events occur. It may be appropriate to let conditioning information play a role outside these periods as well. The natural way to achieve this is to let the decay parameter β be a function of the conditioning variables as well. The time-varying decay parameter β_t is then parameterised as

$$\beta_t = \Phi(z_t\delta), \tag{2.3-7}$$

restricting β_t to be between 0 and 1. The relevant covariates at time t are contained in the $(1 \times q_z)$ vector z_t and δ is the associated $(q_z \times 1)$ vector of parameters. With this parameterization the conditional intensity is reformulated as follows:

$$\lambda_t = \mu_t + \sum_{t_i<t} \alpha_{t_i} \exp\left(-\sum_{\tau=t_i}^{t} \beta_\tau\right) \tag{2.3-8}$$

The parameter vector $\theta = (\zeta' \; \gamma' \; \delta')'$ is estimated by maximising the loglikelihood function of the observed price events at times $t_1, \dots, t_n$:

$$\begin{aligned} &\log L(t_1, \dots, t_n|\theta) \\ &\quad = -\int_0^T \mu_t \, dt - \int_0^T \sum\nolimits_{t_i<t} \alpha_{t_i} e^{-\beta_t(t-t_i)} \, dt \\ &\quad + \int_0^T \log\left(\mu_t + \sum_{t_i<t} \alpha_{t_i} e^{-\beta_t(t-t_i)}\right) dN(t) \end{aligned} \tag{2.3-9}$$

where $dN(t)$ is an indicator function that takes a value of 1 if a price event occurred during period t and a value of 0 if no price event occurred during period t.

2.3.2. The Poisson autoregressive model

Christensen et al. (2009) utilizes a PAR model to model the intensity of price events in the Australian electricity market. For comparison purposes we will apply this model as well. As it turns out, the PAR model will allow for substantially different dynamics of the latent intensity process compared with the Hawkes process. We will therefore compare the suitability of the two different models for forecasting the price spike process.

The full details of the PAR model can be found in Christensen et al. (2009) and in Appendix A. Here we will provide an outline of the model's main components. The model is built around the process of a latent variable, $X_t \in \mathbb{N}^0$, representing the number of stresses in the system. In our example of modeling electricity price spikes, system stresses could represent failures of generators or transmission lines, or events that trigger unanticipated demands. While this variable is latent, it is linked to the observable variable of price events as follows:

$$y_t = \begin{cases} 1 & X_t > 0 \\ 0 & X_t = 0 \end{cases} \qquad (2.3\text{-}10)$$

Hence, a price event is the result of a minimum of one system stress occurring. What remains is the description of the arrival and departure processes (as any existing stress factor is allowed to remain for several periods) of these system stresses. The arrival process for a system stress is modeled as an independent Bernoulli process, and the departure as a binomial thinning process. Both the probability of arrival and the probability of departure can be modeled as either fixed probabilities or as functions of exogenous variables. In Appendix A we provide details of the somewhat involved recursive calculations that are needed to obtain a likelihood function, which is then optimized to obtain maximum likelihood parameter estimates.

2.4. Data of Australian National Electricity Market

During the 1990s, the Australian electricity market started a gradual deregulation by opening up their electricity sectors to competition. By December 1998 a National Electricity Market (NEM) had been formed between the states of New South Wales, Queensland, South Australia and Victoria. The electricity grids of these four states were fully connected by February 2001. According to Christensen et al. (2012), the electricity price data exhibit significant differences before and after the full interconnection, and therefore it seems reasonable to let the investigation period begin on March 1, 2001. The data set is then split into two periods. The estimation period begins from March 1, 2001 and ends on July 31, 2010, and the forecast period begins on August 1, 2010 and ends July 31, 2012.

The available data is half-hourly electricity spot prices, system load[4] and daily maximum and minimum temperatures for each state[5].

(INSERT Table 2.4-1 HERE)

[4] This data is available for download from the Australian Electricity Market Operator (AEMO) website: http://www.aemo.com.au/data/price_demand.html.

[5] Temperature data was obtained from the Australian Government Bureau of Meteorology (URL: www.bom.gov.au). Temperatures were measured in the respective state capitals.

The descriptive statistics for the estimation sample in Table 2.4-1 reveal that electricity spot prices in all regions have high levels of excess kurtosis and are heavily skewed to the right. These are common observations for deregulated electricity prices (Escribano Sáez et al. 2002). It is, of course, the thick right-hand tail that is associated with price events as defined in this chapter (ie, days on which the price exceeds A$100/MWh for at least one half-hour period). The last row in Table 2.4-1 indicates that price spikes occur on a significant number of days.

We observe the daily average spot prices in Victoria seem to have highest skewness and kurtosis compared to other regions. According to the State of the Energy Market report published by Australian Energy Regulator (AER) in 2011, the baseload coal power plant supplies more than 60% of New South Wales and Queensland demand capacity and around 56% of Victoria demand capacity. This means higher cost power plant such as gas fired generation is used to satisfy a larger proportion of demand in Victoria. Given the convexity of the electricity supply curve and almost perfectly inelastic electricity demand curve, when gas fired generation is the marginal generator even a small increase in demand would have caused a large price spikes. This could explain the high skewness and kurtosis in Victoria's daily average spot prices compared to other regions.

While in South Australia, their daily average spot prices produce highest standard deviation and proportion of days (17.17%) with price spikes. These price series come second after Victoria being second highest in skewness and kurtosis. This is likely due to the fact that South Australia relies heavily on peaking gas fired generation and much less on baseload coal power plants (supplies only 20% of South Australia demand capacity) when compared to the other regions.

Any reasonable forecasting model should take advantage of potential exogenous variables that may be related to the occurrence of price events. As the spot price is determined by the interactions between electricity supply and demand, variables that determine these are our prime candidates for potential covariates.

2.4.1. Electricity load

It can be seen in Figure 2.4-1 that the daily load series behaves in a very regular way. We can clearly discern an overall trend, seasonal fluctuations (with load being generally highest in winter and intermediate peaks for the summer period) and a weekly seasonal

pattern[6]. Furthermore, there is a vast literature on load forecasting (see Weron (2007) for an overview), and many regulators and market participants will have their own sophisticated, proprietary load forecasting model. These models will go well beyond modeling seasonal regularities and use information that is very specific to the upcoming day. For the purpose of modeling the probability of a price event occurring at time t it is important to be precise about the load information available at time $t-1$ that should be utilized in such a model. We conjecture that it would be the predicted load for day t that exceeds the normally expected load at that time of the year. We are therefore required to obtain both load forecasts (at day $t-1$ for day t) and a normally expected load.

(INSERT Figure 2.4-1 HERE)

The normally expected load series, l_t^e, is constructed as a weighted average of the previous seven days of load, $l_{t-1}, \dots, l_{t-7}$. The weights in this calculation are specific for each day of the week and are optimized to minimise the in-sample fit of this series. Unfortunately, a series of load forecasts is unavailable at this stage. We therefore use the actual realised load series, l_t, making the assumption that the available proprietary load-forecasting models would produce fairly accurate forecasts of these[7].

In short we calculate

$$\widetilde{el}_t = l_t - l_t^e \qquad (2.4\text{-}1)$$

As our very simple model that calculates l_t^e is likely to miss out on some annual seasonality, $\widetilde{el}_t$ is then deseasonalized from annual seasonality by applying the rolling volatility technique proposed by Weron (2007), resulting in $load_t$. From Figure 2.4-2 it is clear that the deseasonalized time series, $load_t$, has little apparent seasonality and/or trend, making it a plausible candidate to capture the forecast electricity demand that exceeds the demand we would usually expect for that day of the week at that particular time of the year. This will be the load series used to forecast the intensity at time t.

(INSERT Figure 2.4-2 HERE)

[6] The high frequency fluctuations in Figure 2.4-1 reflect that the electricity demand on weekends is significantly reduced.

[7] Admittedly, this is an unattainable (at time $t-1$) proxy to the preferred series of unseasonal load predictions. The results here could be interpreted as the maximum attainable forecast performance (in the context of the relevant model) if we had extremely good load forecasts available. A similar argument can be made for the temperature series in the next subsection.

2.4.2. Maximum and minimum daily temperature

A significant amount of variation in system load is related to variation in the prevailing temperature. For instance, in winter, colder days trigger a higher electricity demand through additional heating needs. Equally, very hot days in summer increase the demand for electricity through additional use of air-conditioning. We would therefore expect to find a close (although nonlinear) relation between temperature and load.

A significant part of that temperature variation is predictable, in particular, the general seasonal pattern. It is expected that in general the generator's bidding pattern allows for these variations and that price spikes will be the result of unexpected or unseasonal temperature extremes.

The maximum, $Tmax_t$ and minimum, $Tmin_t$, temperatures series in this chapter are designed to capture such unexpected and unseasonal hot days (in the warm seasons of the year) and extremely cold days (in the cold season of the year), respectively. Using the $Tmax_t$ series as an example, we therefore consider the absolute difference of the observed maximum daily temperature, $T^{obs}_{max_t}$, from the seasonally expected maximum temperature, $T^{s}_{max_t}$. The $T^{s}_{max_t}$ series is obtained by fitting a seasonal pattern of maximum daily temperatures, modeled by trigonometric functions[8]. As an unseasonably warm winter day is unlikely to have the same impact on the probability of a price spike as a very hot summer day, we want to eliminate the former from the $Tmax_t$ series. This is achieved by applying the following definition:

$$Tmax_t = \begin{cases} |T^{obs}_{max_t} - T^{s}_{max_t}| & if\ T^{obs}_{max_{t_i}} > T^{s}_{max_t}\ and\ T^{s}_{max_t} > \bar{T}_{max} \\ 0 & otherwise \end{cases} \quad (2.4\text{-}2)$$

where $\bar{T}_{max}$ is the average of $T^{obs}_{max_t}$series. The same considerations are applied to $Tmin_t$, ensuring that this series captures unseasonably cold days in the winter season only. As a result of this operation, we do not expect these series to display any significant (linear) correlation with the system load across the entire sample.

As for the load series it is important to carefully consider the timing of the variables to be used for forecasting the intensity λ_t. The relevant temperatures will be those on day t. These are, however, not available on day $t-1$ and are as such not useful for building a

[8] We use the same approach to fitting trigonometric functions as that described in Section 2.3.1. Details are available upon request.

forecast model for intensity λ_t. In practice we would want a temperature forecast for day t available at day $t-1$. No such series are easily available as a long time series and in the context of this chapter we shall assume that very precise and unbiased temperature forecasts will be available at time $t-1$. We therefore use $Tmax_t$ and $Tmin_t$ in our forecast model for intensity λ_t.

2.4.3. Count

The fourth covariate used in our models is the $count_t$ variable. It represents the number of half-hourly spot prices, $s_{t,j}$, that exceed the threshold, τ (i.e. $A\$100/MwH$) on day *t*. The $count_t$ variable is used as a proxy for the number of price spikes during a day and it may well contain valuable information. A high number of such threshold exceedences in a particular day might indicate the presence of a more severe system stress that cannot be resolved quickly. This may well indicate an increased probability for another price event during the following day.

2.4.4. Weekdays

In Section 2.3.1, we argued that the baseline intensity (the probability of a price event occurring) of the Hawkes model may well depend on a deterministic variable such as a weekend / no-weekend dummy variable (through a trigonometric series with an appropriately chosen frequency). However, the deterministic variable may affect the intensity not only through different baseline intensities but perhaps through a time-varying intensity decay rate, β_t. Appropriately defined dummy variables will therefore be considered in the estimation process.

(INSERT Table 2.4-2 HERE)

In Table 2.4-2 we report summary statistics and a correlation matrix for these variables in NSW (information for other regions in Table 2.4-3, Table 2.4-4, and Table 2.4-5). It is apparent that the correlations between these variables are at most moderate, justifying their separate consideration as explanatory variables in our intensity models.

(INSERT Table 2.4-3 TO Table 2.4-5 HERE)

We realise that none of these exogenous variables are incorporating information from the electricity supply. Ideally, we would like to use data on reserve margin forecast for this purpose (Chen & Bunn 2010; Bunn et al. 2012). However, unlike the UK electricity

market where the system operator produces forecasts of the available reserve margin for each half-hourly trading period, we cannot find such data available from AEMO. Augmenting the conditioning covariates with information on drivers of the reserve margins is one of the extensions that we would want to incorporate in the future works on this chapter.

2.5. Estimation

The Hawkes model specifications used in this chapter are the seasonal Hawkes with time-varying α_t (and constant β), which we denote by HAWa, and the seasonal Hawkes with time-varying α_t and β_t, which we denote by HAWab, where their parameters are defined in equation (2.3-5) and (2.3-7). These parameters are estimated by maximising the loglikelihood function in equation (2.3-9), while the standard errors of the parameters are computed using the typical robust sandwich estimator for standard errors[9]. These models are called the seasonal Hawkes models, as they both allow the baseline intensity to vary with deterministic functions as described in equation (2.3-3). Further, we will apply the PAR model as specified in Section 2.3.2 and Appendix A.

In the remainder of this section we will report the estimation results of the models (Section 2.5.1 – Section 2.5.3). This is followed by a brief discussion of a small robustness test in Section 2.5.4.

2.5.1. Hawkes with time-varying α

This is the most basic Hawkes model for which we report results. The intensity process is represented by

$$\lambda_t^{sea,tva} = \mu_t + \sum_{t_i<t} \alpha_{t_i} \, e^{-\beta(t-t_i)} \qquad (2.5\text{-}1)$$

where "sea" and "tva" indicate that the intensity is seasonally adjusted and allows for a time-varying α.

We have considered more basic model specifications, restricting μ and α to be time-invariant. However, these restrictions are comfortably rejected by likelihood ratio tests and are therefore not considered here[10].

[9] We are grateful to Micheal Rockinger for providing a MATLAB code to estimate the standard errors.

[10] Results are available upon request from the authors.

One of the difficult aspects of empirical modeling is determining an appropriate estimation period. In the absence of structural breaks we would prefer to estimate on as long a sample as possible. However, if there are structural breaks in the sample period, then a shorter estimation may be advisable (see Pesaran & Timmermann (2007) for a detailed discussion on the issue of optimal estimation window choice). Here we apply a somewhat simplified approach. We split the estimation period (March 1, 2001 to July 31, 2010) into three sub-periods (March 1, 2001 to July 31, 2004; August 1, 2004 to July 31, 2007; August 1, 2007 to July 31, 2010). We then estimate the model on all these periods and test the null hypothesis (no structural breaks) with a likelihood ratio (LR) test. This null hypothesis is rejected at all standard significant levels. Therefore, we choose to estimate the models on a shorter subperiod beginning with August 1, 2007 to July 31, 2010 for the first forecast.

The explanatory variables used to explain variation in α_{t_i} for equation (2.5-1) are the load ($load_{t_i}$) and temperature variables ($Tmax_{t_i}$ and $Tmin_{t_i}$) discussed in Sections 2.4.1 and 2.4.2 as well as the number of half-hourly periods that the wholesale spot price exceeded the price threshold ($count_{t_i}$; see Section 2.4.3).

The important feature of this model is that the increase in intensity following a price event can vary with the conditions that prevailed on the day, t_i, of the price event. The occurrences of a price event therefore increase the intensity for future events, resulting in event "clustering", which is an obvious feature of our electricity price data (Figure 2.2-2). If α_{t_i} is a constant parameter, each price event carries the same effect on future intensity. However, since α_{t_i} is determined by a set of covariates, each price event contributes differently (depending on the value of covariates at time of the event, t_i) to future intensity, λ_t. Larger values of α_{t_i} result in larger contributions to the intensity of future price events.

(INSERT Figure 2.5-1 HERE)

This is illustrated in Figure 2.5-1, showing how the contribution of an individual price event to the probability of future price events decreases over time. In this example the contribution decreases with a constant rate of decay, β, but different initial contributions α_{t_i} in events 1 - 3.

The result in Table 2.5-1 shows how coefficients of explanatory variables in α_{t_i} determine the relative importance of particular price events toward the probability of future events.

The coefficients to the load ($load_{t_i}$) variable are negative and significant in NSW and QLD and significantly positive in VIC. This inconsistent result is somewhat difficult to explain and may well be due to the fact that the decay is restricted to be constant and the model could be misspecified. We therefore suspect that the inconsistency of the estimation results for the load variable is due to this misspecification. Indeed, once we relax this restriction in the Hawkes model with time-varying beta, the coefficient of load variable is consistently negative and significant for all of the regions (refer to Table 2.5-2).

The variable tracking the number of price spikes during a day ($count_{t_i}$) is estimated consistently positive and significant in QLD and VIC, and thus increases probability of future price events through an increase in α_{t_i}. The coefficients of unpredicted maximum, $Tmax_{t_i}$ and minimum, $Tmin_{t_i}$, temperatures variables are negative and highly significant in all regions except in SA. This means both of these variables have a mainly negative effect on α_{t_i}. Therefore, price events that occurred due to extreme and unpredicted temperatures cause (ceteris paribus) a smaller increase in the intensity for future price events. This result can be interpreted in the following way. Very extreme temperatures (high values for $Tmax_{t_i}$ and $Tmin_{t_i}$) are likely to cause quicker (and possibly more radical) adjustments to the supply schedule such that the next periods are less likely to exhibit price events.

(INSERT Table 2.5-1 HERE)

(INSERT Figure 2.5-2 HERE)

Figure 2.5-2 displays the conditional intensities, $\lambda_t^{Sea,tva}$ against the time-series of price events, y_t, for all the regions. By definition, the price events (y_t) series equal to 0 if there is no event on day, t or 1 if an event occurred, while the conditional intensities series are the probability of an event occurring on day t. The figure illustrates how repeated price events trigger a sustained rise in the intensity, which only decays to the baseline level during a sustained absence of price spikes. Clearly, the seasonality in baseline intensity is dominated by the autoregressive feature of the intensity series. These results are qualitatively similar across all four states.

2.5.2. Hawkes with time-varying α and β

One limiting feature of the model in the previous subsection is that the influence of covariates is restricted to the event times t_i, as the coefficient α_{t_i} is only relevant on the

occasion of a price event. By relaxing the assumption of a constant decay parameter, β, it is possible to introduce the influence of covariates in any period t. The seasonal Hawkes, time-varying α and β model allows for the initial intensity impact of a price event, α and the decay parameter, β, to be time-varying and to be related to covariates. The intensity process is then represented by

$$\lambda_t^{sea,tvab} = \mu_t + \sum_{t_i<t} \alpha_{t_i} \exp\left(-\sum_{\tau=t_i}^{t} \beta_\tau\right) \qquad (2.5\text{-}2)$$

where "tvab" indicates that α and β are allowed to vary through time.

Figure 2.5-3 illustrates the effect of changing betas on the contribution of past events to the intensity λ_t

(INSERT Figure 2.5-3 HERE)

The explanatory variables contributing to the initial contribution, α_{t_i} and time-varying decay parameter, β_τ in equation (2.5-2) are as defined in Section 2.5.1 but with the addition of a weekend / no-weekend dummy variable in the decay parameter, β_τ.

The result in Table 2.5-2 shows that most of the variables in α_{t_i} give consistent results with those in the seasonal Hawkes (HAWa) model but now for a consistently negative coefficient for the load variable. For the time-varying decay parameter, β_τ, the coefficient of electricity load, $load_\tau$, and the number of price spikes in a day, $count_\tau$ are negative for most of the regions, while the unpredicted minimum temperature, $Tmin_\tau$, and maximum temperature, $Tmax_\tau$, are mainly positive, expect for SA. To recall, Figure 2.5-3 illustrates the intensity decays at slower rate with lower value of β_τ. Therefore, negative coefficients from $load_\tau$ and $count_\tau$ decrease the contribution to β_τ, resulting in the intensity decaying at a slower rate. The intensity that reverts slowly carries longer memory, and thus increases probability of future price events. The unexpected increase in load and frequently price spikes occurrences within 48 half hours indicates the possible presence of a more severe system stress, which cannot be resolved quickly.

(INSERT Table 2.5-2 HERE)

As discussed in relation to the previous model, larger unseasonal and unpredicted temperatures have a smaller impact on future intensity, and this is confirmed by increased intensity decay. The $weekdays_\tau$ dummy variable proves not to be statistically significant.

(INSERT Figure 2.5-4 HERE)

Figure 2.5-4 displays the conditional $\lambda_t^{Sea,tvab}$ against the time-series of price events, y_t for all regions. The seasonal Hawkes (HAWab) model is still able to produce a sustained rise in the intensity after episodes of price events, but once the assumption of a constant decay parameter β is relaxed, the intensity tends to rise higher and decays at different rates depending on the effects of the explanatory variables on the time-varying decay parameter, β_τ. When performing LR tests for the validity of the restriction that all the coefficients to the explanatory variables in the decay parameter are equal to zero, this restriction is comfortable rejected at any standard significance level[11].

2.5.3. PAR

The third model implemented is the PAR model. The covariates used to explain the time variation in arrival and survival of system stresses, z_{1t}' and z_{2t}' in the Christensen et al. (2009) modified PAR model are the same set of covariates used in the Hawkes models, excluding the $count_t$ variable[12]. In contrast with Christensen et al. (2009), we are using unseasonable load, $load_t$, and unpredicted maximum, $Tmax_{t_i}$, and minimum, $Tmin_{t_i}$, temperatures as covariates in arrival and survival of system stresses, z_{1t}' and z_{2t}'. We argue that one of the factors that influences system stresses is unexpected or unseasonable variations in these variables. Recall that the variables in z_{1t}' are those that control the probability of a new event occurring in period t, and the role of the variables in z_{2t}' is to set the probability of an existing system stress to disappear in period t.

The result in Table 2.5-3 shows the load ($load_t$) and unpredicted maximum and minimum temperatures ($Tmax_t$ and $Tmin_t$) variables of the arrival probability, π_t remain significant in most of the regions with positive coefficient. This suggests the unanticipated increase in electricity load and temperature motivate system stresses to happen. However, other more applicable explanatory variables are needed to explain variation in the survival probability, ρ_t, as all the chosen covariates are insignificant.

(INSERT Table 2.5-3 HERE)

(INSERT Figure 2.5-5 HERE)

[11] For brevity, these results are not reported but are available upon request.

[12] Including the $count_t$ variable into the PAR model prevented the optimization from obtaining a global maximum.

Figure 2.5-5 shows the conditional intensities of PAR model, p_t, are more volatile than those of both of the Hawkes models. It delivers higher intensities but reverts very quickly (if not instantaneously) to a low level after prior events. Since the arrival, π_t, and survival, ρ_t, probabilities are conditioned on any time t and determined by a set of covariates, the intensity of the PAR model, p_t, shows more variation during times without price events although its persistency would not survive once price events are over.

In Figure 2.5-6 we compare the intensities for each of the models.

(INSERT Figure 2.5-6 HERE)

All three models are able to produce increasing intensity after a series of price events but in some periods the PAR model delivers higher intensity. Although the PAR model produces faster intensity jumps and decays, the intensity of Hawkes models has stronger persistency even after the price events end. One of the reasons is the arrival, π_t, and survival, ρ_t, probabilities of system stresses are conditioned on any time period, t. In the Hawkes models, however, the scale parameter, α_{t_i}, is conditioned at all times on all previous price events. During nonprice events or benign periods the intensity of Hawkes models show less variation compared to the PAR model. The baseline of the Hawkes models only reflect the deterministic, seasonal component, μ_t, while the baseline of the PAR model is conditioned on the probability of arrival and survival of system stresses.

As the PAR and Hawkes models allow for such different intensity dynamics (when estimated on the same data set) this demonstrates that there is a place for both types of models. In other words, one type of model may be better suited for one data set, while the other will perform better on another data set. In Section 2.6 we will evaluate which of the models is better suited for forecasting price events in the context of our data set.

2.5.4. Effects of different threshold levels

In this section we investigate the impact on the model when a different threshold of price exceedance is used. The latest research by Christensen et al. (2012) used an alternative threshold of A$300/MWh to define price events. It is equivalent to the strike price of some derivative products in the NEM. Similarly, the seasonal Hawkes with time-varying α_t (and constant β) (HAWa) is re-estimated using a threshold of A$300/MWh instead of A$100/MWh. Figure 2.5-7 displays the events and the estimated intensity for the HAWa

model applied to the event series based on the two different price thresholds. It is obvious from this figure that the series based on the A$300/MWh threshold has much less structure (in particular seasonality), although some pattern of persistence remains. Apparently, these are series with very different characteristics and in that sense models with different thresholds are difficult to compare. It will be an interesting line of future research to establish how the characteristics of the event series change as the threshold increases.

(INSERT Figure 2.5-7 HERE)

2.6. Forecasting

The practical value of any model is usually determined by its ability to give prior warning of future price events. In this section, the forecast performances of the models described in Section 2.5 are compared. One-day-ahead intensity forecasts are produced for all models. Additionally, the parameters used for forecasting purposes are re-estimated every thirty days with estimation windows of constant size.

(INSERT Figure 2.6-1 HERE)

The one-day-ahead forecast intensities for all models (using NSW as an example) are demonstrated in Figure 2.6-1. The forecast intensities of HAWab and PAR models deliver higher and decay faster after episodes of price events than the HAWa model. The forecast intensities by both Hawkes models give a genuine nonprice-event signal during benign periods. But the forecast intensity of the PAR model reacts faster once price events have happened, even though it also decays much faster once the price events are over. However, there are periods when only one price event occurs in isolation, and none of the models delivers high forecast intensity for that day. This seems to indicate that, in the main, isolated price events cannot be anticipated. Similar features of forecast intensities are observed in other regions (Figure 2.6-2, Figure 2.6-3 and Figure 2.6-4).

(INSERT Figure 2.6-2 TO Figure 2.6-4 HERE)

Following Christensen et al. (2012), the forecast performance of all three models are compared against a "naïve" model. The naïve model predicts that spikes that occurred during weekdays and weekends persist on the next day only if next day is a weekday. Since there are no universally superior measures of forecast accuracy, it is informative to evaluate the forecast performance in different ways. If all these measures were to point to

one superior model, we would have a more robust result on the model forecasting performance.

It would be most useful to implement either a derivatives trading or a hedging strategy based on the model's prediction. A realistic implementation of any such exercise is hindered by the following two aspects. First, most exchange traded-energy derivatives (certainly in Australia) are based on continuous price outcomes rather than a binary event series. Our modelling of a binary series therefore provides insufficient information to fully value any such derivatives. For hedging strategies we would have to devise hedging strategies for the continuous price process on the basis of the binary event predictions. This can be done in many different ways[13] and results could not be generalised beyond the particular implementation. Second, the exchange-traded derivatives market is very thin and this makes it difficult to use its prices for a comprehensive pricing comparison (even in the context of modeling a continuous price process).

The first class of feasible forecast evaluations uses the intensity measure as a predictor to the binary event series. The forecast error ($y_t - \lambda_{t|t-1}$ for Hawkes models and $y_t - p_{t|t-1}$ for the PAR model) are summarised using different forecasting measurements, such as mean absolute error (MAE), root mean square error (RMSE) and the asymmetric loss score (Asym) (see for example, Rudebusch & Williams (2009) and Christensen et al. (2009)).

Based on MAE and asymmetric loss score measurement in Table 2.6-1, the forecast intensity from the naïve model consistently outperforms the forecasts of other models when evaluated by the MAE and the "Asym" loss functions while the RMSE tends to favour the HAWab model.

Comparing the outcome of our event variable with the intensity forecast as done above is somewhat limiting and may not be representative of the way in which these models can be used in practice. Users are likely to look at the forecast intensity and then judge whether the signaled probability of a price event in the next period is significant enough to trigger some action, such as an increase in a hedge ratio. Consider that if the forecast intensity exceeds a certain threshold value, then the user judges that there is a substantial enough probability for an event happening such that they engage in such action. We say that in that

[13] See Christensen et al. (2012) for one example and also a discussion of the limitations of their approach.

case the user "predicted" an event. Such a prediction can then turn out to be correct (if an events occurs) or false.

The use of the term "predicted" needs to be clarified here. This seems to suggest that the user would predict an event if the intensity exceeds a value of 0.5. This is not necessarily the case, as a user may think that even a probability of (say) 30% of an event occurring is sufficiently large to trigger some sort of action. In fact, the value of this trigger intensity is crucial. What should it be? In practice this will depend on how costly it is to miss an event (eg, being exposed to a price hike) compared with the cost of making a false event prediction/alarm (eg, cost of hedging when in hindsight it turned out to be unnecessary).

(INSERT Table 2.6-1 HERE)

Therefore, our second forecast-evaluation measure will be based on evaluating the frequency of false alarms and missed events (see also Christensen et al. (2012)). In particular, we introduce a forecast failure statistic that is defined as

$$forecast\ failure = Tot^{false\ alarm} + \nu(Tot^{missed\ event}) \quad (2.6\text{-}1)$$

where the number of false alarms is defined as $Tot^{false\ alarm}$, and the number missed events as $Tot^{missed\ event}$. We also introduce a scaling factor, ν. This scaling factor reflects the relative cost of missed events as opposed to false alarms. A value of $\nu = 5$ means that missed events are five times more costly than false alarms. Different users will have different views on the value of ν, and in this chapter we will not take a view on this value. However, it is important, as it will have a major impact on the trigger intensity that will induce a user to take, say, additional hedging action. *Ceteris paribus*, the larger the value ν the smaller will be the trigger intensity.

The approach taken here is, for each model, to find that value for the trigger intensity that minimises the forecast failure statistic and to use this value to establish whether or not the predicted intensity of a particular model predicts an event. This is done for a range of values of $\nu\ (= 1, 2, \ldots, 10)$. For $\nu = 1$ this implies that we find the trigger intensity (for each model) that minimizes the total number of missed events and false alarms for any model.

Based on this, Figure 2.6-5 to Figure 2.6-10 report each model's number of missed events and false alarms as a function of ν, as well as the respective numbers for the naïve model. The number of missed events and false alarms depends on the trigger intensity and

therefore varies with the value of the trigger intensity and hence with ν. Generally, larger ν will lead to a lower trigger intensity and hence to more false alarms and fewer missed events.

A number of interesting findings transpire from this exercise. In all states except for SA, it appears feasible to reduce the number of missed events below that produced by the naïve model. However, this can only happen at the expense of producing more false alarms. In the case of NSW (Figure 2.6-5) and QLD (Figure 2.6-6), however, only the Hawkes models manage to produce a smaller number of missed events (for sufficiently large values of ν).

(INSERT Figure 2.6-5 to Figure 2.6-10 HERE)

For both these states the PAR model is unambiguously worse than the naïve model for large ν (greater than 5 in NSW and greater than 2 in QLD) as it produces both more missed events and more false alarms. For SA (Figure 2.6-7) the trade-off between missed events and false alarms is present for all three estimated models. In this case it is indeed the PAR model that (at sufficiently large values of ν) can produce the smallest number of missed events but again at the price of also producing the largest number of false alarms. For VIC (Figure 2.6-11 to Figure 2.6-12), however, the results of the Hawkes models are little different from those of the naïve model at all values of ν, whereas the application of the PAR model appears to be largely ineffective.

The overall finding, in the context of this chapter, is that it is very difficult to predict price events even by application of state-of-the-art econometric models. Even a very simplified and naïve prediction model is difficult to beat with the set of information used in this chapter.

(INSERT Figure 2.6-11 TO Figure 2.6-12)

2.7. Conclusions

The contributions of this chapter are threefold. First, we are able to establish exogenous variables that are relevant to explain variation in the intensity that drives the process of price events. These variables include unseasonal loads, maximum and minimum temperatures and the number of price events that occurred on the previous day. However, it is certain that this list is incomplete. An interesting line of future research will investigate

whether one-day-ahead market scenarios, which are published in near real-time by the Australian Electricity Market Operator, contain any information that is useful in predicting future price events.

Second, the results and findings in this chapter show significantly different characteristics of forecast intensities modeled by Hawkes and PAR models. Although all the models are able to produce sustained rises in the intensity after episodes of price events, the PAR model produces faster intensity increases and decays, while the Hawkes models deliver a more persistent intensity series. During periods with no price events the intensity of the Hawkes models show less variation than the PAR model. These significantly different model characteristics (when applied to the same series) suggest that there is a place for both models, depending on the series modeled.

The last contribution relates to the evaluation of the forecast performance of the models estimated in this chapter. Overall it is apparent that forecasting price events is a difficult task and neither Hawkes nor PAR models are able to predict price events that occurred in isolation. While in some Australian regions some models (in particular, the Hawkes-type models) were able to produce a smaller number of missed events than a reasonable but simplistic forecasting procedure, this comes at a price of an increased number of false predictions.

References

Anderson, E.J., Hu, X. & Winchester, D., 2007. Forward contracts in electricity markets: The Australian experience. *Energy Policy*, 35(5), pp.3089–3103. Available at: http://linkinghub.elsevier.com/retrieve/pii/S0301421506004289 [Accessed August 14, 2013].

Australia Competition & Consumer Conmission, 2002. *Final Determination on Bidding and Rebidding Code Changes*, Available at: http://transition.accc.gov.au/content/index.phtml/itemId/744502/fromItemId/401858/display/acccDecision.

Australian Competition & Consumer Commission, 2002. *Draft Determination on Bidding and Rebidding Code Changes*, Available at: http://transition.accc.gov.au/content/index.phtml/itemId/744502/fromItemId/401858/display/acccDecision.

Australia Electricity Regulator. (2011). State of the energy market , 25–51. Retrieved from http://www.aer.gov.au/sites/default/files/State of the energy market 2011 - complete report.pdf

Australian Energy Market Commission, 2013. National Electricity Rules Version 56 Chapter 3: Market Rules. , p.146. Available at: http://www.aemc.gov.au/Electricity/National-Electricity-Rules/Current-Rules.html.

Australian Energy Market Commission, 2012. Potential Generator Market Power in the NEM. *Rule Determination*. Available at: http://www.aemc.gov.au/electricity/rule-changes/completed/potential-generator-market-power-in-the-nem.html.

Australian Energy Market Operator, 2010. An introduction to australia's national electricity market. *Available at AEMO website http://http://www.aemo.com.au/About-the-Industry/Energy-Markets/National-Electricity-Market*.

Bardak Ventures Pty Ltd, 2005. *The Effect of Industry Structure on Generation Competition and End-User Prices in the National Electricity Market*.

Barlow, M.T., 2002. A diffusion model for electricity prices. *Mathematical Finance*, 12(4), pp.287–298.

Barone-Adesi, G. & Gigli, A., 2002. Electricity derivatives. *Unpublished manuscript, National Centre of Competence in Research*, (May). Available at: http://www.nccr-finrisk.uzh.ch/media/pdf/financeSeminarZH_SS02_giovanniBarone-adesi.pdf [Accessed August 14, 2013].

Barz, G. & Johnson, B., 1998. Modeling the prices of commodities that are costly to store: the case of electricity. In *Proceedings of the Chicago Risk Management Conference*.

Bauwens, L. & Hautsch, N., 2006. Stochastic Conditional Intensity Processes. *Journal of Financial Econometrics*, 4(3), pp.450–493. Available at: http://jfec.oxfordjournals.org/cgi/doi/10.1093/jjfinec/nbj013 [Accessed August 14, 2013].

Becker, R., Enders, W. & Hurn, S., 2004. A general test for time dependence in parameters. *Journal of Applied Econometrics*, 19(7), pp.899–906. Available at: http://doi.wiley.com/10.1002/jae.751 [Accessed August 14, 2013].

Bhanot, K., 2000. Behaviour of power prices: Implications for the valuation and hedging of financial contracts. *Journal of Risk*, 2, pp.43–62.

Borenstein, S., Bushnell, J. & Wolak, F., 2002. Measuring market inefficiencies in California's restructured wholesale electricity market. *American Economic Review*, 92(5), pp.1376–1405. Available at: http://www.jstor.org/stable/10.2307/3083255 [Accessed August 12, 2013].

Bowsher, C.G., 2007. Modelling security market events in continuous time: Intensity based, multivariate point process models. *Journal of Econometrics*, 141(2), pp.876–912. Available at: http://linkinghub.elsevier.com/retrieve/pii/S030440760600251X [Accessed August 14, 2013].

Bunn, D., Andresen, A., Chen, D., & Westgaard, S. (2012). Analysis and Forecasting of Electricity Price Risks with Quantile Factor Models

Burger, M. et al., 2004. A spot market model for pricing derivatives in electricity markets. *Quantitative Finance*, 4(1), pp.109–122. Available at: http://www.tandfonline.com/doi/abs/10.1088/1469-7688/4/1/010 [Accessed August 14, 2013].

Burger, M., Graeber, B. & Schindlmayr, G., 2008. Managing energy risk: An integrated view on power and other energy markets, Wiley. com.

Bushnell, J. & Saravia, C., 2002. An empirical assessment of the competitiveness of the New England electricity market. Available at: http://escholarship.org/uc/item/8913v4bk [Accessed August 12, 2013].

Byström, H.N.E., 2005. Extreme value theory and extremely large electricity price changes. *International Review of Economics & Finance*, 14(1), pp.41–55. Available at: http://linkinghub.elsevier.com/retrieve/pii/S1059056003000327 [Accessed August 14, 2013].

Cartea, Á. & Figueroa, M.G., 2005. Pricing in Electricity Markets: A Mean Reverting Jump Diffusion Model with Seasonality. *Applied Mathematical Finance*, 12(4), pp.313–335. Available at: http://www.tandfonline.com/doi/abs/10.1080/13504860500117503 [Accessed August 14, 2013].

Chen, D. & Bunn, D.W., 2010. Analysis of the Nonlinear Response of Electricity Prices to Fundamental and Strategic Factors. IEEE Transactions on Power Systems, 25(2), pp.595–606. Available at: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5378456.

Christensen, T., Hurn, S. & Lindsay, K., 2009. It never rains but it pours: modeling the persistence of spikes in electricity prices. *Energy Journal*, 30(1), p.25.

Christensen, T.M., Hurn, a. S. & Lindsay, K. a., 2012. Forecasting spikes in electricity prices. *International Journal of Forecasting*, 28(2), pp.400–411. Available at: http://linkinghub.elsevier.com/retrieve/pii/S0169207011000550 [Accessed August 14, 2013].

David, A.K. & Wen, F., 2000. Strategic bidding in competitive electricity markets: a literature survey. In *Power Engineering Society Summer Meeting, 2000. IEEE*. pp. 2168–2173 vol. 4. Available at: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=866982.

Deng, S., 2000. *Stochastic models of energy commodity prices and their applications: Mean-reversion with jumps and spikes*, University of California Energy Institute.

Deng, S.-J. & Jiang, W., 2005. Levy process-driven mean-reverting electricity price model: the marginal distribution analysis. *Decision Support Systems*, 40(3-4), pp.483–494. Available at: http://linkinghub.elsevier.com/retrieve/pii/S016792360400123X [Accessed August 14, 2013].

Eichler, M. et al., 2012. *Modeling spike occurrences in electricity spot prices for forecasting*, METEOR, Maastricht research school of Economics of Technology and Organizations. Available at: http://arno.unimaas.nl/show.cgi?fid=25393.

Electricity Market Performance, 2012. *Factors Contributing to Differences Between Dispatch and Pre-Dispatch Outcomes*, Available at: http://www.aemo.com.au/Electricity/Market-Operations/Dispatch/Factors-Contributing-to-Differences-between-Dispatch-and-Predispatch-Outcomes.

Escribano Sáez, Á., Pena, J.I. & Villaplana, P., 2002. Modeling electricity prices: international evidence. In *EFMA 2002 London Meetings*.

Ethier, R. & Mount, T., 1998. Estimating the volatility of spot prices in restructured electricity markets and the implications for option values. *Cornell University, Ithaca New York*, (1997), pp.1–45. Available at: http://www.pserc.wisc.edu/documents/publications/papers/1998_general_publications/pserc_31.pdf [Accessed August 19, 2013].

Fabra, N. & Toro, J., 2005. Price wars and collusion in the Spanish electricity market. *International Journal of Industrial Organization*, 23(3), pp.155–181. Available at: http://www.sciencedirect.com/science/article/pii/S0167718705000123 [Accessed August 12, 2013].

Gillett, R. & Market Operations, 2010. *Pre-dispatch Process Description*, Available at: http://www.aemo.com.au/Electricity/Market-Operations/Dispatch/Predispatch-Process-Description [Accessed August 12, 2013].

Hall, A.D. & Hautsch, N., 2007. Modelling the buy and sell intensity in a limit order book market. *Journal of Financial Markets*, 10(3), pp.249–286. Available at: http://linkinghub.elsevier.com/retrieve/pii/S1386418106000681 [Accessed August 14, 2013].

Hawkes, A.G., 1971. Spectra of some self-exciting and mutually exciting point processes. *Biometrika*, 58(1), pp.83–90.

Heydari, S. & Siddiqui, A., 2010. Valuing a gas-fired power plant: A comparison of ordinary linear models, regime-switching approaches, and models with stochastic volatility. *Energy Economics*, 32(3), pp.709–725.

Hortacsu, A. & Puller, S., 2005. Understanding strategic bidding in restructured electricity markets: a case study of ERCOT. , (010366). Available at: http://www.nber.org/papers/w11123 [Accessed August 12, 2013].

Huisman, R. & De Jong, C., 2003. Option pricing for power prices with spikes. *Energy power risk management*, 7(11), pp.12–16.

Huisman, R. & Mahieu, R., 2003. Regime jumps in electricity prices. *Energy Economics*, 25(5), pp.425–434. Available at: http://linkinghub.elsevier.com/retrieve/pii/S0140988303000410.

Jong, C. De & Huisman, R., 2002. Option Formulas for Mean-Reverting Power Prices with Spikes This paper can be downloaded without charge from the Social Science Research Network at : E RASMUS R ESEARCH I NSTITUTE OF M ANAGEMENT REPORT SERIES.

Kharbach, S., Fredette, M. & Pineau, P.-O., 2010. *Impacts of Imports and Natural Gas on Electricity Prices: The Case of Ontario*, Groupe d'études et de recherche en analyse des

décisions. Available at: http://www.gerad.ca/fichiers/cahiers/G-2010-70.pdf [Accessed August 12, 2013].

Knittel, C.R. & Metaxoglou, K., 2008. Diagnosing unilateral market power in electricity reserves market. Available at: http://escholarship.org/uc/item/14q6c0mk.pdf [Accessed August 12, 2013].

Knittel, C.R. & Roberts, M.R., 2005. An empirical examination of restructured electricity prices. *Energy Economics*, 27(5), pp.791–817. Available at: http://linkinghub.elsevier.com/retrieve/pii/S0140988305000228 [Accessed August 14, 2013].

Lucia, J.J. & Schwartz, E.S., 2002. Electricity prices and power derivatives: Evidence from the nordic power exchange. *Review of Derivatives Research*, 5(1), pp.5–50.

Mansur, E., 2001. *Pricing behavior in the initial summer of the restructured PJM wholesale electricity market*, University of California Energy Institute. Available at: http://www.ucei.berkeley.edu/ucei/PDF/pwp083.pdf.

Mansur, E.T., 2007. Upstream competition and vertical integration in electricity markets. *Journal of Law and Economics*, 50(1), pp.125–156. Available at: http://www.jstor.org/stable/10.1086/508309.

Martin, V., Hurn, S. & Harris, D., 2012. *Econometric modelling with time series: specification, estimation and testing*, Cambridge University Press.

National Electricity Code Administrator, 2001a. *Generators ' bidding and rebidding strategies and their effect on prices*, Available at: http://www.neca.com.au/TheCode0378.html?CategoryID=34&SubCategoryID=83&ItemID=951.

National Electricity Code Administrator, 2001b. *Generators' bidding and rebidding strategies and their effect on prices Written comments in response to Panel's consultation document*, Available at:

http://www.neca.com.au/Files/C_Bidding_rebidding_strategies_and_effect_on_prices_written_comments_Sep2001.pdf.

Ogata, Y., 1978. The asymptotic behaviour of maximum likelihood estimators for stationary point processes. *Annals of the Institute of Statistical Mathematics*, 30(1), pp.243–261. Available at: http://www.springerlink.com/index/247815U5W2407822.pdf.

Oh, H., 2003. *Simulation methods for modeling offer behavior and spot prices in restructured markets for electricity*, Cornell University.

Park, H., Mjelde, J.W. & Bessler, D. a., 2006. Price dynamics among U.S. electricity spot markets. *Energy Economics*, 28(1), pp.81–101. Available at: http://linkinghub.elsevier.com/retrieve/pii/S0140988305000903 [Accessed August 14, 2013].

Pesaran, M.H. & Timmermann, A., 2007. Selection of estimation window in the presence of breaks. *Journal of Econometrics*, 137(1), pp.134–161. Available at: http://linkinghub.elsevier.com/retrieve/pii/S0304407606000418 [Accessed August 14, 2013].

Pirrong, C. & Jermakyan, M., 2008. The price of power: The valuation of power and weather derivatives. *Journal of Banking & Finance*, 32(12), pp.2520–2529. Available at: http://linkinghub.elsevier.com/retrieve/pii/S0378426608000861 [Accessed August 14, 2013].

Rodriguez, C.P. & Anders, G.J., 2004. Energy price forecasting in the Ontario competitive power system market. *Power Systems, IEEE Transactions on*, 19(1), pp.366–374. Available at: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1266590.

Rudebusch, G.D. & Williams, J.C., 2009. Forecasting Recessions: The Puzzle of the Enduring Power of the Yield Curve. *Journal of Business & Economic Statistics*, 27(4), pp.492–503. Available at: http://www.tandfonline.com/doi/abs/10.1198/jbes.2009.07213 [Accessed August 14, 2013].

Russell, J.R., 1999. Econometric modeling of multivariate irregularly-spaced high-frequency data. *Manuscript, GSB, University of Chicago*, (November). Available at: http://faculty.chicagobooth.edu/jeffrey.russell/research/multi.pdf [Accessed August 19, 2013].

Saravia, C., 2003. Speculative Trading and Market Performance: The Effect of Arbitrageurs on Efficiency and Market Power in the New York Electricity Market. Available at: http://escholarship.org/uc/item/0mx44472#page-4.

Simonsen, I., 2005. Volatility of power markets. *Physica A: Statistical Mechanics and its Applications*, 355(1), pp.10–20. Available at: http://linkinghub.elsevier.com/retrieve/pii/S0378437105002712 [Accessed August 14, 2013].

Sweeting, A., 2007. Market Power In The England And Wales Wholesale Electricity Market 1995-2000. *The Economic Journal*, 117(520), pp.654–685. Available at: http://onlinelibrary.wiley.com/doi/10.1111/j.1468-0297.2007.02045.x/full [Accessed August 12, 2013].

Weron, R., 2009. Heavy-tails and regime-switching in electricity prices. *Mathematical Methods of Operations Research*, 69(3), pp.457–473. Available at: http://link.springer.com/10.1007/s00186-008-0247-4 [Accessed August 14, 2013].

Weron, R., 2008. Market price of risk implied by Asian-style electricity options and futures. *Energy Economics*, 30(3), pp.1098–1115. Available at: http://linkinghub.elsevier.com/retrieve/pii/S014098830700076X [Accessed August 14, 2013].

Weron, R., 2007. *Modeling and forecasting electricity loads and prices: A statistical approach*, Wiley. com.

Wolak, F., 2003. Measuring unilateral market power in wholesale electricity markets: the California market, 1998-2000. *The American economic review*, 93(2), pp.1998–2000. Available at: http://www.jstor.org/stable/10.2307/3132266 [Accessed August 12, 2013].

Wolak, F. & Patrick, R., 2001. *The impact of market rules and market structure on the price determination process in the England and Wales electricity market*, Available at: http://medcontent.metapress.com/index/A65RM03P4874243N.pdf [Accessed August 12, 2013].

Wolak, F.A., 2000. An Empirical Analysis of the Impact of Hedge Contracts on Bidding Behavior in a Competitive Electricity Market. *International Economic Journal*, 14(2), pp.1–39.

Wolfram, C.D., 1999. Measuring duopoly power in the British electricity spot market. *American Economic Review*, 89(4), pp.805–826. Available at: http://www.jstor.org/stable/10.2307/117160 [Accessed August 12, 2013].

Wolfram, C.D., 1998. Strategic Bidding in a Multiunit Auction: An Empirical Analysis of Bids to Supply Electricity in England and Wales. *The RAND Journal of Economics*, 29(4), pp.pp. 703–725. Available at: http://www.jstor.org/stable/2556090.

Wooldridge, J.M., 2002. *Econometric analysis of cross section and panel data*, The MIT press.

Zareipour, H. et al., 2006. Application of public-domain market information to forecast Ontario's wholesale electricity prices. *Power Systems, IEEE Transactions on*, 21(4), pp.1707–1717. Available at: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1717574 [Accessed August 11, 2013].

Zareipour, H., Bhattacharya, K. & Canizares, C.A., 2006. Forecasting the hourly Ontario energy price by multivariate adaptive regression splines. In *Power Engineering Society General Meeting, 2006. IEEE*. p. 7–pp. Available at: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1709474.

Zareipour, H., Caizares, C.A. & Bhattacharya, K., 2007. The operation of Ontario's competitive electricity market: overview, experiences, and lessons. *Power Systems, IEEE Transactions on*, 22(4), pp.1782–1793. Available at: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=4349150.

Appendix A

In Christensen et al. (2009) it is assumed that in any period t there can only be one new stress. The variable e_t takes the value 1 if a new stress arrives and 0 otherwise. e_t is assumed to be a sequence of independent Bernoulli random variables:

$$P(e_t = 1) = \pi \tag{2.7-1}$$

$$P(e_t = 0) = 1 - \pi \tag{2.7-2}$$

The assumption of one new stress per period is necessary to disentangle the arrival and departure of the latent process. The departure of a system stress is modeled as a binomial thinning process, where each stress factor present at time $t-1$ (X_{t-1}) survives with probability ρ. The number of surviving stress factors at time t is therefore determined by $X_t^{surv} = \rho \circ X_{t-1}$ ($0 \leq X_t^{surv} \leq X_{t-1}$), where $\circ$ is the binomial thinning operator. The number of stress factors at time t is therefore represented by

$$X_t = \rho \circ X_{t-1} + e_t \tag{2.7-3}$$

The parameters that describe the dynamics of this model can again be made time-varying using the following specifications:

$$\pi_t = 1 - \exp(-\exp(z_{1t}\beta_1)) \tag{2.7-4}$$

$$\rho_t = 1 - \exp(-\exp(z_{2t}\beta_2)) \tag{2.7-5}$$

in which z_{1t} and z_{2t} are sets of regressors for the arrival and survival of the system stresses, respectively, and β_1 and β_2 are (HAWa) the associated parameter vectors.

As illustrated by Christensen et al. (2009), the procedure of calculating the probabilities for price spikes depends on the history of the price spike process $\Psi_{t-1} = \{y_0, y_1, \dots, y_{t-1}\}$. If $y_{t-1} = 0$, then $X_t = 0$ or $X_t = 1$ (a maximum of one new arrival of system stress at any t). Therefore,

$$\begin{aligned} p_t = P(y_t = 1|\Psi_{t-1}) = P(e_t = 1) = \lambda_t \\ 1 - p_t = P(y_t = 0|\Psi_{t-1}) = P(e_t = 0) = 1 - \lambda_t \end{aligned} \tag{2.7-6}$$

The values of p_t are required in order to calculate the loglikelihood function in equation (2.7-11). The calculation of these values, however, is somewhat more convoluted if there is a history of consecutive price events such as $y_{t-1} = 1$ and potentially $y_{t-2} = 1, \dots, y_{t-k} = 1$.

If $y_{t-2} = 0$ and $y_{t-1} = 1$, then X_t can possibly take the values 0, 1 or 2. In general, if the k previous days exhibited price events ($y_{t-1} = y_{t-2} = \cdots = y_{t-k} = 1$ and $y_{t-(k+1)} = 0$), then there could be anything between a minimum of $X_{t-1} = 1$ and $X_{t-1} = k$ stresses present at time $t-1$. The calculation of $p_t = P(y_t = 1|\Psi_{t-1})$ in such a case depends on $P(X_{t-1} = m|\Psi_{t-1})$ for $m = 1, \dots, k$, the probability of m system stresses being present at time $t-1$ and $P(X_{t-1} = m|\mathrm{X}_{t-2} = n)$ based on the arrival and survival processes of the system stresses.

This is calculated as

$$P(X_{t-1} = m|\Psi_{t-1}) = \sum_{n=1}^{k-1} P(X_{t-1} = m|X_{t-2} = n)P(X_{t-2} = n|\Psi_{t-2}) \quad (2.7\text{-}7)$$

The probabilities $P(X_{t-2} = m|\Psi_{t-2})$ are readily available if $y_{t-2} = 0$ (see above). The calculation of probabilities of the type $P(X_{t-1} = m|X_{t-2} = n)$ will be detailed below. Once probabilities $P(X_{t-1} = m|\Psi_{t-1})$ are obtained we can continue to calculate

$$P(X_t = m|\Psi_{t-1}) = \sum_{n=1}^{k} P(X_t = m|X_{t-1} = n)P(X_{t-1} = n|\Psi_{t-1}) \quad (2.7\text{-}8)$$

for $m = 0, \dots, k+1$, ie, all possible number of stresses at time t. In order to obtain values of $p_t = P(y_t = 1|\Psi_{t-1})$, required for the calculation of the loglikelihood in equation (2.7-11), we need to recognise that $p_t = P(y_t = 1|\Psi_{t-1}) = 1 - P(y_t = 0|\Psi_{t-1})$ and further that $y_t = 0$ only if $X_t = 0$. Hence

$$p_t = 1 - P(X_t = 0|\Psi_{t-1}) \quad (2.7\text{-}9)$$

From the above it is apparent that the calculation of p_t (when $y_{t-1} = y_{t-2} = \cdots = y_{t-k} = 1$ and $y_{t-(k+1)} = 0$) requires a recursive calculation starting at the beginning of the most recent price event episode ($y_{t-k} = 1, y_{t-(k+1)} = 0$).

To perform the above calculations we require probabilities of the type $P(X_t = m|X_{t-1} = n)$, which, following the assumptions of the arrival and survival processes of the system stresses, are calculated as follows:

$$P(X_t = m|X_{t-1} = n)$$
$$= \begin{cases} 0, & m > n+1 \\ \rho_t^n \pi_t, & m = n+1 \\ B(n, m, \rho_t)(1-\pi_t) + B(n, m-1, \rho_t)\pi_t, & 1 \le m \le n \\ (1-\rho_t)^n (1-\pi_t), & m = 0 \end{cases} \quad (2.7\text{-}10)$$

Here π_t and ρ_t are the arrival and survival probabilities and $B(n, m, \rho_t)$ gives the probability of drawing m successes in n repeated drawings of a binomial random variable with success probability ρ_t.

Finally, the loglikelihood function is calculated as

$$\log L = \sum_{t=1}^{T} y_t \log p_t + (1 - y_t) \log(1 - p_t) \quad (2.7\text{-}11)$$

with $y_1, y_2, \ldots, y_T$ being the sequence of observed price events. The loglikelihood function is maximised with respect to the parameter vector $\beta = (\beta_1' \; \beta_2')'$. Recall that the parameter vectors β_1' and β_2' are used to parameterise the time-varying arrival (π_t) and survival (ρ_t) probabilities, respectively. Once the parameters of the modified PAR are estimated by maximising equation (2.7-11), they are used for forecasting purpose in Section 2.6.

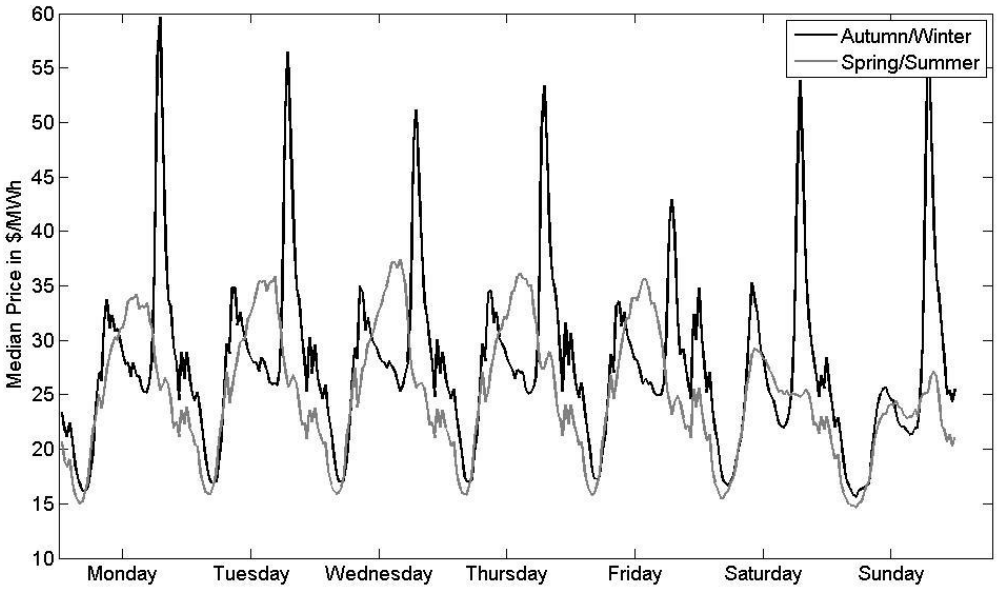

Figure 2.2-1: This figure is showing a weekly profile (Monday to Friday) of median spot prices for autumn and winter months (March 21 to September 20) and for spring and summer months (September 21 to March 20).

The weekly series are calculated based on the half-hourly median prices across the 336 half-hour periods of the week. The figure displays different season of weekly profiles for median spot prices using data between September 21 to March 20 (spring and summer) and March 21 to September 20 (autumn and winter) separately. The data is New South Wales (NSW) half-hourly price series from March 2001 to July 2010 obtained from the Australia Electricity Market Operator (AEMO) website (http://www.aemo.com.au/Electricity/Data/Price-and-Demand/Aggregated-Price-and-Demand-Data-Files).

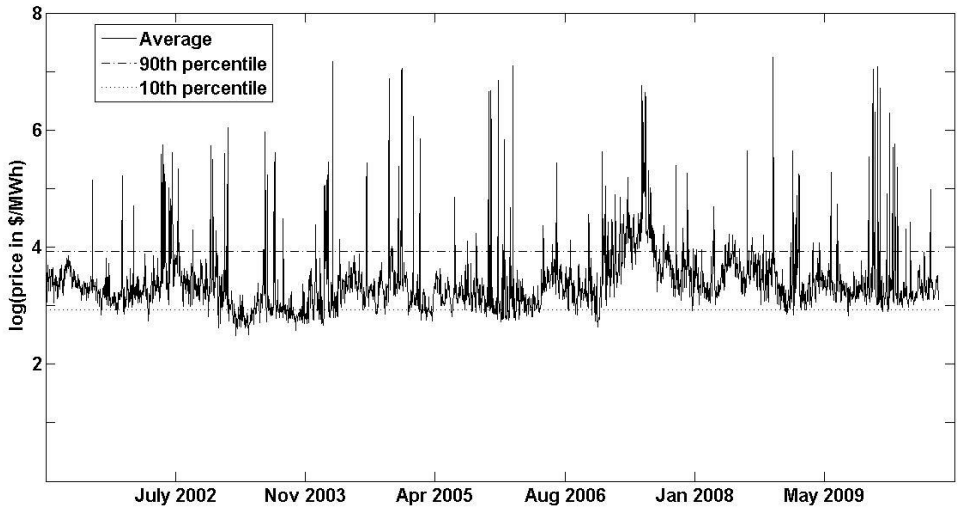

Figure 2.2-2: This figure is showing the time series of (log) daily average spot prices from 2001 to 2010 in NSW.

The daily average spot prices are calculated by taking the average of the 48 half-hour spot prices in a day. Therefore a total of 165,120 half-hourly spot prices are used to calculate the daily average spot prices from March 1, 2001 to July 31, 2010. In order to show the low and high extremity of the daily average spot price series, we calculated their 10^{th} and 90^{th} percentiles of the daily average series. The 10^{th} percentile of the series is A$18.67/MWh and the 90th percentile of the series is $50.49/MWh. The extreme temporary price jumps feature in the price series are made obvious by taking the logarithmic of the daily average spot prices. It can be seen that the logarithmic of the price series exceeded the logarithmic of the high extremity (90th percentile) values frequently.

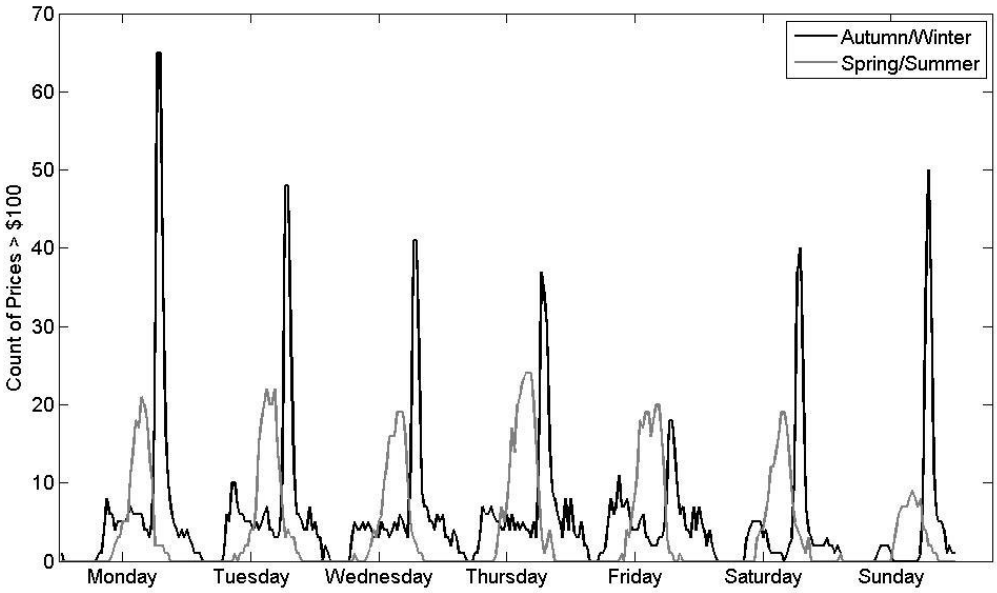

Figure 2.2-3: This figure is showing a weekly pattern for the frequency of spot prices exceeding A$100/MWh.

Using NSW wholesale spot prices, the numbers of hourly prices that exceeds A$100/MwH are counted across day of the week. They are counted for spring and summer season (September 21 to March 20) and for autumn and spring season (March 21 to September 20).

	NSW	QLD	SA	VIC
Mean	38.7906	35.7624	42.3296	34.5638
Median	27.1702	25.2533	29.3967	26.8015
Std. dev.	74.5241	63.6779	103.4392	58.9765
Skewness	11.8879	13.2374	14.9365	23.8209
Kurtosis	170.4310	230.7358	265.4799	774.3291
% spikes days	13.47%	15.25%	17.17%	12.32%

Table 2.4-1: This table shows the descriptive statistics of the daily average spot prices in four regions of the Australia NEM. % spikes days is definite for proportion of days with price spikes. The five regions are New South Wales (NSW), Queensland (QLD), South Australia (SA) and Victoria (VIC).

These descriptive statistics are using half-hourly wholesale spot prices from March 1, 2001 to July 31, 2010.

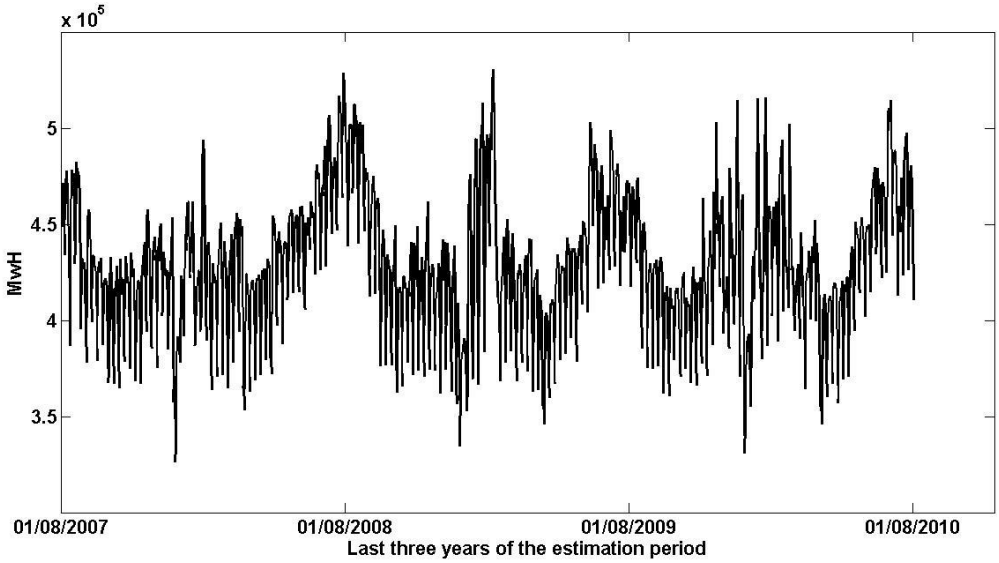

Figure 2.4-1: This figure is showing daily electricity consumption from August 1, 2007 to August 1, 2010 in NSW.

The daily load series are calculated by summing all the 48 half-hour load series in a day. For the purpose of illustration, we are using a total of 52,656 half-hourly load series obtained from AEMO's website in Aggregate Price and Demand Data Files (http://www.aemo.com.au/Electricity/Data/Price-and-Demand/Aggregated-Price-and-Demand-Data-Files). It can be observed from the figure that the daily electricity load series show a clear annual and weekly seasonal pattern.

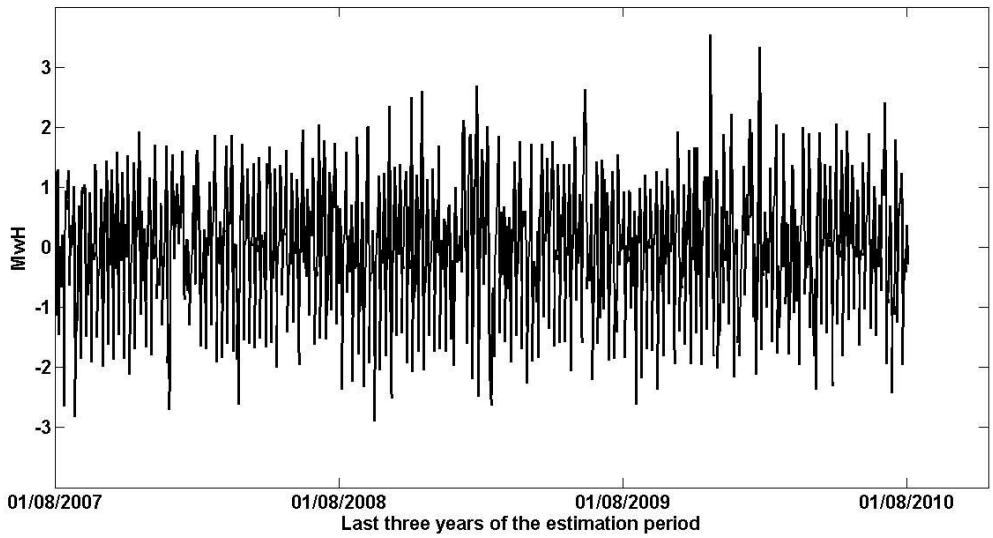

Figure 2.4-2: This figure is showing daily deseasonalized load time series, $load_t$ plotted from August 1, 2007 to August 1, 2010 in NSW.

After we have constructed normally expected load series, l_t^e using weighted average of the previous seven days of loads, $l_{t-1}, \dots, l_{t-7}$, we deducted the actual realised load series, l_t (as a proxy of forecast load available on day, $t-1$ for day, t) from the normally expected load series, l_t^e. In doing so, we are able to obtain the forecast electricity demand that exceeds the demand we would usually expect for that day of the week at that particular time of the year. However, the residual series, $\widetilde{el}_t$ is further deseasonalized using rolling volatility (Weron 2007) to eliminate any remaining annual seasonality in the series. The end-result of these processes is the series shown is the figure. They show no obvious trends or seasonality making the $load_t$ series is a plausible candidate to capture the forecast electricity demand that exceeds the demand we would usually expect for that day of the week at that particular time of the year.

	Correlation Matrix			
Exogenous	$load_t$	$Tmax_t$	$Tmin_t$	$count_t$
$load_t$	1			
$Tmax_t$	-0.0024	1		
$Tmin_t$	0.0530	-0.1416	1	
$count_t$	0.2646	0.1006	-0.0447	1
	Descriptive Analysis			
Mean	-0.0066	0.5507	0.4234	0.5465
Median	0.0649	0	0	0
Std. dev.	1.0204	1.6877	0.9767	2.1015
Skewness	-0.2206	4.6215	2.6301	5.1791
Kurtosis	3.0504	28.6803	9.6966	32.7006

Table 2.4-2: This table shows the correlations matrix and descriptive analysis of the exogenous variables as discuss in Section 2.4 for NSW. They are used as conditioning covariates in the parameters of the Hawkes and PAR models.

The information used to measure daily $load_t$ is half-hourly load series obtained from AEMO's website in Aggregate Price and Demand Data Files (http://www.aemo.com.au/Electricity/Data/Price-and-Demand/Aggregated-Price-and-Demand-Data-Files) and the series are calculated based on the processes discussed in Section 2.4.1. $Tmax_t$ is measured using daily maximum temperature series available from The Bureau of Methodology's website (http://www.bom.gov.au/climate/data/) for NSW based on a station number, 66062. The $Tmax_t$ is calculated based on the processes discussed in Section 2.4.2 in order to capture unexpected and unseasonal hot days (in the warm seasons of the year). Similar details mentioned on $Tmax_t$ apply to $Tmin_t$ except it is used to capture unexpected and unseasonal extremely cold days (in the cold season of the year). $count_t$ is measured from the 48 half-hour spot prices in a day. It represents the number of half-hourly spot prices, $s_{t,j}$, that exceed the threshold, τ (i.e. $A\$100/MwH$) on day *t*. All of the data used to measure the exogenous variables in this table start from August 1, 2007 to July 31, 2010.

	Correlation Matrix			
Exogenous	$load_t$	$Tmax_t$	$Tmin_t$	$count_t$
$load_t$	1			
$Tmax_t$	0.0705	1		
$Tmin_t$	0.0227	-0.1399	1	
$count_t$	0.1769	0.1043	-0.0381	1
	Descriptive Analysis			
Mean	-0.0040	0.3495	0.4799	0.6542
Median	0.0621	0	0	0
Std. dev.	1.0108	0.9289	1.2026	2.2309
Skewness	-0.2707	4.5988	3.0739	4.6777
Kurtosis	3.1356	34.0664	13.5522	28.1624

Table 2.4-3: This table shows the correlations matrix and descriptive analysis of the exogenous variables as discuss in Section 2.4 for QLD. They are used as conditioning covariates in the parameters of the Hawkes and PAR models.

Further explanation on this table is similar to the caption of Table 2.4-2, except the daily maximum and minimum temperature series for QLD is based on a station number, 40842.

	Correlation Matrix			
Exogenous	$load_t$	$Tmax_t$	$Tmin_t$	$count_t$
$load_t$	1			
$Tmax_t$	0.1815	1		
$Tmin_t$	-0.0037	-0.1981	1	
$count_t$	0.3103	0.4447	-0.0596	1
	Descriptive Analysis			
Mean	0.0327	1.2289	0.5630	0.7354
Median	0.0608	0	0	0
Std. dev.	1.1194	2.9840	1.1715	2.8655
Skewness	-0.1484	2.7087	2.2973	6.2571
Kurtosis	3.3962	9.8944	8.2197	55.5421

Table 2.4-4**:** This table shows the correlations matrix and descriptive analysis of the exogenous variables as discuss in Section 2.4 for SA. They are used as conditioning covariates in the parameters of the Hawkes and PAR models.

Further explanation on this table is similar to the caption of Table 2.4-2, except the daily maximum and minimum temperature series for SA is based on a station number, 23090.

	Correlation Matrix			
Exogenous	$load_t$	$Tmax_t$	$Tmin_t$	$count_t$
$load_t$	1			
$Tmax_t$	0.1805	1		
$Tmin_t$	0.0067	-0.1881	1	
$count_t$	0.2775	0.3253	-0.0039	1
	Descriptive Analysis			
Mean	-0.0058	1.3050	0.5497	0.5338
Median	0.0641	0	0	0
Std. dev.	1.0582	3.1260	1.1829	2.1902
Skewness	-0.1912	2.7678	2.4432	7.3709
Kurtosis	3.1774	10.5875	9.0029	81.1905

Table 2.4-5**:** This table shows the correlations matrix and descriptive analysis of the exogenous variables as discuss in Section 2.4 for VIC. They are used as conditioning covariates in the parameters of the Hawkes and PAR models.

Further explanation on this table is similar to the caption of Table 2.4-2, except the daily maximum and minimum temperature series for VIC is based on a station number, 86071.

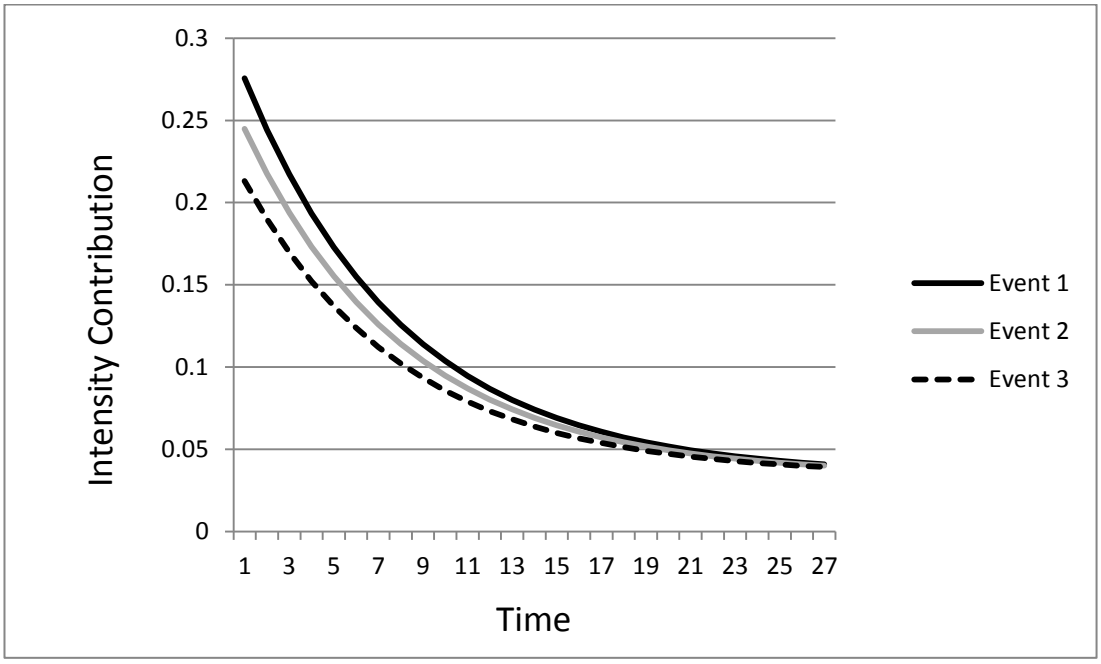

Figure 2.5-1: This figure is used to show the mechanism of α_{t_i} as initial intensity jumps.

Following a price event happened at time 0 (t_i), using the Hawkes model in equation (2.3-4) we are able to calculate future intensity at time 1 (t). The contribution of a price event at time 0 to future intensity starts at time 1 and continue at an increasing lags until time 27 (refer to the x-axis in the figure). If α_{t_i} is conditioned only by $constant$ exogenous variable, this means price event at time 0 carries the same effect on future intensity. However, Event 1, Event 2 and Event 3 are used to illustrate how the contribution of a price event at time 0 changes based on the parameter of α_{t_i} that are now determined by a set of covariates at time t_i. The set of covariates at time t_i follow equation (2.3-5) and (2.3-6) with the conditioning covariates for HAWa model are as specified in Table 2.5-1. Since different event results to different values of α_{t_i}, the initial impact on the future intensity from Event 1 to Event 3 changes from 0.27 to 0.23. It can be observed from the figure that larger values of α_{t_i}(eg. Event 1) gives larger contributions to the probability of future price events since the intensity contribution of Event 1 are consistently superior from Event 2 and Event 3 from time 0 until time 27. For the purpose of this illustration, decay parameter β is set constant ($\beta = 0.4$).

Variable	NSW	QLD	SA	VIC
μ_t^				
α_{t_i}				
γ				
$constant$	-0.2718* (0.0023)	-0.7501* (0.0029)	-1.3597* (0.0084)	-9.9990* (0.0029)
$load_{t_i}$	-0.0989* (0.0093)	-0.3093* (0.0008)	0.2988 (0.2203)	5.4033* (0.6753)
$Tmax_{t_i}$	-0.5683* (0.0088)	-0.6174* (0.0037)	0.0485* (0.0135)	-1.2018* (0.0007)
$Tmin_{t_i}$	-0.1735* (0.0135)	-0.0285* (0.0033)	0.2648* (0.1416)	0.2076 (0.7011)
$count_{t_i}$	0.1362 (0.1092)	0.3970* (0.1552)	0.1847 (0.1493)	3.0635* (0.3691)
β				
δ				
$constant$	0.0003* (0.0001)	0.0002* (0.0001)	0.0003* (0.0001)	0.0003* (0.0001)
log $L^{Sea,tva}$	-387.8466	-457.7136	-411.2108	-377.3038

Table 2.5-1: This table presents the coefficients and robust standard errors (in brackets) of parameter vector θ ((γ' δ')′) used on the conditioning covariates in the initial intensity jumps, α_{t_i} and time-invariant intensity decay, β parameter.

These are the results of the HAWa model for all the regions (NSW, QLD, SA and VIC) during the estimation period (August 1, 2007 to July 31, 2010). ^The coefficients and standard errors of parameter vector of the trigonometric term ζ (allowing for seasonal variations, μ_t) are not reported but are available upon request. All of these parameter vectors are estimated by maximising the loglikelihood function in equation (2.3-9). In the bottom row, we report the optimised loglikelihood. The covariates conditioned in α_{t_i} are an intercept ($constant$), unseasonal load series ($load_{t_i}$) as discussed in Section 2.4.1, unexpected increase/decrease in temperatures ($Tmax_{t_i}$,$Tmin_{t_i}$) as discussed in Section 2.4.2 and the number of price spikes during a day ($count_{t_i}$). While the β remains constant. *Parameters are significant at 10% significant level.

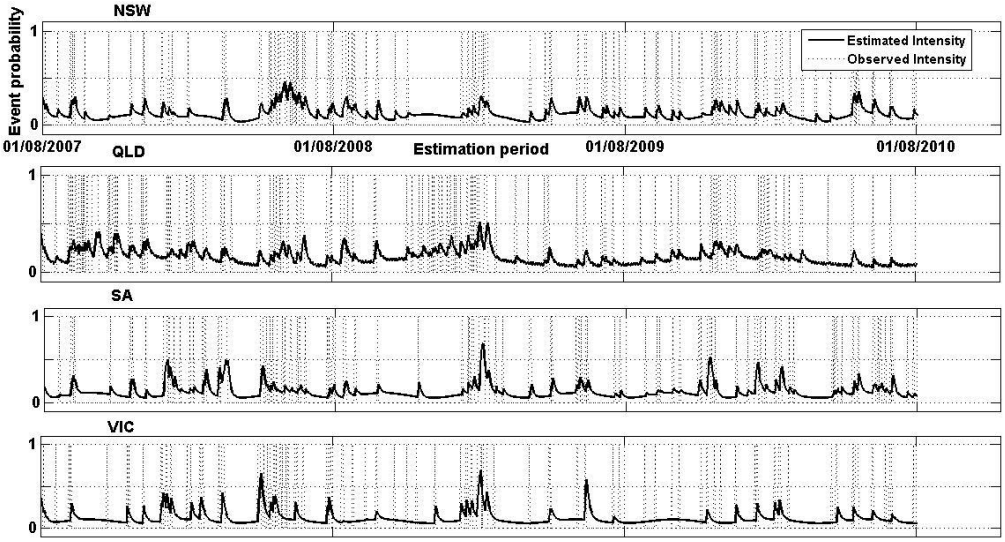

Figure 2.5-2: This figure illustrates the in-sample of one-day-ahead conditional intensity forecast, $\lambda_t^{Sea,tva}$ during the estimation period (August 1, 2007 to July 31, 2010) for each of the regions.

The dependent variables of Hawkes model, y_t is set to 1 if any of the half-hourly spot prices, $s_{t,j}$ exceeds $A\$100/MwH$ on day t and 0 otherwise. y_t is labelled as 'Observed Intensity' in the legend. While, the probability of future price events, $\lambda_t^{Sea,tva}$ is labelled as 'Estimated Intensity' in the legend. Therefore, the y-axis represents event probability. The Estimated Intensity is measured based on equation (2.5-1) where the x-axis represents days with a price event.

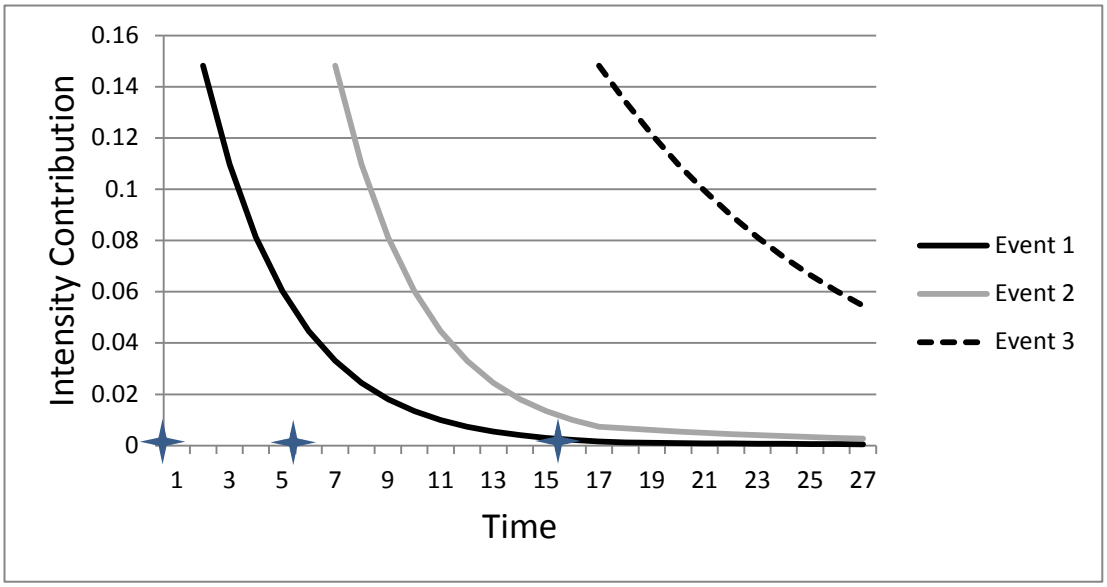

Figure 2.5-3: This figure is used to show the mechanism of time-varying decay parameter, β_τ which are used to introduce the influence of covariates in any period t (since $\tau = t_i, \dots, t$).

For the purpose of this illustration, the initial intensity contribution α_{t_i} is set constant ($\alpha_{t_i} = 0.2$). Therefore each of the events carries the same initial intensity jumps. In contrast to HAWa model, in HAWab model the influence of covariates are allow to play a role outside the times, t_i, at which price events occur. Therefore the probability of future price events under this model, $\lambda_t^{sea,tvab}$ follows equation (2.5-2) and the time-varying decay parameter β_t is parameterised using equation (2.3-7). The added flexibility allows for increased or delayed intensity decay due to prior events. In this figure, price event happened at time 0, time 5 and time 15 (labelled by ✦). In the figure, the contribution of past event (at time 0) to the intensity of Event 1 has a decay parameter β_τ that takes a value 0.3 up to $t = 16$ and 0.1 afterwards. Therefore the contributions to the probability of future price events for Event 1 decays at a smaller rate after $t = 16$. The same situations happened for Event 2, the contribution of past event (at time 5) to the intensity of Event 2 has a decay parameter β_τ that takes a value 0.3 up to $t = 16$ and 0.1 afterwards. Hence it become obvious that the contributions to the probability of future price events for Event 3 decays at slowest rate compared to Event 1 and Event 2. In summary, this means with high value of decay parameter β_τ, the intensity will decay at faster rate compared to lower value of decay parameter β_τ.

Variable	NSW	QLD	SA	VIC
μ_t^				
α_{t_i}				
γ				
$constant$	0.6162* (0.2582)	0.1813 (0.1952)	-0.5811* (0.0128)	-0.9535* (0.0127)
$load_{t_i}$	-0.3941* (0.0013)	-0.5552* (0.0056)	-0.1046* (0.0061)	-0.2272* (0.0030)
$Tmax_{t_i}$	-0.7187* (0.0072)	-0.0657* (0.0056)	0.1028* (0.0518)	-0.3996* (0.0182)
$Tmin_{t_i}$	-0.2495* (0.0020)	-0.0500* (0.0143)	0.2911* (0.1401)	0.1666 (0.1527)
$count_{t_i}$	0.0795* (0.0004)	0.0840* (0.0037)	0.1166 (0.0796)	0.5765* (0.0740)
β_τ				
δ				
$constant$	0.0096* (0.0007)	0.0062 (0.0042)	0.0052 (0.0119)	0.0023 (0.0015)
$load_\tau$	-0.0139* (0.0047)	-0.0010 (0.0110)	-0.0216* (0.0014)	-0.0066 (0.0061)
$Tmax_\tau$	0.0187* (0.0009)	0.0084* (0.0018)	-0.0102* (0.0035)	-0.0018 (0.0075)
$Tmin_\tau$	0.0107* (0.0032)	0.0041 (0.0090)	0.0080 (0.0215)	0.0060 (0.0075)
$count_\tau$	-0.1122* (0.0020)	-0.0600* (0.0022)	-0.0850* (0.0095)	-0.0263* (0.0054)
$Weekdays_\tau$	1.5983 (2.4725)	0.4147 (0.5131)	1.3669 (1.4165)	1.7989 (1.4864)
log $L^{Sea,tvab}$	-374.2704	-443.2520	-395.1150	-359.6421

Table 2.5-2: This table presents the coefficients and robust standard errors (in brackets) of parameter vector θ $((\gamma' \delta')')$ used on the conditioning covariates in the initial intensity jumps, α_{t_i} and time-varying decay parameter, β_τ.

These are the results of the HAWab model for all the regions (NSW, QLD, SA and VIC) during the estimation period (August 1, 2007 to July 31, 2010). ^The coefficients and standard errors of parameter vector of the trigonometric term ζ (allowing for seasonal variations, μ_t) are not reported but are available upon request. All of these parameter vectors are estimated by maximising the loglikelihood function in equation (2.3-9). In the bottom row, we report the optimised loglikelihood. The covariates conditioned in α_{t_i} and β_τ are as defined in Table 2.5-1 but with the addition of a weekend / no-weekend dummy variable, $Weekdays_\tau$ in the decay parameter, β_τ.
*Parameters are significant at 10% significant level.

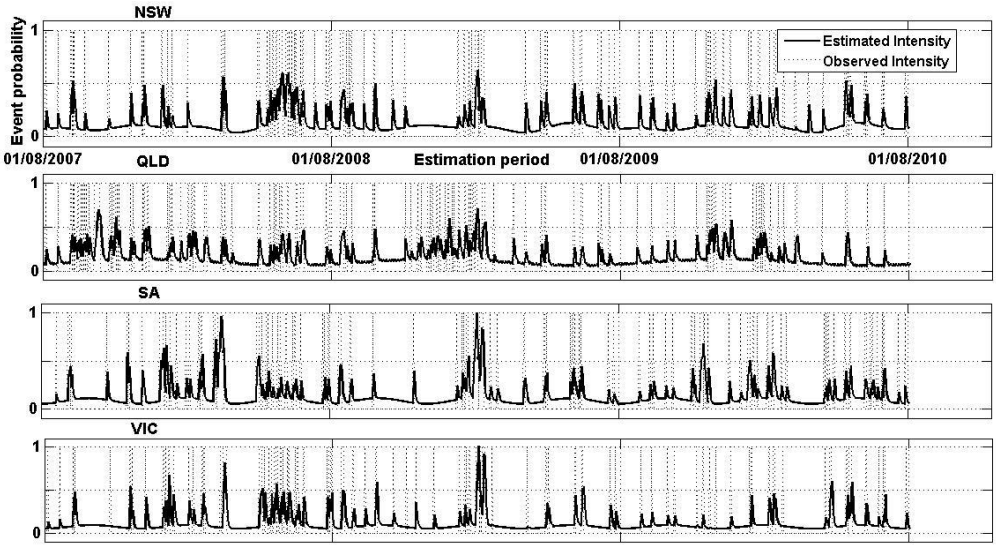

Figure 2.5-4: This figure illustrates the in-sample of one-day-ahead conditional intensity forecast, $\lambda_t^{Sea,tvab}$ during the estimation period (August 1, 2007 to July 31, 2010) for each of the regions.

Further explanation on the y and x-axis are the same with the caption of Figure 2.5-2, except that the estimated intensity is measured based on equation (2.5-2).

Variable	NSW	QLD	SA	VIC
π_t				
β_1				
$constant$	-2.6416*	-2.3216*	-2.9685*	-3.0580*
	(0.1571)	(0.1296)	(0.1854)	(0.1771)
$load_t$	0.7355*	0.4992*	0.7742*	0.7185*
	(0.1683)	(0.1354)	(0.1500)	(0.1753)
$Tmax_t$	-0.0165	0.1619*	0.1341*	0.1368*
	(0.0751)	(0.0601)	(0.0386)	(0.0275)
$Tmin_t$	0.0002	-0.0120	0.1418*	0.1659*
	(0.1365)	(0.1017)	(0.0829)	(0.0884)
ρ_t				
β_2				
$constant$	-1.4804*	-1.7787*	-1.4428*	-1.0978*
	(0.5939)	(0.7052)	(0.5552)	(0.5938)
$load_t$	0.3988	-0.1148	0.2908	-0.0848
	(0.3007)	(0.2779)	(0.2181)	(0.3643)
$Tmax_t$	-0.1200	-0.1399	0.0444	-0.1358
	(0.0955)	(0.1713)	(0.0638)	(0.1763)
$Tmin_t$	-0.0950	0.0944	0.0814	0.1040
	(0.1934)	(0.1007)	(0.1762)	(0.1411)
$Weekdays_t$	0.3862	0.8405	-0.1592	-0.1573
	(0.6733)	(0.6735)	(0.5166)	(0.6475)
log L^{PAR}	-355.0726	-425.4564	-346.7250	-332.9346

Table 2.5-3**:** This table presents the coefficients and robust standard errors (in brackets) of parameter vector, β_1 and β_2 used on the conditioning covariates relating to the variation in arrival probability of system stresses, π_t and in survival probability of system stresses, ρ_t.

These are the results of the PAR model for all the regions (NSW, QLD, SA and VIC) during the estimation period (August 1, 2007 to July 31, 2010). All of these parameter vectors are estimated by maximising the loglikelihood function in equation (2.7-11). In the bottom row, we report the optimised loglikelihood. The covariates conditioned in π_t are as defined for α_{t_i} and the covariates conditioned in ρ_t are as defined for β_τ but without $count_t$ exogenous variables. *Parameters are significant at 10% significant level.

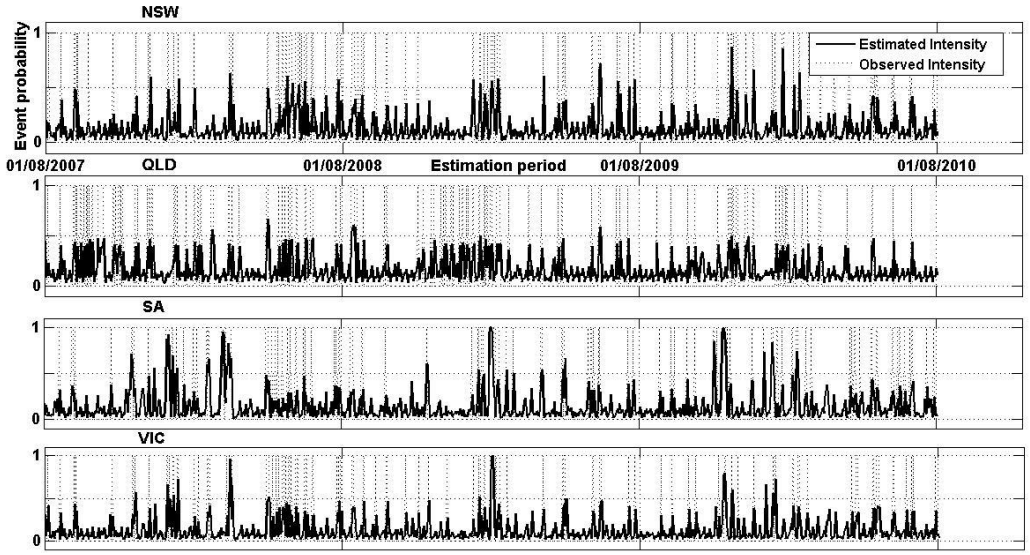

Figure 2.5-5: This figure illustrates the in-sample of one-day-ahead conditional intensity forecast, p_t, during the estimation period (August 1, 2007 to July 31, 2010) for each of the regions.

Further explanation on the y and x-axis are the same with the caption of Figure 2.5-2, except that the estimated intensity is measured based on equation (2.7-10) in Appendix A

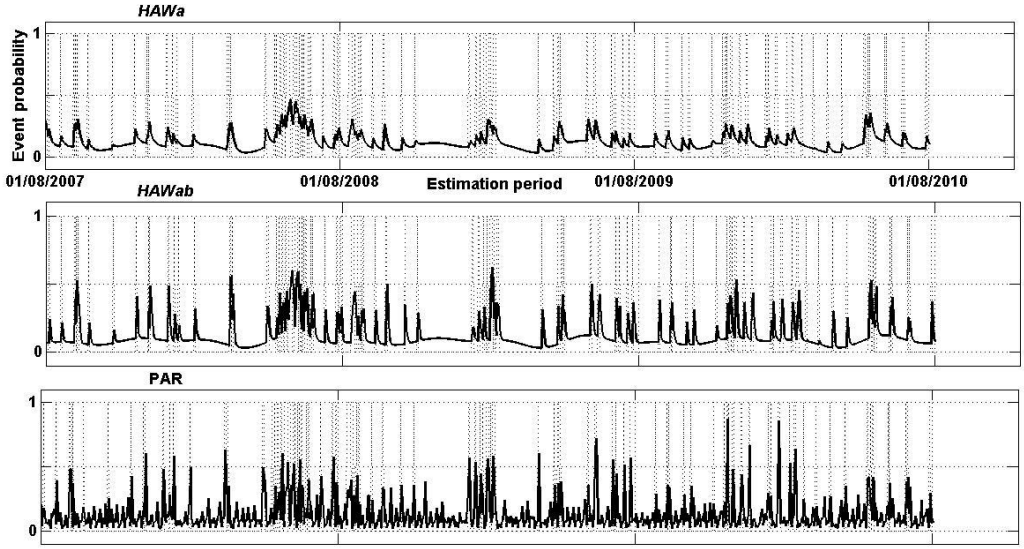

Figure 2.5-6: This figure illustrates the in-sample of one-day-ahead conditional intensity forecast of all the models, $\lambda_t^{Sea,tva}$, $\lambda_t^{Sea,tvab}$ and p_t during the estimation period (August 1, 2007 to July 31, 2010) for NSW. It is apparent from the graph that different models allow for quite different features in the intensity series.

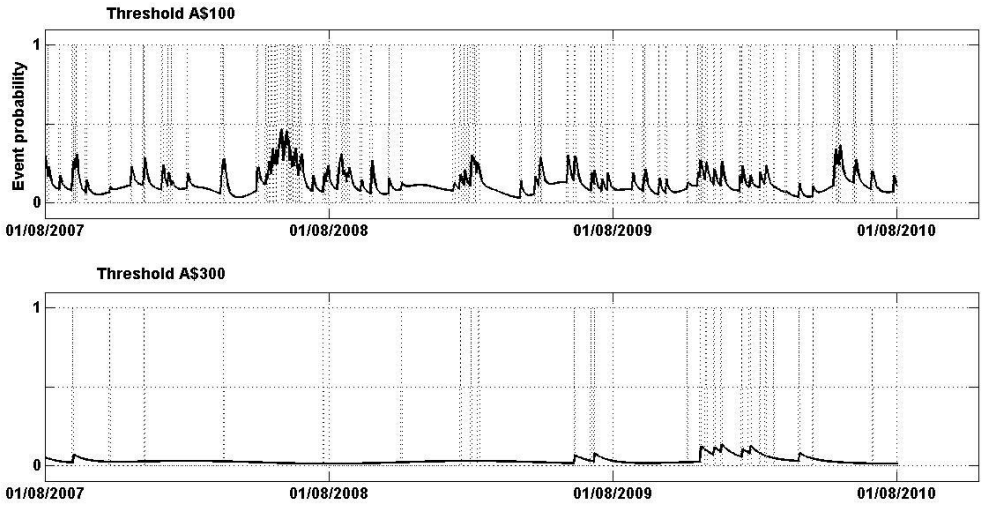

Figure 2.5-7**:** This figure illustrates the in-sample of one-day-ahead conditional intensity forecast, $\lambda_t^{Sea,tva}$ during the estimation period (August 1, 2007 to July 31, 2010) for NSW and based on different threshold value, τ.

In the first panel, the dependent variables, y_t is set to 1 if any of the half-hourly spot prices, $s_{t,j}$ exceeds $A\$100/MwH$ on day t and 0 otherwise while in the second panel, the dependent variables, y_t is set to 1 if any of the half-hourly spot prices, $s_{t,j}$ exceeds $A\$300/MwH$ on day t and 0 otherwise. It can be observed the days with a price event reduces tremendously when the threshold value increases from $A\$100/MwH$ to $A\$300/MwH$.

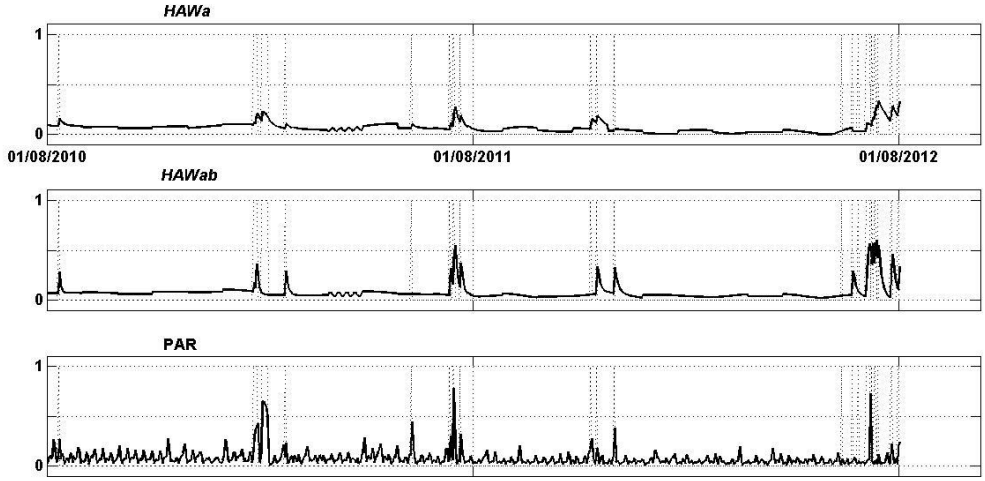

Figure 2.6-1: This figure illustrates the out-of-sample of one-day-ahead conditional intensity forecast for all the models ($\lambda_t^{Sea,tva}$, $\lambda_t^{Sea,tvab}$, p_t) during the forecast period (August 1, 2010 to July 31, 2012) for NSW.

All the parameters used to forecast future price events for these models are re-estimate every 30 days. For example, the parameters vectors of all the models are estimated by maximising the loglikelihood function in equation (2.3-9) and (2.7-11) using information during the initial estimation period (August 1, 2007 to July 31, 2010). These parameters vectors are used to produce one-day-ahead conditional intensity forecast from August 1, 2010 to August 30, 2010 using equation (2.5-1), (2.5-2) and equation (2.7-10) in Appendix A. The parameters are re-estimate every 30 days, meaning that the second estimation window includes data from August 30, 2007 to August 29, 2010. These updated parameter estimates are then used to produce the next 30 one-day-ahead conditional intensity forecast from August 31, 2010 to September 29, 2010. This process continues until we obtain two years' worth of intensity forecasts.

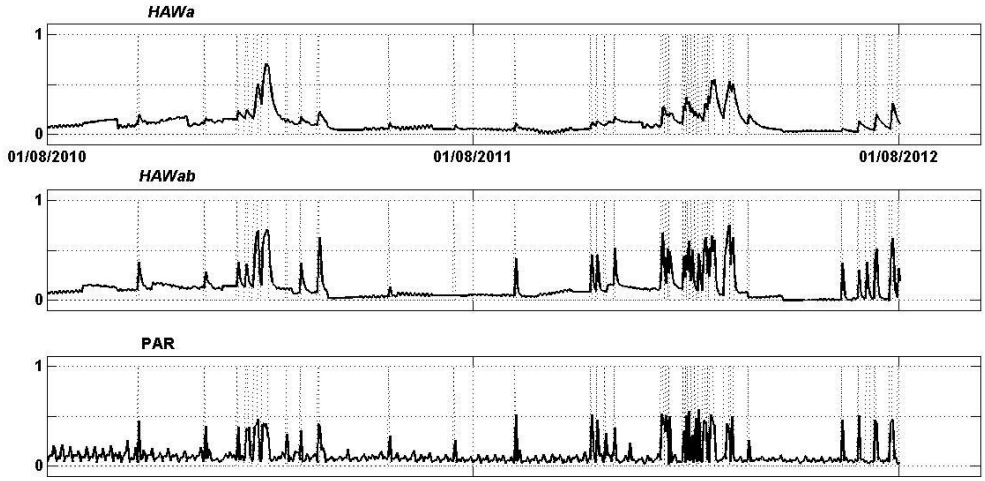

Figure 2.6-2: This figure illustrates the out-of-sample of one-day-ahead conditional intensity forecast for all the models ($\lambda_t^{Sea,tva}$, $\lambda_t^{Sea,tvab}$, p_t) during the forecast period (August 1, 2010 to July 31, 2012) for QLD.

Further explanation on this figure is similar to the caption of Figure 2.6-1

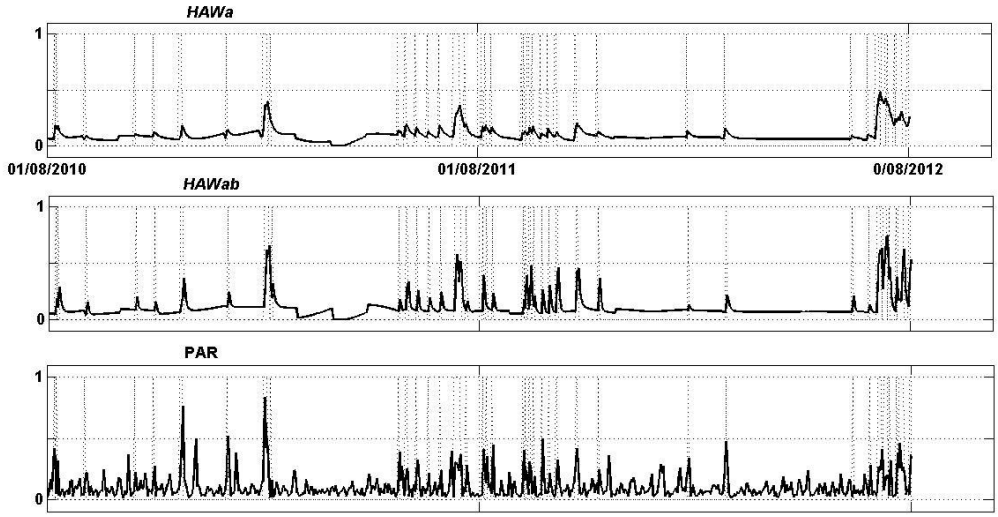

Figure 2.6-3**:** This figure illustrates the out-of -sample of one-day-ahead conditional intensity forecast for all the models ($\lambda_t^{Sea,tva}$, $\lambda_t^{Sea,tvab}$, p_t) during the forecast period (August 1, 2010 to July 31, 2012) for SA.

Further explanation on this figure is similar to the caption of Figure 2.6-1

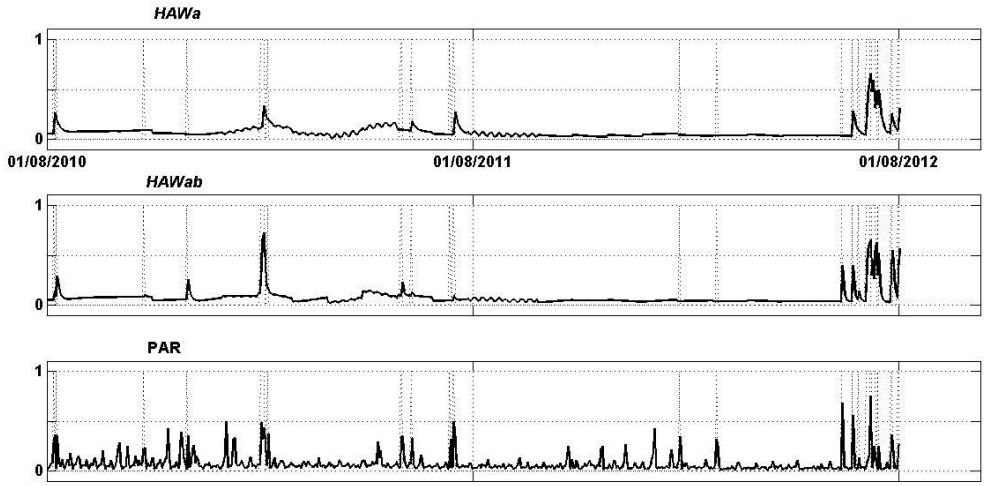

Figure 2.6-4: This figure illustrates the out-of -sample of one-day-ahead conditional intensity forecast for all the models ($\lambda_t^{Sea,tva}$, $\lambda_t^{Sea,tvab}$, p_t) during the forecast period (August 1, 2010 to July 31, 2012) for VIC.

Further explanation on this figure is similar to the caption of Figure 2.6-1

Measure	Model			
	Naive	HAWa	HAWab	PAR
		NSW		
MAE	**0.0506**	0.1007	0.1077	0.1061
RMSE	0.2250	0.2062	0.2081	**0.2048**
Asym	**0.0540**	0.0929	0.0945	0.0926
		QLD		
MAE	**0.0917**	0.1656	0.1590	0.1613
RMSE	0.3027	0.2748	**0.2689**	0.2802
Asym	**0.1019**	0.1572	0.1481	0.1573
		SA		
MAE	**0.0985**	0.1517	0.1529	0.1539
RMSE	0.3138	0.2662	**0.2624**	0.2707
Asym	**0.1081**	0.1488	0.1444	0.1478
		VIC		
MAE	**0.0479**	0.1010	0.0988	0.1024
RMSE	0.2188	0.1969	**0.1916**	0.2091
Asym	**0.0527**	0.0883	0.0851	0.0906

Table 2.6-1**:** This table presents the forecast evaluation statistics (mean absolute error, MAE; root mean square error, RMSE; asymmetric loss score, Asym) for all the regions.

Based Rudebusch & Williams (2009) and Christensen et al. (2009) using forecast error ($y_t - \lambda_{t|t-1}$) of Hawkes models as an example;

$$MAE = \frac{1}{t_1 - t_0 + 1}\sum_{t=t_0}^{t_1}\left|y_t - \lambda_{t|t-1}\right|$$

$$RMSE = \sqrt{\frac{1}{t_1 - t_0 + 1}\sum_{t=t_0}^{t_1}\left(y_t - \lambda_{t|t-1}\right)^2}$$

$$Asym = \frac{1}{t_1 - t_0 + 1}\sum_{t=t_0}^{t_1}\left(y_t(1+\kappa) + (1-y_t)(1-\kappa)\right)\left|y_t - \lambda_{t|t-1}\right|$$

where t_0 and t_1 denote the beginning and the end of the forecast period and κ is equal to 0.5 since the failure to predict an actual price event is penalized by three times the rate of predicting a price event that does not actually occurred.

This table reports the average value of the respective loss function evaluated for Models Naïve, HAWa, HAWab and PAR for the forecast period from August 1, 2010 to August 1, 2012. The bold entry indicates which model produces the lowest average loss.

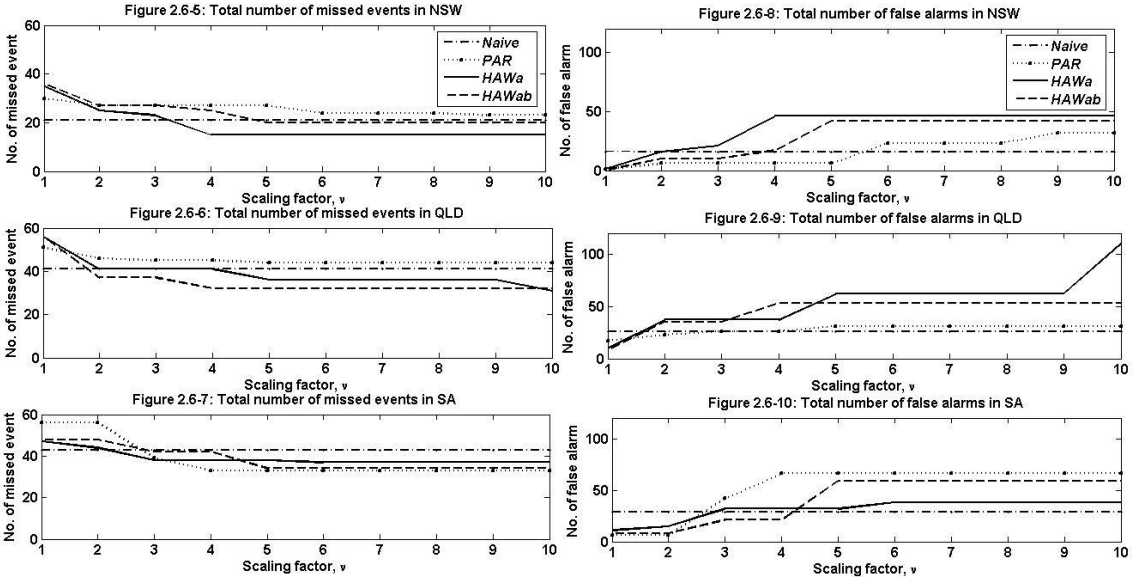

Figure 2.6-5 to Figure 2.6-10: This figure illustrates the total number of missed events, $Tot^{missed\ event}$ and false alarms, $Tot^{false\ alarm}$ in NSW, QLD and SA.

Both the total number of missed events and false alarms are based on the lowest signal threshold that produces lowest forecast failure. Therefore for $\nu = 1$ this implies that we find the trigger intensity (for each model) that minimizes the total number of missed events and false alarms for any model. The scaling factor, ν on the x-axis is used to reflect the relative cost of missed events as opposed to false alarms. For example, when $\nu = 5$ this means the missed events are five times more costly than false alarms. The number of missed events and false alarms depends on the trigger intensity and therefore varies with the value of the trigger intensity and hence with ν. Generally, larger ν will lead to a lower trigger intensity and hence to more false alarms and fewer missed events.

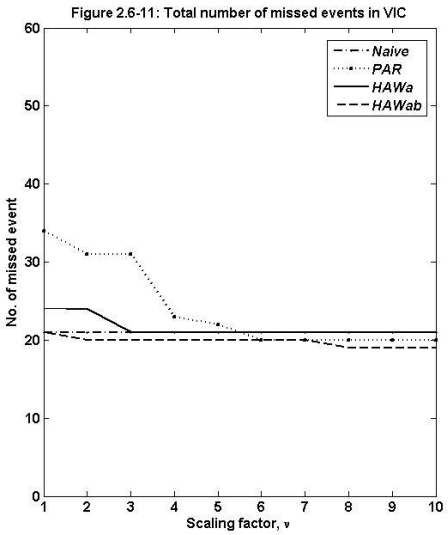

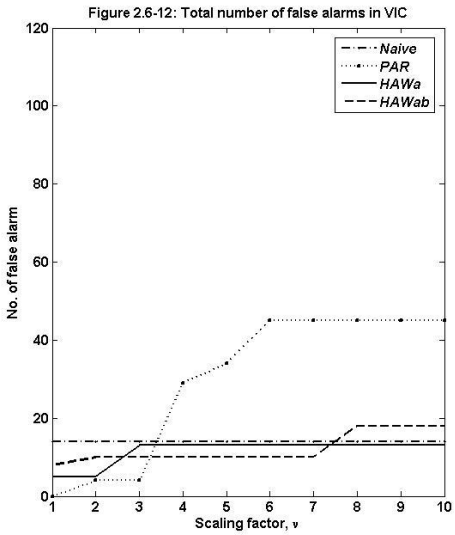

Figure 2.6-11 to Figure 2.6-12: This figure illustrates the total number of missed events, $Tot^{missed\ event}$ and false alarms, $Tot^{false\ alarm}$ in VIC.

Further explanation on this figure is similar to the caption of Figure 2.6-5 to Figure 2.6-10.

Chapter 3

3. The Australian Electricity Market's pre-dispatch process: Some observations on its efficiency

ABSTRACT

In the past doubts have been raised as to whether the pre-dispatch process in Australia Electricity Market is able to give market participants and market operator good and timely quantity and price information. It is the purpose of this chapter to introduce a framework to analyse whether the pre-dispatch process is delivering biased predictions of the actual wholesale spot price outcomes. Here we investigate the bias by comparing the actual wholesale market spot price outcome to pre-dispatch sensitivity prices established the day before dispatch and on the day of dispatch. We observe a significant bias (mainly indicating that the pre-dispatch process tends to underestimate spot price outcomes) and we further establish the seasonality features of the bias across seasons and/or trading periods. We also establish changes in bias across the years in our sample period (1999 to 2007). In the formal setting of an ordered probit model we establish that there are some exogenous variables that are able to explain increased probabilities of over- or under-predictions of the spot price. It transpires that meteorological data, expected pre-dispatch prices and information on past over- and under-predictions contribute significantly to explaining variation in the probabilities for over- and under-predictions. The results allow us to conjecture that some of the bids and re-bids provided by electricity generators are not made in good faith.

3.1. Introduction

The Australian National Electricity Market (NEM) is an electricity pool into which all generators submit bids to supply energy and the Australian Energy Market Operator (AEMO)[14] eventually attempts to meet the demand using the cheapest possible combination of generators. All electricity that is supplied is compensated at the price which was required for the marginal generator.

In order to give market participants (suppliers and electricity customers) as well as the market operator good and timely quantity and price information (Australian Energy Market Operator 2010), AEMO performs a pre-dispatch process. By 12.30pm on the day before dispatch, electricity generators are obliged to submit supply schedules for all 48 hours of the following trading day. These supply schedules are then aggregated and matched with demand forecasts to produce expected wholesale electricity prices for the next day, also known as pre-dispatch prices for the next day. In order to evaluate how sensitive this predicted price is with respect to changes in demand, AEMO repeats this exercise for different demand quantities (demand forecast +/- 200 MWh[15]) and reports the resulting predicted price outcomes, $P^{+\Delta}$ and $P^{-\Delta}$, also known as pre-dispatch sensitivity prices for the next day.

Generators can change their bids (re-bidding) on the quantity but not the price band of their supply schedules up until shortly before dispatch. Any bid or rebid, however, should reflect the generators "genuine intention to honour" (Australian Energy Market Commission, 2013, paragraph 3.8.22A.b) their bid. In order to reflect any such changes, the pre-dispatch matching process of supply bids and demand forecasts is repeated every 30 minutes up until the actual dispatch period.

The availability of these scenarios allows us to place the actual price outcome on the dispatch day relative to these pre-dispatch sensitivity prices. For a sensibly working pre-dispatch system one would expect the price outcomes to mainly fall between the $P^{+\Delta}$ and $P^{-\Delta}$ prices and occasions on which the actual price exceeds $P^{+\Delta}$ not too significantly outnumber the occasions on which the price outcome is below $P^{-\Delta}$.

[14] The Australian Energy Market Operator (AEMO) was established in July 2009 to manage NEM and gas market and undertake the electricity functions previously carried out by National Electricity Market Management Company (NEMMCO)

[15] In fact they will produce more scenarios (+/- 500 and 1000 MWh) but our analysis will be restricted to +/- 200 MWh scenarios.

If the pre-dispatch price was available[16] alternative measures of bias could also be calculated; such as the average deviation between the pre-dispatch price and the spot price. It should be noted, however, that, due to the convexity of the aggregate supply curve we would expect such a measure to be positively skewed. Consequently we would likely to expect the average deviation to be positive. One advantage of using measures of the spot price against the pre-dispatch sensitivity price is that the latter account for this supply curve convexity.

To illustrate this we shall use a simulate example. At the core of this example is a convex aggregate supply curve (Figure 3.1-1, Panel B). We assume that the load forecast is 1000, resulting in a pre-dispatch price, P^{0MwH} and pre-dispatch scenario prices of $P^{-200MwH}$ and $P^{+200MwH}$. We now simulate realised demands to investigate the properties of the resulting price distribution (Panel C) and the distribution of the discrete price outcomes (Panel D).

The results will, of course, crucially depend on the properties of the demand distribution. Here we assume that the demand outcomes are symmetrically distributed around our demand forecast of 1000 (Panel A). Based on a review of the large demand forecasting literature (see Weron, 2007 for an overview), we could not find any evidence to either support or refute this assumption. This literature, almost exclusively, discusses measures like the mean absolute percentage error (MAPE) or root mean squared forecast error (RMSFE), and do not discuss any measures of symmetry. We could identify only one paper (Nowicka-Zagrajek & Weron 2002) in which the actual distribution of demand forecast errors was shown. This distribution (referring to a fairly small sample of California electricity demand forecasts) did not have any obvious asymmetric properties.

(INSERT Figure 3.1-1 HERE)

Given the convex nature of the aggregate supply curve (in panel B, top middle), the distribution of spot prices (in panel C, top right), are heavily skewed to the right, even though we assumed a symmetric demand distribution. If we are, however, merely counting the number of occasions on which the actual price, P^{Spot} exceeds $P^{+200MwH}$ or is smaller than $P^{-200MwH}$ then we obtain (despite the convexity in the supply curve) a symmetric, discrete price outcome distribution (Panel D).

[16]The pre-dispatch prices are available for very recent periods from AEMO website (http://www.aemo.com.au/Electricity/Data), but were not included in the historical dataset that we obtained from AEMO

In the past doubts have been raised as to whether this pre-dispatch process meets its purpose and indeed changes have and are being proposed as there is doubt whether there aren't electricity suppliers which possess and use excessive market power (see a recent proposed rule change to identify generators with market power, Australian Energy Market Commission, 2012).

In this context the objectives of this chapter are threefold: *First*, we will introduce a framework in which we propose to analyse whether there is an apparent bias in the pre-dispatch process in the sense that it systematically under- or overestimates the actual electricity price outcomes. *Second*, we will analyse whether the answer to the first question varies systematically across years, seasons and/or trading periods. *Third*, we will establish whether any exogenous variables are relevant to explain variation in the probability of spot price outcomes to fall above $P^{+\Delta}$ or below $P^{-\Delta}$.

We find that there is a significant bias in the pre-dispatch process in the sense that the pre-dispatch process tends to underestimate the actual price outcomes (i.e. dispatch price > $P^{+\Delta}$). While there is variation in the extent of the bias the general finding is remarkably persistent. We do find that this bias is particularly strong in the morning and evening peak periods of the day. We also find that the probability with which the dispatch price exceeds $P^{+\Delta}$ varies significantly with a number of explanatory variables and indeed that there is a fair degree of persistence.

The analysis in this chapter is essentially a descriptive analysis and it is impossible to use these results to definitively make judgements on the existence of market power of individual generators. The findings in this chapter, however, do allow the conjecture that some of the bids and re-bids provided by generators are not made in good faith as defined by the National Electricity Rules (Australian Energy Market Commission, 2013, paragraph 3.8.22A).

The remainder of the chapter is as follows. In Section 3.2 we will give a more detailed analysis of the institutional settings. In Section 3.2.1 we discuss pre-dispatch rules in NEM and their effect on the bidding behaviour together with the empirical properties of a functioning pre-dispatch process. In Section 3.2.2 we will give detailed descriptive analysis of spot price outcomes using unconditional probabilities on each datasets. We also illustrate the bias we observed in the price outcomes. In order to analyse whether any

exogenous variables are relevant to explain variation in the probability of spot price outcomes (relative to pre-dispatch sensitivities), we use an ordered probit model. This and the explanatory variables used are presented in Section 3.3 and 3.4. A summary of the results and their discussion are given in Section 3.5 and 3.6. We conclude the chapter in Section 3.7.

3.2. The Bidding and Re-Bidding Process in the NEM

In this Section we will provide some more detail on the NEM pre-dispatch process. The electricity pool setup with day-ahead bids and the opportunity for rebids is similar to the setup of the Canadian and New Zealand electricity market[17]. In the context of this chapter the specifics of the Australian NEM are of importance and we will detail these by following an imaginary timeline from 12.30pm on the day prior to dispatch (day $t-1$) to the various dispatch periods on the day on which electricity is actually dispatched (day t). The electricity trading day starts at 4.01am in the morning and ends at 4am the next day. It is separated into 48 half hour periods. Wholesale electricity prices are set every five minutes, by balancing supply and demand, and setting the price of the marginal generator as the price to be paid for all electricity supplied in that period. Prices are then averaged across the 6 five minute period of every half-hour to obtain the wholesale price for every half-hour period.

The pre-dispatch process can be described as below:

- By 12.30pm on the day before dispatch of electricity ($t-1$), generators need to submit their daily bids for all 48 half-hourly trading intervals during the dispatch day (t). Each generator supplies 10 price bands and indicates how much electricity they are willing to supply for every 48 half hour periods at each price band. AEMO will then aggregate all supply bids and evaluate how the forecast demand (at all 48 half hour periods) can be met at the lowest possible price[18]. This will deliver the pre-dispatch price and by varying the demand forecast by +/- 200 MWh it will also deliver the $P^{+\Delta}$ and $P^{-\Delta}$ price sensitivities.

[17] In Ontario, a Region in Canada, the pre-dispatch schedules are published at 11.00am while in New Zealand the schedules are produced before and after 1.00pm a day before dispatch. The market participants in both markets are allowed to rebids until 2 hours ahead of dispatch.

[18] In this process AEMO will have to take a number of restrictions into account. For example, some generators cannot be switched on and off in quick succession. Details are available in (Gillett & Market Operations 2010).

- After 12.30pm on day $t-1$ and up to 3.55am on day t, the generators can change their bid information. These are called rebids. The rebids, just as the initial bids, have to be submitted in "good faith" which is interpreted as the understanding that the bids reflect the "genuine intention to honour that offer [...] if the material conditions and circumstances [...] remain unchanged until the relevant dispatch interval" (Australian Energy Market Commission, 2013, paragraph 3.8.22A). All rebids have to be accompanied by a justification that outlines the change in circumstances that led to the changed bid. Every half hour AEMO will repeat the pre-dispatch process and publish the resulting price sensitivities.

- From 3.55am onwards on the dispatch day the generators can continue to submit rebids as long as they are submitted a minimum of five minutes before the beginning of the 30 minutes dispatch period for which changes in the supply schedule are submitted. AEMO will continue to undertake the pre-dispatch process and will publish the resulting $P^{+\Delta}$ and $P^{-\Delta}$ price sensitivities.

- As the actual dispatch interval arrives AEMO will run an optimisation process very similar to the pre-dispatch process that identifies the cheapest combination of electricity generators to meet the prevalent electricity demand (load), taking into account a large range of generator and transmission constraints.

It is useful to illustrate the optimisation process that happens during pre-dispatch and at dispatch with a graph. In Figure 3.2-1 we show a stylised aggregate supply schedule, the associated demand forecasts and the resulting pre-dispatch prices. It is important to understand that this information needs to be indexed with the relevant dispatch period and the time at which the pre-dispatch was published. In this figure, for instance, we display the supply schedule for the 8.00am dispatch period on day t as available from the pre-dispatch published on day $t-1$ at 15.00pm, $Sup_{t,08:00|t-1,15:00}$. We can also see the equivalent demand forecast, $D_{t,08:00|t-1,15:00}$, and the resulting pre-dispatch price, $P_{t,08:00|t-1,15:00}$. Further we can see the equivalent low and high demand scenarios and the resulting $P^{+\Delta}$ and $P^{-\Delta}$ prices.

(INSERT Figure 3.2-1 HERE)

To measure the performance of the pre-dispatch process we will use this information as follows. First we restrict our attention to two pre-dispatch periods, the day before dispatch ($t-1$, 15:00, *PD1*) and the pre-dispatch on the dispatch day (t) at 6.00am (*PD2*). We will then compare price outcomes for the half hour dispatch periods starting with the 06:30 and ending with the 23:30 period to the respective pre-dispatch prices for these periods.

This is done by comparing the actual wholesale market spot price outcome for the τth sub-period on day t, $P_{t,\tau}$, with the respective pre-dispatch price sensitivity prices from *PD1* and *PD2*. The following ordered categorical variables are created:

$$S_{t,\tau|j} = \begin{cases} 2 & if\ P_{t,\tau} > P_{t,\tau|j}^{+\Delta} \\ 1 & if\ P_{t,\tau|j}^{+\Delta} \geq P_{t,\tau} > P_{t,\tau|j}^{-\Delta} \\ 0 & if\ P_{t,\tau|j}^{-\Delta} \geq P_{t,\tau} \end{cases} \qquad (3.2\text{-}1)$$

$where\ j = PD1\ or\ PD2$ and $\tau = 06{:}30\ to\ 23{:}30$. The empirical analysis in this chapter will be based on $S_{t,\tau|j}$.

3.2.1. Pre-Dispatch Rules and Bidding Behaviour

The value of the pre-dispatch and associated price sensitivities to market participants depends on the quality of AEMO's demand forecasts and on generator's bids being reflective of their actual supply intentions. If the rebids, which are possible up until five minutes before the relevant dispatch period, were used strategically, then the value of pre-dispatch information would be doubtful.

Not long after its formation in 1998 the first doubts about the workings of the pre-dispatch system were raised in a Australia National Electricity Code Administrator (NECA)[19] issue paper (National Electricity Code Administrator 2001b; National Electricity Code Administrator 2001a) that analysed generator's rebidding and identified significant variations between pre-dispatch and dispatch prices. They were conjectured to be the result of rebidding activities. In particular the authors highlight that many rebids happen close to dispatch when there is no opportunity for an effective competitive or demand-side response. A variety of recommendations to change the re-bidding process were made, of which only two were picked up in rule changes in 2002. As a result the "good faith" rule referred to earlier (Australian Energy Market Commission, 2013, paragraph 3.8.22A.a) was

[19] As of 1 July 2005 all roles and functions of NECA have been taken over by the Australian Energy Market Commission (AEMC) and the Australian Energy Regulator (AER).

introduced. Further it is upon the generator that submits a rebid, to provide "a brief, verifiable and specific reason for the rebid" and to "substantiate and verify" the reasons if so requested by the Australian Energy Regulator (AER) (Australian Energy Market Commission, 2013, paragraph 3.8.22).

From this it follows that, as part of our empirical analysis we will investigate whether there are any clear changes in the characteristics of our variable $S_{t,\tau|j}$. Little research has been done in this field, although Bardak Ventures Pty Ltd (2005) note that the rule changes have not been successful, and fail to detect significant differences in bidding behaviour.

Ideally with demand forecasts and supply bids being updated continuously and promptly the pre-dispatch price should be the best possible forecast of the actual dispatch price, at any given time. Borrowing ideas from the efficient market hypothesis literature this should leads us to the proposition that information available at pre-dispatch should not be useful in predicting the outcome of our categorical random variable $S_{t,\tau|j}$.

If, however, there were generators that could exercise market power, in combination with a not strictly enforced or enforceable "good faith" bidding condition, then it may well be possible that we may find systematic variation in the conditional distributions of $S_{t,\tau|j}$. In particular generators that posess market power may then be tempted to use initial supply bids as merely an "opening gambit".

There is a significant literature that investigates the issue of market power in electricity markets. Studies of market power in Australia by Wolak (2000), Pennsylvania–New Jersey– Maryland by Mansur (2007), New England by Bushnell & Saravia (2002), England and Wales by Wolak & Patrick (2001) and Sweeting (2007), New York by Saravia (2003), Spain by Fabra & Toro (2005), California by Knittel & Metaxoglou (2008) and Texas by Hortacsu & Puller (2005) confirm the existence of market power in some restructured electricity markets. David & Wen (2000) note that market power can be created or strengthened by collusion involving frequently repeated auctions for electricity under similar demand and supply conditions and intimate knowledge of a rival's operating costs and almost immediate knowledge of a rival's actions.

Generators could conjecture from pre-dispatch information how their rivals adjust their supply schedules, and this could help generators plan when it is optimal to be marginal

generators and bid accordingly. When a generator has the opportunity to bid at prices higher than their marginal costs, by taking advantage of a poor market design, then they are able to exercise market power. This is called strategic bidding. The literature on strategic bidding is vast stating the existence of rich bidding strategies in restructured electricity markets such as England and Wales (Wolfram 1999; Wolfram 1998), California (Borenstein et al. 2002; Wolak 2003), Pennsylvania–New Jersey– Maryland (Mansur 2001; Oh 2003), and Texas (Hortacsu & Puller 2005).

Previous findings by Rodriguez & Anders (2004); Hamidreza Zareipour et al. (2006a); H Zareipour et al. (2006) and Kharbach et al. (2010) prove the importance of considering pre-dispatch prices as additional information in wholesale market price forecast models. The aim here is slightly different, as we consider the usefulness of the pre-dispatch prices as the forecast itself. To the best of our knowledge, this is the first attempt to study the performance of pre-dispatch scenario prices as estimates of future wholesale market spot prices[20] in Australia's NEM. The results from the analysis of spot price outcomes, $S_{t,\tau|PD1}$ and $S_{t,\tau|PD2}$ will not allow us to directly infer anything about the bidding behaviour of individual generators. However, we will be able to state whether the combination of the collective bidding behaviour with the bidding rules in the NEM make for a useful pre-dispatch procedure.

While it is not expected that the pre-dispatch (at *PD1* or indeed *PD2*) will provide an exact spot price forecast, we would expect a functioning pre-dispatch process to have the following empirical properties:

1. We should expect to see a roughly symmetric distribution for the ordered categorical variables $S_{t,\tau|PD1}$ and $S_{t,\tau|PD2}$.
2. We should expect the distribution of $S_{t,\tau|PD2}$ to be more centred around $S_{t,\tau|PD2} = 1$ than $S_{t,\tau|PD1}$. Both the demand forecast as well as the supply schedules at *PD2* should reflect more precise information for the dispatch period t, τ compared to *PD1*.
3. If all supply schedules reflect all available information and are made in "good faith" we should expect that any information available at *PD1* or *PD2* is irrelevant for explaining the outcome of $S_{t,\tau|PD1}$ and $S_{t,\tau|PD2}$ respectively.

[20] Electricity Market Performance (2012a) discusses on the factors affecting the ability of pre-dispatch as forecast of dispatch outcomes

In the remainder of the chapter we will be able to establish to what extend the pre-dispatch process in Australia is able to match these expectations.

3.2.2. Descriptive analysis of spot price outcomes

We have data for the wholesale spot prices, $P_{t,\tau}$, and the pre-dispatch prices $\left(P^{-\Delta}_{t,\tau|j}, P^{+\Delta}_{t,\tau|j}\right)$ from 1 March 1999 to 31 October 2007[21]. These are used, as in equation (3.2-1) to calculate our new variable $S_{t,\tau|j}$.

We begin our empirical analysis by reporting the empirical distributions for our variable $S_{t,\tau|j}$ for *j = PD1*, *PD2* and $\tau = 06{:}30\ to\ 23{:}30$. These results are reported in Figure 3.2-2 in which the results are broken down into 6 subpanels. In Panel A (top left) we can see the empirical probabilities of $S_{t,\tau|j} = 0$, 1 or 2 as calculated across all days of the year 2000. On the horizontal axis we have the 35 half hour periods from $\tau = 06{:}30\ to\ 23{:}30$.

(INSERT Figure 3.2-2 HERE)

We can see that for the first (06:30) dispatch period, the probability of $S_{t,\tau|PD1} = 2$ exceeds 60%. This means that in more than 60% of days in 2000 the actual spot price outcome exceeded the upper price sensitivity $P^{+\Delta}_{t,06:30|PD1}$ at *PD1*. The probability that the spot price is lower than $P^{-\Delta}_{t,06:30|PD1}$ is less than 10%. In the remaining approximate 25% of days the actual spot price fell between $P^{-\Delta}_{t,06:30|PD1}$ and $P^{+\Delta}_{t,06:30|PD1}$. This indicates that the pre-dispatch process at 15:00 the day before pre-dispatch (*PD1*) on average significantly underestimates the wholesale price for the 06:30 subperiod. From Panel A it is apparent that this bias is obvious for the morning and the evening sub-periods, whereas the pre-dispatch process does not display such a clear bias for the trading periods in the middle of the day. However, even during these mid-day periods $\mathrm{P}\left(S_{t,06:30|PD1} = 2\right)$ is always about 10% points larger than $\mathrm{P}\left(S_{t,06:30|PD1} = 0\right)$. Panels B to F show the equivalent results for a range of different sub-samples in 2000, weekdays (Panel B), weekends (C), Winter days (D), Summer days (E) and Spring and Autumn days (F). By large the results remain unchanged. There are, however, a number of minor variations (for the year 2000).

For instance the bias of the pre-dispatch system towards underestimating the actual spot price outcomes appears stronger on weekends as compared to weekdays, in particular in the mid-day dispatch periods. The same bias is also stronger during winter as compared to

[21] Details on the data sources can be found in Appendix C

other seasons (but particularly summer days) in the morning and mid-day dispatch periods. In fact the summer mid-day periods are the only ones in which the bias is reversed (i.e. the pre-dispatch, *PD1*, seems to overestimate the spot price outcomes, although this bias is rather modest).

Next, we plot the equivalent unconditional probabilities for $S_{t,\tau|PD2}$. As discussed previously we should expect the distribution of $S_{t,\tau|PD2}$ to be more centered around $S_{t,\tau|PD2} = 1$ as the *PD2* uses more up to date (relative to the actual dispatch time) information than *PD1*. Further, this information advantage is larger for the early morning dispatch periods. These results again presented for the year 2000 only and with the same subsampling in Figure 3.2-3.

(INSERT Figure 3.2-3 HERE)

In general it is apparent that the pre-dispatch process at 06:00 on the dispatch day (*PD2*) produces less biased results than *PD1*. The proportion of outcomes of $S_{t,\tau|PD2} = 1$ has increased significantly and now frequently represents the most likely outcome. In general, however, there is still some bias with $\mathrm{P}(S_{t,06:30|PD1} = 2)$ exceeding $\mathrm{P}(S_{t,06:30|PD1} = 0)$ in the morning and evening dispatch periods (see Panel A). In terms of the differences between the different sub-samples nothing much has changed with respect to the outcomes discussed for *PD1*. The relative frequencies of $S_{t,\tau|PD2} = 2$ and $S_{t,\tau|PD2} = 0$ remain basically unchanged and therefore, even at *PD2* we can still see the tendency of the pre-dispatch process to underestimate the price outcome, especially in the morning and evening sub-periods.

There are two more aspects of these results that are worth highlighting. The improvement in terms of stronger centering (higher probability of $S_{t,\tau|PD2} = 1$) is clearest in the morning periods. This seems to imply that the information gained between *PD1* and *PD2* is most significant for the morning periods. In a way this is not surprising as the time between 06:00 (*PD2*) and the actual morning dispatch is very limited and one would expect not much more new information to become available in that time. However, it is noticeable that despite *PD2* being very close in time to the morning dispatch there is still a noticeable bias with price outcomes still being significantly more likely to exceed the higher scenario price than being smaller than the lower scenario price. This seems to suggest that even between

PD2 and the morning dispatch periods there appears to be significant and price relevant re-bidding activity[22].

Next we will investigate whether the above features of the distributions of $S_{t,\tau|j}$ change significantly throughout our sample period from 1999 to 2007. To facilitate this investigation we decided to split the investigated trading interval into morning (06:30 to 09:30 dispatch intervals, $\tau \in m$), daytime (10:00 to 17:30, $\tau \in d$) and evening (18:00 to 23:30, $\tau \in e$) sub-periods. When averaging across the years we now average across all trading intervals in, say, the morning sub-period, rather than only a particular half-hour dispatch period as we did for the previous results. These sub-periods were created by subsuming those sub-periods that appear to have similar characteristics (as judged from the previous two Figure 3.2-2 and Figure 3.2-3).

Figure 3.2-4 presents two characteristics of the distributions of $S_{t,\tau|j}$ for $\tau = m, d, e$ and $j = PD1, PD2$ as they change through the years (horizontal axis). In the left column we show $\mathrm{P}(S_{t,\tau|j} = 1)$ as an indication of how many of the spot price outcomes land between the two pre-dispatch scenario prices. The second feature is $\mathrm{P}(S_{t,\tau|j} = 2) - \mathrm{P}(S_{t,\tau|j} = 0)$ giving an indication of the bias. The larger this value the more does the pre-dispatch process underestimate the actual spot price outcome. The three rows in Figure 3.2-4 represent the morning, daytime and evening sub-periods respectively.

(INSERT Figure 3.2-4 HERE)

As reported for the year 2000, we can see that in all years and in all sub-periods *PD2* produces more price outcomes between the two scenario prices, indicating that *PD2* is indeed better than *PD1*. The improvement is clearly largest for the morning sub-period and barely noticeable for the evening sub-period. The results are very stable across the years although a downward trend for $\mathrm{P}(S_{t,\tau|j} = 1)$ is noticeable.

The biasedness indicator, $\mathrm{P}(S_{t,\tau|j} = 2) - \mathrm{P}(S_{t,\tau|j} = 0)$, takes almost exclusively positive values, indicating that the previously discussed tendency to underestimate the spot price is a persistent feature in our dataset. The only exception to this is the year 2001 during which there is no clear bias for the daytime and evening sub-periods. This measure, as discussed in Figure 3.2-2 and Figure 3.2-3, is smallest for the daytime period. Again, the

[22] Re-bidding can happen until up to 5 minutes before the start of the dispatch activity.

bias does reduce from *PD1* to *PD2*, mostly so for the morning period and hardly for the evening sub-period. However, even at *PD2* the bias in the morning period is still substantial and larger than 20% points (but for 2001). Importantly, and corresponding to the decline in $\mathrm{P}(S_{t,\tau|j} = 1)$, this bias appears to increase with time, indicating that the pre-dispatch process is getting worse towards the end of our sample period. Although not presented here, it should be noted that the general thrust of these results remain unchanged when investigating the (weekend/weekday and seasons sub-samples).

It was the reports by National Electricity Code Administrator (2001a); Australia Competition & Consumer Conmission (2002) and Australian Competition & Consumer Commission (2002) that highlighted the issue of late re-bidding activity meaning that there would be no opportunity for an effective competitive or demand-side response and, so suspected the report's authors, contributed to very short-term price spikes events. None of the substantial proposed changes (e.g. no re-bidding in the last three hours prior to dispatch) were actually implemented, although the previously cited sections that spell out the meaning of "good faith" were clarified. It should be noted that the NECA report used data prior to 2001.

It is interesting to note that the one year in which there was least evidence of a pre-dispatch bias was indeed the year 2001 in which the more wide ranging changes were being discussed. The implementation of the clarifications in 2001 did not result in an "improved" pre-dispatch process. In contrast, and as discussed previously, a clear trend to more severe under-prediction of the spot price is discernible.

3.3. Methodology

3.3.1. Ordered probit model for spot price outcomes

There is a small literature by Rodriguez & Anders (2004); H Zareipour et al. (2006); Hamidreza Zareipour et al. (2006a) and Kharbach et al. (2010) that uses pre-dispatch information as an explanatory variable in models to forecasts spot price outcomes in Ontario, a market in which there are similarities (but also important differences) in the pre-dispatch process. In this literature it transpires that pre-dispatch price information is indeed informative with respect to spot price outcome and that it becomes more informative for more recent pre-dispatch information (Hamidreza Zareipour et al. 2006a).

This is, of course, not surprising given the advanced amount of supply and demand information that is used in the pre-dispatch process. Interestingly, however, there is also evidence that the pre-dispatch prices in Ontario (at least for the early sample period) tends to overestimate the actual spot market outcome, a result that is very different to our finding of under-prediction of the pre-dispatch process in New South Wales. One obvious explanation is that the Ontario electricity market incorporates a financial compensation mechanism that covers profit shortfalls of suppliers (when compared to the predicted profits from pre-dispatch) as discussed in Zareipour et al. (2007).

In contrast to the above literature we do not aim to provide a forecasting model for spot price outcomes. Rather it is our aim to understand the circumstances in which the pre-dispatch process tends to produce pre-dispatch information that tends to underestimate the market outcomes. To this end we will estimate an ordered probit model (Wooldridge 2002) for the ordered categorical variables, $S_{t,\tau|j}$ used in our previous empirical analysis. We will establishe whether there are any (weakly) exogenous variables that are useful in predicting the outcomes of $S_{t,\tau|j}$.

The time-varying ordered probit model for the τth sub-period on day t is based on a latent variable $S^*_{t,\tau|j}$ which is related to a $(k \times 1)$ vector of explanatory variables $x_{t,\tau|j}$ via the following linear relationship

$$S^*_{t,\tau|j} = \mu_{\tau|j} + \gamma_{\tau|j}{}' x_{t,\tau|j} + \varepsilon_{t,\tau|j}$$

where $\mu_{\tau|j}$ is a scalar constant, $\gamma_{\tau|j}$ is a $(k \times 1)$ coefficient vector and $\varepsilon_{t,\tau|j}$ is assumed to be normally distributed. For ease of notation we subsume the constant into $\boldsymbol{x}_{t,\tau|j} = (1\ x_{t,\tau|j}{}')'$ and define $\boldsymbol{\beta}_{\tau|j} = (\mu_{\tau|j}\ \gamma_{\tau|j}{}')'$ and hence

$$S^*_{t,\tau|j} = \boldsymbol{\beta}_{\tau|j}{}' \boldsymbol{x}_{t,\tau|j} + \varepsilon_{t,\tau|j} \qquad (3.3\text{-}1)$$

This model is estimated for any combination of dispatch period ($\tau \in m, d, e$) and pre-dispatch period ($j = PD1, PD2$). The latent index $S^*_{t,\tau|j}$ and $S_{t,\tau|j}$ in turn are related via

$$S_{t,\tau|j} = \begin{cases} 2 & if\ S^*_{t,\tau|j} > \rho_{1\tau|j} \\ 1 & if\ \rho_1 \geq S^*_{t,\tau|j} > \rho_{0\tau|j} \\ 0 & if\ \rho_0 \geq S^*_{t,\tau|j} \end{cases} \qquad (3.3\text{-}2)$$

with ρ_0 and ρ_1 being threshold parameters, which are collected in $\boldsymbol{\rho}_{\tau|j} = (\rho_0\ \rho_1)'$. To complete the model it is necessary to specify the distribution of the error term $\varepsilon_{t,\tau|j}$. In the

context of the ordered probit model this is defined to follow a standard normal ($N(0,1)$) distribution. This is a standard modelling technique for ordered categorical variables. The specification allows the calculation of conditional category probabilities $P(S_{t,\tau|j} = i|\boldsymbol{x}_{t,\tau|j})$, which in turn allow us to specify a log-likelihood function (see Section 3.5). The model parameters $(\boldsymbol{\beta}_{\tau|j}, \boldsymbol{\rho}_{\tau|j})$ that maximise that log-likelihood function are the Maximum Likelihood (ML) parameter estimates.

3.4. Explanatory variables

The explanatory variables considered for inclusion in $x_{t,\tau|j}$ fall into two categories. Those variables that remain the same for any sub-period τ or pre-dispatch period j. Amongst these are seasonal variables and weather characteristics for day t. The second set of variables will differ with τ and j. As argued before, it is the aim of these models to explain why the pre-dispatch process frequently delivers under-estimated spot prices and to identify the market conditions under which this is more likely. Therefore, some of the variables will relate to information only available on the dispatch day t. They will therefore not be available at the time of PD1 (15:00 on day $t-1$) or PD2 (06:00 on day t) and consequently the estimated models should not be understood to be forecast models. We will now introduce the explanatory variables.

3.4.1. Seasonal variables

In Section 3.2.2, we discussed that $\Pr(S_{t,\tau|j} = 1)$ and $\Pr(S_{t,\tau|j} = 2) - \Pr(S_{t,\tau|j} = 0)$ may well depend on the weekday and the annual seasons. Hence, it is appropriate to include $D_t^{wk}, D_t^{win}, D_t^{sum}$ in the model. The base case is therefore a weekday during either spring or autumn.

3.4.2. Load Error and Forecast

As previously discussed, the pre-dispatch process is the result of combining the most up to date demand and supply information at *PD1* and *PD2*. If one wants to build a forecasting model for wholesale spot electricity prices, and the pre-dispatch process was delivering unbiased spot price forecasts, one would not expect any additional demand (load) information to be relevant in our ordered probit model[23].

[23] Also see Hamidreza Zareipour et al. (2006a) who establish that the inclusion of pre-dispatch information renders many other variables irrelevant.

Here, however, we start from the presumption (supported by our results in Section 3.2.2) that the pre-dispatch process is not delivering such unbiased outcomes and it is our desire to characterise when this process is more likely to deliver under-predicted spot prices. We therefore include two load related variables. First, we include a simple load forecast available using $t-1$ information, l_t^e. Ideally we would use the load forecast used in the pre-dispatch process, but these data are not available in our dataset. We therefore build an, admittedly simplistic, but useful load forecast along the lines of Weron's ((Weron 2007), Section 3.4.1 similar day method, this was also used in Chapter 2). We are attempting to establish whether high load days are more likely to deliver biased pre-dispatch outcomes.

The second load related variable included is the difference between the actual load on day t, l_t, and the above load prediction, l_t^e. This series, $\widetilde{el_t}$, is de-seasonalised using the rolling volatility method described in ((Weron 2007), Section 2.4.5). Using this variable will help to identify the days on which the actual load outcomes were unexpectedly high or low. We would expect that on days with unexpectedly high load we can see an increased probability for $S_{t,\tau|j} = 2$.

3.4.3. Temperature Data

It is well known that electricity load is closely related to ambient temperatures, mainly through heating and air-conditioning activities. Given the availability of fairly accurate weather forecasts, it is expected that the load forecasts used in the pre-dispatch process will be reflective of that information. However, given the existing bias of the pre-dispatch process it is interesting to establish whether the probability to underestimate the spot price does depend on any meteorological information. We use the same $Tmax_t$ and $Tmin_t$ variables as in by Chapter 2. $Tmax_t$ is defined as follows:

$$Tmax_t = \begin{cases} |T^{obs}_{max_t} - T^{s}_{max_t}| & if\ T^{obs}_{max_{t_i}} > T^{s}_{max_t}\ and\ T^{s}_{max_t} > \bar{T}_{max} \\ 0 & otherwise \end{cases} \quad (3.4\text{-}1)$$

where $T^{obs}_{max_t}$ is the realised maximum temperature series on day t

$T^{s}_{max_t}$ is the seasonal pattern of maximum temperature series, modelled by a trigonometric function

$\bar{T}_{max}$ is the average of the $T^{obs}_{max_t}$series

Hence $Tmax_t$ is non-zero only on unseasonably hot days in summer while $Tmin_t$ is defined similarly and only captures unseasonably cold days in winter.

3.4.4. Spot price outcome persistency

It is our conjecture that the outcomes for $S_{t,\tau|j}$ may largely be driven by generator's re-bidding behaviour. A paper on modelling spike occurrences by Eichler et al. (2012) finds that price spikes in the NEM are often followed by further spikes. This suggests that some aspects of generator's bidding behaviour are persistent. It seems therefore reasonable to establish whether the outcomes of $S_{t,\tau|j}$ also display persistent characteristics. We therefore include the following explanatory variables

$$D^i_{t,\tau-1|j} = \begin{cases} 1 & if\ S_{t,\tau-1|j} = i \\ 0 & otherwise \end{cases} \quad (3.4\text{-}2)$$

$$D^i_{t-1,\tau|j} = \begin{cases} 1 & if\ S_{t-1,\tau|j} = i \\ 0 & otherwise \end{cases} \quad (3.4\text{-}3)$$

These are dummy variables which indicate whether one trading period prior ($D^i_{t,\tau-1|j}$) or one day prior ($D^i_{t-1,\tau|j}$) the spot price outcome was placed either above the upper scenario ($i = 2$) or below the lower scenario price ($i = 0$).

We expect there to be significant persistence with respect to the one trading period (30 minutes) lag, as it is likely that any unexpected events (e.g. supply shortfall) may persist beyond one dispatch period. More interestingly we will examine whether there is persistence across days. A properly working pre-dispatch process should be expected to adjust for any information available the day before dispatch and hence there is no reason to expect there to be persistence that extends across days.

3.4.5. Expected pre-dispatch price

The pre-dispatch prices will give us an advanced knowledge of spot price fluctuations and signal the upcoming demand and supply response; for example when pre-dispatch price published 30 minutes before dispatch increased, industrial consumers with real-time meters would reduce their consumption while the generators would either bids at the lower supply schedule to ensure their power plants are called for dispatch or withdraw some of the lower supply schedule and bid at higher marginal price to increase the price further. As stated in Section 3.2.1, if the pre-dispatch prices deliver good forecasts of actual dispatch price, the above variables should be insignificant in explaining variation in $S_{t,\tau|j}$.

Since the price series is not easily available, we proxy the pre-dispatch price using the following equation:

$$P_{t,\tau|j} = \left(P_{t,\tau|j}^{-200} + P_{t,\tau|j}^{-100} + P_{t,\tau|j}^{+100} + P_{t,\tau|j}^{+200}\right)/4 \qquad (3.4\text{-}4)$$

This concludes the list of explanatory variables used in the ordered probit model.

As discussed in Section 2.4.4 it would be desirable to have further supply side information available. In particular a measure of the reserve margin, as this was found to be an important factor in Chen & Bunn (2010) and Bunn et al. (2012). Such a series is, however, not available at the necessary frequency.

3.5. Estimation

The model parameter vector $(\boldsymbol{\beta}_{\tau|j}, \boldsymbol{\rho}_{\tau|j})$ in the model described by equation (3.3-1) and (3.3-2) can be estimated by maximum likelihood. The log-likelihood function is

$$\begin{aligned} & llf_{\tau|j}\left(\boldsymbol{\beta}_{\tau|j}, \boldsymbol{\rho}_{\tau|j}\right) \\ & = I_2 \log\left[P\left(S_{t,\tau|j} = 2 \middle| \boldsymbol{x}_{t,\tau|j}; \boldsymbol{\beta}_{\tau|j}, \boldsymbol{\rho}_{\tau|j}\right)\right] \\ & + I_1 \log\left[P\left(S_{t,\tau|j} = 1 \middle| \boldsymbol{x}_{t,\tau|j}; \boldsymbol{\beta}_{\tau|j}, \boldsymbol{\rho}_{\tau|j}\right)\right] \\ & + I_0 \log\left[P\left(S_{t,\tau|j} = 0 \middle| \boldsymbol{x}_{t,\tau|j}; \boldsymbol{\beta}_{\tau|j}, \boldsymbol{\rho}_{\tau|j}\right)\right] \end{aligned} \qquad (3.5\text{-}1)$$

where $I_i = \mathrm{I}\left(S_{t,\tau|j} = i\right)$ is an indicator function which equals one in case $S_{t,\tau|j} = i$ is true and 0 otherwise. At the heart of the likelihood calculation is the calculation of the category probabilities $P\left(S_{t,\tau|j} = i \middle| \boldsymbol{x}_{t,\tau|j}; \boldsymbol{\beta}_{\tau|j}, \boldsymbol{\rho}_{\tau|j}\right)$, which are calculated as

$$\begin{aligned} & P\left(S_{t,\tau|j} = 2 \middle| \boldsymbol{x}_{t,\tau|j}; \boldsymbol{\beta}_{\tau|j}, \boldsymbol{\rho}_{\tau|j}\right) \\ & \qquad = P\left(S_{t,\tau|j}^{*} > \rho_{1\tau|j} \middle| \boldsymbol{x}_{t,\tau|j}; \boldsymbol{\beta}_{\tau|j}, \boldsymbol{\rho}_{\tau|j}\right) \\ & \qquad = 1 - \Phi(\rho_{1\tau|j} - \boldsymbol{\beta}_{\tau|j}{}' \boldsymbol{x}_{t,\tau|j}) \\ & P\left(S_{t,\tau|j} = 1 \middle| \boldsymbol{x}_{t,\tau|j}; \boldsymbol{\beta}_{\tau|j}, \boldsymbol{\rho}_{\tau|j}\right) \\ & = P\left(\rho_1 \geq S_{t,\tau|j}^{*} > \rho_0 \middle| \boldsymbol{x}_{t,\tau|j}; \boldsymbol{\beta}_{\tau|j}, \boldsymbol{\rho}_{\tau|j}\right) \\ & = \Phi\left(\rho_{1\tau|j} - \boldsymbol{\beta}_{\tau|j}{}' \boldsymbol{x}_{t,\tau|j}\right) - \Phi\left(\rho_{0\tau|j} - \boldsymbol{\beta}_{\tau|j}{}' \boldsymbol{x}_{t,\tau|j}\right) \\ & P\left(S_{t,\tau|j} = 0 \middle| \boldsymbol{x}_{t,\tau|j}; \boldsymbol{\beta}_{\tau|j}, \boldsymbol{\rho}_{\tau|j}\right) \\ & \qquad = P\left(\rho_0 \geq S_{t,\tau|j}^{*} \middle| \boldsymbol{x}_{t,\tau|j}; \boldsymbol{\beta}_{\tau|j}, \boldsymbol{\rho}_{\tau|j}\right) \\ & \qquad = \Phi(\rho_{0\tau|j} - \boldsymbol{\beta}_{\tau|j}{}' \boldsymbol{x}_{t,\tau|j}) \end{aligned} \qquad (3.5\text{-}2)$$

where Φ represents the standard normal CDF. The log-likelihood function can be optimised by any standard nonlinear optimiser as long as $\rho_{0\tau|j}$ is constrained to be smaller

than $\rho_{1\tau|j}$. Asymptotic standard errors for the estimated ML parameter estimates $(\widehat{\boldsymbol{\beta}}_{\tau|j}, \widehat{\boldsymbol{\rho}}_{\tau|j})$ are calculated based on the basis of the expected inverse Hessian Matrix of the variance covariance matrix (see e.g. Martin et al. (2012), Chapter 3)[24].

3.5.1. Marginal effects

In binary response model, as the one estimated here, the primary interest lies in the effects of a change in an explanatory variable on the category probabilities $P\left(S_{t,\tau|j} = i \middle| \boldsymbol{x}_{t,\tau|j}; \boldsymbol{\beta}_{\tau|j}, \boldsymbol{\rho}_{\tau|j}\right)$, $\partial P_i/\partial x_{tk}$, where x_{tk} is the *k* element of $\boldsymbol{x}_{t,\tau|j}$. In what follows conditionality on $\tau|j$ is suppressed for notational simplicity only. $\partial P_i/\partial x_{tk}$ is a nonlinear function of all parameters $(\boldsymbol{\beta}, \boldsymbol{\rho})$ and the values of $\boldsymbol{x}_t$. The parameter coefficients values themselves are not easily interpreted and will only be looked at in the context of $\partial P_i/\partial x_{tk}$. The sign of the *k*th coefficient in $\boldsymbol{\beta}$, β_k, however, will determine how the probabilities for the extreme categories changes. In particular, if $\beta_k > 0$ then $\partial P_2/\partial x_{tk} > 0$ and $\partial P_0/\partial x_{tk} < 0$. A value of $\beta_k = 0$ indicates that the *k*th explanatory variables is irrelevant.

The conditionality on $\boldsymbol{x}_t$ is typically overcome in one of two ways (Wooldridge 2002). Either the marginal effects are calculated for all observed values of $\boldsymbol{x}_t$ and one subsequently calculates the average value across these effects (*mean marginal effect*), or alternatively the marginal effect is calculated at the mean value $\bar{\boldsymbol{x}}$ (*marginal effect at mean*). Here we apply the former and present the *mean marginal effects*. The calculation of these marginal effects $\partial P_i/\partial x_{tk}$ also depends on whether x_{tk} is a continuous or discrete variable.

For continues exogenous variables, the mean marginal effect of x_{tk} on the response probabilities for the central category is derived as follows:

$$\begin{aligned} &\mathrm{E}\left(\frac{\partial[P(S_t = 1|x_{tk}, \boldsymbol{\beta}, \boldsymbol{\rho})]}{\partial x_{tk}}\right) \\ &= \mathrm{E}\left(\frac{\partial[P(\rho_1 \geq S_t^* > \rho_0|x_{tk}, \boldsymbol{\beta}, \boldsymbol{\rho})]}{\partial x_{tk}}\right) \\ &= \mathrm{E}\left(\frac{\partial[\Phi(\rho_1 - \boldsymbol{\beta}'\boldsymbol{x}_t) - \Phi(\rho_0 - \boldsymbol{\beta}'\boldsymbol{x}_t)]}{\partial x_{tk}}\right) \\ &= E(\beta_k[\phi(\rho_0 - \boldsymbol{\beta}'\boldsymbol{x}_t) - \phi(\rho_1 - \boldsymbol{\beta}'\boldsymbol{x}_t)]) \end{aligned} \tag{3.5-3}$$

[24] The models were estimated in MATLAB using the ORDER_PROBIT function written by Hang Qian, Iowa State University available from http://www.public.iastate.edu/~hqi/

Here ϕ represents the standard normal density and hence the term in square brackets can be either positive or negative resulting in the sign of the mean marginal effect being potentially different to the sign of β_k. The mean marginal effects for the extreme categories ($S_t = 0$ and $S_t = 2$) can be derived accordingly. In these cases only one density term will survive and as the density is always positive the respective mean marginal effects will be unambiguously determined by the sign of β_k. If β_k is positive then $\mathrm{E}(\partial[P(S_t = 2|x_{tk}, \boldsymbol{\beta}, \boldsymbol{\rho})/\partial x_{tk}]) > 0$ and $\mathrm{E}(\partial[P(S_t = 0|x_{tk}, \boldsymbol{\beta}, \boldsymbol{\rho})/\partial x_{tk}]) < 0$. The mean marginal effects are then estimated by calculating the term inside the expectations operator for all observed values of $\boldsymbol{x}_t$ and subsequently computing the average.

For binary exogenous variables, the calculation of the mean marginal effect is somewhat more straightforward as one merely has to consider the change in category probabilities as x_{tk} changes from 0 to 1:

$$\mathrm{E}\left(\frac{\partial[P(S_t = i|x_{tk}, \boldsymbol{\beta}, \boldsymbol{\rho})]}{\partial x_{tk}}\right) = \mathrm{E}[P(S_t = i|x_{tk} = 1) - P(S_t = i|x_{tk} = 0)] \quad (3.5\text{-}4)$$

where the category probabilities $P(S_t = i|x_{tk} = 1)$ are used as defined in equation (3.5-2). As for the case of continuous exogenous variables the sign of β_k will clearly determine the sign of the changes in the category probabilities of the extreme categories: $\mathrm{E}(\partial[P(S_t = 2|x_{tk}, \boldsymbol{\beta}, \boldsymbol{\rho})/\partial x_{tk}])$ will have the same sign as β_k and $\mathrm{E}(\partial[P(S_t = 0|x_{tk}, \boldsymbol{\beta}, \boldsymbol{\rho})/\partial x_{tk}])$ will have the opposite sign of β_k. Again, the sign of the effect on the intermediate category probability ($P(S_t = 1|x_{tk}, \boldsymbol{\beta}, \boldsymbol{\rho})$) is undetermined and depends on the particular set of observations at hand.

3.6. Results

In this Section we will present the results obtained from estimating the ordered probit model parameters in equation (3.3-1) and (3.3-2) by ML. The dependent variable of the models, $S_{t,\tau|j}$, is indexed by the trading sub-period τ , $\tau = 06{:}30\ to\ 23{:}30$, for which the spot price is compared against the pre-dispatch prices of the jth pre-dispatch, $j = PD1, PD2$. This implies that we can estimate a model for each (τ, j) combination. In order to economise on the number of models to be estimated and analysed we grouped the data into three daily sub-periods as mentioned before with $\tau = m, d, e$ for morning (06:30 to

09:30), day evening (10:00 to 17:30) and evening (18:00 to 23:30) sub-periods[25]. We shall therefore refer to $M_{\tau,j}$ as the model estimated for sub-period τ using pre-dispatch j.

Next, we need to determine an optimal estimation period. Our nine year sample period is from 1 January 1999 to 31 December 2007. In order to allow for structural changes in the relationships we are interested in, we estimate the ordered probit model for yearly sub-samples, $M_{y,\tau,j},\ y = 1997,\ldots,2007$. As it turns out the restriction that the models are identical across the years can be rejected for all different (τ, j) combination. All p-values of the relevant Likelihood-Ratio (LR) tests are smaller than 0.01

While it would be, in principle, possible to apply a search procedure to identify the number and locations of possible structural breaks, this is not done here. The computational effort required to achieve this in nonlinear models, such as the ordered probit model, is significant and stands in no relation to the limited additional insights.
We therefore estimate 54 ($= 9$ years $\times$ 3 daily sub-periods $\times$ 2 pre-dispatch periods) models. The number of observations used in each model is equal to the number of half hours included in the particular sub-period times the number of days in the particular year. By way of example, $M_{1999,m,j}$ uses 7 (seven half hours $06{:}30,\ldots,09{:}30$) times 365 (=2555) observations.

In particular we will be interested in how changes in explanatory variables affect the category probabilities, e.g. under what conditions is it more likely to find spot price outcomes, $P_{t,\tau}$, exceeding the relevant high demand pre-dispatch scenario price, $P_{t,\tau|j}^{+\Delta}$ (e.g. $S_{t,\tau|j} = 2$). Any such findings will be presented in Section 3.6.1 by summarising the relevant mean marginal effects. As discussed previously these are (nonlinear) functions of the estimated model parameters, which by itself are of limited interest.

We will initially present the results for two of these models (1999, *m*, PD1 & PD2). This will serve to explain the available results. As presenting the results for all 54 models is not particularly informative we will, here, merely present results for two of these models. We will then continue to summarise and discuss the substantial results that arise from the estimation of all the 54 models[26] in Section 3.6.1.

[25] These categories were established to encompass broadly similar periods of pre-dispatch price behaviour (relative to spot price outcomes)
[26] Detailed results for all 54 models are available in Appendix B

(INSERT Table 3.6-1 HERE)

The results for $M_{1999,m,PD1}$ and $M_{1999,m,PD2}$ are presented in Table 3.6-1. The first three results columns refer to the model that uses spot price outcomes evaluated against the pre-dispatch scenarios at *PD1* as the dependent variable ($M_{1999,m,PD1}$), whereas the final three columns represent results for $M_{1999,m,PD2}$. In the left hand column we find the exogenous variables that are used to explain variation in the category probabilities of $S_{t,m|j} = i$, $i = 0, 1\ or\ 2$. In the second row we can see the unconditional category probabilities. When the spot prices are compared to the pre-dispatch scenarios at *PD1* in the morning sub-period in year 1999, we find that in 20% of cases the wholesale spot price ends up being smaller than the low scenario pre-dispatch price, in 37% of cases it lies between the low and high pre-dispatch scenario price and in the remaining 44% of cases it exceeds the high scenario price[27].

In each model we estimate one coefficient for each of the specified explanatory variables. Only estimated coefficients that are statistically significant from 0 at 5% level are reported in the respective columns ($i = 0,1,2$) in Table 3.6-1. These columns show the calculated mean marginal effects from equations (3.5-3) and (3.5-4) for continuous and binary variables respectively. Mean marginal effects based on parameters that were not found to be significantly different from 0 are not shown (blank cell entries).

Let's consider the effect of the weekend dummy variable, D_t^{wk}. The largest positive mean marginal effect falls to the probability of the upper category $P(S_{t,m|PD1} = 2)$. On average, the probability of spot prices exceeding the high pre-dispatch (*PD1*) scenario increases by 8.5% (points) on weekends when compared to weekdays (largest increase is shown in bold). Note that the corresponding decreases in probabilities for the central and lower categories (-1.1% and -7.4%) ensure that the sum of all these three changes is 0. By construction the sum of all marginal effects equal to 0 as at any stage the sum of all the conditional probabilities, $P(S_{t,m|PD1} = i|\boldsymbol{x}_{t,\tau|j}; \boldsymbol{\beta}_{\tau|j}, \boldsymbol{\rho}_{\tau|j})$ have to sum to 1.

From these results we can infer that the estimated parameters for the D_t^{wk} explanatory variable is positive, showing the same sign as the probability change for the higher category (or opposite sign of the probability change in the lower category).

[27] These statistics correspond to the values reported for year 1999 in Panels A and B of Figure 3.2-4

The entries for the other explanatory variables can be interpreted in a similar manner, e.g. a one unit increase in the load forecast, l_t^e, has a statistically significant effect on the conditional outcome probabilities. More specifically, a higher load forecast (ceteris paribus) increases the probability of spot price outcomes exceeding the high scenario pre-dispatch price (*PD1*) by 12.7% (points)[28].

From the results in Table 3.6-1, we can draw a number of conclusions. As discussed previously, it is apparent that the outcome distribution is more focused on the central category for PD2 compared to PD1 (49% of $S_{t,m|PD2} = 1$ as compared to 37% of $S_{t,m|PD1} = 1$). Also we can see that there are fewer significant explanatory variables that can explain variation in conditional probabilities. In particular the expected pre-dispatch price and the load variables are no longer significant in explaining response probabilities at PD2. It is also apparent from these results that there is a fair degree of persistence in the spot price outcomes. The probability of $S_{t,\tau|j} = 2$ at PD1 (PD2) increases by 27.2% (31%) if the spot price outcomes in the previous half-hour is $S_{t,\tau-1|j} = 2$ and by 7.6% (11.3%) if 48 half-hours prior $S_{t-1,\tau|j} = 2$. This clearly indicates that high spot price outcomes are likely to be followed by further high spot price outcomes. The significant effect of $S_{t,\tau-1|j} = 2$ is not surprising as any unanticipated (at *PD1*) event that may cause the spot price at 08:30 to exceed the high scenario, is likely to still be in place at 09:00. However, the fact that the previous days spot price exceeded the high scenario (for sub-period τ) is known at the time of the pre-dispatch (certainly for *PD2* and also for all morning periods at *PD1*). As such this information should have been worked into the pre-dispatch. The fact that this variable still has a sizeable and statistically significant effect on the category probabilities is one indication for the pre-dispatch process not incorporating all relevant information.

Interestingly, the equivalent persistence effect of $S^0_{t,\tau-1|j}$ and $S^0_{t-1,\tau|j}$ on the conditional category probability $P(S_{t,\tau|j} = 0)$ is also statistically significant but not quite as strong as the persistence for the highest category. The effect is weaker than the high category persistence at both lags (half-hour and 48 half-hours) and at both *PD1* and *PD2*.

[28] For readability we will omit the "points" clarification in what follows.

3.6.1. Findings

When evaluating the results of all 54 estimated models a number of general findings arise. These are now discussed in turn.

3.6.1.1. Stronger persistence for $S_{t,\tau|j} = 0$ and $S_{t,\tau|j} = 2$

(INSERT Figure 3.6-1 HERE)

There is strong positive persistence for the spot price outcomes to be placed in the same outcome categories as it was in previous half hour in m, d and e periods for both *PD1* and *PD2*. As these results do not qualitatively differ either across years or the periods of the day (m, d and e) we present, in Figure 3.6-1, results for mean marginal effects that are averaged across these two dimensions.

As for the results for $M_{1999,m,PD1}$ and $M_{1999,m,PD2}$ we can clearly discern that the persistence is much stronger for the half hour lag (Panels A and B) as compared to the 48 half hour lag (Panels C and D). For instance, knowing that the spot price outcome is in the low category increases the probability of ending up in the low category, in the next half hour, (ceteris paribus) by about 30% (panel A). Similarly, if the current spot price is in the high category this will increase the probability of next half hour's spot price being in the high category by about 35% (panel B). The 48 half hour persistence is much lower at about 3.5% (for the low category, in panel C) and 4% (high category, in panel D) respectively. We can still detect a slight asymmetry in the sense that the persistence is somewhat larger for the high category.

3.6.1.2. Strategic bidding

(INSERT Figure 3.6-2 HERE)

Next we discuss mean marginal effects when the demand one would expect for that day of the week, $\widetilde{el}_t$ increases by 1MwH. A 1MWh change is in fact a very small variation in load, in particular when considering that the low and high pre-dispatch scenarios are generated by changing the load forecast by +/- 200MWh. As such it is unlikely that any variations in category probabilities reported, as a result of this marginal change, is actually due to the change in the load. They are more likely to be the result of re-bidding activity triggered by the conditions that make the actual load exceed the load forecast.

In the rows of Figure 3.6-2 we can see the different results for the morning (m), day time (d) and evening (e) sub-periods. In the left column we present the results for *PD1* and those for *PD2* on the right. We can see that the increase in the category probability for $S_{t,\tau|j} = 2$, as a result of the small increase in $\widetilde{el}_t$, is in the order of 10% (points). Taking into account the size of the change considered this seems very high. In fact this value increases to an increase of about 15% (points) for the *PD2* model for the morning period. This seems to indicate that conditions in which actual load exceeds the load forecast trigger more significant re-bids for the morning dispatch period.

This may be due to the fact that the time left to the morning dispatch periods (recall that *PD2* represents the 6am pre-dispatch on the dispatch day) is very short. This is a scenario in which strategic rebidding appears quite fertile as it leaves little to no opportunity for an effective competitive response to any rebids, either by generators or indeed from the demand side. Indeed it was highlighted by National Electricity Code Administrator (2001b); Australia Competition & Consumer Conmission (2002) and Australian Competition & Consumer Commission (2002) that such late re-bidding activity plays an important role in causing the short-term price spikes.

3.6.1.3. Temperature

Of the temperature variables, only the unseasonably cold days in winter variable, $Tmin_t$, could explain any of the variation in the spot price outcomes during the morning (m) period. For the daytime (d) period, however, both the $Tmax_t$ and $Tmin_t$ variables were significant. The spot price outcomes in the evening (e) period, were independent of the two temperature variables. This pattern illustrates that the pre-dispatch process seems to anticipate and incorporate all the temperature information relevant to spot price outcomes in the evening periods into the pre-dispatch process, even at *PD1*. Logically, it is the daytime period that is most exposed to extreme temperature variations.

Where these variables are significant it appears as if more extreme temperatures, sensibly, increase the probability of the wholesale spot price being larger than the high pre-dispatch scenario price. This is illustrated in Figure 3.6-3 in which we can see that more extreme values in $Tmax_t$ and $Tmin_t$ result in an increase in the probability for the high category.

(INSERT Figure 3.6-3 HERE)

3.6.1.4. Expected pre-dispatch price

(INSERT Figure 3.6-4 HERE)

Last we investigate whether the category probabilities depend on the value of the expected pre-dispatch price level as measured by $P_{t,\tau|j}$. In Figure 3.6-4, we observe that if $P_{t,\tau|j}$ increases, there is a significantly lower probability for $S_{t,\tau|j} = 2$ and a higher probability of the spot price being below the *PD1* low scenario price. The coefficients to the $P_{t,\tau|j}$ variable are significant during the morning (*m*) period on *PD1* and the evening (*e*) period on *PD1* and *PD2*. Notably, when the pre-dispatch is close to dispatch (*m*, *PD2*) the variation in the spot price outcomes to be placed in any of the categories does not vary with $P_{t,\tau|j}$ anymore. This indicates that supply and demand adjustment for the morning trading period have taken place by 6.00am (the time of *PD2*). This could be the result of demand elasticity (consumers decreasing their consumption) as a response to the high prices signal or rebidding mechanism by the generators resulting to cheaper power plants being made available or supply schedules being adjusted downwards.

3.7. Conclusions

The aim of this chapter was threefold. *First*, to introduce a framework in which we can analyse whether the pre-dispatch process, as operated by the AEMO, delivers unbiased predictions of the actual wholesale spot price outcomes. This was achieved by investigating the statistical properties of a categorical, ordinal variable based on high and low pre-dispatch scenarios. It transpired that in general the pre-dispatch process is downward biased, i.e. underestimates the actual price outcomes of the wholesale spot prices.

Second, we investigate whether this bias varies systematically across years, seasons and/or trading periods. There is strong evidence that the pre-dispatch process seems to underestimate the actual price outcomes more severely for the morning and evening peak periods of the day. This finding is remarkably persistent throughout our sample period. If at all this bias becomes stronger for the later sample. When comparing the properties of the pre-dispatch process undertaken on the afternoon of the day prior to dispatch (*PD1*) to that undertaken at 6am on the day of dispatch (*PD2*) it becomes apparent that the latter is indeed more precise in the sense that we find a higher proportion of spot price outcomes

between the pre-dispatch's high and low scenario price. Having said that, at *PD2* there is still a strong downward bias even for the morning trading period.

Third, by modelling our categorical, ordered outcome variable with an ordered probit model, we are in a position to investigate whether there are any exogenous variables that are able to explain the variation in the probability of spot price outcomes to fall above (below) the high (low) pre-dispatch scenario prices. Exogenous variables included are seasonal variables (to allow for systematic variations observed), load forecasts and load forecast errors, temperature data and expected pre-dispatch prices. We also include variables designed to investigate whether spot price outcomes (relative to pre-dispatch scenario prices) show any persistence. While such persistence is clearly expected (and indeed observed) in the very short run it is somewhat surprising to find this persistence to still be statistically significant across a full day. This clearly indicates that the pre-dispatch process, as operated in New South Wales, fails to fully accommodate all information available on the day prior to dispatch. Further results demonstrate that there is some indication of market participants adjusting their behaviour to the price signal delivered by the pre-dispatch process. Higher price signals result in increased probabilities for price outcomes below the low pre-dispatch scenario price. From the data available here it is impossible to say whether these adjustments are driven by demand or supply side adjustments.

These results improve our understanding of the strengths and weaknesses of the pre-dispatch process in the Australian Electricity Market. We identified circumstances when the pre-dispatch process tends to produce biased price forecasts but also found evidence that it does fulfil some of its signalling function and allows market participants to adjust their behaviour. The results obtained (in particular the persistence of the overall bias) allow us to conjecture that some of the bids and re-bids provided by generators are not made in good faith as required by the National Electricity Rules (Australian Energy Market Commission, 2013, paragraph 3.8.22A).

In order to substantiate such a conjecture the analysis would have to analyse firm and bid specific information. This information is in principle available and we anticipate that it will provide a fertile ground of future research.

References

Australia Competition & Consumer Commission, 2002. *Final Determination on Bidding and Rebidding Code Changes*, Available at: http://transition.accc.gov.au/content/index.phtml/itemId/744502/fromItemId/401858/display/acccDecision.

Australian Competition & Consumer Commission, 2002. *Draft Determination on Bidding and Rebidding Code Changes*, Available at: http://transition.accc.gov.au/content/index.phtml/itemId/744502/fromItemId/401858/display/acccDecision.

Australian Energy Market Commission, 2013. National Electricity Rules Version 56 Chapter 3: Market Rules. , p.146. Available at: http://www.aemc.gov.au/Electricity/National-Electricity-Rules/Current-Rules.html.

Australian Energy Market Commission, 2012. Potential Generator Market Power in the NEM. *Rule Determination*. Available at: http://www.aemc.gov.au/electricity/rule-changes/completed/potential-generator-market-power-in-the-nem.html.

Australian Energy Market Operator, 2010. An introduction to australia's national electricity market. *Available at AEMO website http://http://www.aemo.com.au/About-the-Industry/Energy-Markets/National-Electricity-Market*.

Bardak Ventures Pty Ltd, 2005. The Effect of Industry Structure on Generation Competition and End-User Prices in the National Electricity Market.

Borenstein, S., Bushnell, J. & Wolak, F., 2002. Measuring market inefficiencies in California's restructured wholesale electricity market. *American Economic Review*, 92(5), pp.1376–1405. Available at: http://www.jstor.org/stable/10.2307/3083255 [Accessed August 12, 2013].

Bunn, D., Andresen, A., Chen, D., & Westgaard, S. (2012). Analysis and Forecasting of Electricity Price Risks with Quantile Factor Models.

Bushnell, J. & Saravia, C., 2002. An empirical assessment of the competitiveness of the New England electricity market. Available at: http://escholarship.org/uc/item/8913v4bk [Accessed August 12, 2013].

Chen, D. & Bunn, D.W., 2010. Analysis of the Nonlinear Response of Electricity Prices to Fundamental and Strategic Factors. IEEE Transactions on Power Systems, 25(2), pp.595–606. Available at: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5378456.

David, A.K. & Wen, F., 2000. Strategic bidding in competitive electricity markets: a literature survey. In *Power Engineering Society Summer Meeting, 2000. IEEE*. pp. 2168–2173 vol. 4. Available at: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=866982.

Eichler, M. et al., 2012. *Modeling spike occurrences in electricity spot prices for forecasting*, METEOR, Maastricht research school of Economics of Technology and Organizations. Available at: http://arno.unimaas.nl/show.cgi?fid=25393.

Electricity Market Performance, 2012. *Factors Contributing to Differences Between Dispatch and Pre-Dispatch Outcomes*, Available at: http://www.aemo.com.au/Electricity/Market-Operations/Dispatch/Factors-Contributing-to-Differences-between-Dispatch-and-Predispatch-Outcomes.

Fabra, N. & Toro, J., 2005. Price wars and collusion in the Spanish electricity market. *International Journal of Industrial Organization*, 23(3), pp.155–181. Available at: http://www.sciencedirect.com/science/article/pii/S0167718705000123 [Accessed August 12, 2013].

Gillett, R. & Market Operations, 2010. *Pre-dispatch Process Description*, Available at: http://www.aemo.com.au/Electricity/Market-Operations/Dispatch/Predispatch-Process-Description [Accessed August 12, 2013].

Hortacsu, A. & Puller, S., 2005. Understanding strategic bidding in restructured electricity markets: a case study of ERCOT. , (010366). Available at: http://www.nber.org/papers/w11123 [Accessed August 12, 2013].

Kharbach, S., Fredette, M. & Pineau, P.-O., 2010. *Impacts of Imports and Natural Gas on Electricity Prices: The Case of Ontario*, Groupe d'études et de recherche en analyse des décisions. Available at: http://www.gerad.ca/fichiers/cahiers/G-2010-70.pdf [Accessed August 12, 2013].

Knittel, C.R. & Metaxoglou, K., 2008. Diagnosing unilateral market power in electricity reserves market. Available at: http://escholarship.org/uc/item/14q6c0mk.pdf [Accessed August 12, 2013].

Mansur, E., 2001. *Pricing behavior in the initial summer of the restructured PJM wholesale electricity market*, University of California Energy Institute. Available at: http://www.ucei.berkeley.edu/ucei/PDF/pwp083.pdf.

Mansur, E.T., 2007. Upstream competition and vertical integration in electricity markets. *Journal of Law and Economics*, 50(1), pp.125–156. Available at: http://www.jstor.org/stable/10.1086/508309.

Martin, V., Hurn, S. & Harris, D., 2012. *Econometric modelling with time series: specification, estimation and testing*, Cambridge University Press.

National Electricity Code Administrator, 2001a. *Generators ' bidding and rebidding strategies and their effect on prices*, Available at: http://www.neca.com.au/TheCode0378.html?CategoryID=34&SubCategoryID=83&ItemID=951.

National Electricity Code Administrator, 2001b. *Generators' bidding and rebidding strategies and their effect on prices Written comments in response to Panel's consultation document*, Available at: http://www.neca.com.au/Files/C_Bidding_rebidding_strategies_and_effect_on_prices_written_comments_Sep2001.pdf.

Nowicka-Zagrajek, J. & Weron, R., 2002. Modeling electricity loads in California: ARMA models with hyperbolic noise. Signal Processing, 82(12), pp.1903–1915.

Oh, H., 2003. *Simulation methods for modeling offer behavior and spot prices in restructured markets for electricity*, Cornell University.

Rodriguez, C.P. & Anders, G.J., 2004. Energy price forecasting in the Ontario competitive power system market. *Power Systems, IEEE Transactions on*, 19(1), pp.366–374. Available at: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1266590.

Saravia, C., 2003. Speculative Trading and Market Performance: The Effect of Arbitrageurs on Efficiency and Market Power in the New York Electricity Market. Available at: http://escholarship.org/uc/item/0mx44472#page-4.

Sweeting, A., 2007. Market Power In The England And Wales Wholesale Electricity Market 1995-2000. *The Economic Journal*, 117(520), pp.654–685. Available at: http://onlinelibrary.wiley.com/doi/10.1111/j.1468-0297.2007.02045.x/full [Accessed August 12, 2013].

Weron, R., 2007. *Modeling and forecasting electricity loads and prices: A statistical approach*, Wiley. com.

Wolak, F., 2003. Measuring unilateral market power in wholesale electricity markets: the California market, 1998-2000. *The American economic review*, 93(2), pp.1998–2000. Available at: http://www.jstor.org/stable/10.2307/3132266 [Accessed August 12, 2013].

Wolak, F. & Patrick, R., 2001. *The impact of market rules and market structure on the price determination process in the England and Wales electricity market*, Available at: http://medcontent.metapress.com/index/A65RM03P4874243N.pdf [Accessed August 12, 2013].

Wolak, F.A., 2000. An Empirical Analysis of the Impact of Hedge Contracts on Bidding Behavior in a Competitive Electricity Market. *International Economic Journal*, 14(2), pp.1–39.

Wolfram, C.D., 1999. Measuring duopoly power in the British electricity spot market. *American Economic Review*, 89(4), pp.805–826. Available at: http://www.jstor.org/stable/10.2307/117160 [Accessed August 12, 2013].

Wolfram, C.D., 1998. Strategic Bidding in a Multiunit Auction: An Empirical Analysis of Bids to Supply Electricity in England and Wales. *The RAND Journal of Economics*, 29(4), pp.pp. 703–725. Available at: http://www.jstor.org/stable/2556090.

Wooldridge, J.M., 2002. *Econometric analysis of cross section and panel data*, The MIT press.

Zareipour, H. et al., 2006. Application of public-domain market information to forecast Ontario's wholesale electricity prices. *Power Systems, IEEE Transactions on*, 21(4), pp.1707–1717. Available at: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1717574 [Accessed August 11, 2013].

Zareipour, H., Bhattacharya, K. & Canizares, C.A., 2006. Forecasting the hourly Ontario energy price by multivariate adaptive regression splines. In *Power Engineering Society General Meeting, 2006. IEEE*. p. 7–pp. Available at: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1709474.

Zareipour, H., Caizares, C.A. & Bhattacharya, K., 2007. The operation of Ontario's competitive electricity market: overview, experiences, and lessons. *Power Systems, IEEE Transactions on*, 22(4), pp.1782–1793. Available at: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=4349150.

Appendix B

	$M_{1999,m,j}$						$M_{1999,d,j}$						$M_{1999,e,j}$					
	$j = PD1$			$j = PD2$			$j = PD1$			$j = PD2$			$j = PD1$			$j = PD2$		
i	0	1	2	0	1	2	0	1	2	0	1	2	0	1	2	0	1	2
$\Pr(S_{t,\tau\|j} = i)$	20%	37%	**44%**	12%	**49%**	39%	27%	**36%**	**36%**	29%	**41%**	31%	31%	34%	**35%**	28%	**39%**	33%
D_t^{wk}	-7%	-1%	**8%**	-3%	-2%	**5%**	-4%	-1%	**5%**	-3%	0%	**3%**	-4%	-1%	**5%**	-3%	0%	**3%**
D_t^{win}	**8%**	1%	-9%	**2%**	1%	-3%	**5%**	1%	-6%	**2%**	0%	-2%	**9%**	1%	-10%	**6%**	0%	-6%
D_t^{sum}																		
$D_{t,\tau-1\|j}^0$	**23%**	3%	-26%	**21%**	-1%	-20%	**49%**	**-7%**	-42%	**50%**	-30%	-20%	**38%**	-3%	-35%	**39%**	-14%	-25%
$D_{t-1,\tau\|j}^0$	**5%**	1%	-6%	**5%**	2%	-7%				**1%**	0%	-1%	**5%**	1%	-6%	**2%**	0%	-2%
$D_{t,\tau-1\|j}^2$	-21%	-6%	**27%**	-13%	-18%	**31%**	-27%	-23%	**50%**	-28%	-12%	**40%**	-25%	-12%	**37%**	-27%	-14%	**41%**
$D_{t-1,\tau\|j}^2$	-7%	-1%	**8%**	-6%	-5%	**11%**	-3%	-1%	**4%**	-3%	0%	**3%**	-5%	-1%	**6%**	-4%	0%	**4%**
$P_{t,\tau\|j}$	**2%**	0%	-2%										**1%**	0%	-1%	**2%**	0%	-2%
$\widetilde{el}_t$	-9%	-1%	**10%**				-10%	-3%	**13%**	-7%	0%	**7%**	-13%	-2%	**15%**	-6%	0%	**6%**
l_t^e	-11%	-2%	**13%**				-7%	-1%	**8%**				-8%	-1%	**9%**	-5%	0%	**5%**
$Tmin_t$	-2%	0%	**2%**	-1%	-1%	**2%**												
$Tmax_t$							-1%	0%	**1%**	-1%	0%	**1%**						

Table B-1: This table is used to show the mean marginal effect, $M_{y,\tau,j}$ in year (1999), y during morning (m), daytime (d) and evening (e) sub-periods, τ for both dispatch periods, j; the day before dispatch ($t-1$, 15:00, *PD1*) and on the dispatch day (t) at 6am (*PD2*). The unconditional probability of *PD1* and *PD2* for each of the spot price outcomes; $\Pr(S_{t,\tau|j} = \{0,1,2\})$ are reported in the third row. While the mean marginal effects of the explanatory variables listed in first column are display in the remaining rows. $D_t^{wk}, D_t^{win}, D_t^{sum}$ are the dummy variables giving value of 1 when day t is weekends, winter, summer or 0 otherwise. $\widetilde{el}_t$, is de-seasonalised load forecast error and l_t^e is load forecast as discussed in Section 3.4.2. $Tmax_t$ and $Tmin_t$ are unseasonably temperature as discussed in Section 3.4.3. $D_{t,\tau-1|j}^2$,$D_{t-1,\tau|j}^2$, $D_{t,\tau-1|j}^0$ and $D_{t-1,\tau|j}^0$ are dummy variables which indicate whether one trading period prior ($D_{t,\tau-1|j}^i$) or one day prior ($D_{t-1,\tau|j}^i$) the spot price outcome was placed either above the upper scenario ($i = 2$) or below the lower scenario price ($i = 0$). And lastly, $P_{t,\tau|j}$ is the expected pre-dispatch price calculated using equation (3.4-4). The mean marginal effects for continuous and binary variables are calculated based on equation (3.5-3) and (3.5-4) respectively. All the mean marginal effects displayed are based on the significant estimated parameters, β and α at 5% level. The insignificant ones are not shown (blank cell entries). Only the highest positive mean marginal effects are bolded.

	$M_{2000,m,j}$						$M_{2000,d,j}$						$M_{2000,e,j}$					
	$j = PD1$			$j = PD2$			$j = PD1$			$j = PD2$			$j = PD1$			$j = PD2$		
i	0	1	2	0	1	2	0	1	2	0	1	2	0	1	2	0	1	2
$\Pr(S_{t,\tau\|j} = i)$	19%	34%	**47%**	15%	**47%**	38%	30%	33%	**38%**	31%	**37%**	32%	24%	33%	**43%**	22%	37%	**41%**
D_t^{wk}	-6%	-3%	**9%**	-7%	-2%	**9%**	-2%	-1%	**3%**	-4%	1%	**3%**						
D_t^{win}							-2%	0%	**2%**	-2%	0%	2%						
D_t^{sum}				**3%**	0%	-3%	**2%**	0%	-2%	2%	0%	-2%						
$D_{t,\tau-1\|j}^0$	**28%**	-2%	-26%	**20%**	-7%	-13%	**50%**	-27%	-23%	**51%**	-38%	-13%	**36%**	-17%	-19%	**37%**	-28%	-9%
$D_{t-1,\tau\|j}^0$	**6%**	2%	-8%	**5%**	0%	-5%	1%	1%	-2%									
$D_{t,\tau-1\|j}^2$	-21%	-15%	**36%**	-15%	-6%	**21%**	-29%	-9%	**38%**	-30%	4%	**26%**	-26%	-6%	**32%**	-25%	8%	**17%**
$D_{t-1,\tau\|j}^2$	-3%	-2%	**5%**	-5%	-1%	**6%**	-2%	0%	2%	-1%	0%	1%	-5%	0%	5%	-3%	1%	**2%**
$P_{t,\tau\|j}$													**0.4%**	0%	-0.4%	**0.3%**	-0.1%	-0.2%
$\widetilde{el}_t$	-10%	-5%	**15%**	-13%	-2%	**15%**	-5%	-1%	**6%**	-5%	1%	**4%**						
l_t^e																**3%**	-1%	-2%
$Tmin_t$	-3%	-1%	**4%**	-3%	0%	**3%**	**1%**	0%	-1%				-1%	0%	1%	-1%	0%	1%
$Tmax_t$							-1%	0%	**1%**	-0.5%	0.1%	**0.4%**						

Table B-2: This table is used to show the mean marginal effect, $M_{y,\tau,j}$ in year (2000), y during morning (m), daytime (d) and evening (e) sub-periods, τ for both dispatch periods, j; the day before dispatch ($t-1$, 15:00, *PD1*) and on the dispatch day (t) at 6am (*PD2*).

Further explanation on this table is similar to the caption of Table B-1

	$M_{2001,m,j}$						$M_{2001,d,j}$						$M_{2001,e,j}$					
	$j = PD1$			$j = PD2$			$j = PD1$			$j = PD2$			$j = PD1$			$j = PD2$		
i	0	1	2	0	1	2	0	1	2	0	1	2	0	1	2	0	1	2
$\Pr(S_{t,\tau\|j} = i)$	21%	30%	**49%**	22%	**43%**	35%	**39%**	26%	35%	**42%**	32%	27%	**38%**	24%	**38%**	**39%**	26%	34%
D_t^{wk}	-9%	-5%	**14%**	-5%	-1%	**6%**	-2%	0%	2%									
D_t^{win}	**7%**	3%	-10%	**5%**	1%	-6%							**6%**	-1%	-5%	**6%**	-2%	-4%
D_t^{sum}							2%	0%	-2%	**3%**	-2%	-1%						
$D_{t,\tau-1\|j}^0$	**24%**	5%	-29%	**22%**	-2%	-20%	**47%**	-32%	-15%	**52%**	-45%	-7%	**41%**	-17%	-24%	**43%**	-28%	-15%
$D_{t-1,\tau\|j}^0$	**4%**	2%	-6%	**9%**	1%	-10%							**5%**	0%	-5%	**5%**	-1%	-4%
$D_{t,\tau-1\|j}^2$	-23%	-15%	**38%**	-19%	-8%	**27%**	-39%	14%	**25%**	-37%	**22%**	15%	-31%	1%	**30%**	-33%	11%	**22%**
$D_{t-1,\tau\|j}^2$	-4%	-2%	**6%**	-5%	-1%	**6%**	-3%	0%	3%	-4%	2%	2%	-4%	0%	4%	-3%	1%	**2%**
$P_{t,\tau\|j}$																**0.5%**	-0.1%	-0.4%
$\widetilde{el_t}$	-11%	-5%	**16%**	-8%	-2%	**10%**	-3%	0%	3%									
l_t^e							3%	0%	-3%	**4%**	-2%	-2%						
$Tmin_t$	-1%	-1%	**2%**	-2%	0%	2%												
$Tmax_t$				-1%	0%	1%												

Table B-3: This table is used to show the mean marginal effect, $M_{y,\tau,j}$ in year (2001), y during morning (m), daytime (d) and evening (e) sub-periods, τ for both dispatch periods, j; the day before dispatch ($t-1$, 15:00, *PD1*) and on the dispatch day (t) at 6am (*PD2*).

Further explanation on this table is similar to the caption of Table B-1

	$M_{2002,m,j}$						$M_{2002,d,j}$						$M_{2002,e,j}$					
	$j = PD1$			$j = PD2$			$j = PD1$			$j = PD2$			$j = PD1$			$j = PD2$		
i	0	1	2	0	1	2	0	1	2	0	1	2	0	1	2	0	1	2
$\Pr(S_{t,\tau\|j} = i)$	12%	24%	**64%**	17%	**46%**	37%	26%	25%	**49%**	32%	31%	**37%**	26%	26%	**48%**	29%	28%	**44%**
D_t^{wk}	-3%	-1%	**4%**	-6%	-4%	**10%**												
D_t^{win}				**5%**	3%	-8%	-2%	-2%	**4%**	-2%	0%	**2%**	**4%**	1%	-5%	**4%**	1%	-5%
D_t^{sum}				**2%**	1%	-3%	**3%**	2%	-5%	**3%**	0%	-3%						
$D_{t,\tau-1\|j}^0$	**20%**	5%	-25%	**32%**	4%	-36%	**39%**	-10%	-29%	**48%**	-27%	-21%	**35%**	3%	-38%	**41%**	2%	-43%
$D_{t-1,\tau\|j}^0$				**2%**	1%	-3%	**2%**	1%	-3%	**3%**	0%	-3%	**5%**	1%	-6%	**3%**	1%	-4%
$D_{t,\tau-1\|j}^2$	-21%	-8%	**29%**	-16%	-18%	**34%**	-30%	-14%	**44%**	-31%	-4%	**35%**	-27%	-16%	**43%**	-27%	-13%	**40%**
$D_{t-1,\tau\|j}^2$	-2%	-1%	**3%**	-3%	-2%	**5%**	-1%	-1%	**2%**				-4%	-1%	**5%**	-3%	0%	**3%**
$P_{t,\tau\|j}$	**2%**	0%	-2%	1%	1%	-2%							**2%**	1%	-3%	**2%**	0%	-2%
$\widetilde{el_t}$	-7%	-2%	**9%**	-13%	-9%	**22%**	-5%	-4%	**9%**	-6%	0%	**6%**	-9%	-2%	**11%**	-9%	-2%	**11%**
l_t^e	-6%	-2%	**8%**	-7%	-5%	**12%**							-4%	-1%	**5%**	-6%	-1%	**7%**
$Tmin_t$				-2%	-1%	**3%**												
$Tmax_t$							-0.4%	-0.3%	**0.7%**	-0.4%	0.0%	**0.4%**				**0.5%**	0.1%	-0.6%

Table B-4: This table is used to show the mean marginal effect, $M_{y,\tau,j}$ in year (2002), y during morning (m), daytime (d) and evening (e) sub-periods, τ for both dispatch periods, j; the day before dispatch ($t-1$, 15:00, *PD1*) and on the dispatch day (t) at 6am (*PD2*).

Further explanation on this table is similar to the caption of Table B-1

	$M_{2003,m,j}$						$M_{2003,d,j}$						$M_{2003,e,j}$					
	$j = PD1$			$j = PD2$			$j = PD1$			$j = PD2$			$j = PD1$			$j = PD2$		
i	0	1	2	0	1	2	0	1	2	0	1	2	0	1	2	0	1	2
$\Pr(S_{t,\tau\|j} = i)$	12%	22%	**66%**	15%	40%	**45%**	24%	25%	**51%**	29%	29%	**42%**	20%	22%	**57%**	21%	27%	**53%**
D_t^{wk}	-5%	**3%**	2%	-6%	0%	**6%**	-2%	-1%	**3%**	-2%	-1%	**3%**						
D_t^{win}	**6%**	-3%	-3%	**7%**	0%	-7%												
D_t^{sum}							**3%**	2%	-5%	**3%**	1%	-4%				**2%**	1%	-3%
$D_{t,\tau-1\|j}^0$	**20%**	-12%	-8%	**26%**	1%	-27%	**37%**	-1%	-36%	**44%**	-13%	-31%	**31%**	9%	-40%	**27%**	3%	-30%
$D_{t-1,\tau\|j}^0$	**2%**	-1%	-1%	**4%**	0%	-4%				**2%**	0%	-2%	**2%**	1%	-3%	**2%**	1%	-3%
$D_{t,\tau-1\|j}^2$	-19%	**14%**	5%	-16%	-1%	**17%**	-28%	-21%	**49%**	-30%	-15%	**45%**	-23%	-21%	**44%**	-25%	-20%	**45%**
$D_{t-1,\tau\|j}^2$	-2%	1%	1%	-2%	0%	**2%**	-3%	-2%	**5%**	-2%	0%	**2%**	-4%	-2%	**6%**	-3%	-2%	**5%**
$P_{t,\tau\|j}$	**2%**	-1%	-1%				1%	1%	-2%	**1%**	0%	-1%						
$\widetilde{el_t}$	-9%	**5%**	4%	-16%	-1%	**17%**	-5%	-3%	**8%**	-6%	-1%	**7%**	-8%	-4%	**12%**	-7%	-3%	**10%**
l_t^e	-10%	**6%**	4%	-10%	-1%	**11%**												
$Tmin_t$	-1.4%	**0.8%**	0.6%	-1%	0%	**1%**							-1%	0%	**1%**	-1%	0%	**1%**
$Tmax_t$				-1%	0%	**1%**	-1%	0%	**1%**	-1%	0%	**1%**						

Table B-5: This table is used to show the mean marginal effect, $M_{y,\tau,j}$ in year (2003), y during morning (m), daytime (d) and evening (e) sub-periods, τ for both dispatch periods, j; the day before dispatch ($t-1$, 15:00, *PD1*) and on the dispatch day (t) at 6am (*PD2*).

Further explanation on this table is similar to the caption of Table B-1

	$M_{2004,m,j}$						$M_{2004,d,j}$						$M_{2004,e,j}$					
	$j = PD1$			$j = PD2$			$j = PD1$			$j = PD2$			$j = PD1$			$j = PD2$		
i	0	1	2	0	1	2	0	1	2	0	1	2	0	1	2	0	1	2
$\Pr(S_{t,\tau\|j} = i)$	10%	18%	**72%**	14%	35%	**50%**	22%	22%	**56%**	25%	29%	**46%**	20%	21%	**59%**	23%	25%	**53%**
D_t^{wk}	-4%	-4%	**8%**	-3%	-3%	**6%**	-2%	-2%	**4%**	-3%	-1%	**4%**						
D_t^{win}	4%	4%	-8%	**3%**	2%	-5%	1%	1%	-2%	**2%**	1%	-3%	**3%**	1%	-4%	**3%**	1%	-4%
D_t^{sum}	2%	2%	-4%				**3%**	2%	-5%	**3%**	1%	-4%	**3%**	1%	-4%	**2%**	1%	-3%
$D_{t,\tau-1\|j}^0$	**16%**	12%	-28%	**19%**	6%	-25%	**34%**	5%	-39%	**38%**	-3%	-35%	**30%**	10%	-40%	**33%**	7%	-40%
$D_{t-1,\tau\|j}^0$	4%	4%	-8%	**3%**	2%	-5%	2%	2%	-4%	**2%**	1%	-3%	**2%**	1%	-3%	**3%**	1%	-4%
$D_{t,\tau-1\|j}^2$	-18%	-20%	**38%**	-17%	-17%	**34%**	-28%	-25%	**53%**	-28%	-24%	**52%**	-23%	-18%	**41%**	-24%	-16%	**40%**
$D_{t-1,\tau\|j}^2$	-3%	-3%	**6%**	-4%	-3%	**7%**	-1%	-1%	**2%**				-4%	-2%	**6%**	-3%	-1%	**4%**
$P_{t,\tau\|j}$	1%	1%	-2%							**3%**	1%	-4%	**1%**	0%	-1%	**1%**	0%	-1%
$\widetilde{el}_t$	-5%	-5%	**10%**	-12%	-10%	**22%**	-5%	-4%	**9%**	-8%	-2%	**10%**	-8%	-3%	**11%**	-8%	-3%	**11%**
l_t^e	-3%	-3%	**6%**							-3%	-1%	**4%**	-4%	-2%	**6%**	-5%	-2%	**7%**
$Tmin_t$							**0.5%**	0.4%	-0.9%	**1%**	0%	-1%						
$Tmax_t$																		

Table B-6: This table is used to show the mean marginal effect, $M_{y,\tau,j}$ in year (2004), y during morning (m), daytime (d) and evening (e) sub-periods, τ for both dispatch periods, j; the day before dispatch ($t-1$, 15:00, *PD1*) and on the dispatch day (t) at 6am (*PD2*).

Further explanation on this table is similar to the caption of Table B-1

	$M_{2005,m,j}$						$M_{2005,d,j}$						$M_{2005,e,j}$					
	$j = PD1$			$j = PD2$			$j = PD1$			$j = PD2$			$j = PD1$			$j = PD2$		
i	0	1	2	0	1	2	0	1	2	0	1	2	0	1	2	0	1	2
$\Pr(S_{t,\tau\|j} = i)$	11%	14%	**74%**	15%	32%	**54%**	18%	17%	**65%**	22%	27%	**51%**	18%	18%	**64%**	22%	23%	**55%**
D_t^{wk}	-2%	1%	1%				-2%	-1%	**3%**	-3%	-1%	**4%**						
D_t^{win}	**3%**	-1%	-2%															
D_t^{sum}	**2%**	-1%	-1%	**3%**	2%	-5%	**2%**	1%	-3%	**3%**	0%	-3%						
$D_{t,\tau-1\|j}^{0}$	**12%**	-5%	-7%	**18%**	8%	-26%	**30%**	14%	-44%	**35%**	6%	-41%	**23%**	10%	-33%	**28%**	1%	-29%
$D_{t-1,\tau\|j}^{0}$				**4%**	3%	-7%							**3%**	2%	-5%	**4%**	2%	-6%
$D_{t,\tau-1\|j}^{2}$	-23%	11%	**12%**	-17%	-14%	**31%**	-26%	-26%	**52%**	-28%	-18%	**46%**	-25%	-18%	**43%**	-26%	-13%	**39%**
$D_{t-1,\tau\|j}^{2}$	-3%	1%	**2%**	-3%	-2%	**5%**	-1%	-1%	**2%**				-5%	-3%	**8%**	-2%	-1%	**3%**
$P_{t,\tau\|j}$													**0.3%**	0.2%	-0.5%			
$\widetilde{el_t}$	-8%	3%	**5%**	-6%	-4%	**10%**	-7%	-4%	**11%**	-10%	-2%	**12%**	-6%	-4%	**10%**	-6%	-3%	**9%**
l_t^e	-6%	2%	**4%**	-6%	-4%	**10%**	-2%	-1%	**3%**	-6%	-1%	**7%**						
$Tmin_t$							**1%**	0%	-1%	**1%**	0%	-1%						
$Tmax_t$																		

Table B-7: This table is used to show the mean marginal effect, $M_{y,\tau,j}$ in year (2005), y during morning (m), daytime (d) and evening (e) sub-periods, τ for both dispatch periods, j; the day before dispatch ($t-1$, 15:00, *PD1*) and on the dispatch day (t) at 6am (*PD2*).

Further explanation on this table is similar to the caption of Table B-1

	$M_{2006,m,j}$						$M_{2006,d,j}$						$M_{2006,e,j}$					
	$j = PD1$			$j = PD2$			$j = PD1$			$j = PD2$			$j = PD1$			$j = PD2$		
i	0	1	2	0	1	2	0	1	2	0	1	2	0	1	2	0	1	2
$\Pr(S_{t,\tau\|j} = i)$	10%	21%	**70%**	11%	37%	**52%**	21%	23%	**56%**	20%	27%	**53%**	20%	21%	**59%**	20%	23%	**58%**
D_t^{wk}	-1%	-2%	**3%**				-3%	-1%	**4%**	-2%	-2%	**4%**						
D_t^{win}													**2%**	1%	-3%	**2%**	1%	-3%
D_t^{sum}							**2%**	0%	-2%	1%	1%	-2%	**3%**	1%	-4%	**3%**	2%	-5%
$D_{t,\tau-1\|j}^0$	**19%**	14%	-33%	**17%**	6%	-23%	**30%**	7%	-37%	**36%**	0%	-36%	**27%**	7%	-34%	**27%**	6%	-33%
$D_{t-1,\tau\|j}^0$	2%	2%	-4%				**2%**	0%	-2%				**3%**	1%	-3%	**3%**	2%	-5%
$D_{t,\tau-1\|j}^2$	-18%	-22%	**40%**	-16%	-20%	**36%**	-29%	-28%	**57%**	-25%	-27%	**52%**	-24%	-13%	**37%**	-25%	-20%	**45%**
$D_{t-1,\tau\|j}^2$	-2%	-2%	**4%**	-3%	-3%	**6%**	-1%	-1%	**2%**	-2%	-1%	**3%**	-4%	-1%	**5%**	-3%	-2%	**5%**
$P_{t,\tau\|j}$	1%	1%	-2%										**1%**	0%	-1%	**0.1%**	0%	-0.1%
$\widetilde{el}_t$	-3%	-3%	**6%**	-4%	-5%	**9%**	-9%	-2%	**11%**	-8%	-6%	**14%**	-5%	-1%	6%			
l_t^e	-3%	-3%	**6%**				-4%	-1%	**5%**				-5%	-1%	6%			
$Tmin_t$	-1%	-1%	**2%**	-2%	-2%	**4%**	**1%**	0%	-1%									
$Tmax_t$							**0.4%**	0.1%	-0.5%	**0.3%**	0.2%	-0.5%						

Table B-8: This table is used to show the mean marginal effect, $M_{y,\tau,j}$ in year (2006), y during morning (m), daytime (d) and evening (e) sub-periods, τ for both dispatch periods, j; the day before dispatch ($t-1$, 15:00, *PD1*) and on the dispatch day (t) at 6am (*PD2*).

Further explanation on this table is similar to the caption of Table B-1

	$M_{2007,m,j}$						$M_{2007,d,j}$						$M_{2007,e,j}$					
	$j = PD1$			$j = PD2$			$j = PD1$			$j = PD2$			$j = PD1$			$j = PD2$		
i	0	1	2	0	1	2	0	1	2	0	1	2	0	1	2	0	1	2
$\Pr(S_{t,\tau\|j} = i)$	10%	17%	**73%**	17%	37%	**46%**	25%	19%	**56%**	25%	28%	**48%**	27%	20%	**53%**	24%	23%	**53%**
D_t^{wk}	-5%	**3%**	2%	-3%	-2%	**5%**	-3%	0%	**3%**	-3%	0%	**3%**	-2%	0%	**2%**			
D_t^{win}	**5%**	-3%	-2%				**2%**	0%	-2%									
D_t^{sum}							**6%**	0%	-6%	**7%**	0%	-7%	**2%**	0%	-2%			
$D_{t,\tau-1\|j}^0$	**10%**	-6%	-4%	**17%**	6%	-23%	**32%**	1%	-33%	**35%**	2%	-37%	**31%**	4%	-35%	**32%**	3%	-35%
$D_{t-1,\tau\|j}^0$	**3%**	-2%	-1%	**3%**	2%	-5%							**5%**	0%	-5%	**2%**	1%	-3%
$D_{t,\tau-1\|j}^2$	-22%	**16%**	6%	-20%	-15%	**35%**	-31%	-2%	**33%**	-27%	-5%	**32%**	-31%	-10%	**41%**	-28%	-16%	**44%**
$D_{t-1,\tau\|j}^2$	-3%	**2%**	1%	-3%	-2%	**5%**	-3%	0%	**3%**	-3%	0%	**3%**				-2%	-1%	**3%**
$P_{t,\tau\|j}$	**0.4%**	-0.2%	-0.2%				**0.4%**	0%	-0.4%	**0.3%**	0%	-0.3%	**0.1%**	0%	-0.1%			
$\widetilde{el_t}$	-5%	**3%**	2%	-10%	-5%	**15%**	-8%	0%	**8%**	-11%	-1%	**12%**	-6%	-1%	**7%**			
l_t^e	-7%	**4%**	3%	-5%	-3%	**8%**	-6%	0%	**6%**	-8%	0%	**8%**	-5%	0%	**5%**			
$Tmin_t$				-1%	-1%	**2%**												
$Tmax_t$																		

Table B-9: This table is used to show the mean marginal effect, $M_{y,\tau,j}$ in year (2007), y during morning (m), daytime (d) and evening (e) sub-periods, τ for both dispatch periods, j; the day before dispatch ($t-1$, 15:00, *PD1*) and on the dispatch day (t) at 6am (*PD2*).

Further explanation on this table is similar to the caption of Table B-1

Appendix C

Variable	Description	Source
$P_{t,\tau}$	The actual wholesale market spot price outcome for the τth sub-period on day t	http://www.aemo.com.au/Electricity/Data/Price-and-Demand/Aggregated-Price-and-Demand-Data-Files
$P_{t,\tau\|j}^{+\Delta}$, $P_{t,\tau\|j}^{-\Delta}$	The pre-dispatch price sensitivity prices between +/- 200 MWh scenarios on the day before dispatch ($t-1$, 15:00, *PD1*) and on the dispatch day (t) at 6am (*PD2*)	http://www.aemo.com.au/Electricity/Market-Operations/Dispatch/PreDispatch-Sensitivities http://www.aemo.com.au/Electricity/Data/Market-Management-System-MMS#Predispatch
$D_t^{wk}, D_t^{win}, D_t^{sum}$	These are dummy variables that takes a value of 1 when day t is weekends, winter or summer and 0 otherwise	
l_t^e, $\widetilde{el_t}$	Both of these variables are calculated from the actual load on day t, l_t	http://www.aemo.com.au/Electricity/Data/Price-and-Demand/Aggregated-Price-and-Demand-Data-Files
$Tmax_t$ and $Tmin_t$	Both of these variables are calculated from the actual maximum and minimum temperature on day t, $T_{max_t}^{obs}$ and $T_{min_t}^{obs}$	http://www.bom.gov.au/climate/data/ Station number for NSW is 66062

Table C-1: This table is used to show details on the source of data used in Chapter 3

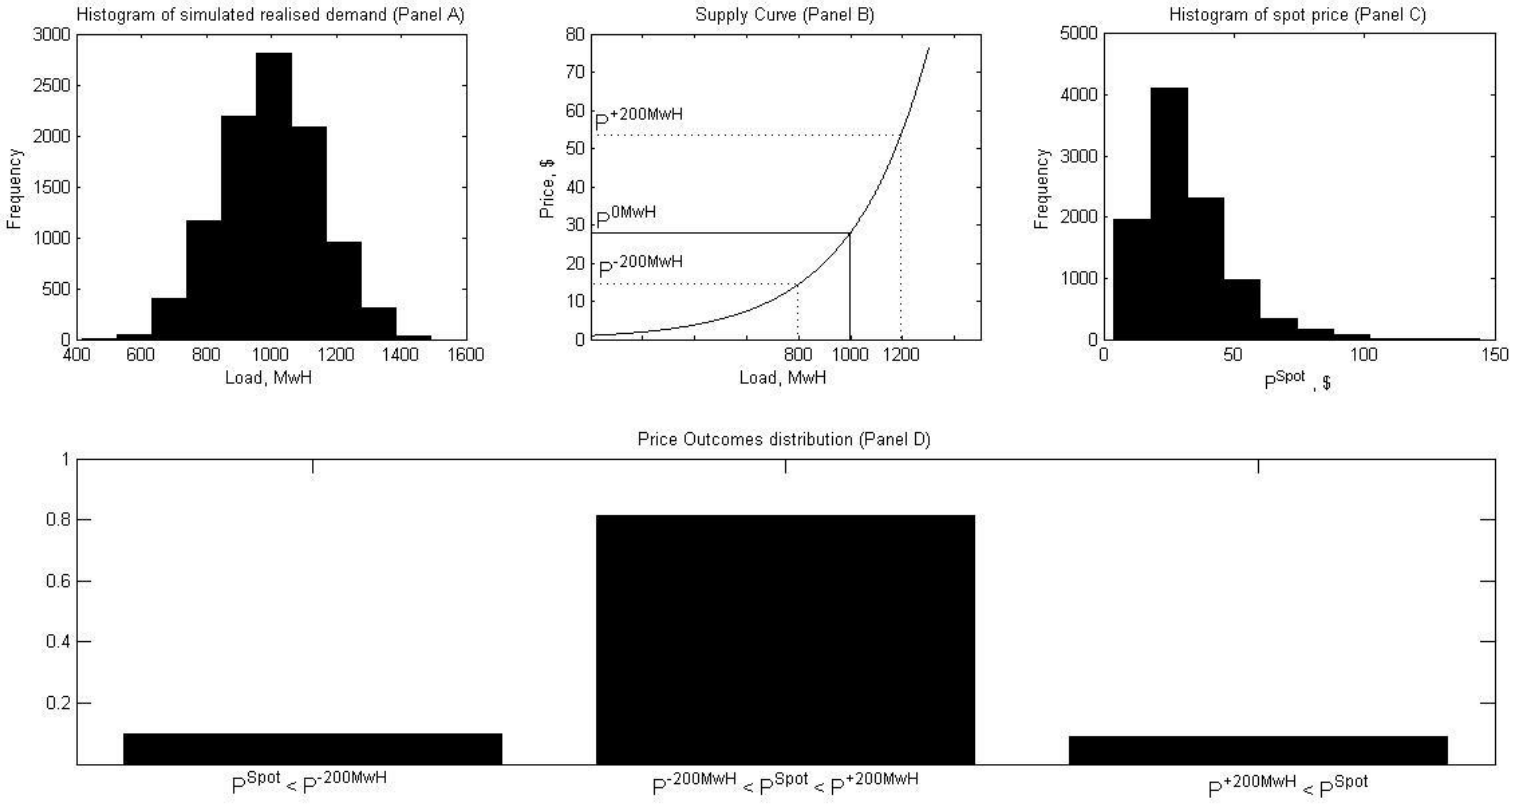

Figure 3.1-1: In panel A to illustrate the deviations of actual demand from the demand forecast is symmetric, we add a variation drawn from normally distributed pseudorandom numbers to demand forecast which is set to be 1000MwH and plot their histogram. An electricity supply curve is illustrated in panel B to show their influence on the skewness of the spot price distribution in panel C, even though these spot prices are based on symmetric demand distribution assumption. Hence, with this assumption, we expect an equal numbers of these spot prices (P^{Spot}) from the symmetric distributed demand to appear below $P^{-200MwH}$ and above $P^{+200MwH}$ resulting to the distribution of the discrete price outcomes to be symmetric (in panel D).

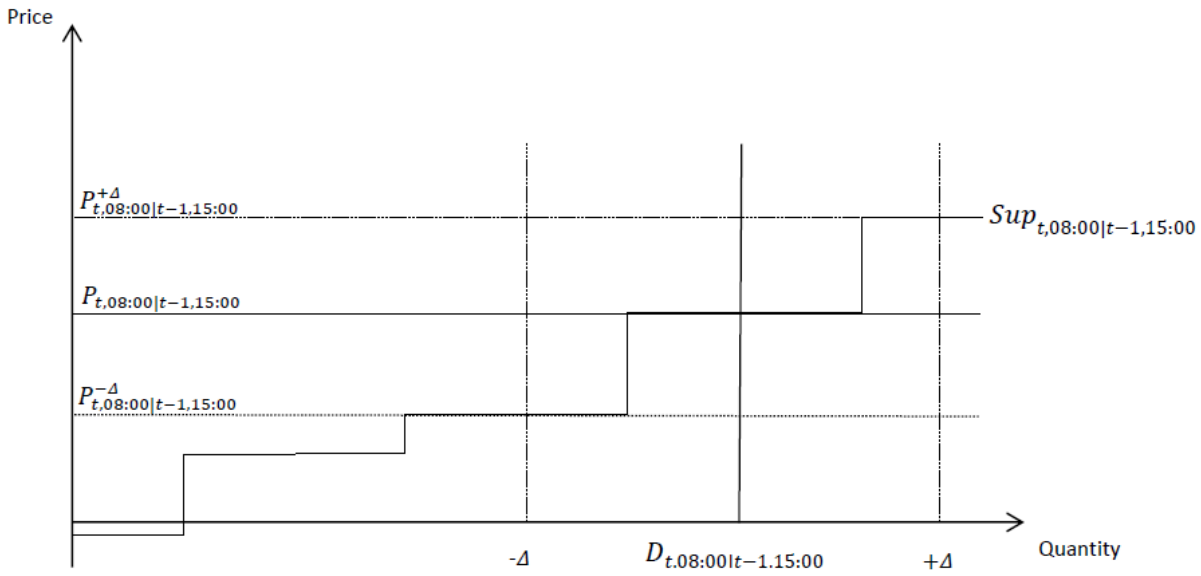

Figure 3.2-1: This figure is used to show the optimisation process during pre-dispatch and at dispatch.

In the figure, $Sup_{t,08:00|t-1,15:00}$ is a stylised supply schedule for dispatch period 08:00 on day t as produced from the 15:00 pre-dispatch on day $t-1$ and $D_{t,08:00|t-1,15:00}$ is the demand forecast (for the 08:00 period on day t) as of 15:00 on day $t-1$. The intersection of the supply and demand curves resulting to pre-dispatch price denoted by $P_{t,08:00|t-1,15:00}$ on the y-axis.

Δ on the x-axis is used to represent the demand variation considered in the sensitivity scenarios. In the case of the NSW electricity market, Δ is equivalent to 200 MWh. Therefore, if the demand forecast increases/decreases by Δ, this means the pre-dispatch price changes to $P^{+\Delta}_{t,08:00|t-1,15:00}$ or $P^{-\Delta}_{t,08:00|t-1,15:00}$ (as shown in the x-axis). $P^{+\Delta}_{t,08:00|t-1,15:00}$and $P^{-\Delta}_{t,08:00|t-1,15:00}$ are also known as pre-dispatch scenarios prices.

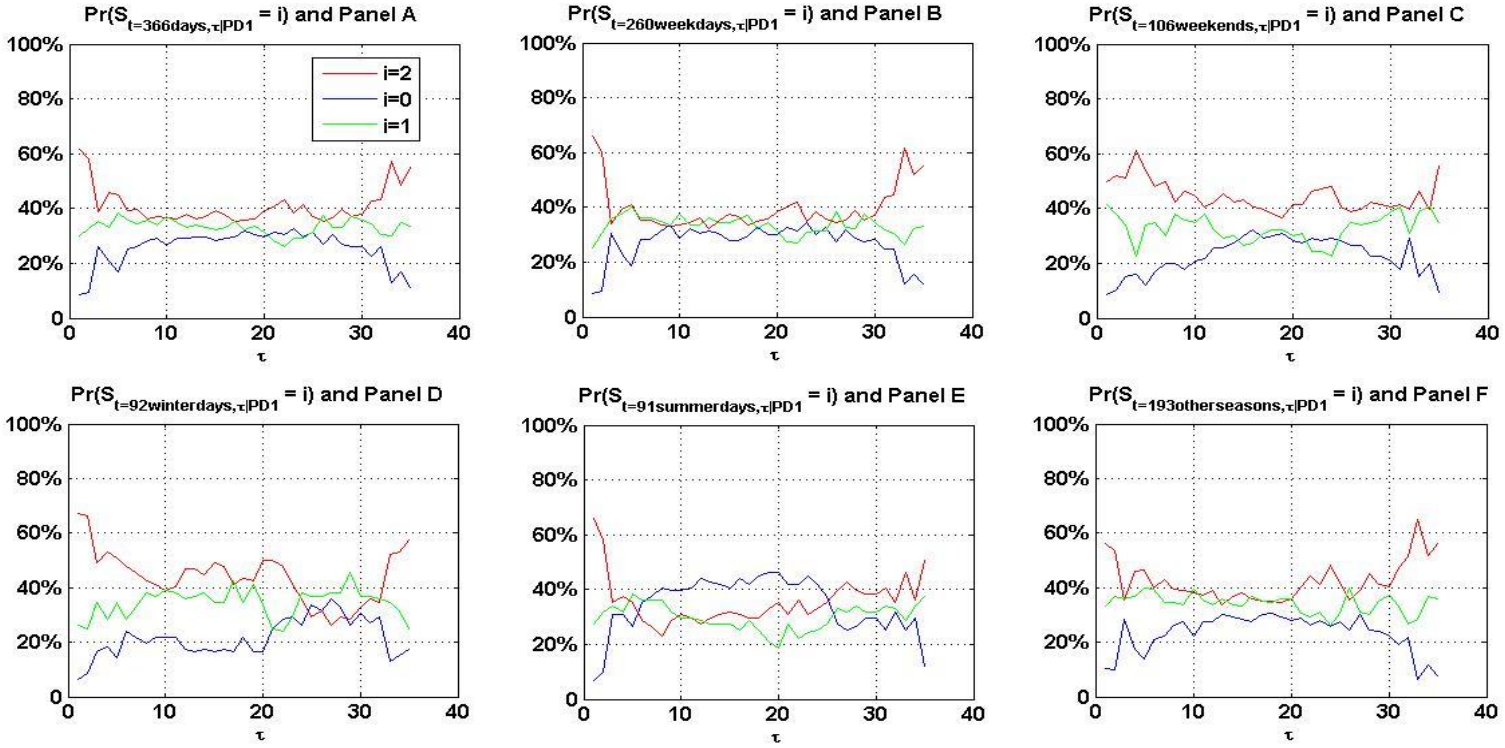

Figure 3.2-2: In Panel A, we report the empirical distribution for $S_{366days,\tau|PD1}$ in year 2000 by plotting their empirical probabilities, $\Pr(S_{366days,\tau|PD1} = i)$ in the y-axis as a percentage with respect to 35 half-hour in a day, $\tau = 06{:}30, \ldots, 23{:}30$ (1st half-hour,…,35th half-hour). The $\Pr(S_{366days,\tau|PD1} = \{0,1,2\})$ with respect to 35 half-hour in a day are as labelled by the legend of the graphs. To calculate $\Pr(S_{366days,06:30|PD1} = 0)$

$$\Pr(S_{366days,06:30|PD1} = 0) = \left(\frac{\sum_{t=1}^{366} I_{S_{t,06:30|PD1}=0}}{366}\right) \times 100\%$$

$$I_{S_{1day,06:30|PD1}=0} = \begin{cases} 1 & if\ S_{t,06:30|PD1} = 0 \\ 0 & otherwise \end{cases}$$

In Panel B, we are only looking at the empirical distributions of $S_{t,\tau|PD1}$ during weekdays, $\Pr(S_{260\ weekdays,\tau|PD1} = i)$, panel C during weekends, $\Pr(S_{106\ weekends,\tau|PD1} = i)$, in panel D during winter, $\Pr(S_{92\ winter,\tau|PD1} = i)$, in Panel E during summer $\Pr(S_{91summer,\tau|PD1} = i)$ and in Panel F during other seasons $\Pr(S_{192\ other\ seasons,\tau|PD1} = i)$.

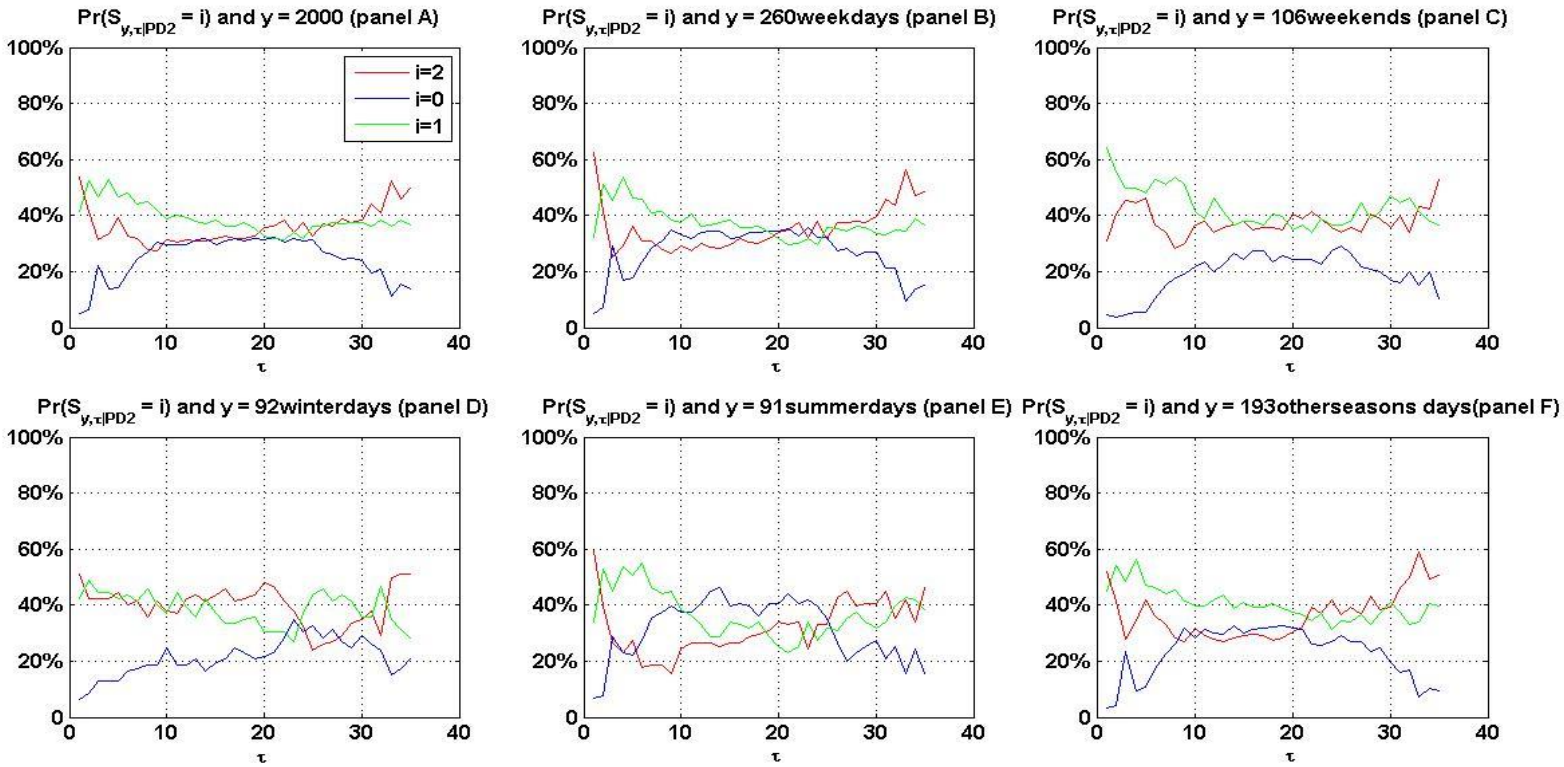

Figure 3.2-3: In Panel A, we report the empirical distribution for $S_{366days,\tau|PD2}$ in year 2000 by plotting their empirical probabilities, $\Pr(S_{366days,\tau|PD2} = i)$ in the y-axis as a percentage with respect to 35 half-hour in a day, $\tau = 06{:}30, \ldots, 23{:}30$ (1st half-hour,…,35th half-hour). The $\Pr(S_{366days,\tau|PD2} = \{0,1,2\})$ with respect to 35 half-hour in a day are as labelled by the legend of the graphs. To calculate $\Pr(S_{366days,06:30|PD2} = 0)$

$$\Pr(S_{366days,06:30|PD2} = 0) = \left(\frac{\sum_{t=1}^{366} I_{S_{t,06:30|PD2}=0}}{366}\right) \times 100\%$$

$$I_{S_{1day,06:30|PD2}=0} = \begin{cases} 1 & if\ S_{t,06:30|PD2} = 0 \\ 0 & otherwise \end{cases}$$

In Panel B, we are only looking at the empirical distributions of $S_{t,\tau|PD2}$ during weekdays, $\Pr(S_{260\ weekdays,\tau|PD2} = i)$, panel C during weekends, $\Pr(S_{106\ weekends,\tau|PD2} = i)$, in panel D during winter, $\Pr(S_{92\ winter,\tau|PD2} = i)$, in Panel E during summer $\Pr(S_{91summer,\tau|PD2} = i)$ and in Panel F during other seasons $\Pr(S_{192\ other\ seasons,\tau|PD2} = i)$.

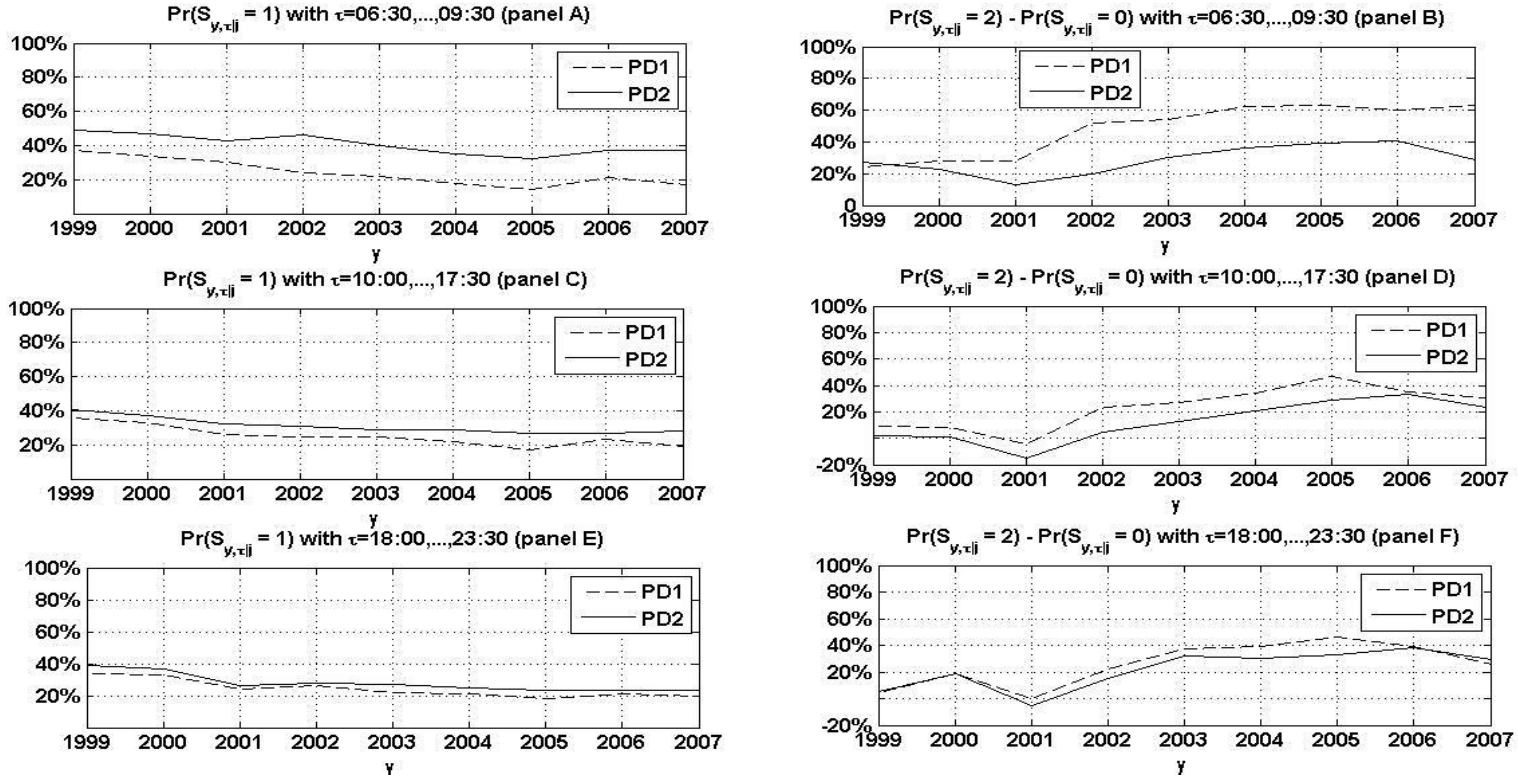

Figure 3.2-4: In panel A, we plotted a yearly unconditional probability of the spot price outcomes that land between the two pre-dispatch scenario prices, $\Pr(S_{t,\tau|j} = 1)$ from 1 March 1999 to 31 October 2007 in morning sub-period (06:30 to 09:30 dispatch intervals, $\tau \in m$). In this figure the yearly unconditional probability are reported for both dispatch periods; the day before dispatch ($t-1$, 15:00, *PD1*) and on the dispatch day (t) at 6am (*PD2*). The respective yearly unconditional probability for *PD1* and *PD2* are as labelled in the legend. As follows; panel C show the $\Pr(S_{t,\tau|j} = 1)$ in daytime sub-period (10:00 to 17:30, $\tau \in d$) and panel E in evening sub-period (18:00 to 23:30, $\tau \in e$).
While Panel B, D and F, the yearly unconditional probability, $\mathrm{P}(S_{t,\tau|j} = 2) - \mathrm{P}(S_{t,\tau|j} = 0)$ are giving an indication of the bias to be placed in the $S_{t,\tau|j} = 2$ according to the sequence of sub-periods. The larger this value the more does the pre-dispatch process underestimate the actual spot price outcome.

	$M_{1999,m,PD1}$			$M_{1999,m,PD2}$		
i	0	1	2	0	1	2
$\Pr(S_{t,\tau\|j} = i)$	20%	37%	**44%**	12%	**49%**	39%
D_t^{wk}	-7%	-1%	**8%**	-3%	-2%	**5%**
D_t^{win}	**8%**	1%	-9%	**2%**	1%	-3%
D_t^{sum}						
$D_{t,\tau-1\|j}^{0}$	**23%**	3%	-26%	**21%**	-1%	-20%
$D_{t-1,\tau\|j}^{0}$	**5%**	1%	-6%	**5%**	2%	-7%
$D_{t,\tau-1\|j}^{2}$	-21%	-6%	**27%**	-13%	-18%	**31%**
$D_{t-1,\tau\|j}^{2}$	-7%	-1%	**8%**	-6%	-5%	**11%**
$P_{t,\tau\|j}$	**2%**	0%	-2%			
$\widetilde{el_t}$	-9%	-1%	**10%**			
l_t^e	-11%	-2%	**13%**			
$Tmin_t$	-2%	0%	**2%**	-1%	-1%	**2%**
$Tmax_t$						

Table 3.6-1: This table is used to show the mean marginal effect, $M_{y,\tau,j}$ in year (1999), y during morning (m) sub-period, τ for both dispatch periods, j; the day before dispatch ($t-1$, 15:00, *PD1*) and on the dispatch day (t) at 6am (*PD2*). The unconditional probability of *PD1* and *PD2* for each of the spot price outcomes; $\Pr(S_{t,\tau|j} = \{0,1,2\})$ are reported in the second row. While the mean marginal effects of the explanatory variables listed in first column are display in the remaining rows. $D_t^{wk}, D_t^{win}, D_t^{sum}$ are the dummy variables giving value of 1 when day t is weekends, winter, summer or 0 otherwise. $\widetilde{el_t}$, is de-seasonalised load forecast error and l_t^e is load forecast as discussed in Section 3.4.2. $Tmax_t$ and $Tmin_t$ are unseasonably temperature as discussed in Section 3.4.3. $D_{t,\tau-1|j}^{2}, D_{t-1,\tau|j}^{2}, D_{t,\tau-1|j}^{0}$ and $D_{t-1,\tau|j}^{0}$ are dummy variables which indicate whether one trading period prior ($D_{t,\tau-1|j}^{i}$) or one day prior ($D_{t-1,\tau|j}^{i}$) the spot price outcome was placed either above the upper scenario ($i = 2$) or below the lower scenario price ($i = 0$). And lastly, $P_{t,\tau|j}$ is the expected pre-dispatch price calculated using equation (3.4-4). The mean marginal effects for continuous and binary variables are calculated based on equation (3.5-3) and (3.5-4) respectively. All the mean marginal effects displayed are based on the significant estimated parameters, β and α at 5% level. The insignificant ones are not shown (blank cell entries). Only the highest positive mean marginal effects are bolded.

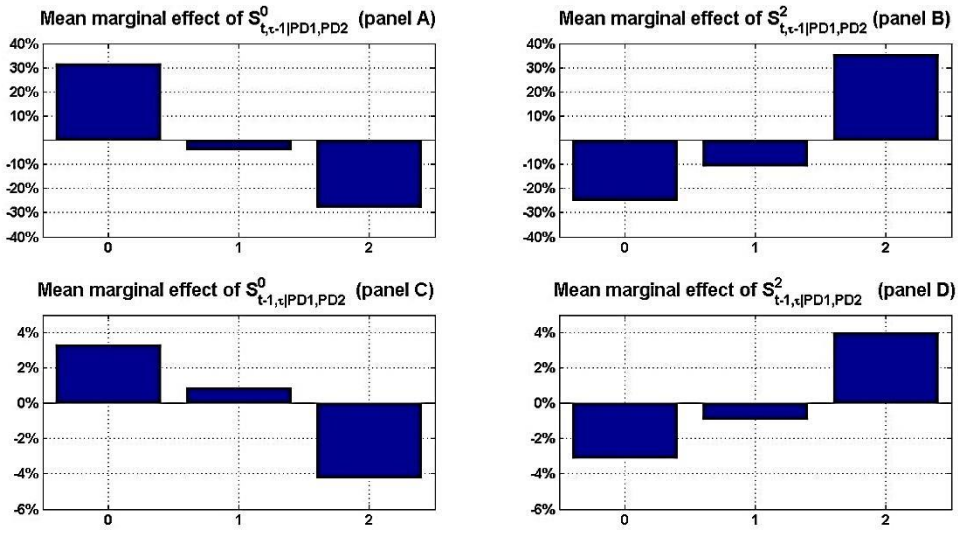

Figure 3.6-1: These are the mean marginal effects of $S^i_{t,\tau-1|j}$ (panel A and B) and $S^i_{t-1,\tau|j}$ (panel C and D) averaged across two dimensions; the years and all trading intervals and for both *PD1* and *PD2*.

On average, the probability of $S_{t,\tau|j} = 0$ increases by 31% if the spot price outcomes in the previous half-hour is $S_{t,\tau-1|j} = 0$ (in panel A). To calculate the probability, we take an average of all the mean marginal effects of $S^0_{t,\tau-1|PD1}$ and $S^0_{t,\tau-1|PD2}$ from 1999 to 2007 for all the subperiods. Note that the corresponding decreases in probabilities for the central and upper categories (-3% and -28%) ensure that the sum of all these three changes is 0. By construction the sum of all marginal effects equal to 0 as at any stage the sum of all the conditional probabilities, $P\left(S_{t,\tau|j} = i \middle| \boldsymbol{x}_{t,\tau|j}; \boldsymbol{\beta}_{\tau|j}, \boldsymbol{\rho}_{\tau|j}\right)$ have to sum to 1.

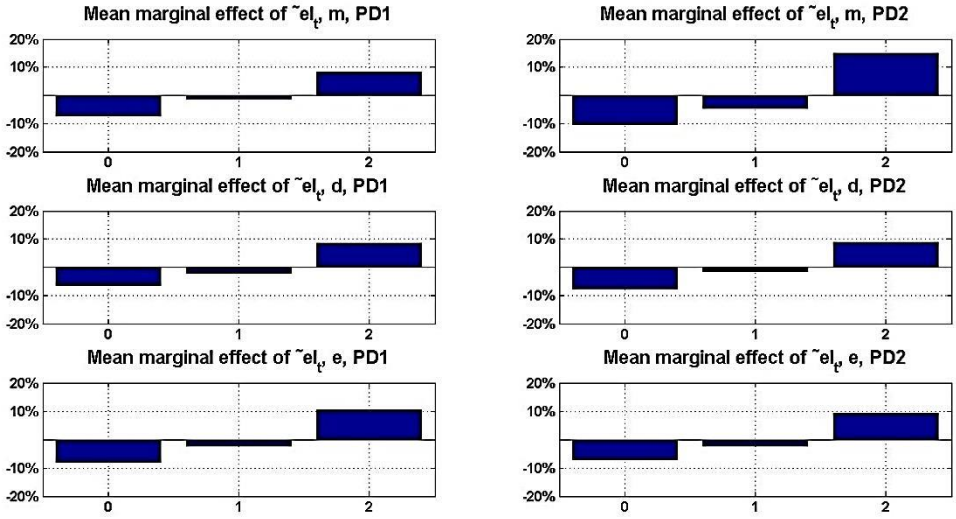

Figure 3.6-2: These are the mean marginal effects of $\widetilde{el_t}$ averaged across the years. The left-side is based on *PD1* while the right side is based on *PD2*. For mean marginal effect of $\widetilde{el_t}, m, PD1$ in the first panel, we take an average of all the mean marginal effects of $\widetilde{el_t}, m, PD1$ from 1999 to 2007.

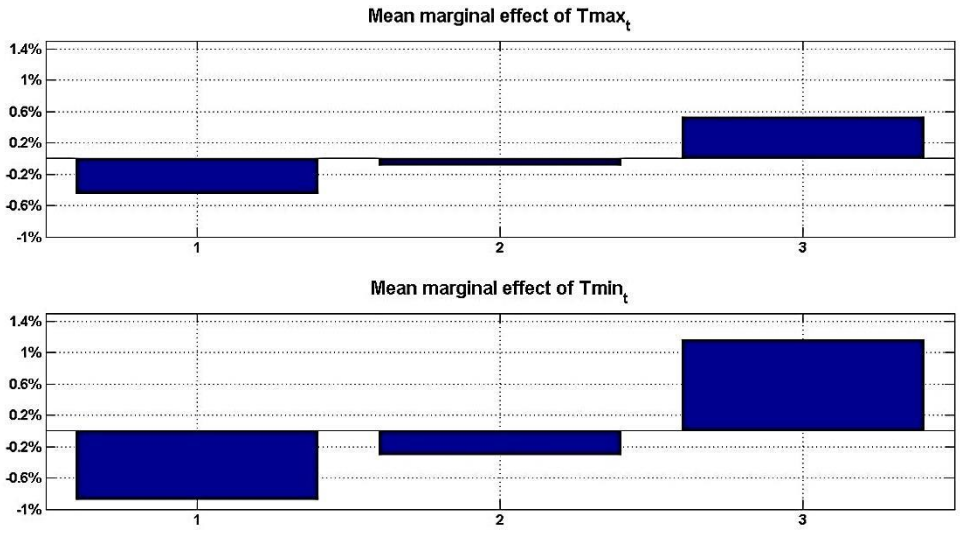

Figure 3.6-3: These are the mean marginal effects of $Tmax_t$ and $Tmin_t$ averaged across the years and all trading intervals for both *PD1* and *PD2*. For these mean marginal effects, we take an average of all the mean marginal effects of $Tmax_t$ and $Tmin_t$ from 1999 to 2007 for all the sub-periods.

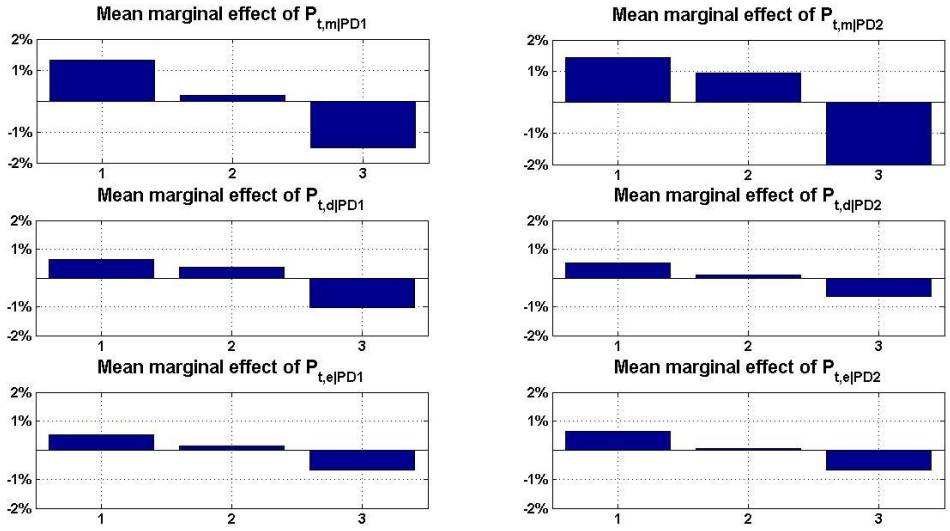

Figure 3.6-4: These are the mean marginal effects of $P_{t,\tau|j}$ averaged across the years. The left-side is based on *PD1* while the right side is based on *PD2*. For mean marginal effect of $P_{t,\tau|j}, m, PD1$ in the first panel, we take an average of all the mean marginal effects of $P_{t,\tau|j}, m, PD1$ from 1999 to 2007.

Chapter 4

4. Does information from the pre-dispatch process help predicting price spikes in the Australian Electricity Market?

ABSTRACT

The pre-dispatch process is an important element of Australia Electricity Market operation, intended to provide market participants with useful and accurate information for next-day's trading period. It is the purpose of this chapter to investigate whether information from this pre-dispatch process can be useful when predicting next-day price spikes. In a preliminary analysis we establish the effect of pre-dispatch prices on the quantiles of the spot price distribution. A Quantile regression approach reveals that higher pre-dispatch prices signal only to a certain extent an increased probability of higher spot price outcomes. They also signal a higher uncertainty about the resulting spot price outcomes. We further establish whether the inclusion of information from the pre-dispatch process can significantly improve the predictability of price spikes when these are modelled as a point process (as in Chapter 2). The models used here are Hawkes and Poisson Autoregressive Models which allow for time variation (correlated to exogenous information) in the intensity process that governs the occurrence of price spikes. It transpires that the pre-dispatch process of the Australian Electricity Market does not provide any information that can be used in a systematic manner to help predicting on what days price spikes are more likely to occur.

4.1. Introduction

Until the early 1990s the electricity sector has been a vertically integrated industry, often under state control, in virtually all countries. Regulators were responsible to fix prices and did so considering the costs of generation, transmission and distribution costs. This central price setting process resulted in fairly stable electricity prices.

In recent years electricity markets have been deregulated in many countries. Deregulation aims to introduce competition in generation and supply activities, with transmission and distribution often still regarded as a natural monopoly. This is done by unbundling the generation activity from transmission and distribution as a separate business. While this unbundling is common to almost all deregulation processes, the process which determines the market price of electricity differs substantially between countries[29]. Most deregulated markets are organised by power pools and managed by an independent electricity operator (see eg. (Weron 2007)). In such markets electricity suppliers are required to submit offers to sell electricity while the wholesale buyers submit bids to buy electricity.

An independent electricity operator receives these bids and the market prices will be determined through some pricing algorithm that matches the resulting supply with the system demand in the most cost efficient way. The market price is then the cost of the marginal (most expensive) unit of electricity used in this matching process.

Electricity is a homogenous commodity with physical characteristics unlike other commodities. In order to maintain the stability of the electricity network the market needs to clear at any moment in time. The 'non-storability' of electricity means that inventories cannot be used to 'smooth' out supply[30]. This causes substantial intra-day price seasonality, mainly driven by the typical daily electricity demand pattern. As retail prices are regulated it is typical to find a very low price elasticity of demand. This, in combination with a highly convex supply curve, results in a very volatile price process.

In fact, the electricity price time-series exhibits more extreme and more frequent price spikes than even financial data. In the electricity market price spikes are large upward jumps of electricity prices which will usually return to the normal price level very quickly.

[29] In most countries the retail market is still heavily regulated. In what follows "electricity market", if not made explicit otherwise, refers to the wholesale electricity market.
[30] Pump-storage reservoirs can, to a certain extent, be seen as inventories.

Due to the complex empirical properties of electricity prices (seasonality, time-varying volatility, jumps) in deregulated markets the process of modelling these in an econometric model has proven challenging. Consequently forecasting electricity prices is at least equally as complex a task (Bunn 2000; Song & Wang 2003). In this chapter we are interested in a particular aspect of the price process, the sudden and very large jumps of the spot price, also known as "price spikes". The Federal Energy Regulatory Commission (1998) concluded that these price events can mainly be attributed to unusually hot weather, equipment outages, transmission congestion and retail price inflexibility.

Figure 4.1-1 shows the time-series of daily average prices (on a logarithmic scale) on the New South Wales[31] wholesale electricity market. It is immediately apparent that price spikes are a prominent feature of this time-series. Price spikes in this particular market feature prices that are several orders of magnitude higher than normal price levels, jumping back down swiftly to a normal price range. While the above features of the electricity market can be made partly responsible for the prevalence of price spikes in the Australian National Electricity Market (NEM), Kwoka (2012) conjectures that a fair proportion of these price spikes are due to strategic bidding behaviour.

(INSERT Figure 4.1-1 HERE)

Kwoka (2012) discusses situations in which opportunities for strategic bidding arise. The most obvious is when an individual supplier is critical in order to ensure that the electricity demand can be met by the supply side. In such situations this suppliers has an incentive and the ability to manipulate the market clearing price by potentially withholding generation capacity. These situations are more likely to arise when the market demand gets close to the generation capacity. In the absence of effective behavioural controls on electricity supplier's bidding behaviour this is likely to contribute to the high frequency and size of price spikes in the NEM.

Extraordinary price events represent a significant source of risk to retailers. In deregulated market, retailers buy electricity from a grid at spot price but are required to sell them to end user at a heavily regulated price. Hence, it is important for retailers to predict (if and where this is possible) these extreme events and hedge them accordingly.

[31] The New South Wales region is one of several, connected regions in the Australian Electricity Market.

One element in the setup of the NEM that is meant to increase market transparency and therefore to facilitate the prediction of potential price events is the pre-dispatch process. In this process electricity suppliers provide the Australian Energy Market Operator (AEMO) with their day-ahead supply curves. The market operator will aggregate these to a market supply curve and match it with its best possible demand forecast. The result of this exercise is the market operator's prediction of the next day's wholesale market price.

It is the main aim of this chapter to incorporate information from this pre-dispatch process into price spike models. To the best of our knowledge no such attempt has been made in the literature. Methodologically our chapter follows those of Christensen et al. (2012), Christensen et al. (2009) and Chapter 2 in which price events are defined as a binary series or a point process. These papers investigate whether the occurrence of the price events can be predicted through variation in exogenous variables (such as meteorological or load information) and the history of the binary event series features itself. The contribution of this chapter is therefore the inclusion of pre-dispatch information into the relevant information sets.

In a preliminary analysis we investigate some fundamental features of the relationship between pre-dispatch prices and the wholesale market price outcomes. In this preliminary analysis, which consists of a careful analysis of correlations and quantile regressions, we find evidence that pre-dispatch prices are a surprisingly week predictor of price outcomes. This allows the conjecture that electricity generators have the tendency to change their supply bids substantially between their initial submission and the actual dispatch. In the light of these findings it is not surprising to find that the pre-dispatch information does not have any substantial explanatory power for price spikes.

The remainder of this chapter is structured as follows; in Section 4.2 we discuss the state of the literature in modelling the time-series of wholesale market electricity prices, putting a particular emphasis on those papers that focus on price spikes. The methodology used to investigate the relationship between pre-dispatch prices and the wholesale market price outcomes and the point process models used to predict price spikes are discussed in Section 4.3. A brief description on the structure of Australia National Electricity Market (NEM), pre-dispatch process and the conditioning covariates is presented in Section 4.4. Later in Section 4.5, we discuss the result from our preliminary analysis, the estimation

and forecasting exercise using the point process models and conditioning covariates discuss previously. Finally in Section 4.6, we conclude this chapter.

4.2. Literature review

Since wholesale electricity markets have been deregulated modelling and forecasting wholesale electricity prices has attracted a considerable amount of attention. Understanding the behaviour of future electricity prices is particularly important for market participants which are, due to the market design, exposed to large price variations. In the case of the NEM these are electricity retailers as they have to purchase electricity from the deregulated wholesale market but sell this electricity on a regulated retail market. Price distribution forecasts will inform the hedging strategies of these retailers. On the other side, Arciniegas & Arciniegas Rueda (2008) argue that electricity suppliers with a superior understanding of future market conditions and prices can implement more profitable bidding strategies.

Due to the complexity of the electricity price process (such as seasonality, skewness, high and clustering volatility as well as the presence of jumps[32]) the modelling and forecasting process is not straightforward and many different methodologies have been proposed. Most of the studies that proposed approaches to model electricity prices focused on modelling the trajectory of the spot price or its return across time and by so doing attempt to deal with most or all of the above properties (Amjady & Keynia 2008; Geman & Roncoroni 2006; M. Barlow 2002; Mount et al. 2006; Bunn 2000; De Jong & Huisman 2002; Amjady & Keynia 2009). The methodology adopted by these studies include time series based models such as; traditional autoregressive time series models, nonlinear time series models such as Markov-Switching models, continuous-time diffusion or jump-diffusion models and artificial intelligence models such as neural networks, structural models, machine learning models, and hybrid models (see Weron (2007) for discussions on a range of these models).

In this chapter we focus on one particular aspect of the electricity price distribution, the so called price spikes. There is no particular definition of what constitutes a price spike, but for the purpose of this chapter we follow the lead of Christensen et al. (2009) who define price spikes as a binary variable defined by whether the wholesale price exceeds some price threshold. Christensen et al. (2009) found evidence of significant persistence in the

[32] These features have been documented in a number of papers. See inter alia (Wu & Shahidehpour 2010; Aggarwal et al. 2008; Catalão et al. 2007; Nogales & Contreras 2002; Duan n.d.; Burger, Klar, M ller, et al. 2004; Conejo et al. 2005).

occurrence of such price spikes. This feature is commonly neglected by most traditional price series models since they, if they allow for such a feature, treat these as a memory-less jump process with an intensity which is independent of its own history.

In order to rectify this shortcoming more recent studies turn their attention to the feature of price spikes and attempt to model these either directly or indirectly. The latter category consists of a range of papers who recognise that the electricity price process can come from distinctly different distributions. One of these distributions is seen as generating extreme price events. Kanamura & Ōhashi (2008), Mount et al. (2006) and Becker et al. (2004) all propose to allow for time-varying transition probabilities in a regime switching model that allows for electricity prices to be drawn from different distribution. These papers identify variables which can explain systematic variations in the implied transition probabilities and therefore allow the practitioner to improve her ability to predict whether next day's electricity price is likely to come from the regime that is more likely to produce price spikes.

More recently, the issue of price spikes has been tackled in a more direct manner by defining a binary series that identifies the instances in which the electricity price exceeds a certain price threshold[33] and subsequently modelling the dynamics of this series. Research papers in this vein are in Chapter 2, Christensen et al. (2012), Christensen et al. (2009), Clements et al. (2012) and Eichler et al. (2012). It is the essence of these papers that they model and forecast the probability of price spikes (as defined through the binary series). The papers differ in the precise way in which they do model this binary series. To the best of our knowledge the first study using this approach was Christensen et al. (2009) who used a modified Poisson autoregressive (PAR) framework to forecast next day's price spike occurrence. In Chapter 2 we tackle the same task using Hawkes models. Both approaches essentially allow for the probability of a price spike occurring to vary with (weakly) exogenous covariates and to display persistence. The latter has been recognised as an important stylized fact when analysing the occurrence of electricity price spikes on a daily basis[34]. Any sensible model should be able to replicate this feature of the data.

[33] A related paper is by Zhao et al. (2005). They also define price spikes as a binary series, but do allow for a time varying threshold price. While there is some appeal to their logic, this approach complicates the modelling and forecasting exercise significantly, as one needs to forecast the threshold price and the probability of the actual price exceeding that level. It is our view that fixing the threshold allows for a better focus on the dynamics of the probability of exceeding that threshold.

[34] See Section 4.3 for details on how these papers derived a daily binary series from the half-hourly series of electricity prices.

While the papers above focus on a derived daily price spike series, the papers by Christensen et al. (2012) and Eichler et al. (2012) model a half-hourly, binary price spike series. The approaches chosen in this context are the Autoregressive conditional hazard model (Christensen et al. 2012) and a dynamic logistic regression model (Eichler et al. 2012). By modelling the half-hourly series these papers need to tackle the thorny issue of significant intra-daily seasonality. This makes the two problems of forecasting electricity price spikes half an hour ahead or a day ahead two significantly different problems. In this chapter we will focus on the one day ahead forecasting.

Our intention is not to introduce a new forecasting method for the daily price spike series, but to analyse whether information from the pre-dispatch process is able to improve on forecast of price spikes. The type of information that was used in the previously discussed papers is either information on the system load or meteorological information like temperatures. To the best of our knowledge, the existing literature (apart from two papers by Zareipour & Bhattacharya (2006) and Hamidreza Zareipour et al. (2006b)) does not consider information from any pre-dispatch process in their forecasting models. Details of the pre-dispatch process and its importance for price spike modelling are discussed in Section 4.4.

4.3. Methodology

The wholesale electricity price in the NEM is settled as the price of the most expensive unit of electricity that is needed to satisfy the current level of demand (load). This is done every five minutes, but AEMO calculates an average of all five minutes intervals in a half hour to arrive at a half-hourly series of wholesale electricity prices. This is the price paid for all units of electricity supplied in that half hour.

As we are interested in price spikes only we need to define such events. Let $P_{t,\tau}$ be the electricity spot price for the τ-th half hour ($\tau = 1, \dots, 48$) on the t-th day. We define a price event on day t as having occurred whenever any of the half-hour spot prices exceeds a threshold price, th of A$100/MwH.

$$y_t = \begin{cases} 1 & \text{if any } P_{t,\tau} > th \text{ for any } \tau = 1, \dots, 48 \\ 0 & \text{if any } P_{t,\tau} \leq th \text{ for all } \tau = 1, \dots, 48 \end{cases} \qquad (4.3\text{-}1)$$

The new series y_t is a binary series which takes the value 1 for all days in which $P_{t,\tau}$ exceeded th for at least one half hour. The threshold has been set to A$100/MwH in order to facilitate a comparison of these results to those in Chapter 2. It is, of course an arbitrary value although it seems widely used[35]. The timing of the price spikes is an important element in our modelling strategy and for this purpose we define the time index $\{t_i\}, i = 1, \dots, m$ as the time when the i-th price event occurred. Here we assume that in our sample we observed m price events.

Both approaches used, the PAR and the Hawkes models, model the probability or intensity of a price event occurring at time t, $P(y_t = 1|I_{t-1}) = \lambda_t$. The intensity λ_t is modelled based on information available at time $t-1$, I_{t-1}. This information consists of all the price events occurring up to time $t-1$ and exogenous variables that are available at time, $t-1$. Both these approaches are introduced in subsections 4.3.2.

4.3.1. Relationship between Pre-dispatch Prices and Price Outcomes

It is the main contribution of this chapter to introduce a novel set of explanatory variables into models for electricity price spikes. As a pre-cursor to this analysis we will investigate the relationship between pre-dispatch prices, $P_{t,\tau|j}$, and actual price outcomes, $P_{t,\tau}$, more generally[36]. This is done initially with some simple correlation analysis. However, as we are mainly interested in the more extreme, higher, percentiles of the price distribution, a correlation analysis may not reveal all. Acknowledging that the relationship between pre-dispatch prices, $P_{t,\tau|j}$, and price outcomes, $P_{t,\tau}$, may be different in the tails of the price distribution, naturally leads to Quantile Regressions (refer to Bunn et al. (2012) and Koenker & Jr (1978) for an application in the context of predicting electricity prices) as a methodology of choice. Here we will relate the logarithmic (log) of price outcomes, $\log(P_{t,\tau})$, and the logarithmic (log) of the pre-dispatch price, $\log(P_{t,\tau|j})$, linearly,

$$\log(P_{t,\tau}) = \alpha_p + \beta_p \log(P_{t,\tau|j}) + errors \qquad (4.3\text{-}2)$$

but acknowledge that the parameters describing this relationship may differ at different quantiles, p. This setup will allow us to examine the effect of pre-dispatch prices on the entire distribution of spot prices, by establishing the form of conditional quantiles of $\log(P_{t,\tau})$.

[35] Eichler et al. (2012) also provide a brief discussion on this threshold level.

[36] The exact definition of pre-dispatch prices will be discussed in Section 4.4.2.5. The subscript j will index the different times at which the pre-dispatch prices for period t, subperiod τ is formed.

Here we will briefly outline the estimation strategy used to obtain the sample estimates for $\hat{\alpha}_p$ and $\hat{\beta}_p$[37]. Let $c_p\big(\log(P_{t,\tau})|\log(P_{t,\tau|j})\big)$ be the p-th ($p = 0.1,02, \dots ,0.9$) quantile of the conditional distribution of $\log(P_{t,\tau})$ given $\log(P_{t,\tau|j})$. Using the linear specification we know that

$$c_p\big(\log(P_{t,\tau})|\log(P_{t,\tau|j})\big) = \alpha_p + \beta_p \log\big(P_{t,\tau|j}\big) \qquad (4.3\text{-}3)$$

The quantile loss function for the p-th quantile can be defined as follows:

$$L_p\left(\log(P_{t,\tau}), c_p\big(\log(P_{t,\tau})|\log(P_{t,\tau|j})\big)\right) = \begin{cases} p\left(\log(P_{t,\tau}) - c_p\big(\log(P_{t,\tau})|\log(P_{t,\tau|j})\big)\right) & if\ \log(P_{t,\tau}) \geq c_p\big(\log(P_{t,\tau})|\log(P_{t,\tau|j})\big) \\ (1-p)\big(c_p\big(\log(P_{t,\tau})|\log(P_{t,\tau|j})\big) - \log(P_{t,\tau})\big) & otherwise \end{cases} \qquad (4.3\text{-}4)$$

The estimates $\hat{\alpha}_p$ and $\hat{\beta}_p$ are obtained by minimising this loss function (evaluated at all n observations) at every p:

$$\underset{(\alpha_p, \beta_p)}{argmin} \left(n^{-1} \sum_{i=1}^{n} L_p\left(\log(P_{t,\tau})_i, c_p\big(\log(P_{t,\tau})|\log(P_{t,\tau|j})_i\big)\right)\right) \qquad (4.3\text{-}5)$$

If $p = 0.5$, then we are minimising the expected absolute loss. The resulting $c_{0.5}\big(\log(P_{t,\tau})|\log(P_{t,\tau|j})\big)$ is the predicted median of the conditional distribution. Using $p = 0.9$, we obtain the estimated 90th percentile of the conditional distribution, $c_{0.9}\big(\log(P_{t,\tau})|\log(P_{t,\tau|j})\big)$.

The results from this analysis will establish the effect of changes in the pre-dispatch prices on the respective conditional quantiles and therefore will shed light to what extend high pre-dispatch prices actually signal subsequent high price outcomes.

4.3.2. Hawkes point process model

In order to model the point process of actual price spikes we follow an approach that was first introduced by Hawkes (1971). The attraction of Hawkes models is that they provide a natural modelling framework in which the intensity of events can depend on the timing of previous occurrences. Such models have recently attracted renewed attention in the financial econometrics literature where they have been used to model the arrival of trades or orders (see e.g. Bauwens & Hautsch (2006) for an overview). The Hawkes model is particularly appropriate when events, such as the electricity price spikes, cluster in time. This implies that there is a degree of persistence in the process intensity, λ_t, and a univariate Hawkes process captures such persistence by specifying λ_t as follows

[37] For a detailed discussion of estimation strategies one can refer to Perlich et al. (2007). The estimation results shown in our chapter are produced by STATA

$$\lambda_t = \mu_t + \sum_{t_i<t} w(t - t_i) \qquad (4.3\text{-}6)$$

Both, μ_t and the function $w(\cdot)$ need to be non-negative to avoid negative intensities. The intensity, λ_t, is decomposed into two components, a deterministic (seasonal) component, μ_t, and the stochastic component, $\sum_{t_i<t} w(t - t_i)$, the roles of which will now be explained.

The electricity price is known to have distinct seasonal patterns. In the main they are a weekly pattern (with weekends usually displaying lower load and prices) and an annual pattern where loads and prices vary with the different climatic conditions throughout the year. Therefore, the deterministic component is used to capture any variations which can be ascribed to these seasonalities. We use trigonometric functions (as in Lucia & Schwartz (2002), Heydari & Siddiqui (2010) and Becker et al. (2004)) to model the weekly and annual seasonality in μ_t. As in Chapter 2 a range of trigonometric functions with frequencies chosen to fit weekly and annual variations is fit to the observed series of price spikes, y_t. Only the trigonometric functions that are statistically significant are then included in the further analysis. More precisely, we use $\mu_t = \varphi' \, trig_t$, where $trig_t$ is a row vector that contains the relevant values of these significant trigonometric functions at time t and φ is an appropriately sized column vector of parameters. The parameter vector φ is estimated alongside all other model parameters as discussed below[38].

The intensity at time t, λ_t, depends on all the occurrences of price events that occurred at times, t_i, before the current time, $t(t_i < t)$. The stochastic component of a Hawkes process is then defined as;

$$\lambda_t = \mu_t + \sum_{t_i<t} w(t - t_i) \qquad (4.3\text{-}7)$$

$$w(t - t_i) = \alpha_{t_i} e^{-\beta(t-t_i)}$$

Each price event that occurred before t contributes to the intensity at time t. A price spike at time t_i has an initial contribution to the intensity of α_{t_i}, but its contribution to the intensity at time t is dampened by the decay factor $e^{-\beta(t-t_i)}$.

[38] In order to simplify the exposition below we shall keep the term μ_t in the discussion below.

In this specification a price spike's initial contribution is indexed by t_i, the time at which the price spike occurred. This is to indicate that we allow this initial contribution[39] to vary with covariates available at time t_i, x_{t_i}. This will, e.g. allow the intensity contribution to be different for weekend spikes as opposed to middle of the week price spikes. To ensure the point process complies with stationary conditions (see; Ogata (1978)) α_{t_i} is defined as follows;

$$\alpha_{t_i} = \Phi(\omega_{t_i}) \cdot \beta \tag{4.3-8}$$

$$\omega_{t_i} = x_{t_i}\gamma \tag{4.3-9}$$

since $\omega_{t_i} \in \mathbb{R}$, and $\Phi(\cdot)$ being the standard normal cumulative distribution, $\Phi(\omega_{t_i}) \cdot \beta$ ensures that $\alpha_{t_i} \leq \beta$. This model allows for the impact of covariates on the intensity for price spikes, but only of covariates at price spike times t_i. This model will be labelled HAWa.

In order to allow covariates at any time influence the intensity for price spikes, in Chapter 2 we make the decay parameter, β, dependent on exogenous variables as well. Therefore in the second model, the decay parameter is parameterised to be time varying, β_t, as follows[40]:

$$\beta_t = \Phi(z_t\delta) \tag{4.3-10}$$

where β_t is restricted to be between 0 and 1 via $\Phi(\cdot)$. The $(1 \times q_z)$ vector, z_t, contains the relevant covariates conditioned at time t and δ is the appropriate $(q_z \times 1)$ parameter vector. This modelled will be called HAWab.

The parameter vector, $\theta = (\zeta'\ \gamma'\ \delta')'$ is estimated by maximising the log-likelihood function (LLF) in equation (4.3-11) based on the observed price events at time $t_1, \dots, t_n$;

$$\begin{aligned}\log L(t_1, \dots, t_n|\theta) &= -\int_0^T \mu_t\ dt - \int_0^T \sum\nolimits_{t_i<t} \alpha_{t_i} e^{-\beta_t(t-t_i)}\, dt \\ &+ \int_0^T \log\left(\mu_t + \sum_{t_i<t} \alpha_{t_i}\, e^{-\beta_t(t-t_i)}\right) dN(t)\end{aligned} \tag{4.3-11}$$

where $dN(t)$ denoted as an indicator function taking a value of 1 if a price event occurred during period t and a value of 0 if no price event occurred during period, t.[41]

[39] Refer to Figure 2.5-1 on the the mechanism of α_{t_i}

[40] Refer to Figure 2.5-3 on the the mechanism of β_t

[41] For the first model in which only the initial intensity increase parameter α is covariant dependent, $\theta = (\zeta'\ \gamma')'$. The changes to the LLF are straightforward.

4.3.3. The Poisson Autoregressive Model

The last model to be considered here is the model that was introduced in Christensen et al. (2009) as their model of choice for the intensity of electricity price events. In contrast to the Hawkes model, the PAR model allows for different dynamics in the latent intensity process[42]. The PAR model is based on the process of a latent variable, $X_t \in \mathbb{N}^0$ defined as the number of system stresses. An example of system stress could be any events that could trigger unanticipated demand or events that cause a drop in supply, e.g. a generator failure. The latent variable is related to the observed price events as follows:

$$y_t = \begin{cases} 1 & if\ X_t > 0 \\ 0 & if\ X_t = 0 \end{cases} \qquad (4.3\text{-}12)$$

In the PAR model we infer from the presence of a price event that at least one system stress occurred. The model is built around the arrival and departure process of system stresses. Importantly any existing system stress is allowed to persist for several periods and at any period there could be more than one stress present ($X_t > 1$). The arrival of the system stress is modelled by an independent Bernoulli process while the departure of the system stress is modelled by a binomial thinning process. The probabilities of the arrival and departure for the system stresses can be modelled as functions of the conditioning covariates. These covariates are set to be the same as in the Hawkes models.

The arrival of a system stress is modelled by assuming that at any period, t, there can only be one new stress, defined by $(e_t = 1)$ and this arrival process is assumed to follow an independent Bernoulli process:

$$P(e_t = 1) = \pi \quad new\ stress\ arrives \qquad (4.3\text{-}13)$$

$$P(e_t = 0) = 1 - \pi \quad no\ new\ stress \qquad (4.3\text{-}14)$$

Limiting this arrival process to only one new stress is necessary to disentangle the arrival and departure of the latent process.

The departure of the system stress is modelled as a binomial thinning process whereby the number of stress factors surviving from time $t-1$ to time t, is equal to $X_t^{surv} = \rho \circ X_{t-1}$ $(0 \leq X_t^{surv} \leq X_{t-1})$ where $\circ$ is the binomial thinning operator and ρ is the probability of each stress factor present at time $t-1$ (X_{t-1}) surviving. The number of stress factors at time t is therefore represented by

[42] It is apparent from the empirical work in Chapter 2 that the intensities implied by the PAR model exhibit less persistence than the Hawkes models.

$$X_t = \rho \circ X_{t-1} + e_t \qquad (4.3\text{-}15)$$

The time-varying arrival and survival of the system stresses is conditioned by the exogenous variables as follows;

$$\pi_t = 1 - \exp(-\exp(z_{1t}\beta_1)) \qquad (4.3\text{-}16)$$

$$\rho_t = 1 - \exp(-\exp(z_{2t}\beta_2)) \qquad (4.3\text{-}17)$$

in which z_{1t} and z_{2t} are sets of regressors for the arrival and survival of the system stresses respectively and β_1 and β_2 are the associated parameter vectors that produce the scalar weighted values $z_{1t}\beta_1$ and $z_{2t}\beta_2$.

This model is suitable to model the probability of a price spike occurring $P(y_t = 1|\Psi_{t-1})$ as that is equivalent to $P(X_t > 0|\Psi_{t-1})$. Similar to the Hawkes model, the probability of the price spikes depend on the history of price spikes process, $\Psi_{t-1} = \{y_0, y_1, \dots, y_{t-1}\}$. For example, if $y_{t-1} = 0$ then we know that $X_{t-1} = 0$ and therefore $X_t = 0$ or 1 as we limited system stress arrival to one new arrival at any time. This feature has an interesting implication. In the case that $y_{t-1} = 0$, the relevant information set reduces to $\Psi_{t-1} = y_{t-1}$ and the probability of a price event is simply the probability of a new system stress arriving,

$$p_t = P(y_t = 1|y_{t-1} = 0) = P(e_t = 1) = \lambda_t \qquad (4.3\text{-}18)$$

In fact, the relevant information set always reaches back to the last period in which we did not have any price spike. Consider the following history of price events, $y_{t-1} = y_{t-2} = \cdots = y_{t-k} = 1$ and $y_{t-(k+1)} = 0$. The relevance of this history is that at time $t-1$, there was certainly at least one system stress present, but it could have been up to k stresses in the case in which the last k periods each produced a new stress and none of the existing ones disappeared.

The probability required for the calculation of the likelihood function, and hence to estimate parameters, is $p_t = P(y_t = 1|\Psi_{t-1}) = 1 - P(y_t = 0|\Psi_{t-1})$. It is the calculation of the latter that is somewhat involved. Consider the above case in which during the last k periods price spikes were observed. Then we know that at time $t-1$ there was anything between 1 and k stresses present. If we knew that there were k stresses the probability of $P(y_t = 0|\Psi_{t-1}, X_{t-1} = k) = P(X_t = 0|\Psi_{t-1}, X_{t-1} = k)$ could easily be calculated as $(1-\rho_t)^k(1-\pi_t)$, i.e. the probability that all existing stresses disappear, $(1-\rho_t)^k$, times the probability of not seeing a new stress in time t, $(1-\pi_t)$. But we need to allow for the

possibility that X_{t-1} could take any value between 1 and k. This leads to the following approach

$$P(X_t = 0|\Psi_{t-1}) = \sum_{n=1}^{k} P(X_{t-1} = 0|X_{t-1} = n)P(X_{t-1} = n|\Psi_{t-1}) \qquad (4.3\text{-}19)$$

Finally, to calculate $P(X_{t-1} = n|\Psi_{t-1})$ we require a recursive calculation starting from k periods back, which is the last period where we were certain about the number of system stresses ($y_{t-k} = 1, y_{t-(k+1)} = 0 \Rightarrow X_{t-k} = 1$). The details of this recursive calculation are outlined in Appendix A.

The result of this recursive calculation is $P(X_t = 0|\Psi_{t-1})$ which is then used to calculate $p_t = P(y_t = 1|\Psi_{t-1}) = 1 - P(y_t = 0|\Psi_{t-1})$. This allows the formulation of the following log likelihood function (LLF)

$$\log L\,(\beta; y_1, y_2, \dots, y_T) = \sum_{t=1}^{T} y_t \log p_t + (1 - y_t)\log(1 - p_t) \qquad (4.3\text{-}20)$$

where $\{y_1, y_2, \dots, y_T\}$ is the sequence of observed price events and as mentioned previously the parameter vector $\beta = (\beta_1' \beta_2')'$ is used to parameterise the time-varying arrival (π_t) and survival (ρ_t) probabilities respectively(as outlined in equations (4.3-16) and (4.3-17)). The parameters are estimated by maximising the LLF.

4.4. Overview of the Australian New Electricity Market (NEM) and the Data

In this Section we will give a short overview of the Australian National Electricity Market (NEM) with a particular focus on the pre-dispatch process. This is followed by the description of the data used.

4.4.1. The Australian NEM and the Pre-Dispatch Process

The restructuring of the electricity market in the regions of Australia happened in different periods of time. The Australia NEM wholesale market was first organised into two separate electricity pools in the states of Victoria and New South Wales. In the mid-1990s this was followed by two more regions, Queensland and South Australia, which were then connected to the Australia National Electricity Market (NEM). The NEM began operating on December 13, 1998 as a pooled market in which the market operator, AEMO, is

responsible to ensure that all the available supplies are aggregated to meet the demand as cost efficient as possible[43].

On the dispatch day, AEMO is responsible to give dispatch instruction to generators to meet the prevailing demand as cost effective as possible in every five minute interval. Hence, the dispatch price is determined by the bid price of the marginal generator dispatched into production. In the NEM, there are 48 half-hourly intervals in a day with each half-hour is divided into six five-minute dispatch intervals. The half-hourly spot price is calculated as an arithmetic mean of the six five-minute dispatch prices within half hour.

As part of the restructuring process, AEMO is responsible to provide generators with sufficient information to indicate whether their capacity would be called for dispatch. One of these pieces of information is the pre-dispatch price. The daily supply bids, submitted by generators at 12.30pm a day before dispatch, are matched against the regional demand forecast to produce pre-dispatch prices for all half hour periods in the next day. Pre-dispatch sensitivities are published together with pre-dispatch prices. These indicate how prices are expected to change if the demand was somewhat higher or lower than currently expected. These are denoted by $P^{+\Delta}$ and $P^{-\Delta}$ respectively[44].

The day-ahead market (pre-dispatch) is based on current supply bids and demand forecasts, and it can be seen as a short-term forecast of the wholesale electricity price for the respective dispatch periods on the next trading day (Andalib & Atry 2009) and will help to plan the volume expected to be supplied through the interconnectors between regions (Australian Energy Market Operator 2010). The quality of these pre-dispatch prices relies on the quality of the matching mechanism[45], the accuracy of demand forecasts (which are commonly accepted to be good and unbiased, (Australian Energy Market Operator 2013; Electricity Market Performance 2012c; System Operations 2011a; System Operations 2011b)) and the reliability of the supply bids made. Acknowledging the critical nature of this last point, the Australian Energy Market Commission (AEMC) has mandated that the

[43] In May 1996, the government of Australia formed two companies to operate the NEM. They are the National Electricity Market Management Company Limited (NEMMCO) and National Electricity Code Administrator Limited (NECA). However, on 1 July 2009 the operation of NEMMCO was ceased and all their roles and responsibilities are now being replaced by the Australian Energy Market Operator (AEMO); (see; Chevallier (2010)).

[44] AEMO produces +/- 200, 500 and 1000 MWh scenario prices but our analysis uses +/- 200 MWh scenario prices only.

[45] The task of finding the most cost effective combination of suppliers to meet the demand under a range of constraints is delegated to a computer algorithm. While there are slight differences between the algorithm used in the pre-dispatch periods and that used for the actual dispatch period (Electricity Market Performance 2012b) these differences are minor and are not expected to lead to any systematic biases.

initial bids and subsequent rebids have to be submitted in "good faith" which means the bids have to reflect the "genuine intention to honour that offer [...] if the material conditions and circumstances [...] remain unchanged until the relevant dispatch interval" (Australian Energy Market Commission, 2013, paragraph 3.8.22A).

The pre-dispatch information will help generators to anticipate whether and if so how much of their offered generation will be dispatched. Market participants will also use this information as a basis of rebids which can be submitted until approximately five minutes before dispatch. But as indicated above, any such bids have to be justified on the basis of changed circumstances. The pre-dispatch price and its sensitivities prices are updated and published every half-hour for the rest of the current day and throughout the dispatch day. The pre-dispatch prices and sensitivities will change as the demand forecasts and supply bids change.

In earlier work in Chapter 3 some aspects of the relationship between pre-dispatch prices and dispatch price outcomes on the dispatch day are investigated. In Chapter 3 our aim is to analyse whether there is an apparent bias in the pre-dispatch process; such that the pre-dispatch prices are systematically under- or overestimating the actual electricity price outcomes. We also analyse whether there is any systematic variation in the bias across years, seasons and/or trading periods and whether there are exogenous variables that could explain the systematic variation. We find that there is a significant bias in the pre-dispatch process. Overall, the pre-dispatch process tends to underestimate the actual price outcomes (i.e. price outcomes tend to exceed $P^{+\Delta}$). In addition, the variation in the probabilities of price outcomes to be more than $P^{+\Delta}$ or less than $P^{-\Delta}$ (high or low boundary) are remarkably persistent across the years. Further, the bias is particularly strong in the morning and evening peak periods of the day.

Following the findings in Chapter 3, in this chapter we wish to investigate whether any pre-dispatch information can be used in modelling and predicting price spikes. There are several studies in the deregulated electricity market of Ontario (Canada) that included pre-dispatch information in their forecast models producing favourable forecasting results. In a pair of papers by Zareipour & Bhattacharya (2006) and Hamidreza Zareipour et al. (2006b) the authors use pre-dispatch price and demand information as exogenous variables and find an improvement in the accuracy of short-term hourly Ontario electricity price (HOEP) forecast. This pre-dispatch information is used as exogenous information in nonlinear time-

series models of the HOEP. It is conjectured that the value of this information lies in its complex information content, as it is based on actions of market participants. Interestingly, however, the authors also state that this information does not seem to foreshadow extreme price movements.

4.4.2. The Data

In this study we are using New South Wales (NSW) regional market data.

4.4.2.1. Wholesale Electricity Prices

The half-hourly spot electricity prices used in this study are obtained from the AEMO website. The price is defined in Australian Dollars per Megawatt Hour ($/MWh). The data cover the period from 2 March 1999 to 31 October 2007. Since we have a missing observation from January 1 to 31, 2003 and July 2 to August 1, 2005 this delivers a sample size of 148,992 half-hourly observations or 3135 daily observation. This price series is used to define the dependent price spike variable as in equation (4.3-1). While it is the aim of this chapter to predict the probability of a price spike occurring on a particular day, the series of price spikes is a derivative of the actual price series. Descriptive statistics for the daily spot prices (averages of the 48 half hour periods in a day) are therefore shown in Table 4.4-1.

(INSERT Table 4.4-1 HERE)

From Table 4.4-1, we can observe the median of daily time series of spot prices are consistent throughout the years except for 2007 where it increases to $51.14/Mwh. While high standard deviation can be observed particularly from 2004 onwards. The price spikes seem to jump highest on 2004 with based on its 99^{th} percentile value that is more than $800/Mwh. There are also a high proportion of days with price spikes (45.14%) in 2004. Although the magnitude of the price spikes jump ($589.90/Mwh) on 2007 is lower than on 2004, they have highest proportion of days (72.88%) and half-hour (51.14%) with price spikes and mean values (4.924) of Count Data among all the years.

We will now discuss the potential exogenous variables that maybe related to the occurrence of price events. These variables are used as they will all potentially contribute to extreme price events.

4.4.2.2. The System Load

It is known that electricity load is a main driver on the movement of the electricity spot price in deregulated markets (Moghram & Rahman 1989; Wu & Shahidehpour 2010; Vucetic 2001). Therefore it is important to use them as an input variable in building a price model. We illustrate the behaviour of NEM load series in Figure 4.4-1. The daily load series shows day-of-week and month-of-year seasonality and a clear upwards trend over the years.

(INSERT Figure 4.4-1 HERE)

As in Eichler et al. (2012) we aim to produce a forecast of electricity demand that exceeds the demand that one would normally expect for that day of the week at that particular time of the year. That is the demand measure we believe is likely to contain most information for price spikes. We therefore need two elements. First, the normally expected load series, which is constructed by weighting the average load from the previous seven days (see; Chapter 2). This series as available at time $t-1$ for time t. Second, we would ideally use the load forecast for day t available on day $t-1$. However we do not have that series available. Instead we use the actual realised load series on day t as a proxy of the forecast load series on day $t-1$ for day t. This approach comes with the caveat that the information extracted from this series is an upper-limit of the information actually achievable[46].

Before we include the difference between realised load and normally expected load in the price model, the series needs to be de-seasonalised from annual seasonality by applying a rolling volatility technique proposed by Weron (2007). The resulting series is illustrated in Figure 4.4-2 and label as L_t.

(INSERT Figure 4.4-2 HERE)

4.4.2.3. Temperature Data

A significant part (but not all) of the variation in system load is due to variation in the weather. Extreme weather conditions could result to high electricity demand which translates into high electricity price since high-cost generators need to be activated to satisfy the escalating demand. The most obvious cases of extreme weather are very high and low temperatures as these may trigger extra demand for cooling or heating respectively.

[46] We feel that this approach, despite its obvious shortcomings is appropriate as it is well known that AEMO (and generators) have excellent and detailed load forecasts available.

Temperature data are available from Australian Bureau of Methodology on daily basis. There is, of course, a large degree of seasonality in the temperature data. We are not really interested in the main seasonal pattern as this is unlikely to be correlated to finding any extreme price events. What we want to identify are temperature extremes. Here we are using the same approach as in Chapter 2. We consider the absolute difference of the observed temperature series from the seasonally expected temperature which is modelled by a trigonometric function. Furthermore we differentiate between un-seasonally hot and un-seasonally cold days by creating two time-series, $Tmax_t$ and $Tmin_t$. The details of the calculations are discussed in Chapter 2, but it is important to note that these series only capture un-seasonally hot days in summer and un-seasonally cold days in winter. This is in acknowledgement that a very warm winter day is unlikely to produce the same demand pressure on the wholesale electricity market as would a hot day in summer.

In our models, which use information available at time $t-1$ to predict price events at time t, we would of course want to have forecasts for these variables for day t. We have no series of historical temperature forecasts available. This leaves us with one of three options. We could ignore the temperature information, but we think that the potential importance of this information rules this out as an option. We could also attempt to build a basic temperature forecasting model based on historical information. However, it is unlikely that any such attempt would get even close to the accuracy of actual temperature forecasts. We therefore feel that the only viable option is to use actual realised temperature data instead, acknowledging that any significant effect of such a variable would likely appear somewhat more significant than that which could be realised with forecasts.

4.4.2.4. Count Data

From Figure 4.1-1 it was apparent that price spikes, even at a daily level, are highly persistent. This is, of course the feature that is exploited by both the Hawkes and PAR models to predict the occurrence of price spikes. The models build in this chapter, however, work on daily data, despite the price series being a half-hourly series. The aggregation process from equation (4.3-1) used here masks some variation in the price process. Indeed y_t could take a value of 1 if $P_{t,\tau} > A\$100$ for only one half-hour. It will equally take a value of one if the price exceeded the threshold in all 48 half hours of a day.

In order to capture this intra-daily information, that may capture, to some extent, the severity of the price spike, we define a new variable counting the number of times the half-hourly spot prices exceeds the threshold th during day t, C_t. A high value of this variable indicates the presence of more severe system stress and may be more likely to increase the probability of another price event.

4.4.2.5. Pre-dispatch information

This section outlines the new covariates derived from pre-dispatch information. These covariates are for use in the information set Ψ_t used to predict price spikes at time $t+1$. As discussed previously, the pre-dispatch process for day t starts on the day before dispatch, i.e. day $t-1$, but continues all through the actual dispatch. Therefore, the pre-dispatch information used in this chapter includes information from the pre-dispatch at 15:00 on $t-1$, (denoted as *PD1*) but also of the pre-dispatch on day t at 06:00 (*PD2*) and the pre-dispatch 30-minutes prior to the actual dispatch (*PD3*)[47]. *PD1* and *PD2* gather information for all dispatch sub-periods τ at one point in time (15:00 on $t-1$ and 06:00 on day t respectively). *PD3* however, collects pre-dispatch information at a different time for each sub-periods τ.

The first piece of information that uses the pre-dispatch is the relative position of the actual spot price outcome relative to the respective pre-dispatch information. In our study on the performance of the pre-dispatch process in Chapter 3 we placed the spot price outcome relative to the pre-dispatch scenarios. Following that approach the spot price outcome on day t, sub-period τ, relative to pre-dispatch j (where $j = PD1, PD2 \; or \; PD3$) is defined as follows:

$$S_{t,\tau|j} = \begin{cases} 2 & if \; P_{t,\tau} > P_{t,\tau|j}^{+\Delta} \\ 1 & if \; P_{t,\tau|j}^{+\Delta} \geq P_{t,\tau} > P_{t,\tau|j}^{-\Delta} \\ 0 & if \; P_{t,\tau|j}^{-\Delta} \geq P_{t,\tau} \end{cases} \qquad (4.4\text{-}1)$$

We use this measure for sub-periods $\tau = 1, \ldots, 35$ $(06{:}30 \text{ to } 23{:}30)$ only to ensure that we have an equal amount of information available for all j.

We then generate a daily measure by comparing the number of sub-periods for which the actual wholesale spot price outcomes, $P_{t,\tau}$, on the dispatch day are higher than $P_{t,\tau|j}^{+\Delta}$ and lower than $P_{t,\tau|j}^{-\Delta}$, for $j = PD1, PD2 \; and \; PD3$:

[47] It should be noted that the PAR and Hawkes models that are designed to produce intensity forecasts for time t, conditional on information at time $t-1$, only used PD1.

$$B_{t,p|j} = \sum_{\tau \in p} I(S_{t,\tau|j} = 2) - \sum_{\tau \in p} I(S_{t,\tau|j} = 0) \qquad (4.4\text{-}2)$$

We interpret this measure as a measure of bias. From the work in Chapter 3 it is apparent that the relative performance of the pre-dispatch process differs for different periods in the day. We therefore produce this measure for the following three distinctive sub-periods. The morning (06:30 to 09:30 dispatch intervals - 7 trading intervals, $p = m$), the daytime (10:00 to 17:30 - 16 trading intervals, $p = d$) and evening (18:00 to 23:30 - 12 trading intervals , $p = e$).

In Figure 4.4-3, we display the time-series of $B_{t,d|PD1}$. If the total of $B_{t,d\,|PD1}$ is equivalent to 16, then this implies that on that particular dispatch day,t, all actual wholesale spot price outcomes, $P_{t,\tau}$, from 10:00 to 17:30 are larger than the corresponding than $P^{+\Delta}_{t,d|PD1}$. However, if $B_{t,d\,|PD1}$ is equal to -16, this implies that on the particular dispatch day, t, all price outcomes, $P_{t,\tau}$ from 10:00 to 17:30 are less than $P^{-\Delta}_{t,d|PD1}$.

(INSERT Figure 4.4-3 HERE)

Despite the apparent volatility in the series, there is also a fair degree of underlying persistence, which is best seen from the 50 day moving average that is superimposed. While this series fluctuates around 0 in the early part of the sample, there is a clearly discernible upwards drift in the latter part, indicating that there is a tendency for the pre-dispatch process on the day before dispatch (*PD1*), to underestimate the actual spot price outcomes. Similar observation as in Figure 4.4-3 can be observed for the time-series of $B_{t,d|PD2}$ (hence not shown here) indicating not many significant adjustment have been made on the supply bids and demand forecast from a day before until 06:00 on the day of dispatch.

However, from Figure 4.4-4, which shows $B_{t,d|PD3}$, it becomes apparent that this bias has disappeared 30 minutes before dispatch. Not surprisingly we get much lower variation in the series, as 30 minutes to dispatch most relevant information has been worked into the dispatch via adjusted supply bids and updated demand forecasts. It is perhaps somewhat surprising that we still get a significant number of dates on which $\left|B_{t,d|PD3}\right|$ exceeds 5 which is indicative of days during which we get a significant amount of sluggishness in supply bid adjustments[48].

[48] This assumes that the centrally administered demand forecasts do quickly adjust to any new information.

(INSERT Figure 4.4-4 HERE)

A similar summary can be applied to time-series of $B_{t,m\,|j}$ and $B_{t,e\,|j}$ whereby the pre-dispatch process on the day before dispatch (*PD1*) and on dispatch day at 06:00 (*PD2*) tend to underestimate the actual spot price outcomes and this bias only disappeared 30 minutes before dispatch (*PD3*). Later in Section 4.5.3.2, we will use the $B_{t,p|PD1}$ time-series as conditioning variables in our point process models.

The second set of pre-dispatch information used is the pre-dispatch price. While we only have data on the pre-dispatch scenarios, we decided to proxy the actual pre-dispatch prices from a supply curve inferred from the available pre-dispatch scenarios (sensitivities), $P^{+\Delta}$ and $P^{-\Delta}$. In Figure 4.4-5 you can see an example of six available pre-dispatch prices ($P^{-500}, P^{-200}, P^{-100}, P^{+100}, P^{+200}, P^{+500}$) at their respective relative load offset values (-500, -200, -100, 100, 200, 500). These load offset values are relative to the load forecast for a particular period. A cubic spline connects these 6 points with piece-wise cubic functions where the parameters of these functions are chosen in such a way to ensure that the two neighbouring functions which meet in one of the four inner points (at offsets -200, -100, 100, 200) have equal first and second order derivatives (Brandimarte 2013). Once these parameters have been calculated they can be used to infer any pre-dispatch price at any relative load value. We are interested in the pre-dispatch price at the actual load forecast, E(Dd) (i.e. at a load offset value of 0), which, in the example displayed in Figure 4.4-5 is A$16.80. For illustration purpose in Figure 4.4-5, the estimated pre-dispatch price $\hat{P}_{t,09:00|PD1}$ is specific for the fifth sub-period (09:00) on one particular dispatch day t dated March 2, 1999 and conditioned during pre-dispatch process on *PD1*.

(INSERT Figure 4.4-5 HERE)

To show estimated pre-dispatch price on *PD1, PD2* and *PD3* on the fifth sub-period for all 3166 days, we plot the logarithmic series of the $\hat{P}_{t,09:00|PD1}$, $\hat{P}_{t,09:00|PD2}$, $\hat{P}_{t,09:00|PD3}$ in Figure 4.4-6. Except for one large jumps in the logarithmic series of the $\hat{P}_{t,09:00|PD1}$ there is not much difference between the estimated pre-dispatch price on *PD1, PD2* and *PD3* on the fifth sub-period.

(INSERT Figure 4.4-6 HERE)

In Table 4.4-2 to Table 4.4-4 we report summary statistics for the estimated pre-dispatch prices and, for comparison reason, the equivalent dispatch spot prices. We can see that the descriptive statistics display a significant amount of variation across the years.

(INSERT Table 4.4-2 to Table 4.4-4 HERE)

Figure 4.4-7 compares the median of the spot prices (broken down by years and for trading sub-periods $p = m$, d and e) with those from the three pre-dispatch processes. We can see that on average all pre-dispatch processes underestimate the spot price outcomes. The tendency to underestimate is somewhat stronger in the later years of our sample. It also apparent that in general, the pre-dispatch information improves from *PD1* to *PD2* and to *PD3*. For the morning periods *PD2* and *PD3* have very similar characteristics, whereas from 2003 onwards for the daytime and evening sub-periods *PD2* looks more like *PD1* than *PD3*.

(INSERT Figure 4.4-7 HERE)

The cubic spline interpolation is clearly an approximation to the prevalent supply curve only, as the latter is known to be a step function. Apart from being used to derive an estimate of the pre-dispatch price level we also use it to derive an estimate of the slope or first derivative of the supply curve at $\widehat{P}_{t,\tau|j}$, $\widehat{D}_{t,\tau|j}$. The reason for using this variable as a potential explanatory variable in our point process models is that the slope is a potential indicator for how sensitive the equilibrium price is to an event that either shifts the electricity supply or demand.

The above calculations resulted in variables $\widehat{P}_{t,\tau|j}$ and $\widehat{D}_{t,\tau|j}$ which were available for any trading interval τ and for the different pre-dispatches $j = PD1, PD2\ or\ PD3$. As the Hawkes and PAR models are models estimated on the basis of daily observations, this information has to be distilled into daily variables. Here we are interested in modelling extreme price events and therefore we decided to use the average of the 5 largest values of $\widehat{P}_{t,\tau|j}$ and $\widehat{D}_{t,\tau|j}$ across all $\tau = 1, \dots, 35$ on any day t. The resulting daily series are labelled $Pm5_{t|j}$ and $Dm5_{t|j}$.

In order to get a first impression of how these series relate to the spot price outcomes we present six scatterplot in Figure 4.4-8. In order to show this relation better, the scatterplot is using the logarithmic series of daily spot prices, P_t and logarithmic series of daily $Pm5_{t|j}$ and $Dm5_{t|j}$ on *PD1*, *PD2* and *PD3*. From our understanding of the convex nature of

supply curves, it is not surprising to find that the estimates for its slope, $Dm5_{t|j}$, are highly correlated to those of the expected price, $Pm5_{t|j}$ and therefore its positive correlation to the actual spot price outcomes is equally intuitive.

(INSERT Figure 4.4-8 HERE)

4.5. Empirical Analysis

In this Section we will present the empirical analysis in which we attempt to establish whether information from the pre-dispatch process can be utilised to inform any prediction model of price spikes. The initial analysis will focus on the relationship between the pre-dispatch prices and actual price outcomes (Sections 4.5.1 and 4.5.2). Finally we will incorporate pre-dispatch information into the Hawkes and PAR point process models.

4.5.1. Correlation Analysis

To understand the relationship between pre-dispatch and spot prices, we decided to calculate the correlation between half-hourly logarithmic series of pre-dispatch, $\ln(\hat{P}_{t,\tau|j})$, and spot prices, $\ln(P_{t,\tau})$. We use the logarithmic series to mitigate the extreme skewness observed in the raw series.

(INSERT Figure 4.5-1 HERE)

First we calculate the correlation of the logarithmic series of pre-dispatch prices with respect to logarithmic series of spot prices separately for every half hour trading period across all 3166 days in the sample. The resulting correlation profile (across the final 35 half hour trading periods of a trading day) can be seen in Figure 4.5-1 for the correlations between spot price outcomes and PD1 pre-dispatch prices. Two observations can be made. Overall we find a fairly high level of correlation between $\ln(\hat{P}_{t,\tau|PD1})$ and $\ln(P_{t,\tau})$ (refer to the blue line in the first panel of Figure 4.5-1). The level of correlation is lowest for the trading periods in the middle of the day ($\tau = 13{:}30\ to\ 16{:}00$), dropping to around 0.7.

As it is our final aim to investigate whether pre-dispatch information has any explanatory power on extreme price events, we also split our data according to whether the spot price is above or below A\$50[49], and subsequently calculate the corresponding correlation profile for both these datasets. These correlation profiles can be seen along the correlation profile

[49] The sample split according to the spot price has no real economic meaning and is merely done for illustrative reasons.

calculated for the full dataset in Figure 4.5-1 for all three pre-dispatch periods, *PD1* (refer to first panel of Figure 4.5-1), *PD2* (refer to second panel of Figure 4.5-1) and *PD3* (refer to third panel of Figure 4.5-1).

When comparing the two correlation profiles for the split data-sets it is immediately obvious that the correlation between $\ln(\hat{P}_{t,\tau|PD1})$ and $\ln(P_{t,\tau})$ is fairly strong across all 35 subperiods of the day when the eventual spot price remains low (refer to the green line in the first panel of Figure 4.5-1). For high spot price outcomes, however, the (linear) relationship between $\ln(\hat{P}_{t,\tau|PD1})$ and $\ln(P_{t,\tau})$ is much less obvious and the correlation coefficient drops to values mainly between 0.3 and 0.6 (refer to the red line in the first panel of Figure 4.5-1). Only for the early morning periods (07:30 to 08:30) does this correlation exceed 0.6. But notably, the lowest correlation (conditional on $P_{t,\tau} > A\$50$) is observed for the first period of the day analysed, the 06:30 period. This seems to indicate that price events at the very beginning of the day are the predicted the least reliably.

This main pattern of results can also be observed when analysing the correlation profiles at *PD2*. That is the pre-dispatch at 6am in the morning of the dispatch day. While the overall level of correlation has increased somewhat (by about 0.1 correlation units) (refer to the blue line in the second panel of Figure 4.5-1), we can still see a very significant correlation drop-off for high spot price outcomes (refer to the red line in the second panel of Figure 4.5-1). Again we see a general U-shape in the correlation profile and a correlation peak in the morning periods even for the correlation conditional on high price outcomes. Directly following this peak, however, is a very steep drop in correlations (conditional on high prices for the late morning periods (10:00 to 11:00 sub-periods)). It is worth noting that the time between the pre-dispatch (6am) and the actual dispatch varies for the correlations in this Figure. There is significantly more time left between the pre-dispatch and the 18:00 dispatch period than there is between pre-dispatch and the, say 10:00 dispatch. One would therefore expect that, ceteris paribus, the pre-dispatch process for the 10:00 dispatch period should be more precise than that for the 18:00 dispatch period. That, however, does not seem to be the case, as all correlations are larger for the 18:00 dispatch period.

The *PD3* pre-dispatch uses pre-dispatch information recorded 30 minutes before the beginning of every dispatch period. These results are therefore to be interpreted in a slightly different way than the previous correlations as they use a fixed time difference between pre-dispatch and dispatch. In fact the data for the 06:00 dispatch period are

identical to these in *PD2*, but all others differ. The correlation across all data exceeds 0.9 in all dispatch intervals (refer to the blue line in the third panel of Figure 4.5-1). The correlation profile based on low price outcomes mimics that profile very closely (refer to the green line in the third panel of Figure 4.5-1). As far as the correlation for the high price outcomes (refer to the red line in the third panel of Figure 4.5-1) are concerned, we can see that 30 minutes before dispatch most of these prices seem to be well anticipated by the pre-dispatch, with correlations generally being between 0.7 and 0.9, but still markedly lower than for lower price outcomes. Exceptions here are those that have already been pointed out for the *PD2* data. High prices in the 06:30 dispatch period are badly anticipated with the correlation being smaller than 0.3. This seems to indicate that the time period between 06:00 and 06:30 is characterised by either the emergence of significant new information and/or the prevalence of significant re-bidding information. One may explain this with the fact that the dawn of a new day may indeed bring to light new situations (not least weather outcomes) that could not be observed previously. A similar explanation however seems rather implausible to explain the dip in correlation (for high price outcomes) that occurs for the late morning periods (10:00 to 11:00). On the occasions in which the price outcomes are on the high side, significant activity appears to occur in the 30 minutes prior to dispatch.

4.5.2. Quantile regression

Before applying the quantile regression approach to the relation between spot and pre-dispatch data it is worth discussing what one would expect.

In this section, we propose a simple linear model that relates spot to pre-dispatch prices

$$\log(P_{t,\tau}) = \alpha + \beta \log(\hat{P}_{t,\tau|j}) + errors \qquad (4.5\text{-}1)$$

where $\log(P_{t,\tau})$ is the logarithm of the spot price for the τ-th half hour ($\tau = 06{:}30, \dots, 23{:}30$) on the t-th day and $\log(\hat{P}_{t,\tau|j})$ is the corresponding logarithm of the pre-dispatch price at $j = PD1, PD2 \; or \; PD3$.

As $\log(\hat{P}_{t,\tau|j})$ are essentially system predictions of the spot prices, $\log(P_{t,\tau})$, equation (4.5-1) can be interpreted as a Mincer-Zarnowitz type forecast evaluation regression. If the pre-dispatch process delivered an unbiased and efficient forecast (see; Mincer (1969) and Newbold & Harvey (2002)) we would expect $E(\log(P_{t,\tau})|\log(\hat{P}_{t,\tau|j})) = \log(\hat{P}_{t,\tau|j})$ and

hence the above regression relation should deliver the following estimated ordinary least square regression model:

$$E(\log(P_{t,\tau})|\log(\hat{P}_{t,\tau|j})) = 0 + 1 \times \log(\hat{P}_{t,\tau|j}) \qquad (4.5\text{-}2)$$

It will now be interesting to see in what sense the data deviate from this optimal forecast relationship and whether the deviations differ with respect to the quantile we are looking at.

Below we will report estimated quantile regression coefficients, but to start with we will provide a brief graphical analysis of the pre-dispatch and dispatch prices relationship for the 09:00 and 15:00 dispatch periods. First we concentrate on the 09:00 dispatch period using the log of spot price, $\log(P_{t,09:00})$ and the log of pre-dispatch price, $\log(\hat{P}_{t,09:00|PD1})$. If $\beta = 1$ and $\alpha = 0$, the log of pre-dispatch price was an unbiased forecast of the log of spot prices. This idealistic relationship is represented by the solid line in in the first panel of Figure 4.5-2. That panel also features a scatter plot of all pairs of ($\log(P_{t,09:00})$, $\log(\hat{P}_{t,09:00|PD1})$) observations.

(INSERT Figure 4.5-2 HERE)

The data points of the scatter plot should be scattered close to that line if the $\log(\hat{P}_{t,09:00|PD1})$ delivered efficient and unbiased forecasts for the $\log(P_{t,09:00})$. Although we observe a clear and strong positive correlation between pre-dispatch and spot prices, the $\log(\hat{P}_{t,09:00|PD1})$ series mostly seems to underestimate the $\log(P_{t,09:00})$. The scatter points have an increased density above the line representing forecast unbiasedness.

In the second and third panel of Figure 4.5-2 we show the respective scatter plots for pre-dispatches at *PD2* and *PD3*. The scatter plot for *PD2* shows some changes to that at *PD1*. The dispersion has reduced slightly and the underestimation bias is less clearly visible. When considering the scatter plot for *PD3* (the third panel of Figure 4.5-2) we can see another tightening of the scatterplot around the no-bias regression line. It therefore appears as if, when considering the 09:00 dispatch period, both intervals, between 15:00 on the day before dispatch and 06:00 on the dispatch day, as well as that between 06:00 and 08:30 (*PD3*) on the dispatch day, see the arrival of relevant new pre-dispatch information[50]. As

[50] What we mean here is the new arrival of either information that changes the demand forecasts or the arrival of new supply bids, whether they arise because of genuine new information or as a result of strategic re-bidding not backed by new information.

we get closer to dispatch, the pre-dispatch process delivers more accurate dispatch forecasts.

Next we consider the 15:00 dispatch period, and evaluate the relationship between the $\log(P_{t,15:00})$ and $\log(\hat{P}_{t,15:00|PD1})$ for *PD1* in the first panel of Figure 4.5-3. Compared to 09:00 morning dispatch period, the ability of the pre-dispatch prices to predict the next day's spot price outcomes at 15:00 deteriorates further. The most obvious indication for this is the vastly increased dispersion of the scatter diagram around the no-bias line. There is, however, no clear bias visible from the scatter plot. Another marked difference to the 09:00 dispatch is that the pre-dispatch hardly improves between *PD1* and *PD2* (in the second panel of Figure 4.5-3). There is little visible change in the dispersion. Only when we consider the *PD3* data (in the third panel of Figure 4.5-3) can we see a very clear tightening around the no-bias line. This indicates that a significant amount of new information enters into the pre-dispatch process between 06:00 and 14:30 (the time of the *PD3* pre-dispatch for the 15:00 dispatch period) on the dispatch day.

(INSERT Figure 4.5-3 HERE)

This may be seen as an early indication for pre-dispatch information on the day before pre-dispatch to have only limited informational value in the context of price spike prediction. In order to shed more light on this issue we will now estimate the quantile regression parameters α_p and β_p in

$$\log(P_{t,\tau}) = \alpha_p + \beta_p \log(\hat{P}_{t,\tau|j}) + errors \qquad (4.5\text{-}3)$$

Following the previous example, we first explain the relation between the $\log(P_{t,\tau})$ and the $\log(\hat{P}_{t,\tau|j})$ on all of dispatch days t (1st to 3166th) at $\tau = 09{:}00$ (upper panel of Figure 4.5-4) and 15:00 (lower panel of Figure 4.5-4). The estimated values for $\hat{\beta}_p$, for $p = 0.1\ to\ 0.9$, are displayed in Figure 4.5-4. Their placement on the horizontal axis is determined by the corresponding values of the unconditional deciles for $P_{t,\tau}$. For comparison purposes, the OLS coefficient estimate of β is shown with its 95% confidence interval as a horizontal line.

(INSERT Figure 4.5-4 HERE)

From these figures we can draw a number of interesting conclusions. At almost all deciles p we find $\hat{\beta}_p$ to be smaller than one. This suggests that a 1 unit increase in the logarithm of the pre-dispatch price does not result in an equivalently large increase in the resulting

conditional quantiles. Only the 90th conditional percentile (9th decile) has a $\hat{\beta}_{0.9}$ close to or exceeding 1. Also, the estimated coefficients $\hat{\beta}_p$ increase in p. This indicates that the conditional confidence intervals tend to increase with increasing pre-dispatch price.

(INSERT Figure 4.5-5 HERE)

This can best be illustrated referring to Figure 4.5-5. In that Figure we display percentiles for the dispatch price conditional on the relevant pre-dispatch price. The percentiles displayed are the 10th, 25th, 50th, 75th and 90th. It is clear that all these conditional percentiles increase with the pre-dispatch price, but the increase is much slower for the lower percentiles. This, of course, corresponds to smaller values of $\hat{\beta}_p$ for smaller values of p. A result of this is that any confidence interval for the dispatch price becomes larger with increasing values of the pre-dispatch price.

This, by itself, indicates that larger pre-dispatch values are associated with a larger degree of uncertainty with respect to the actual dispatch price outcome. As we know that electricity prices are positively skewed we would certainly expect that a certain degree of heteroscedasticity may be caused by the possibility of large price outcomes, even for the logarithm of prices[51]. This should show in $\hat{\beta}_p$ for large p that are much larger than 1. However, the type of heteroskedasticity we can see here is mainly due to the fact that the lower percentiles stay very low (values of $\hat{\beta}_p$ for small p that are much lower than 1). It is the $\hat{\beta}_p$ for the large values of p that are close to 1 which indicates that the upper percentiles of the confidence interval increase more or less in step with increases in the logarithm of the pre-dispatch prices.

As could be seen from the lower panel of Figure 4.5-4 the $\hat{\beta}_p$ values for the 15:00 dispatch period have a larger range (from 0.5 to 1.1) than those for the 09:00 dispatch period. This results in the conditional confidence interval for $\log(P_{t,\tau})$ fanning out even more as $\log(\hat{P}_{t,\tau|j})$ increases. This is clearly demonstrated in the lower panel of Figure 4.5-5. It is very clear that high pre-dispatch prices signal only to a certain extend an increased probability of higher spot prices, but they also signal a higher uncertainty about the resulting spot prices.

[51] The logarithm of the spot prices is still positively skewed (see Table 4.4-1) although the degree of skewness has decreases substantially.

(INSERT Figure 4.5-6 TO Figure 4.5-7 HERE)

We will now discuss how these findings change as we move from the pre-dispatch at PD1 to that at *PD2* (Figure 4.5-6) and PD3 (Figure 4.5-7). In Figure 4.5-6 we can see the estimated coefficients $\hat{\beta}_p$ for the 09:00 dispatch period (upper panel) and 15:00 dispatch period (lower panel) when pre-dispatch prices from *PD2* are used in the quantile regression. For the 09:00 dispatch period we can see that, compared to the results using the *PD1* information (refer to Figure 4.5-4), the estimated coefficients move somewhat closer to 1 which in turn results in a tightening of the conditional confidence intervals (as shown in the first panel of Figure 4.5-5). However, as for the *PD1* results, all coefficients stay below the value of 1. As this is true even for the highest quantile this indicates that even large pre-dispatch prices cannot be taken to give rise to significant price spike probabilities at the 09:00 dispatch period. When considering the *PD3* results for the 09:00 dispatch (Figure 4.5-7, upper panel) we can see another slight move of these coefficients towards 1 with the 90^{th} percentile coefficient now being very close to 1. In fact the quantile regression coefficients have now moved fairly close to the OLS coefficient value of approximately 0.94 for most quantiles. The most obvious deviation is for the 10^{th} percentile where the coefficient is close to 0.8. Altogether it appears as if both periods (between *PD1* – 15:00 on $t-1$ – and *PD2* – 06:00 on t – as well as between *PD2* and *PD3* – 08:30 on t) deliver new information to the pre-dispatch process for the 09:00 dispatch interval.

The same conclusion cannot be drawn when considering the results for the 15:00 dispatch period. There are hardly any changes in the estimated quantile regression coefficients when moving from *PD1* (Figure 4.5-4, lower panel) to (Figure 4.5-6, lower panel) *PD2*. This mirrors the absence of any qualitative change in the scatter plots in Figure 4.5-3. Even at *PD2* the quantile regression coefficients vary from about 0.53 to 1.18 (Figure 4.5-6, lower panel). As discussed previously this implies that there is a large degree of uncertainty about the actual dispatch price outcomes when pre-dispatch prices (at *PD2*) are large. It is therefore obvious that, as far as the 15:00 dispatch is concerned, the pre-dispatch process will still have to absorb a significant amount of information and very little new information has arrived between *PD1* and *PD2*. By 14:30 (the time of *PD3* for the 15:00 dispatch period) the situation has changed very significantly. The range of quantile coefficients has now tightened to a range from 0.82 to 1.08 (Figure 4.5-7, lower panel). In fact, at the Median the coefficient is $\hat{\beta}_{0.5} \approx 1.0$. But even at this stage the range of the coefficients

suggests that larger pre-dispatch prices are also indicative of larger conditional confidence intervals.

To illustrate the large qualitative change of conditional confidence intervals we plot the conditional quantiles for the 15:00 dispatch period for *PD1* and *PD3* in Figure 4.5-8[52]. We can see a substantial tightening of the conditional confidence interval which suggests that the pre-dispatch information available on day $t-1$ may only be of limited value as far as price event forecasting is concerned.

(INSERT Figure 4.5-8 HERE)

Given the convex nature of the electricity supply curves it is not necessarily surprising to find the type of heteroskedasticity observed in the above results. Ceteris paribus, higher pre-dispatch prices imply that we are operating in a steeper part of the supply curve and hence the same change in demand will result in larger absolute changes of the wholesale electricity price.

It is one of the aims of this chapter to establish whether pre-dispatch information can be used to predict extreme price events. In the context of our point process models this would require that information available on day $t-1$ has useful informational content relating to extreme price events on day t. Two aspects of the results presented in this Section make this seem unlikely. First, large pre-dispatch prices (in particular at *PD1*) seem to signal a large amount of uncertainty about dispatch price outcomes. The increased conditional confidence intervals are not particularly skewed towards high price outcomes. This makes it impossible to interpret high pre-dispatch prices as a clear signal for possible large price outcomes. Second, the pre-dispatch prices at *PD1* seem to be fairly imprecise (best illustrated by the wide spread observed in the scatter diagrams of Figure 4.5-2 and Figure 4.5-3 in the first panel). The closer we move to the dispatch time the more relevant information is worked in to the pre-dispatch price. This makes the *PD1* fairly uninformative.

Despite this pessimistic outlook it is, of course, necessary to test whether the inclusion of any pre-dispatch information does make a significant contribution in the Hawkes or PAR models used to predict extreme price events. The information from the pre-dispatch process that we use are the estimated pre-dispatch price (used as the explanatory variable

[52] The conditional quantiles for *PD2* are very similar to those from *PD1* in the upper panel of Figure 4.5-8 and therefore are not shown here.

in the above quantile regressions), $\widehat{P}_{t,\tau|j}$, as well as the estimated slope $\widehat{D}_{t,\tau|j}$ of the supply curve. As discussed in Section 4.4.2.5 the characteristics of the wholesale electricity supply curve mean that these will be highly correlated[53]. In addition to these pieces of information we shall also use the variable $B_{t,p|j}$ as described in Section 4.2.5.

4.5.3. Hawkes and PAR Models

In this Section we will establish whether extending the Hawkes and PAR models estimated in Chapter 2 with information from the pre-dispatch process delivers any insights. Here we will re-estimate the Hawkes models (HAWa and HAWab). The former allows the initial spike in intensity due to a price event to vary with covariates while the latter also allows the intensity decay to vary with covariates. In addition to the Hawkes models we will also re-estimate the PAR model. The basic models will use L_t, $Tmin_t$, $Tmax_t$ and C_t as covariates to control the time variation of intensity in the Hawkes models and the probability of event arrival and disappearance in the PAR model.

4.5.3.1. Model Specification

In this chapter we will extend these models with variables derived from the pre-dispatch process. The variables considered are those discussed previously. First we use the variable that measures bias in the pre-dispatch prices, $B_{t,p|j}$. As discussed in Section 4.4.2.5 this gives a measure of how biased the pre-dispatch process for period p ($p = d$ for day times or e for evenings[54]) of day t was, where $j = PD1, PD2 \; or \; PD3$. These measures are available on day t.

Further we use measures of the expected price (the value of the supply curve at the expected load) and the slope of the supply curve (at the expected load). In Section 4.4.2.5 we described how we estimated these measures on the basis of the available pre-dispatch scenarios. This resulted in the variables $\widehat{P}_{t,\tau|j}$ and $\widehat{D}_{t,\tau|j}$ which were available for any trading interval τ and for the different pre-dispatches $j = PD1, PD2 \; or \; PD3$. As the Hawkes and PAR models are models estimated on the basis of daily observations, this information had to be distilled into daily variables. Here we are interested in modelling extreme price events and therefore we decided to use the average of the 5 largest values of

[53] This is confirmed in the summary statistics reported in Table 4.5-1

[54] Since the pre-dispatch information is closely related (multicollinearity), we decided to apply a principal component analysis (PCA) to explain the variation amongst these variables. Based on the results, we decided to eliminate the measures of bias in the pre-dispatch prices during morning, $B_{t,m|j}$, as it does not provide any information in addition to the respective daytime and evening measures.

$\widehat{P}_{t,\tau|j}$ and $\widehat{D}_{t,\tau|j}$ across all $\tau = 1, \dots, 48$ on any day t and then take the logarithmic of the daily average series. The resulting daily series are labelled $\ln(Pm5_{t|j})$ and $\ln(Dm5_{t|j})$.

The timing of these variables is important, as we need to distinguish between the time at which the pre-dispatch information is available and dispatch day to which the pre-dispatch relates. Our convention here is that we will refer to the dispatch on day $t + 1$ (as we will want to forecast the probability of a price event on day $t + 1$). Of course, the *PD1* pre-dispatch for $t + 1$ is available on day t and that is what we will use in the subsequent estimating exercise as it is focused on forecasting performance.

In Table 4.5-1 we report the correlation matrix of the explanatory variables derived from pre-dispatch information. These variables are uncorrelated to any of the variables in our Base Model (L_t, $Tmin_t$, $Tmax_t$ and C_t) and hence it is well worth considering this extra information[55]. There are a number of variables amongst the pre-dispatch variables that are correlated with each other. The bias variables $B_{t,d|j}$ and $B_{t,e|j}$ are correlated at pre-dispatches *PD1* (correlated at 0.6016) and *PD2* (correlated at 0.5482). This is not the case at *PD3* (correlated at 0.1064), at which these bias variables appear fairly uncorrelated. This leads us to never include both bias variables for $= PD1 \; or \; PD2$. Further we find that our variables capturing the expected pre-dispatch price information, $Pm5_{t|j}$, and the slope of the supply curve at the expected load, $Dm5_{t|j}$, are strongly correlated at all $j = PD1, PD2 \; and \; PD3$ (correlated at 0.8913, 0.8796, 0.8752). This is not surprising as we know the supply curve to be convex. This then implies that we will include only one of these two variables.

(INSERT Table 4.5-1 HERE)

The Hawkes and PAR models used in this chapter are models that are highly nonlinear in their parameters. It is therefore difficult to implement a comprehensive general to specific search strategy. Hence we will implement the following estimation strategy. Each of our three point process models (HAWa, HAWab and PAR) will be estimated in 4 different versions. These models use the basic covariates L_t, $Tmin_t$, $Tmax_t$ and C_t and then 4 additional sets of variables derived from pre-dispatch information. The details of the covariate choice are provided in Table 4.5-2.

[55] To conserve space these correlations are not reported here but are available on request. None of these correlations is larger than 0.1 (in absolute terms).

(INSERT Table 4.5-2 HERE)

At this stage we limit our covariates to those that are useful in forecasting models. While we will introduce the forecasting setup below, it is important to be clear about the timing of the variables. We are building models that build a probability for the occurrence of price spikes at time $t+1$ using information available at time t. Therefore we need to ensure that all variables are available at time *t*, and therefore we restrict ourselves to $j = PD1$.

By way of example, Model 1 uses $B_{t,d|PD1}$ and $\ln(Pm5_{t|PD1})$ as additional covariates. For the HAWa this implies that these variables are used to explain variation in α_{t_i} as per equation (4.3-8). In the case of the HAWab model the covariates are used to allow variation in α_{t_i} and β_t as outlined in equations (4.3-8) and (4.3-10). In the PAR model the covariates control variation in the probability of stress arrival π_t (equation (4.3-16)) and the binomial thinning parameter ρ_t (equation (4.3-17)).

4.5.3.2. Estimation Results

In this section, we turn our attention to the estimation results for the models outlined above. We are not really interested in identifying one best model, but rather in whether the inclusion of the pre-dispatch information is valuable. We will therefore analyse whether the inclusion of the pre-dispatch information into our models delivers significant findings that are robust across our different models.

In the first instance we shall analyse the signs of the estimated parameters. As it turns out, the estimated signs are quite consistent across the different specifications for the HAWab and the PAR models. This is not the case, however, for the HAWa model (Table 4.5-3). We are including the extra coefficients into the information set that controls by how much a new price event increases (initially) the intensity of another price event for subsequent days.

(INSERT Table 4.5-3 HERE)

The inclusion of the bias variables results in a mixture of negative and positive coefficients for the day-time bias variable, $B_{t,d|PD1}$ and the evening sub-periods bias variable, $B_{t,e|PD1}$. The coefficients are clearly not significantly different from zero and hence we judge that the bias variables do not contribute significantly to explaining variation in the intensity. The inclusion of the expected pre-dispatch price variable $\ln(Pm5_{t+1|PD1})$ and the supply

curve slope variable, $\ln(Dm5_{t+1|PD1})$, also result in a mixture of positive and negative coefficients, and again they are all statistically insignificant.

In order to analyse the collective importance of pre-dispatch information in each of the models, we use the likelihood ratio (LR) test. The optimised log-likelihood and LR test results for each model are displayed in the bottom rows of the respective results table. Models 1 to 4 are compared to the basic model to establish whether the inclusion of the extra pre-dispatch variables add a significant amount of information. As it turns out Model 3 (at the 5% significance level) and Model 4 (at 1%) prove to be significant improvements to the Basic model despite the individual parameters not being significant.

It is possible that any correlation between the unseasonal load (L_{t+1}) and unexpected increase/decrease in temperatures ($Tmax_{t+1}, Tmin_{t+1}$) variables and the pre-dispatch variables may mask the significance of the pre-dispatch information in these models. We therefore re-estimate the models without L_{t+1}, $Tmax_{t+1}$ and $Tmin_{t+1}$ variables in order to establish whether, the pre-dispatch variables alone prove to be significant. The resulting models are, predictably, worse, and moreover even the removal of these variables do not render the pre-dispatch variables to be statistically significant.

The next set of results in Table 4.5-4 shows the results for the extended HAWab models. The pre-dispatch variables are included into both the α_{τ_i} and β_t time-varying parameters. Here the inclusion of the pre-dispatch variables results in a very consistent signs but also insignificant parameters. In general the pre-dispatch variables are particularly insignificant in those models (Models 1 and 2) in which the expected pre-dispatch price variable, $\ln(Pm5_{t+1|PD1})$ are included. In terms of estimated signs, we find that both the variables based on the expected pre-dispatch price, $\ln(Pm5_{t+1|PD1})$ and the supply curve slope, $\ln(Dm5_{t+1|PD1})$ enter α_{τ_i} with a negative and β_t with a positive sign.

This seems to indicate that events that occur on days for which high prices were expected (and the expected load intersected the supply curve in a relatively steeper part) have a smaller impact on subsequent intensity; and that on such days the intensity decays at a higher rate (larger β_t). The bias variables ($B_{t,d|PD1}$and $B_{t,e|PD1}$) also enter β_t positively as do the evening sub-periods bias variable, $B_{t,e|PD1}$ into α_{τ_i}. The day-time bias variable, $B_{t,d|PD1}$ are negative in the α_{τ_i} specification, but are also very close to 0.

(INSERT Table 4.5-4 HERE)

The likelihood ratio tests are insignificant at any conventional significance level and we therefore conclude that the inclusion of the pre-dispatch variables does not significantly improve the in-sample fit for the HAWab models.

(INSERT Table 4.5-5 HERE)

When including the pre-dispatch variables into the PAR model we find that their inclusion produces some individually significant coefficients and models which are significantly improved compared to the basic model. In the PAR model the time varying parameters π_t and ρ_t control the arrival and the survival probability of latent system stresses. First we will evaluate the effect of the pre-dispatch variables on the arrival probability π_t. The variables based on the expected pre-dispatch price, $\ln(Pm5_{t+1|PD1})$ and the supply curve slope, $\ln(Dm5_{t+1|PD1})$ enter negatively into the arrival probability[56]. This seems to indicate that stresses are less likely to arrive in days in which high prices are anticipated. One interpretation would be that extreme price events are mainly due to stresses that are not anticipated in the pre-dispatch, or, expressed somewhat differently, that a market that already expects high prices may be in a better position to absorb any further adverse events without extreme price events. However, these coefficients are at most marginally significant. The inclusion of the bias variables ($B_{t,d|PD1}$ and $B_{t,e|PD1}$) proves more significant. They all enter positively implying that if the pre-dispatch process underestimated the dispatch price outcomes for day *t*, then it is more likely that we will see a new system stress on day *t*+1.

The inclusion of the pre-dispatch variables in the specification for the stress survival probability ρ_t results in negative coefficients for bias, expected pre-dispatch price and slope variables. The only coefficients that are statistically significant, however, are those of the day-time bias, $B_{t,d|PD1}$. This implies that, ceteris paribus, a larger under-estimation of the dispatch prices on day *t* (i.e. larger values for $B_{t,d|PD1}$), results in a smaller survival probability for stresses at time *t*+1. There is no immediately obvious interpretation of this result.

[56] Given the particular specification for π_t (see equation (4.3-16)) and ρ_t (see equation (4.3-17)), the partial derivative of these terms with respect to a particular variable, has the same sign as the respective coefficient.

$$\frac{\partial(1-\exp(-\exp(x\beta)))}{\partial x} = \beta\exp(-\exp(x\beta)+x\beta)$$

Which will take the same sign as β as $\exp(\cdot)$ is positive.

When evaluating the statistical significance of the inclusion of the four pre-dispatch variables into Models 1 to 4 of PAR model, we find all LR statistics to be significant at 5% but not at 1%. We therefore conclude that while the inclusions seem to be marginally significant, these effects do not appear dramatic.

4.5.3.3. Forecasting

In this section we want to investigate whether pre-dispatch information is able to improve the ability of Hawkes and PAR models in predicting next-day price event.

In the forecast exercise, the intensities of next-day price events are calculated during the forecast period from October 31, 2005 to October 31, 2007. The parameters used for the forecast are re-estimated every 30 days. We use a sliding estimation window of constant size. For example, the parameters vectors of all the models are estimated by maximising the log-likelihood function in equations (4.3-11) and (4.3-20) using information from the initial estimation period (March 2, 1999 to October 30, 2005). These parameters vectors are then used to produce one-day-ahead conditional intensity forecast from October 31, 2005 to November 29, 2005 using equation (4.3-7) and equation (2.7-10). The one-day-ahead conditional intensity forecast for the Hawkes model is denote by $\lambda_{t|t-1}$ while for PAR model is denote by $p_{t|t-1}$. This first set of estimated model parameters is used in the calculation of the first 30 one-step-ahead intensity forecasts. The parameters are re-estimated every 30 days, meaning that the second estimation window includes data from March 31, 1999 to November 29, 2005. These updated parameter estimates are then used to produce the next 30 one-day-ahead conditional intensity forecast from November 30, 2005 to December 29, 2005. This process continues until we obtain two years' worth of intensity forecasts.

We are also using a 'naïve' model to compare the forecast performance of Hawkes and PAR models. As mentioned in Chapter 2, the naïve model predicts spikes on day, t only if there is a spike on day $t-1$ and the next day, t falls on a weekday. In this case the one-day-ahead conditional intensity forecast of the naïve model takes a value of 1. If the condition is not satisfied, the intensity forecast is equal to 0. The forecast model specification follows Table 4.5-2. By comparing the forecast performance of the Basic

models with those models which include pre-dispatch information in the conditioning covariates set, we are able to determine whether pre-dispatch information carry valuable information that could improve forecast of the next day price spikes intensities.

The forecast performance of the models is measured using *Mean Absolute Error* (*MAE*), *Root Mean Square Error* (*RMSE*) and the *Asymmetric loss score* (*ASYM*) (see; Rudebusch & Williams (2009) and Christensen et al. (2009)). All of these measurements are based on the forecast error ($y_t - \lambda_{t|t-1}$) of Hawkes and ($y_t - p_{t|t-1}$) PAR models.

(INSERT Table 4.5-6 HERE)

When comparing the forecast performance of the models that include pre-dispatch information with the basic models, it transpires that it is only the HAWab model that responds positively to the inclusion of the pre-dispatch variable. This result mirrors the previous finding that indicated the collective significance of the pre-dispatch variables was largest for the HAWab class of models. For neither the PAR nor the HAWa models can we see any discernible difference in forecasting performance.

When including the naïve model into the forecasting performance comparison, the results are even more disappointing as the basic price spike prediction rule outlined above produces the lowest forecast MAE criterion. The fact that the naïve model does worst for the RMSE criterion indicates, however, that it does tend to make more larger mistakes which are more heavily penalised in a criterion using a squared measure. That is due to the fact that the naïve model, by design takes extreme intensity predictions.

In the light of our findings in Chapter 3, it is it is possibly not surprising to find that the pre-dispatch process of the Australian Electricity Market does not provide any information that can be used in a systematic manner to help predicting on what days price spikes are more likely to occur.

4.6. Conclusions

Following our work in Chapter 3, where we analyse whether there is an apparent bias in the pre-dispatch process; in this Chapter we investigate whether information from the pre-dispatch process can be useful in forecasting next-day price spikes. In our quest to uncover the potential of pre-dispatch information in a forecasting context, we conducted several preliminary analyses. These analyses consist of a careful analysis of correlations and quantile regressions. From the results of these analyses we find evidence that pre-dispatch prices are a surprisingly week predictor of price outcomes.

Despite the finding, we test whether the inclusion of any pre-dispatch information does make a significant contribution in the Hawkes or PAR models used to predict extreme price events. In doing so, we are not really interested in identifying one best model, but whether the inclusion of the pre-dispatch information into the models is valuable. For the HAWa models, the pre-dispatch variables are included into α_{τ_i} time-varying parameter which control by how much a new price event increases (initially) the intensity of another price event for subsequent days. While for the HAWab models, the pre-dispatch variables are included both into the α_{τ_i} and β_t time-varying parameters. Finally in the PAR model we evaluate the effect of the pre-dispatch variables into the time varying parameters π_t and ρ_t that control the arrival and the survival probability of latent system stresses.

As it turns out the best forecast performance comes from models without inclusion of pre-dispatch information in their conditioning covariates set. This means the inclusions of the pre-dispatch information into these point process models do not deliver significant findings that are robust across different models. Hence, it transpires that the pre-dispatch process of the Australian New Electricity Market does not provide any information that can be used in a systematic manner to help predicting price spikes on the next day. These findings are, of course, conditional on the specific information extracted from the pre-dispatch process. It is possible that different pieces of information embodied in that rich data-set may be more useful. It is, however, not obvious what that information should be.

References

Aggarwal, S. K., Saini, L. M., & Kumar, A. (2008). Electricity price forecasting in Ontario electricity market using wavelet transform in artificial neural network based model. *International Journal of Control, Automation, and Systems*, *6*(5), 639–650.

Amjady, N., & Keynia, F. (2008). Day ahead price forecasting of electricity markets by a mixed data model and hybrid forecast method. *International Journal of Electrical Power & Energy Systems*, *30*(9), 533–546. doi:10.1016/j.ijepes.2008.06.001

Amjady, N., & Keynia, F. (2009). Day-Ahead Price Forecasting of Electricity Markets by Mutual Information Technique and Cascaded neuro-evolutionary algorithm. *Power*, *24*(1), 306–318.

Andalib, A., & Atry, F. (2009). Multi-step ahead forecasts for electricity prices using NARX: a new approach, a critical analysis of one-step ahead forecasts. *Energy Conversion and Management*, *50*(3), 739–747. doi:10.1016/j.enconman.2008.09.040

Arciniegas, A. I., & Arciniegas Rueda, I. E. (2008). Forecasting short-term power prices in the Ontario Electricity Market (OEM) with a fuzzy logic based inference system. *Utilities Policy*, *16*(1), 39–48.

Australian Energy Market Commission. (2013). National Electricity Rules Version 56 Chapter 3: Market Rules. Retrieved from http://www.aemc.gov.au/Electricity/National-Electricity-Rules/Current-Rules.html

Australian Energy Market Operator. (2010). An introduction to australia's national electricity market. *Available at AEMO website http://http://www.aemo.com.au/About-the-Industry/Energy-Markets/National-Electricity-Market*.

Australian Energy Market Operator. (2013). *Executive Summary of National Electricity Forecasting Report*. *New directions for youth development* (Vol. 2013, pp. 9–13). doi:10.1002/yd.20054

Barlow, M. (2002). A diffusion model for electricity prices. *Mathematical Finance*, *12*(4), 287–298. Retrieved from http://onlinelibrary.wiley.com/doi/10.1111/j.1467-9965.2002.tb00125.x/abstract

Bauwens, L., & Hautsch, N. (2006). Stochastic Conditional Intensity Processes. *Journal of Financial Econometrics*, *4*(3), 450–493. doi:10.1093/jjfinec/nbj013

Becker, R., Enders, W., & Hurn, S. (2004). A general test for time dependence in parameters. *Journal of Applied Econometrics*, *19*(7), 899–906. doi:10.1002/jae.751

Brandimarte, P. (2013). *Numerical methods in finance and economics: a MATLAB-based introduction*. John Wiley & Sons.

Bunn, D. (2000). Forecasting loads and prices in competitive power markets. *Proceedings of the IEEE*, *88*(2), 163–169. Retrieved from http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=823996

Bunn, D., Andresen, A., Chen, D., & Westgaard, S. (2012). Analysis and Forecasting of Electricity Price Risks with Quantile Factor Models.

Burger, M., Klar, B., M ller, A., & Schindlmayr, G. (2004). A spot market model for pricing derivatives in electricity markets. *Quantitative Finance*, *4*(1), 109–122.

Catalão, Jps., Mariano, S., Mendes, V. M. F., & Ferreira, L. (2007). Short-term electricity prices forecasting in a competitive market: a neural network approach. *Electric Power Systems Research*, *77*(10), 1297–1304.

Chevallier, J. (2010). The impact of Australian ETS news on wholesale spot electricity prices: an exploratory analysis. *Energy Policy*, *38*(8), 3910–3921. doi:10.1016/j.enpol.2010.03.014

Christensen, T., Hurn, S., & Lindsay, K. (2009). It never rains but it pours: modeling the persistence of spikes in electricity prices. *Energy Journal*, *30*(1), 25.

Christensen, T. M., Hurn, a. S., & Lindsay, K. a. (2012). Forecasting spikes in electricity prices. *International Journal of Forecasting*, *28*(2), 400–411. doi:10.1016/j.ijforecast.2011.02.019

Clements, A., Fuller, J., & Hurn, S. (2012). *Semi-parametric forecasting of Spikes in Electricity Prices*.

Conejo, A. J., Contreras, J., Espínola, R., & Plazas, M. a. (2005). Forecasting electricity prices for a day-ahead pool-based electric energy market. *International Journal of Forecasting*, *21*(3), 435–462. doi:10.1016/j.ijforecast.2004.12.005

De Jong, C., & Huisman, R. (2002). Option formulas for mean-reverting power prices with spikes. *Energy Global Research Paper*.

Duan, J. (n.d.). Electricity Price Forecasting by Autoregressive Model and Artificial Neutral Network. *ee.columbia.edu*, 1–11. Retrieved from http://www.ee.columbia.edu/~lavaei/Projects/Jinrui-Duan_Course.pdf

Eichler, M., Grothe, O., Manner, H., & Tuerk, D. (2012). *Modeling spike occurrences in electricity spot prices for forecasting*. METEOR, Maastricht research school of Economics of Technology and Organizations. Retrieved from http://arno.unimaas.nl/show.cgi?fid=25393

Electricity Market Performance. (2012a). *Regional Demand Definition*.

Electricity Market Performance. (2012b). *Factors Contributing to Differences Between Dispatch and Pre-Dispatch Outcomes*.

Federal Energy Regulatory Commission. (1998). *Staff report to the Federal Energy Regulatory Commission on the causes of wholesale electric pricing abnormalities in the Midwest during June 1998*. *Report. Federal Energy Regulatory Commission, Washington DC*.

Geman, H., & Roncoroni, A. (2006). Understanding the Fine Structure of Electricity Prices. *The Journal of Business*. Retrieved from http://www.jstor.org/stable/10.1086/500675

Hawkes, A. G. (1971). Spectra of some self-exciting and mutually exciting point processes. *Biometrika*, *58*(1), 83–90.

Heydari, S., & Siddiqui, A. (2010). Valuing a gas-fired power plant: A comparison of ordinary linear models, regime-switching approaches, and models with stochastic volatility. *Energy Economics*, *32*(3), 709–725.

Kanamura, T., & Ōhashi, K. (2008). On transition probabilities of regime switching in electricity prices. *Energy Economics*, *30*, 1158–1172. doi:10.1016/j.eneco.2007.07.011

Koenker, R., & Jr, G. B. (1978). Regression quantiles. *Econometrica: journal of the Econometric Society*, *46*(1), 33–50. Retrieved from http://www.jstor.org/stable/10.2307/1913643

Kwoka, J. (2012). Price Spikes in Wholesale Electricity Markets: An Update on Problems and Policies, (March). Retrieved from http://warrington.ufl.edu/centers/purc/purcdocs/PAPERS/TRAINING/events/PPRC/0312_PPRC/adddocs/Conference_Paper-Price_Spikes.pdf

Lucia, J. J., & Schwartz, E. S. (2002). Electricity prices and power derivatives: Evidence from the nordic power exchange. *Review of Derivatives Research*, *5*(1), 5–50.

Mincer, J. (1969). *Economic forecasts and expectations: analyses of forecasting behavior and performance*. National Bureau of Economic Research.

Moghram, I., & Rahman, S. (1989). Analysis and evaluation of five short-term load forecasting techniques. *IEEE Transactions on Power Systems*, *4*(4), 1484–1491.

Mount, T., Ning, Y., & Cai, X. (2006). Predicting price spikes in electricity markets using a regime-switching model with time-varying parameters. *Energy Economics*, *28*, 62–80. doi:10.1016/j.eneco.2005.09.008

Newbold, P., & Harvey, D. I. (2002). Forecast combination and encompassing. *A companion to economic forecasting*, 268–283.

Nogales, F., & Contreras, J. (2002). Forecasting next-day electricity prices by time series models. *IEEE Transactions on Power Systems*, *17*(2), 342–348. Retrieved from http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1007902

Ogata, Y. (1978). The asymptotic behaviour of maximum likelihood estimators for stationary point processes. *Annals of the Institute of Statistical Mathematics*, *30*(1), 243–261. Retrieved from http://www.springerlink.com/index/247815U5W2407822.pdf

Perlich, C., Rosset, S., Lawrence, R. D., & Zadrozny, B. (2007). High-quantile modeling for customer wallet estimation and other applications. *Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '07*, 977. doi:10.1145/1281192.1281297

Rudebusch, G. D., & Williams, J. C. (2009). Forecasting Recessions: The Puzzle of the Enduring Power of the Yield Curve. *Journal of Business & Economic Statistics*, *27*(4), 492–503. doi:10.1198/jbes.2009.07213

Song, Y.-H., & Wang, X.-F. (2003). *Operation of market-oriented power systems*. Springer.

System Operations. (2011a). *Power System Operating Procedure Load Forecasting*.

System Operations. (2011b). *Power System Operating Procedure Load Forecasting*.

Vucetic, S. (2001). Discovering price-load relationships in California's electricity market. *IEEE Transactions on Power Systems*, *16*(2), 280–286. Retrieved from http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=918299

Weron, R. (2007). *Modeling and forecasting electricity loads and prices: A statistical approach* (Vol. 403, p. 45,79). Wiley. com.

Wu, L., & Shahidehpour, M. (2010). A hybrid model for day-ahead price forecasting. *IEEE Transactions on Power Systems*, *25*(3), 1519–1530. Retrieved from http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5409501

Zareipour, H, & Bhattacharya, K. (2006). Forecasting the Hourly Ontario Energy Price by Multivariate Adaptive Regression Splines, 1–7.

Zareipour, Hamidreza, Cañizares, C. A., Bhattacharya, K., & Thomson, J. (2006). Application of public-domain market information to forecast Ontario's wholesale electricity prices. *IEEE Transactions on Power Systems*, *21*(4), 1707–1717.

Zhao, J., Dong, Z. Y., Li, X., & Wong, K. P. (2005). A general method for electricity market price spike analysis. In *Power Engineering Society General Meeting, 2005. IEEE* (pp. 286–293).

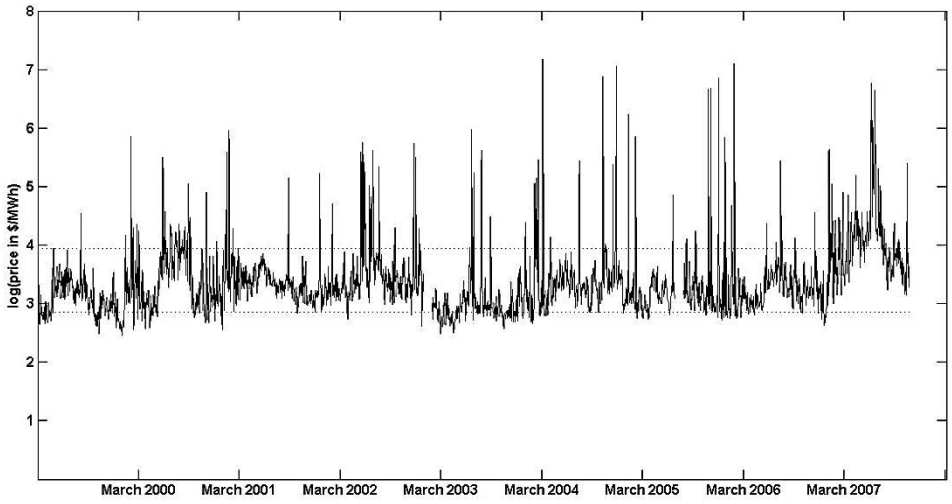

Figure 4.1-1: This figure is showing the time series of (log) daily average spot prices from 1999 to 2007 in NSW.

The daily average spot prices are calculated by taking the average of the 48 half-hour spot prices in a day. Since we have a missing observation from January 1 to 31, 2003 and July 2 to August 1, 2005 a total of 148,992 half-hourly spot prices are used to calculate the daily average spot prices from March 2, 1999 to October 31, 2007. In order to show the low and high extremity of the daily average spot price series, we calculated their 10^{th} and 90^{th} percentiles of the daily average series. The 10^{th} percentile of the series is A$17.39/MWh and the 90th percentile of the series is $51.45/MWh. The extreme temporary price jumps feature in the price series are made obvious by taking the logarithmic of the daily average spot prices.
It can be seen that the logarithmic of the price series exceeded the logarithmic of the 90^{th} percentile value frequently.

	1999	2000	2001	2002	2003	2004	2005	2006	2007
Mean	22.92(3.08)	35.63(3.40)	33.27(3.40)	39.76(3.48)	24.37(3.00)	45.14(3.41)	35.83(3.24)	31.01(3.25)	72.88(4.01)
Median	21.16(3.05)	26.31(3.27)	28.15(3.34)	28.29(3.34)	18.45(2.92)	28.08(3.34)	23.02(3.14)	24.72(3.21)	51.14(3.93)
Std.Dev	7.99(0.30)	29.49(0.53)	30.47(0.36)	41.62(0.51)	35.35(0.42)	118.38(0.57)	83.09(0.52)	64.58(0.40)	92.40(0.61)
Skewness	2.96(0.70)	5.28(0.97)	8.19(3.33)	4.40(2.45)	7.96(4.35)	8.90(3.61)	8.45(4.36)	17.03(3.69)	5.46(1.57)
90^{th} prctile	32.99(3.50)	60.27(4.10)	39.76(3.68)	48.56(3.88)	25.71(3.25)	45.96(3.83)	33.67(3.52)	37.76(3.63)	102.37(4.63)
95^{th} prctile	36.32(3.59)	70.38(4.25)	44.98(3.81)	112.62(4.72)	34.91(3.55)	62.36(4.13)	53.56(3.98)	43.44(3.77)	185.44(5.22)
99^{th} prctile	47.61(3.86)	154.18(5.04)	184.36(5.22)	265.11(5.58)	240.03(5.48)	855.12(6.65)	487.10(6.18)	104.81(4.65)	589.90(6.37)
% days prcspikes	22.92	35.63	33.27	39.76	24.37	45.14	35.83	31.01	72.88
% hh prcspikes	21.16	26.31	28.15	28.29	18.45	28.08	23.02	24.72	51.14
$Mean^{CountData}$	0.049	0.809	0.384	0.753	0.171	0.915	0.540	0.430	4.924

Table 4.4-1: This table shows the descriptive analysis of a daily average spot prices and its logarithmic series (in bracket) for each year throughout the investigation period (02/03/1999 to 31/10/2007)[57] in $/Mwh unit.

The daily prices are calculated by taking an average of 48 half-hour spot prices in a day. Once the daily series is obtained the descriptive analysis shown in the table are calculated. We then take a logarithmic on the daily series to produce the descriptive analysis shown in bracket.

Std.Dev is standard deviation and 90^{th} prctile, 95^{th} prctile and 99^{th} prctile is the 90^{th}, 95^{th} and 99^{th} percentile of the daily average spot prices.
Then the ninth and tenth rows used to display the proportion of days and half-hour with price spikes (% days prcspikes, % hh prcspikes) is in percentage unit. Using year 2000 as an example, they are calculated as follows:

$$\%\ \text{days prcspikes in year } 2000 = \left(\frac{\sum_{t=1}^{366} y_t}{366}\right) \times 100\%$$

$$\%\ \text{hh prcspikes in year } 2000 = \left(\frac{\sum_{t=1}^{366}\sum_{\tau=1}^{48} I_{P_{t,\tau}>th}}{366 \times 48}\right) \times 100\%$$

In the last row, $Mean^{CountData}$ is a mean of Count Data for each year. The Count Data is measured by counting the number of times the half-hourly spot prices exceeds A$100/MwH during a day. Further discussion on the Count Data is available in Section 4.4.2.4

[57] The calculated descriptive analysis in the year 1999, 2003, 2005 and 2007 are based on less than a year daily spot prices since in year 1999, the observation starts on March 2 and in year 2007, the observation ends on October 31. While in 2005 and 2007 we have a missing observation from January 1 to 31, 2003 and July 2 to August 1, 2005

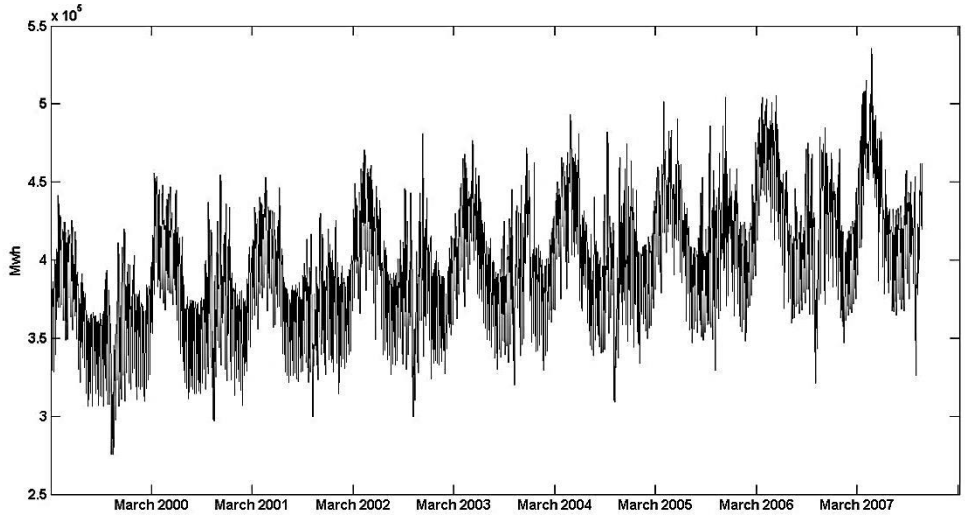

Figure 4.4-1: This figure is showing 3166 daily electricity consumption from March 2, 1999 to October 31, 2007 in NSW

The daily load series are calculated by summing all the 48 half-hour load series in a day. For the purpose of illustration, we are using a total of 151,968 half-hourly load series obtained from AEMO's website in Aggregate Price and Demand Data Files (http://www.aemo.com.au/Electricity/Data/Price-and-Demand/Aggregated-Price-and-Demand-Data-Files).

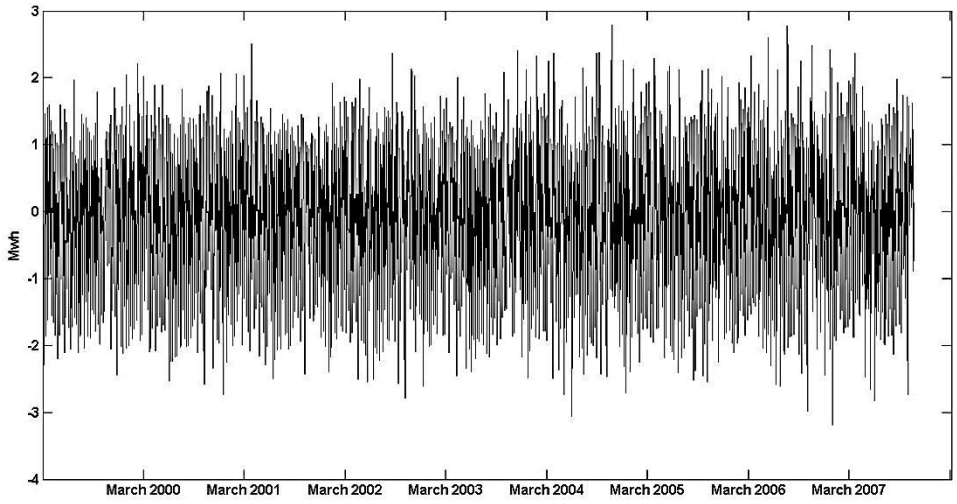

Figure 4.4-2: This figure is showing daily deseasonalized load time series, L_t plotted from March 2, 1999 to October 31, 2007 in NSW.

Details explanation on the construction of the series is as explained in the caption of Figure 2.4-2. The daily L_t series show no obvious trends or seasonality making the series is a plausible candidate to capture the forecast electricity demand that exceeds the demand we would usually expect for that day of the week at that particular time of the year.

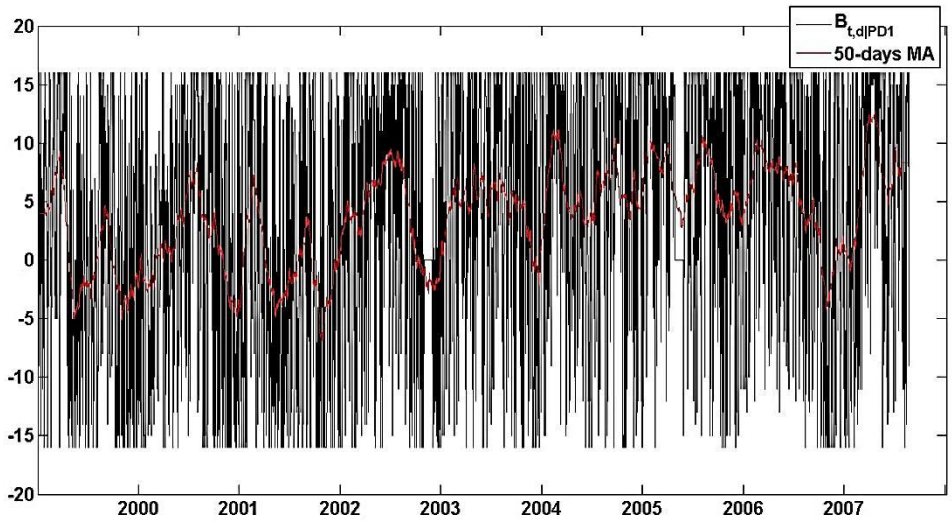

Figure 4.4-3: This figure is showing a daily time-series of $B_{t,d|PD1}$ and its 50 day moving average series.

The daily time-series of $B_{t,d|PD1}$ are calculated as in equation (4.4-2) during daytime period or equivalent to 16 trading interval. In constructing this series we are using 16 half-hourly observations from 10:00 to 17:30 of electricity spot prices, $P_{t,d}$ and the pre-dispatch +/- 200 MWh scenario prices on *PD1*, $P_{t,d|PD1}^{\pm\Delta}$. A 50 day moving average series are superimposed on the graph to highlight longer term trends of $B_{t,d|PD1}$ series.

A gap that we observed in the figure is caused by missing observation as mention in the caption of Figure 4.1-1.

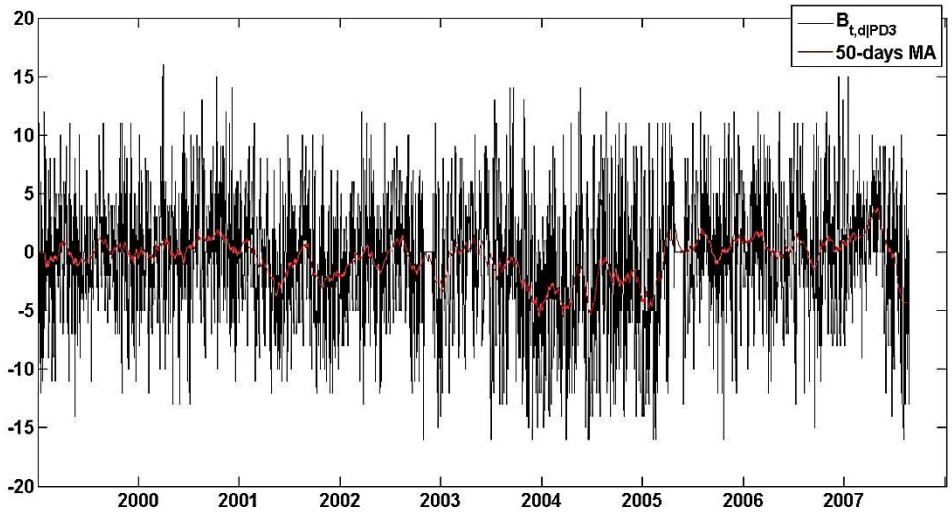

Figure 4.4-4: This figure is showing a daily time-series of $B_{t,d|PD3}$ and its 50 day moving average series.

Details explanation on the construction of the series is as explained in the caption of Figure 4.4-3 except that scenario prices are on *PD3*. A gap that we observed in the figure is caused by missing observation as mention in the caption of Figure 4.1-1.

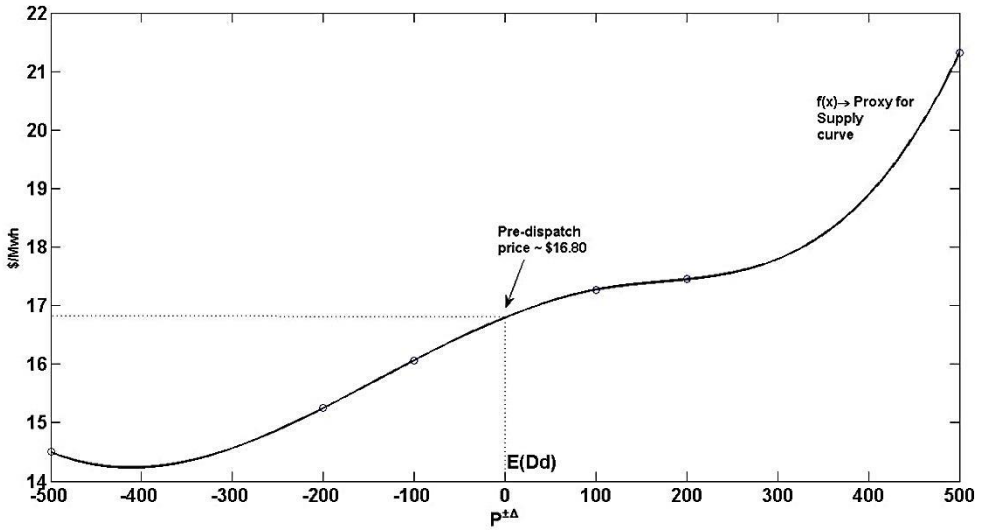

Figure 4.4-5: This figure is showing a cubic spline function, $f(x)$ connecting six available pre-dispatch prices ($P^{-500}, P^{-200}, P^{-100}, P^{+100}, P^{+200}, P^{+500}$) at their respective relative load offset values (-500, -200, -100, 100, 200, 500).

In the figure, $f(x)$ is used to proxy the actual pre-dispatch prices from a supply curve inferred from $P^{\pm\Delta}$. The pre-dispatch price is interpolated at the intersection of the supply curve and a load offset value of 0. For illustration purpose in this figure, the estimated pre-dispatch price $\hat{P}_{t,\tau|j}$ is specific for the fifth sub-period (09:00) on dispatch day t and formed during pre-dispatch on *PD1*.

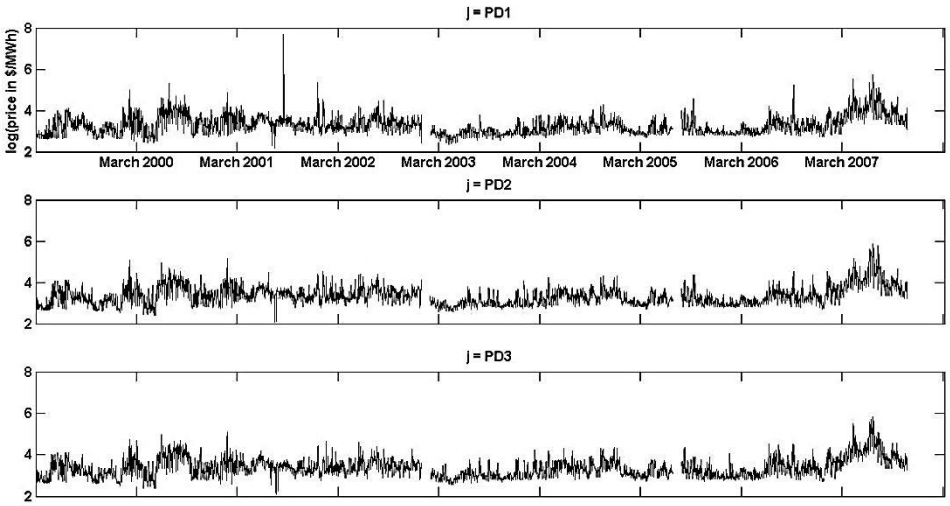

Figure 4.4-6: This figure is showing a logarithmic series of the $\hat{P}_{t,09:00|PD1}$, $\hat{P}_{t,09:00|PD2}$, $\hat{P}_{t,09:00|PD3}$.

To construct the series in the first panel, we take the fifth sub-period (09:00) of the estimated pre-dispatch price, $\hat{P}_{t,09:00|PD1}$ from March 2, 1999 to October 31, 2007. The estimated pre-dispatch price of the fifth sub-period (09:00) on one particular dispatch day t is calculated based on the approach describe in Figure 4.4-5. In order to make the features of the time-series more prominent, we takes the logarithmic of the daily $\hat{P}_{t,09:00|PD1}$. A gap that we observed in the figure is caused by missing observation mentioned in the caption of Figure 4.1-1. Same description applies for the construction of the series in the second and third panels except they are based on pre-dispatch process during *PD2* and *PD3*.

$P_{t,m}$	1999	2000	2001	2002	2003	2004	2005	2006	2007
Mean	26.28	43.18	31.00	29.50	19.59	28.86	23.29	25.28	60.56
Median	22.55	28.75	29.71	28.16	18.66	27.19	21.78	23.18	47.53
Std.Dev	29.12	71.84	10.06	8.71	5.46	9.48	6.68	8.34	43.64
90th prctile	38.71	66.02	41.83	40.69	25.63	41.67	30.63	35.91	99.10
95th prctile	43.05	86.90	47.09	46.12	31.36	46.84	34.79	40.52	151.98
99th prctile	53.50	366.33	63.85	58.27	37.34	57.83	54.79	54.38	252.93
$\widehat{P}_{t,m\|PD1}$									
Mean	22.31	31.71	31.46	25.91	17.36	23.82	20.55	22.46	47.06
Median	20.40	26.43	27.71	24.37	17.05	22.64	19.08	20.38	40.40
Std.Dev	7.84	23.50	45.07	7.41	3.84	6.83	6.94	8.04	27.03
90th prctile	33.37	53.19	39.24	35.53	21.60	33.40	26.69	30.90	78.97
95th prctile	38.52	62.04	43.10	38.28	24.13	36.16	30.06	33.15	94.55
99th prctile	44.37	92.49	66.44	52.11	28.91	44.72	42.29	43.65	155.81
$\widehat{P}_{t,m\|PD2}$									
Mean	22.37	34.24	30.45	28.45	18.65	26.52	22.11	23.50	55.67
Median	20.31	27.34	29.20	27.03	17.87	24.72	20.49	21.54	44.70
Std.Dev	8.79	41.26	10.48	8.10	4.64	8.14	6.53	7.20	37.06
90th prctile	35.18	55.85	40.76	38.71	23.67	37.37	29.45	32.49	96.33
95th prctile	40.65	63.59	47.40	43.28	28.21	40.82	33.41	35.99	124.00
99th prctile	52.08	82.60	57.94	54.82	34.31	54.07	53.66	45.50	216.96
$\widehat{P}_{t,m\|PD3}$									
Mean	22.70	36.32	30.18	28.66	19.07	27.43	22.38	24.65	56.30
Median	20.36	27.94	29.11	27.48	18.08	25.69	20.84	22.41	45.27
Std.Dev	8.89	48.35	10.61	8.12	5.05	8.98	6.31	8.53	36.95
90th prctile	35.59	59.34	40.89	39.01	24.29	38.82	30.62	34.07	95.35
95th prctile	40.56	67.42	45.54	42.35	30.05	45.09	33.09	41.30	128.81
99th prctile	49.61	110.39	67.77	56.15	36.02	53.21	49.72	55.49	222.98

Table 4.4-2: This table shows the descriptive analysis of a daily average spot prices for each year in the morning sub-period, $P_{t,m}$ (06:30 to 09:30 dispatch intervals - 7 trading intervals, $\tau = m$) and the corresponding estimated pre-dispatch prices at $j = PD1, PD2\ or\ PD3$ in the morning sub-period, $\hat{P}_{t,m|j}$.
The daily prices are calculated by taking an average of 7 half-hour spot prices in a day. Once the daily series is obtained the descriptive analysis shown in the table are calculated.
Std.Dev is standard deviation and 90th prctile, 95th prctile and 99th prctile is the 90th, 95th and 99th percentile of the daily average spot prices. Then the ninth and tenth rows used to display the proportion of days and half-hour with price spikes (% days prcspikes, % hh prcspikes) is in percentage unit.

Further explanation on this table is similar to the caption of Table 4.4-1, except $\tau = 1, \ldots, 7 (06{:}30 \text{ to } 09{:}30)$

$P_{t,d}$	1999	2000	2001	2002	2003	2004	2005	2006	2007
Mean	23.81	39.37	43.15	37.38	20.73	78.64	60.30	40.33	79.37
Median	21.99	30.98	29.75	30.89	18.80	28.38	23.88	26.65	56.47
Std.Dev	7.80	54.67	79.46	40.87	7.13	347.08	247.72	187.20	83.84
90^{th} prctile	34.71	61.40	51.38	46.70	27.33	67.00	41.12	41.45	127.27
95^{th} prctile	39.89	74.94	70.75	59.93	35.10	94.13	57.82	58.32	209.45
99^{th} prctile	47.41	152.46	478.60	152.26	55.80	2426.06	1407.47	141.28	534.67
$\widehat{P}_{t,d\|PD1}$									
Mean	22.85	42.85	56.84	36.56	19.75	82.37	25.71	26.36	63.28
Median	20.96	30.40	29.18	27.96	17.42	25.05	20.78	23.64	51.92
Std.Dev	7.73	153.86	245.14	59.70	11.94	393.33	19.25	14.77	61.89
90^{th} prctile	32.63	55.93	61.12	47.02	25.64	45.70	35.52	33.97	95.51
95^{th} prctile	38.82	70.23	82.72	66.14	30.01	91.76	43.63	45.65	128.65
99^{th} prctile	49.53	155.33	224.68	173.50	46.55	2031.46	157.42	73.46	263.17
$\widehat{P}_{t,d\|PD2}$									
Mean	23.51	45.73	53.54	37.95	19.99	66.31	29.68	26.81	78.76
Median	21.24	31.22	30.16	30.15	17.89	25.93	22.26	24.16	53.89
Std.Dev	8.23	186.36	215.85	48.59	7.06	298.33	45.48	19.12	188.41
90^{th} prctile	33.82	55.91	62.13	52.26	26.44	53.44	37.63	34.69	109.58
95^{th} prctile	42.09	67.51	86.74	77.50	30.33	103.61	43.87	41.87	138.93
99^{th} prctile	51.97	158.67	334.63	132.44	48.48	1003.93	174.11	84.92	313.57
$\widehat{P}_{t,d\|PD3}$									
Mean	23.69	39.41	42.07	36.70	20.70	95.89	60.96	49.10	75.00
Median	21.88	30.92	29.97	31.55	18.58	28.64	24.11	26.25	55.72
Std.Dev	7.88	69.80	78.75	39.12	7.11	458.53	224.90	243.96	94.75
90^{th} prctile	35.24	56.50	51.69	48.17	26.46	64.50	41.29	41.71	116.48
95^{th} prctile	38.87	69.69	64.32	58.38	33.61	126.09	55.81	52.61	164.02
99^{th} prctile	49.44	145.88	260.53	122.07	52.79	2403.60	1387.04	392.50	341.33

Table 4.4-3: This table shows the descriptive analysis of a daily average spot prices for each year in the daytime sub-period, $P_{t,d}$ (10:00 to 17:30 dispatch intervals - 16 trading intervals, $\tau = d$) and the corresponding estimated pre-dispatch prices at $j = PD1, PD2\ or\ PD3$ in the daytime sub-period, $\hat{P}_{t,d|j}$.
The daily prices are calculated by taking an average of 16 half-hour spot prices in a day. Once the daily series is obtained the descriptive analysis shown in the table are calculated.
Std.Dev is standard deviation and 90th prctile, 95th prctile and 99th prctile is the 90th, 95th and 99th percentile of the daily average spot prices. Then the ninth and tenth rows used to display the proportion of days and half-hour with price spikes (% days prcspikes, % hh prcspikes) is in percentage unit.

Further explanation on this table is similar to the caption of Table 4.4-1, except $\tau = 8, \ldots, 23 (10{:}00 \text{ to } 17{:}30)$

$P_{t,e}$	1999	2000	2001	2002	2003	2004	2005	2006	2007
Mean	25.09	38.10	33.29	70.07	42.58	37.10	30.11	35.77	111.59
Median	22.25	26.30	30.51	31.76	20.18	31.49	25.15	24.99	51.85
Std.Dev	10.58	47.67	31.10	141.17	137.94	50.24	26.77	59.93	283.23
90th prctile	38.16	63.95	42.81	92.87	33.17	50.27	40.49	52.69	130.83
95th prctile	43.02	73.41	46.87	231.29	54.41	59.25	48.72	67.35	247.95
99th prctile	61.65	181.21	55.33	860.10	896.61	123.25	120.41	133.88	1784.63
$\widehat{P}_{t,e\|PD1}$									
Mean	24.30	35.42	31.92	70.55	28.28	29.35	25.14	29.54	64.78
Median	21.22	24.21	30.80	29.61	18.57	26.87	20.88	22.26	47.44
Std.Dev	9.65	39.00	10.22	163.39	70.65	11.14	14.35	17.17	78.48
90th prctile	39.30	56.96	43.11	104.46	30.22	43.98	32.89	50.58	113.78
95th prctile	42.19	66.04	46.67	255.12	38.12	48.44	39.49	66.98	145.26
99th prctile	50.06	178.10	66.35	898.34	394.62	71.99	93.12	90.93	230.53
$\widehat{P}_{t,e\|PD2}$									
Mean	24.40	35.07	32.45	73.29	35.32	30.38	26.81	29.77	87.11
Median	20.72	24.80	30.83	31.32	18.86	27.94	21.93	22.18	49.03
Std.Dev	9.73	48.23	13.46	180.52	149.53	12.16	18.02	29.64	241.91
90th prctile	38.34	59.12	43.69	92.28	31.21	43.68	34.86	48.25	119.02
95th prctile	44.73	70.73	48.42	141.16	38.10	49.62	42.15	54.45	167.94
99th prctile	52.62	134.46	57.65	916.11	662.62	78.20	136.55	75.77	1498.58
$\widehat{P}_{t,e\|PD3}$									
Mean	24.87	35.13	35.40	76.14	57.61	35.86	28.30	31.07	104.08
Median	21.91	26.25	30.38	31.72	19.94	31.27	23.87	24.36	52.43
Std.Dev	10.17	29.58	80.11	182.57	219.29	34.06	18.51	18.33	266.80
90th prctile	37.49	59.16	42.34	97.08	31.75	50.35	37.89	51.87	122.69
95th prctile	44.18	69.60	46.45	225.93	44.67	57.01	45.85	66.62	218.68
99th prctile	65.10	138.60	58.90	1066.38	1436.26	117.86	114.28	101.61	1769.08

Table 4.4-4: This table shows the descriptive analysis of a daily average spot prices for each year in the evening sub-period, $P_{t,e}$ (18:00 to 23:30 dispatch intervals - 12 trading intervals, $\tau = e$) and the corresponding estimated pre-dispatch prices at $j = PD1, PD2\ or\ PD3$ in the evening sub-period, $\hat{P}_{t,e|j}$.
The daily prices are calculated by taking an average of 12 half-hour spot prices in a day. Once the daily series is obtained the descriptive analysis shown in the table are calculated.
Std.Dev is standard deviation and 90th prctile, 95th prctile and 99th prctile is the 90th, 95th and 99th percentile of the daily average spot prices. Then the ninth and tenth rows used to display the proportion of days and half-hour with price spikes (% days prcspikes, % hh prcspikes) is in percentage unit.

Further explanation on this table is similar to the caption of Table 4.4-1, except $\tau = 24, \ldots, 35 (18{:}00 \text{ to } 23{:}30)$

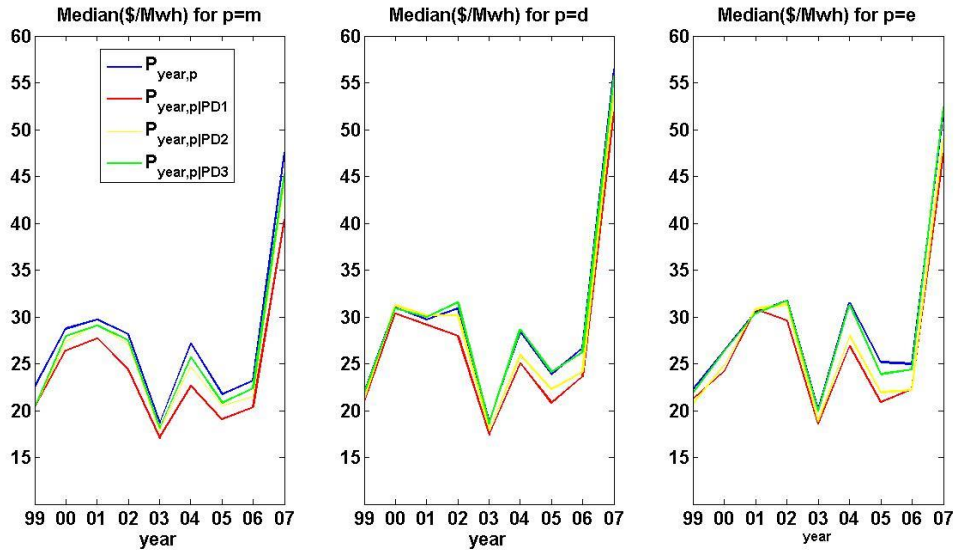

Figure 4.4-7: This figure is showing the median of the daily average spot prices, $P_{year,p}$ and estimated pre-dispatch prices, $\hat{P}_{year,p|j}$ for each year and trading sub-periods ($p = m, d, e$) during each pre-dispatch processes ($j = PD1, PD2, PD3$). The first panel is the median series for the morning sub-period, the second panel is the median series for the daytime sub-period and the third panel is the median series for the evening sub-period.

The $P_{year,p}$ series in the first panel is the median of spot prices calculated for each year from 1999 to 2007 for the morning sub-period, $p = m$. Using an example of $year = 2000$, to calculate median of $P_{2000,m}$, we first take the average of the 7 half-hour (06:30 to 09:30) spot prices in a day to construct 366-days daily time-series. Then by calculating the median of the 366-days daily time-series we obtain value for median of $P_{2000,p}$. The same process applies for the median of other estimated pre-dispatch prices, $\hat{P}_{2000,m|PD1}$, $\hat{P}_{2000,m|PD2}$, $\hat{P}_{2000,m|PD3}$ except they are using the estimated pre-dispatch prices as discuss in Section 4.4.2.5 instead of spot prices.

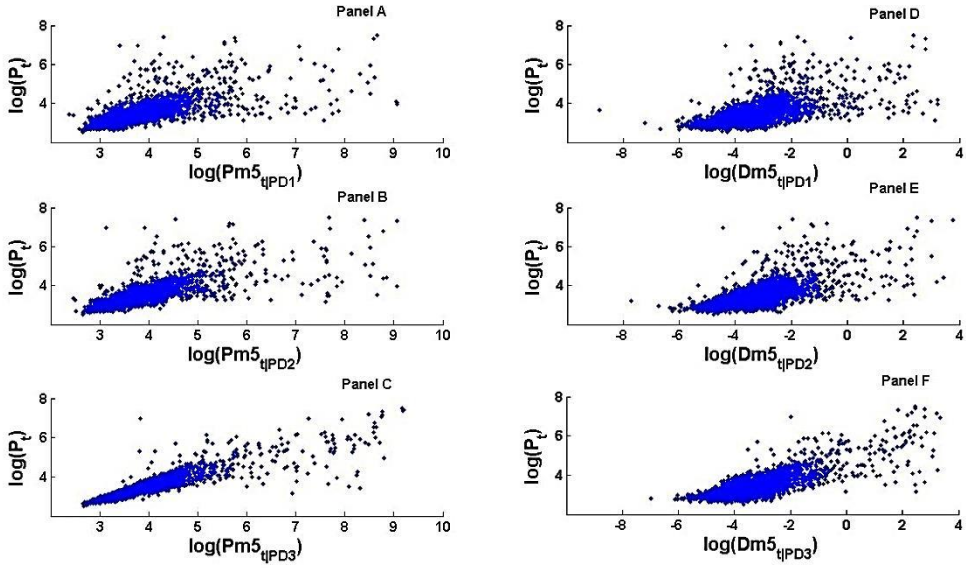

Figure 4.4-8: This figure is showing a scatter plot between $\log(P_t)$ and $\log(Pm5_{t|j})$ in Panel A, B and C and a scatter plot between $\log(P_t)$ and $\log(Dm5_{t|j})$ in Panel D, E and F.

The logarithmic series of daily average spot prices, $\log(P_t)$ used for the data of the y-axis are calculated by taking the average of the 35 half-hour spot prices in a day and then taking the logarithmic of the average series. While the logarithmic series of daily $Pm5_{t|j}$ and $Dm5_{t|j}$ on PD1, PD2 and PD3, $\log(Pm5_{t|j})$ and $\log(Dm5_{t|j})$ used for the data of the x-axis on the respective sub-plot are calculated by taking the average of the 5 largest values of $\hat{P}_{t,\tau|j}$ and $\hat{D}_{t,\tau|j}$ across all $\tau = 1, \dots, 35$ on any day t and the logarithmic of the average series.

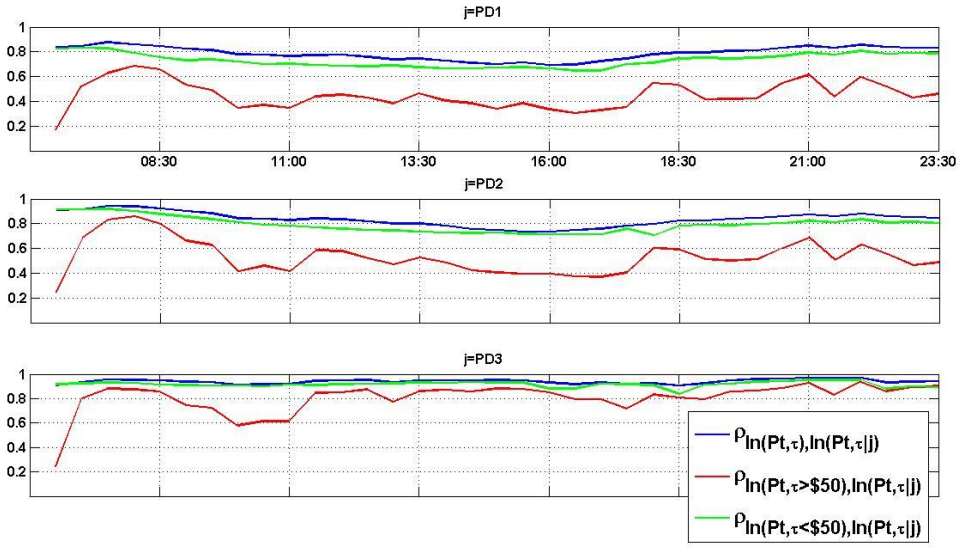

Figure 4.5-1: This figure is showing the 35 half hour correlation between logarithmic series of pre-dispatch and spot prices from March 2, 1999 to October 31, 2007.

In the first panel of this figure, the three distinctive line (with different colours) is used to illustrate the 35 half hour correlation between logarithmic series of pre-dispatch, $\ln(\hat{P}_{t,\tau|j})$, and spot prices, $\ln(P_{t,\tau})$; $\rho_{\ln(P_{t,\tau}),\ln(\hat{P}_{t,\tau|j})}$ plotted using blue line, 35 half hour correlation between logarithmic series of pre-dispatch, $\ln(\hat{P}_{t,\tau|j})$, and spot prices that is above A\$50, $\ln(P_{t,\tau} > \$50)$; $\rho_{\ln(P_{t,\tau}>\$50),\ln(\hat{P}_{t,\tau|j})}$ plotted using red line and lastly 35 half hour correlation between logarithmic series of pre-dispatch, $\ln(\hat{P}_{t,\tau|j})$, and spot prices that is below A\$50, $\ln(P_{t,\tau} < \$50)$; $\rho_{\ln(P_{t,\tau}<\$50),\ln(\hat{P}_{t,\tau|j})}$ plotted using green line.
The second and third panel are used to display the 35 half hour correlation on *PD2* and *PD3*.
Using $\tau = 19$ (at 09:00) as an example, the correlation of one half-hour is calculated as follows:

$$\rho_{\ln(P_{t,19}),\ln(\hat{P}_{t,19|PD1})} = \rho\big(\ln(P_{t,19}), \ln(\hat{P}_{t,19|PD1})\big), t = 1, \dots, 3166\ days$$

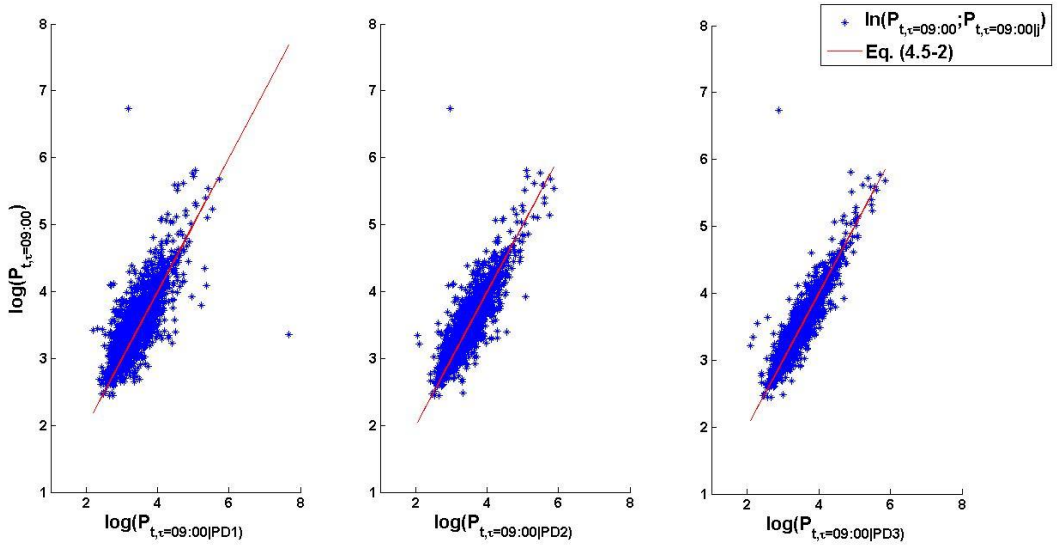

Figure 4.5-2: This figure is used to demonstrate the relationship between $\log(P_{t,09:00})$ and $\log(P_{t,09:00|j})$ during PD1, PD2 and PD3 at $\tau = 09{:}00$.

In the first panel, the scatter points represent all pairs of $\log(P_{t,09.00})$ and $\log(P_{t,09:00|PD1})$ in our sample period (from March 2, 1999 to October 31, 2007). The solid line represents the idealistic relationship between of $\log(P_{t,09.00})$ and $\log(P_{t,09:00|PD1})$ in case the latter provides an unbiased forecast for the logarithm of the spot price. This relationship is represented in equation (4.5-2). The same description applies to the second and third panel except that they are related to the pre-dispatch process at *PD2* and *PD3*.

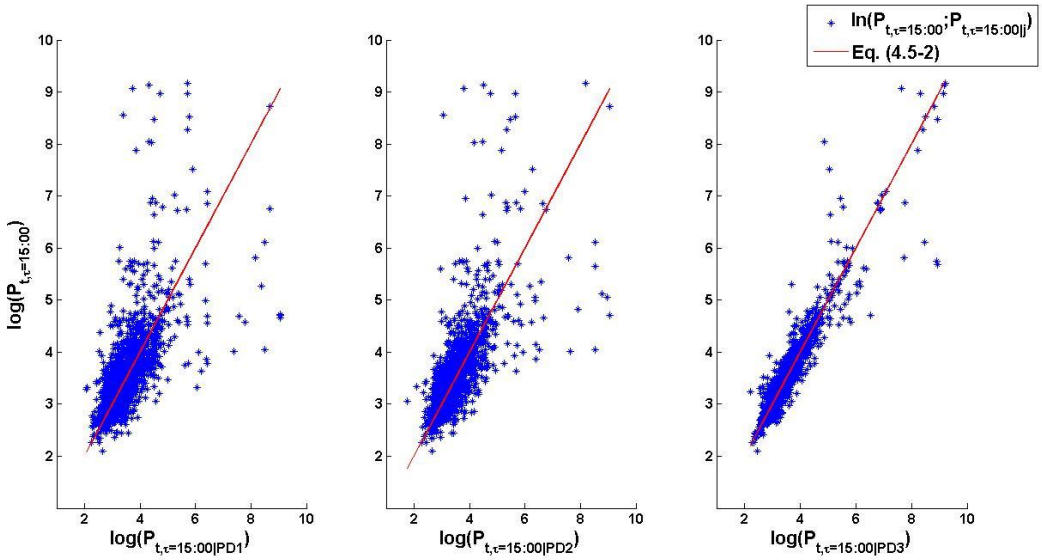

Figure 4.5-3: This figure is used to demonstrate the relationship between $\log(P_{t,15:00})$ and $\log(P_{t,15:00|j})$ during *PD1*, *PD2* and *PD3* at $\tau = 15:00$.

In the first panel, the scatter points represent all pairs of $\log(P_{t,15.00})$ and $\log(P_{t,15:00|PD1})$ in our sample period (from March 2, 1999 to October 31, 2007). The solid line represents the idealistic relationship between of $\log(P_{t,15.00})$ and $\log(P_{t,15:00|PD1})$ in case the latter provides an unbiased forecast for the logarithm of the spot price. This relationship is represented in equation (4.5-2). The same description applies to the second and third panel except that they are related to the pre-dispatch process at *PD2* and *PD3*.

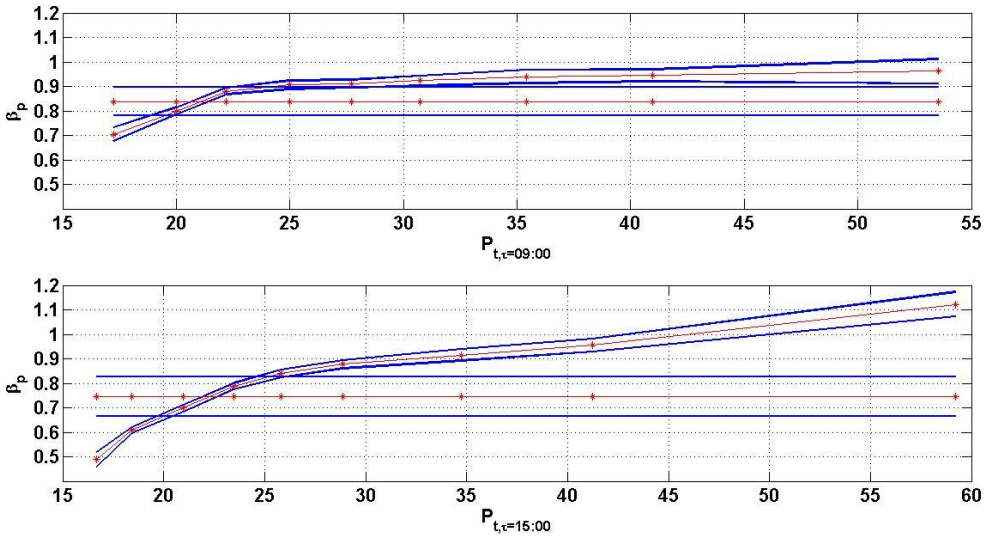

Figure 4.5-4: This figure is showing the quantile regression coefficients β_p calculated using equation (4.3-5) for *PD1* at τ = 09:00 and 15:00 are accompanied by 95% confidence intervals for the each quantile ($p = 0.1\ to\ 0.9$ in increments of 0.1). The estimated coefficients are plotted against the unconditional p deciles of $P_{t,\tau}$ (as indicated by the positions of the *).
The horizontal line represents the least squares (conditional mean) estimate and its 95% confidence interval.

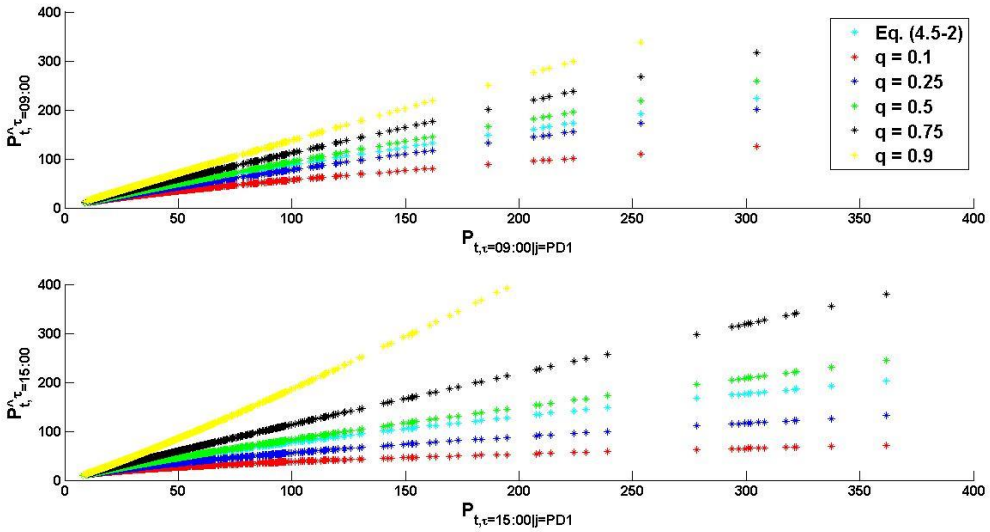

Figure 4.5-5: This figure is showing the estimated series of $\hat{P}_{t,\tau}$.

The raw series of $\hat{P}_{t,\tau}$ is constructed by taking an exponential of $\log(\hat{P}_{t,\tau})$ series. The $\log(\hat{P}_{t,\tau})$ series is estimated using log of pre-dispatch prices at 09:00 and 15:00 during *PD1* with respect to equation (4.5-2) and quantile regression for $p = \{0.1,0.25,0.5,0.75,0.9\}$ as discussed in Section 4.3.1.

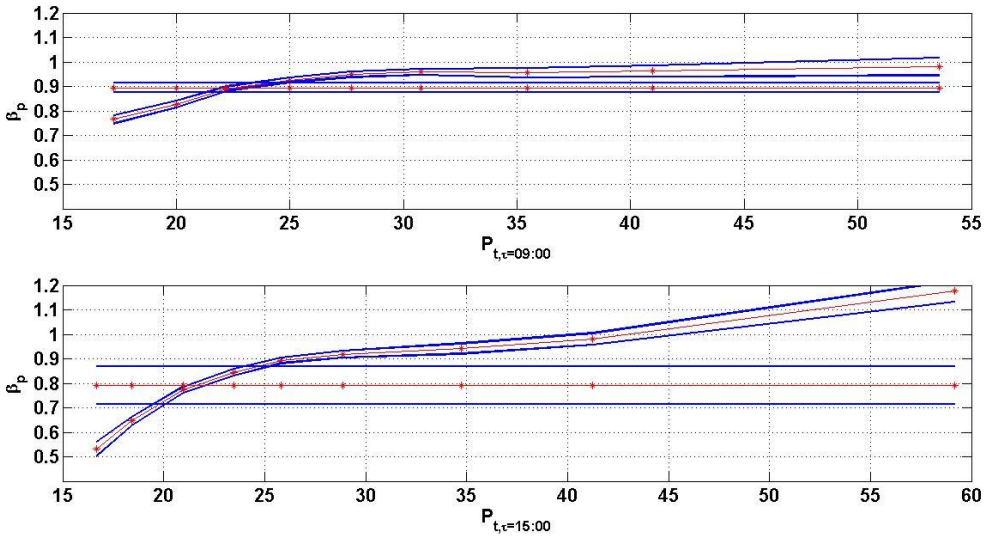

Figure 4.5-6: This figure is showing the quantile regression coefficients β_p calculated using equation (4.3-5) for *PD2* at τ = 09:00 and 15:00 are accompanied by 95% confidence intervals for the each quantile ($p = 0.1\ to\ 0.9$ in increments of 0.1). The estimated coefficients are plotted against the unconditional p deciles of $P_{t,\tau}$ (as indicated by the positions of the *).
The horizontal line represents the least squares (conditional mean) estimate and its 95% confidence interval.

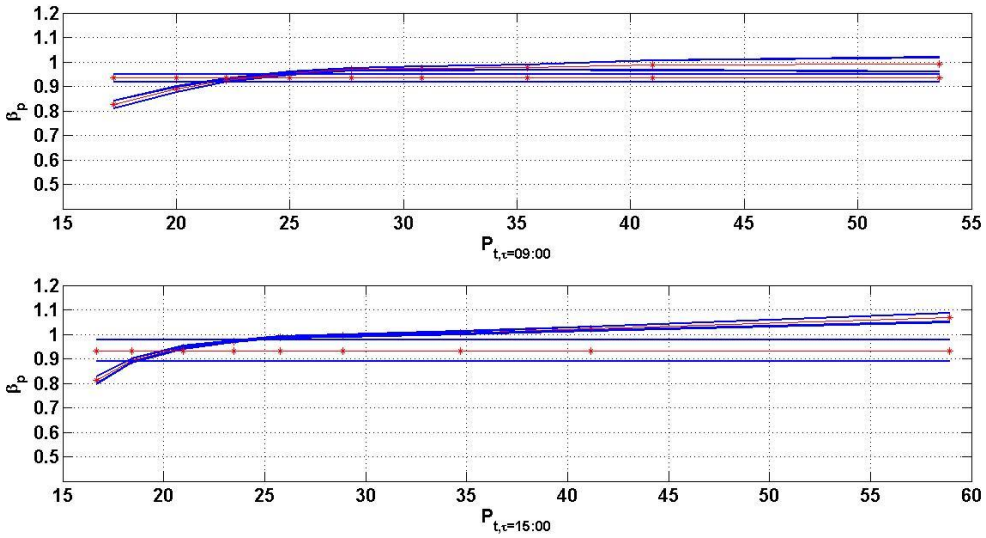

Figure 4.5-7: This figure is showing the quantile regression coefficients β_p calculated using equation (4.3-5) for *PD3* at τ = 09:00 and 15:00 are accompanied by 95% confidence intervals for the each quantile ($p = 0.1\ to\ 0.9$ in increments of 0.1). The estimated coefficients are plotted against the unconditional p deciles of $P_{t,\tau}$ (as indicated by the positions of the *).
The horizontal line represents the least squares (conditional mean) estimate and its 95% confidence interval.

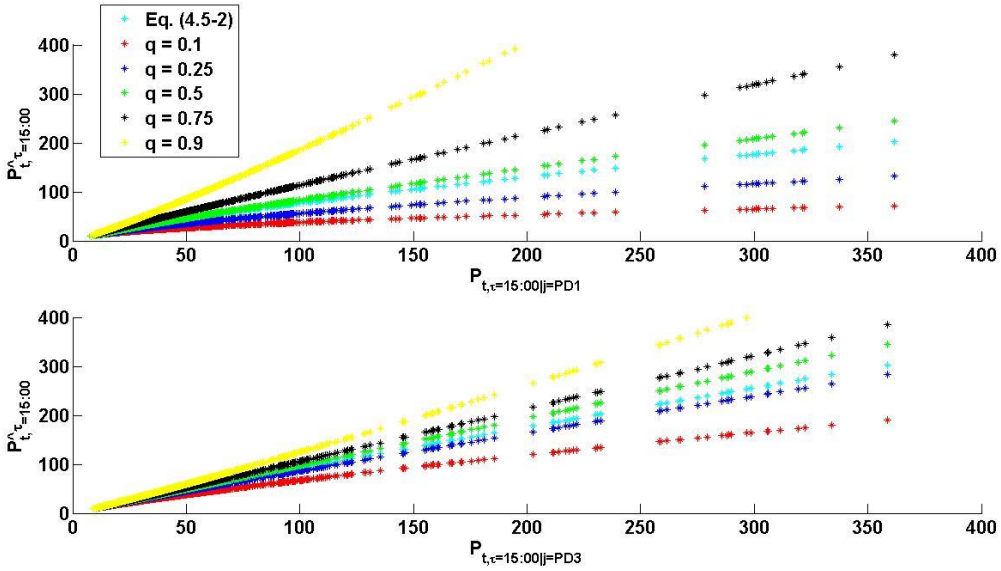

Figure 4.5-8: This figure is showing the estimated series of $\hat{P}_{t,\tau}$.

The raw series of $\hat{P}_{t,\tau}$ is constructed by taking an exponential of $\log(\hat{P}_{t,\tau})$ series. The $\log(\hat{P}_{t,\tau})$ series is estimated using log of pre-dispatch prices at 15:00 during *PD1* and *PD3* with respect to equation (4.5-2) and quantile regression for $p = \{0.1, 0.25, 0.5, 0.75, 0.9\}$ as discussed in Section 4.3.1.

	$\ln(Pm5_{t\|PD1})$	$\ln(Dm5_{t\|PD1})$	$\ln(Pm5_{t\|PD2})$	$\ln(Dm5_{t\|PD2})$	$\ln(Pm5_{t\|PD3})$	$\ln(Dm5_{t\|PD3})$	$B_{t,d\|PD1}$	$B_{t,e\|PD1}$	$B_{t,d\|PD2}$	$B_{t,e\|PD2}$	$B_{t,d\|PD3}$	$B_{t,e\|PD3}$
$\ln(Pm5_{t\|PD1})$	1											
$\ln(Dm5_{t\|PD1)}$	**0.8913**	1										
$\ln(Pm5_{t\|PD2)}$	0.8781	0.7869	1									
$\ln(Dm5_{t\|PD2})$	0.7788	0.7996	**0.8796**	1								
$\ln(Pm5_{t\|PD3})$	0.7125	0.6381	0.7852	0.7099	1							
$\ln(Dm5_{t\|PD3})$	0.6434	0.6556	0.7053	0.7331	**0.8752**	1						
$B_{t,d\|PD1}$	-0.1549	-0.1742	-0.0186	-0.0412	0.1989	0.1651	1					
$B_{t,e\|PD1}$	-0.1812	-0.1741	-0.0746	-0.0719	0.1284	0.1205	**0.6016**	1				
$B_{t,d\|PD2}$	-0.0885	-0.1088	-0.1005	-0.1193	0.1751	0.1453	0.7748	0.4677	1			
$B_{t,e\|PD2}$	-0.1194	-0.1125	-0.1450	-0.1289	0.0948	0.1076	0.4487	0.8019	**0.5482**	1		
$B_{t,d\|PD3}$	-0.0442	-0.0408	-0.0563	-0.0404	0.0656	0.1150	0.4219	0.2201	0.5356	0.2802	1	
$B_{t,e\|PD3}$	0.0095	0.0339	-0.0140	0.0140	-0.0285	0.0738	-0.0079	0.2778	-0.0070	0.3325	**0.1064**	1

Table 4.5-1: This is a correlation matrix for pre-dispatch variables as discussed in Section 4.4.2.5.
This correlation matrix is calculated using daily information throughout the investigation period (02/03/1999 to 31/10/2007). $\ln(Pm5_{t|j})$ and $\ln(Dm5_{t|j})$ are the logarithmic series of the average of the 5 largest values of $\hat{P}_{t,\tau|j}$ and $\widehat{D}_{t,\tau|j}$ at the respective pre-dispatch processes (j). While $B_{t,p|j}$ is a daily measure on the respective trading periods (p) and pre-dispatch processes (j).

	Models			
Variables	1	2	3	4
$Basic$	Y	Y	Y	Y
$B_{t,d\|j}$	PD1		PD1	
$B_{t,e\|j}$		PD1		PD1
$\ln(Pm5_{t\|j})$	PD1	PD1		
$\ln(Dm5_{t\|j})$			PD1	PD1

Table 4.5-2: This is a covariate specification for Models 1 to 4.
The variables in the "Basic" category are L_t, $Tmin_t$, $Tmax_t$ and C_t. The cell entries in the rows for $B_{t,d|j}$, $B_{t,e|j}$, $\ln(Pm5_{t|j})$ and $\ln(Dm5_{t|j})$ indicate from which pre-dispatch (j) the information is taken. Empty cells indicate that the respective variable is not included in the model.

Variable	Basic Model	Model 1	Model 2	Model 3	Model 4
μ_t^					
α_{t_i}					
γ					
Constant	1.1484 (1.7075)	4 (12.8103)	4 (21.9613)	6 (7.8319)	3.5487 (6.7298)
L_{t+1}	-0.0806 (1.5453)	-0.4543 (3.3691)	-0.5651 (4.1289)	2.7963 (3.8443)	-6 (8.8080)
$T_{max_{t+1}}$	-0.0590 (0.3285)	0.1201 (0.9184)	0.1287 (1.3995)	-1.4489 (2.1208)	1.9086 (3.0865)
$T_{min_{t+1}}$	-0.1983 (0.4452)	-0.3475 (0.4947)	-0.3272 (0.4970)	4.4537 (3.3444)	-2.5220 (4.2436)
C_t	-0.0613 (0.2641)	-0.2448 (1.2916)	-0.2638 (1.8843)	-0.3915 (0.5773)	-3.1681 (4.7011)
$B_{t,d\|PD1}$		-0.0125 (0.0966)		0.4906 (0.5936)	
$B_{t,e\|PD1}$			0.0039 (0.1720)		0.2113 (0.3814)
$\ln(Pm5_{t+1\|PD1})$		-0.5542 (1.9876)	-0.5337 (3.8001)		
$\ln(Dm5_{t+1\|PD1})$				0.5210 (0.7307)	-4.9207 (7.2668)
β					
δ					
Constant	0.1352* (0.0387)	0.1306* (0.0345)	0.1317* (0.0328)	0.1180* (0.0217)	0.1246* (0.0246)
Log$L^{Sea,tva}$	-792.004	-790.43	-790.463	-788.799	-786.373
LR test		3.148	3.082	6.410**	11.262***

Table 4.5-3: This table presents the coefficients and robust standard errors (in brackets) of parameter vector θ $((\gamma'\ \delta')')$ used on the conditioning covariates in the initial intensity jumps, α_{t_i} and time-invariant intensity decay, β parameters.

These are the results of the HAWa model for NSW during the estimation period (March 2, 1999 to October 30, 2005). ^The coefficients and standard errors of parameter vector of the trigonometric term ζ (allowing for seasonal variations, μ_t) are not reported but are available upon request. All of these parameter vectors are estimated by maximising the log-likelihood function in equation (4.3-11). In the bottom row, we report the optimised log-likelihood and the results of the likelihood ratio test. The covariates conditioned in α_{t_i} are an intercept ($constant$), unseasonal load series (L_{t+1}) as discussed in Section 4.4.2.2, unexpected increase/decrease in temperatures ($Tmax_{t+1}, Tmin_{t+1}$) as discussed in Section 4.4.2.3, the number of price spikes during a day (C_t) as discussed in Section 4.4.2.4, daily measure in the daytime and evening sub-periods of $B_{t,d|PD1}$ and $B_{t,e|PD1}$ and the logarithmic series of the average of the 5 largest values of $\widehat{P}_{t,\tau|PD1}$ and $\widehat{D}_{t,\tau|PD1}$ during PD1 ($\ln(Pm5_{t|j})$ and $\ln(Dm5_{t|j})$) as discussed in Section 4.4.2.5. While the β remains constant.

*Parameters are significant at 10% significant level.

**Significant at 5% level

***Significant at 1% level

Variable	Basic	Model 1	Model 2	Model 3	Model 4
μ_t					
α_{t_i}					
γ					
Constant	0.6038* (0.3487)	1.5380 (0.9766)	1.4460 (1.1641)	-0.2935 (0.4606)	-0.2466 (0.5477)
L_{t+1}	0.4405 (0.3699)	0.4023 (0.2596)	0.6500 (0.5042)	0.4072* (0.2490)	0.6415 (0.5271)
$T_{max_{t+1}}$	-0.1100 (0.1046)	-0.1442 (0.0929)	-0.1787 (0.1466)	-0.1569* (0.0936)	-0.1828 (0.1637)
$T_{min_{t+1}}$	1.6925 (1.4291)	0.0834 (0.1791)	2.0513 (1.3480)	0.1032 (0.1819)	1.9754 (1.5647)
C_t	0.2446* (0.1145)	0.2118 (0.1401)	0.2266 (0.2654)	0.2052 (0.1644)	0.2179 (0.3271)
$B_{t,d\|PD1}$		-0.0039 (0.0219)		-0.0020 (0.0211)	
$B_{t,e\|PD1}$			0.0672 (0.0554)		0.0675 (0.0532)
$\ln(Pm5_{t+1\|PD1})$		-0.2982 (0.2080)	-0.2685 (0.2558)		
$\ln(Dm5_{t+1\|PD1})$				-0.2167 (0.1404)	-0.2131 (0.1632)
β_t					
δ					
Constant	1.4850* (0.5690)	-0.0320 (1.0796)	0.9284 (1.3628)	1.2659 (0.8073)	1.5711* (0.6584)
L_{t+1}	-0.9692* (0.2155)	-0.7013* (0.2065)	-0.8976* (0.2291)	-0.6898* (0.2145)	-0.8575* (0.2223)
$T_{max_{t+1}}$	0.0334 (0.0511)	0.2901* (0.1693)	0.0551 (0.0673)	0.3186 (0.2311)	0.0673 (0.0707)
$T_{min_{t+1}}$	0.2025 (0.1629)	0.1178 (0.1341)	0.2227 (0.1702)	0.1008 (0.1240)	0.1868 (0.1729)
C_t	-0.0922* (0.0258)	-0.3373* (0.1723)	-0.0949* (0.0331)	-0.3570* (0.2092)	-0.1010* (0.0383)
$B_{t,d\|PD1}$		0.0227 (0.0155)		0.0225 (0.0157)	
$B_{t,e\|PD1}$			0.0060 (0.0276)		0.0050 (0.0268)
$\ln(Pm5_{t+1\|PD1})$		0.1706 (0.2532)	0.0838 (0.2318)		
$\ln(Dm5_{t+1\|PD1})$				0.1917 (0.1846)	0.1198 (0.1243)
$\text{Log}L^{Sea,tvab}$	-747.234	-746.239	-746.156	-745.936	-746.077
LR test		1.990	2.156	2.596	2.314

Table 4.5-4: This table presents the coefficients and robust standard errors (in brackets) of parameter vector θ $((\gamma' \delta')')$ used on the conditioning covariates in the initial intensity jumps, α_{t_i} and time-varying decay parameter, β_τ.

These are the results of the HAWab model for NSW during the estimation period (March 2, 1999 to October 30, 2005). ^The coefficients and standard errors of parameter vector of the trigonometric term ζ (allowing for seasonal variations, μ_t) are not reported but are available upon request. All of these parameter vectors are estimated by maximising the log-likelihood function in equation (4.3-11). In the bottom row, we report the optimised log-likelihood and the results of the likelihood ratio test. The covariates conditioned in α_{t_i} and β_t are as defined in the caption of Table 4.5-3

*Parameters are significant at 10% significant level.

Variable	Basic	Model 1	Model 2	Model 3	Model 4
π_t					
β_1					
Constant	-3.1834* (0.1535)	-2.2681* (0.6520)	-2.4031* (0.6060)	-3.7509* (0.3705)	-3.7601* (0.3630)
Load_{t+1}	0.7914* (0.1496)	0.8266* (0.1484)	0.7997* (0.1452)	0.8218* (0.1468)	0.7967* (0.1437)
$T_{max_{t+1}}$	0.1297* (0.0331)	0.1308* (0.0312)	0.1335* (0.0339)	0.1308* (0.0309)	0.1339* (0.0335)
$T_{min_{t+1}}$	0.1549 (0.1073)	0.1645 (0.1033)	0.1718* (0.1033)	0.1619 (0.1036)	0.1685* (0.1036)
$B_{t,d\|PD1}$		0.0172* (0.0088)		0.0169* (0.0088)	
$B_{t,e\|PD1}$			0.0346* (0.0119)		0.0349* (0.0118)
$\ln(Pm5_{t+1\|PD1})$		-0.2752 (0.1812)	-0.2472 (0.1679)		
$\ln(Dm5_{t+1\|PD1})$				-0.1404 (0.0882)	-0.1323 (0.0864)
ρ_t					
β_2					
Constant	-0.2549* (0.1426)	0.0791 (0.5028)	0.1362 (0.5295)	-0.4045 (0.3309)	-0.4116 (0.3306)
Load_{t+1}	0.1371 (0.1098)	0.1262 (0.1119)	0.1139 (0.1152)	0.1288 (0.1118)	0.1172 (0.1147)
$T_{max_{t+1}}$	-0.2008* (0.0729)	-0.2196* (0.0804)	-0.2115* (0.0822)	-0.2211* (0.0802)	-0.2131* (0.0815)
$T_{min_{t+1}}$	0.0016 (0.0778)	-0.0089 (0.0793)	-0.0078 (0.0800)	-0.0078 (0.0790)	-0.0067 (0.0797)
$B_{t,d\|PD1}$		-0.0143* (0.0088)		-0.0144* (0.0088)	
$B_{t,e\|PD1}$			-0.0140 (0.0146)		-0.0133 (0.0144)
$\ln(Pm5_{t+1\|PD1})$		-0.0747 (0.1337)	-0.0880 (0.1369)		
$\ln(Dm5_{t+1\|PD1})$				-0.0622 (0.0852)	-0.0659 (0.0845)
log L^{PAR}	-680.431	-675.441	-674.28	-675.58	-674.313
LR test		9.980**	12.302**	9.702**	12.236**

Table 4.5-5: This table presents the coefficients and robust standard errors (in brackets) of parameter vector, β_1 and β_2 used on the conditioning covariates relating to the variation in arrival probability of system stresses, π_t and in survival probability of system stresses, ρ_t.

These are the results of the PAR model for NSW during the estimation period (March 2, 1999 to October 30, 2005). All of these parameter vectors are estimated by maximising the log-likelihood function in equation (4.3-20). In the bottom row, we report the optimised log-likelihood and the results of the likelihood ratio test. The covariates conditioned in π_t are as defined for α_{t_i} and the covariates conditioned in ρ_t are as defined for β_t but without C_t exogenous variable.
*Parameters are significant at 10% significant level.
**Significant at 5% level

Measure	Naive	Basic	M1	HAWa M2	M3	M4
MAE	**0.1929**	0.2290	0.2392	0.2327	0.2481	0.2502
RMSE	0.4392	**0.3383**	0.3468	0.3414	0.3476	0.3608
Asym	0.2319	**0.2306**	0.2461	0.2334	0.2658	0.2719
				HAWab		
MAE	**0.1929**	0.2641	0.2608	0.2379	0.2443	0.2371
RMSE	0.4392	0.3747	0.3621	0.3420	0.3444	**0.3411**
Asym	**0.2319**	0.2934	0.2743	0.2470	0.2503	0.2469
				PAR		
MAE	**0.1929**	0.2298	0.2299	0.2313	0.2317	0.2331
RMSE	0.4392	**0.3506**	0.3518	0.3525	0.3522	0.3530
Asym	**0.2319**	0.2526	0.2523	0.2543	0.2543	0.2563

Table 4.5-6: This table presents the forecast evaluation statistics (mean absolute error, MAE; root mean square error, RMSE; asymmetric loss score, Asym) for all models.

Based on Rudebusch & Williams (2009) and Christensen et al. (2009) using forecast error ($y_t - \lambda_{t|t-1}$) of Hawkes models as an example;

$$MAE = \frac{1}{t_1 - t_0 + 1} \sum_{t=t_0}^{t_1} \left| y_t - \lambda_{t|t-1} \right|$$

$$RMSE = \sqrt{\frac{1}{t_1 - t_0 + 1} \sum_{t=t_0}^{t_1} \left(y_t - \lambda_{t|t-1} \right)^2}$$

$$Asym = \frac{1}{t_1 - t_0 + 1} \sum_{t=t_0}^{t_1} \left(y_t (1 + \kappa) + (1 - y_t)(1 - \kappa) \right) \left| y_t - \lambda_{t|t-1} \right|$$

where t_0 and t_1 denote the beginning and the end of the forecast period and κ is equal to 0.5 since the failure to predict an actual price event is penalized by three times the rate of predicting a price event that does not actually occurred.

This table reports the average value of the respective loss function evaluated for Models Naïve, HAWa, HAWab and PAR for the forecast period from October 31, 2005 to October 31, 2007. The bold entry indicates which model produces the lowest average loss.

Chapter 5

Summary and suggestion for a future research

In this thesis, I investigate whether information from the NEM pre-dispatch process can be useful when predicting spike events. This is done in the context of based on the methodology framework of PAR and Hawkes models, which allow for covariate driven time-variation in the probability of price spikes occuring. The purpose is to find model specifications and covariates that are able to predict on what days price spikes are more likely to occur. The overall conclusion is that forecasting is a difficult task, neither PAR nor Hawkes models are able to predict price events that occurred in isolation. Even AEMO price forecast information from pre-dispatch process does not add much useful information to the price spikes forecast.

Chapter 2 examines whether the occurrence of these extreme price events displays any regularities that can be captured using an econometric model. Here I treat these price events as point processes and apply Hawkes and Poisson autoregressive models to model the dynamics in the intensity of this process. I use load and meteorological information to model the time variation in the intensity of the process. The models are applied to data from the Australian wholesale electricity market, and a forecasting exercise illustrates both the usefulness of these models and their limitations when attempting to forecast the occurrence of extreme price events.

Chapter 3 introduces a framework to analyse whether the pre-dispatch process delivers biased predictions of the actual wholesale spot price outcomes. Here I investigate the bias by comparing the actual wholesale market spot price outcomes to pre-dispatch sensitivity prices established the day before dispatch and on the day of dispatch. I observe a significant bias (mainly indicating that the pre-dispatch process tends to underestimate spot price outcomes) and I further establish the seasonality features of the bias across seasons and/or trading periods. I also establish changes in bias across the years in my sample period (1999 to 2007). In the formal setting of an ordered probit model I establish that there are some exogenous variables that are able to explain increased probabilities of over- or under-predictions of the spot price. It transpires that meteorological data, expected pre-dispatch prices and information on past over- and under-predictions contribute significantly to explain variation in the probabilities for over- and under-predictions. The

results allow me to conjecture that some of the bids and re-bids provided by electricity generators are not made in good faith.

Finally, Chapter 4 investigates whether information from this pre-dispatch process can be useful when predicting next-day price spikes. In a preliminary analysis I establish the effect of pre-dispatch prices on the quantiles of the spot price distribution. A Quantile regression approach reveals that higher pre-dispatch prices signal only to a certain extend an increased probability of higher spot price outcomes. They also signal a higher uncertainty about the resulting spot price outcomes. I further establish whether the inclusion of information from the pre-dispatch process can significantly improve the predictability of price spikes when these are modelled as a point process (as in Chapter 2). The models used here are Hawkes and Poisson Autoregressive Models which allow for time variation (correlated to exogenous information) in the intensity process that governs the occurrence of price spikes. It transpires that the pre-dispatch process of the Australian Electricity Market does not provide any information that can be used in a systematic manner to help predicting on what days price spikes are more likely to occur.

Even though I find that pre-dispatch information does not have much use for the purpose of forecasting price spikes, I believe the information may be useful for more generic price forecasting such as price volatility forecasting. Apart from pre-dispatch information, AEMO also publishes individual bids data which contain offers/bids and rebids made by generators and dispatchable loads for each half-hour on the dispatch day. Together, these individual bids and pre-dispatch information could provide a potential avenue for research in the bidding behaviour of generators in the Australian Electricity Market.