Econometric analysis of wholesale electricity markets

Jieyang Chong

Master of Business (Research), Queensland University of Technology, Australia
Bachelor of Business (Hons), Queensland University of Technology, Australia
Bachelor of Commerce, University of Queensland, Australia

A thesis submitted for the degree of Doctor of Philosophy at Monash University in 2021

Department of Econometrics and Business Statistics

Contents

List of Tables

List of Figures

Copyright notice

Abstract

Wholesale electricity price forecasting has been a subject of interest, particularly in response to deregulation of electricity markets around the world. This thesis examines some economic quantities in the German wholesale electricity market and the econometric techniques involved in their analysis.

Various innovative econometric methods have been adopted and adapted for the electricity price setting. However, limited attention has been given to developing models which accommodate changes in the structure of price trajectories over time. The first study in this thesis proposes a model which accounts for mean shifts through the use of smooth time-varying parameter models, which are estimated using semiparametric kernel estimation methods. The findings in this study show that for daily price forecasts, the semiparametric models outperform a constant-parameter benchmark.

Trades in wholesale electricity markets are locked in at hourly intervals. Complete supply and demand curves can be constructed for each hour. These curves may be treated as functional data, which can be modelled and forecasted. The intersection of the forecasted curves then gives a forecast of prices. Research into the use of functional data analysis tools for forecasting wholesale electricity price in this manner is new and extremely limited. This thesis proposes the application of a simpler and less time-consuming approach of producing forecasts of these functional data for this application than that which has been explored in the literature. The findings suggest that this approach is able to outperform a benchmark model based on existing electricity price forecasting literature which uses price data instead of market curves.

Besides prices, price elasticity of demand has been of interest. Estimates from existing regression approaches are typically a single value representing an average over the sample. This method relies on some implicit assumptions which are called into question in this thesis. An alternative way of calculating intra-day electricity demand response in wholesale markets is presented.

This method permits estimation of price elasticity for every hour using fundamental economic principles by exploiting the observed supply and demand curves. Seasonal and intra-day patterns in price elasticity are revealed. The value of price elasticity of demand which is computed using this method is also compared against the estimate from a relevant regression-based model, with a discussion of the relative accuracies.

The studies in this thesis all use relatively simple techniques, but are able to provide improvements on existing methods in the literature. These findings suggest that there is room for improvement from a modelling perspective, and that these innovations do not need to be complicated. Finally, some other research ideas which have not made it into the thesis are discussed briefly, with some scope for future research.

Declaration

This thesis is an original work of my research and contains no material which has been accepted for the award of any other degree or diploma at any university or equivalent institution and that, to the best of my knowledge and belief, this thesis contains no material previously published or written by another person, except where due reference is made in the text of the thesis.

JIEYANG CHONG

19 September 2021

Acknowledgments

I would like to formally express my gratitude to those who have had a direct hand in my PhD journey. From my advisors, Jiti Gao, Wei Wei, and Heather Anderson, I received patience and guidance throughout my PhD candidacy. Their approach to supervision has given me the independence to challenge myself and explore my ideas while always feeling supported and confident. My milestone panels, consisting of George Athanasopoulos, Gael Martin, Benjamin Wong, and Xibin Zhang, have been immensely supportive and encouraging, which has helped me stay on track with a positive attitude. Support from Andrea Meyer, Clare Livesey, and Deborah Fitton has enabled me to focus on my research with little distraction from administrative details. As a whole, the academic and professional staff and students in the Department of Econometrics and Business Statistics create a social and intellectual environment which is conducive to academic research, helping my time here be both enjoyable and productive. I am also grateful to Asger Lunde for permitting the use of the proprietary German wholesale electricity market data in this thesis.

Many others have been a part of my life throughout this PhD journey. The list is too long, and it includes family, friends, fellow students, participants at conferences and workshops, and experts in related fields with whom I have had discussions and from whose founts of knowledge I have benefitted.

This research is supported by an Australian Government Research Training Program (RTP) Scholarship.

Chapter 1

Introduction

Electricity is a necessity for modern human life. We rely on it for various uses including lighting, temperature regulation, entertainment, communication, transportation, medical equipment, manufacturing, delivery of water supply, and national security. As such, any improvements to the efficiency of electricity supply can be expected to have potential benefits to the population at large. This thesis focuses on efficiency from an economic perspective, namely, gaining a better understanding of electricity prices.

Wholesale market data is rich with possibilities. Raw data include transacted volumes and prices for intra-day periods. These series can be modelled or forecasted at the highest granularity, or aggregated to lower frequencies. For example, wholesale prices are forecasted both as volume-weighted daily prices (see, for example, Bierbrauer et al. (2007)) or as daily collections of intra-day prices[1]. Some data include purchase and sale bids made by wholesale market participants for each intra-day bidding period. The intersection of the sale and purchase curves constructed from these bids is determined to be the price in that period. The shape, location, and slope of each of these curves also contain information which may be useful for comprehending and predicting wholesale electricity markets.

The aim of this thesis, just like in any other empirical work, is to propose improved methods of analysing the data. For all of the analysis conducted in this thesis, the emphasis is always to ensure that the data and underlying research questions are adequately addressed by any methods which are employed, without allowing the technical details to overshadow or sideline

[1]For example, in one of their empirical applications, Weron and Misiorek (2008) forecast electricity prices in California for all 24 hours in the following day.

the problem at hand. While other works do a laudable job of introducing more complex or modern methods to better understand electricity markets, the place in the literature which this thesis aims to occupy is to ensure that existing methodologies are appropriate for their applications. In saying that, the following chapters do not criticise any other work. Instead, the projects contained in this thesis identify certain areas which may benefit from some adjustments, and then show that those adjustments provide good empirical performance. Justification of the proposed methods are always made in reference either to economic, regulatory, and market conditions, or to theoretical concerns which are relevant to the data. In this instance, the data is from wholesale markets, therefore the all explanations and interpretations are restricted to the setting of wholesale electricity markets, not the universe of consumers.

To introduce the work contained within this thesis, it is necessary to first discuss wholesale electricity prices and the relevant interesting or distinctive features.

1.1 Wholesale electricity prices

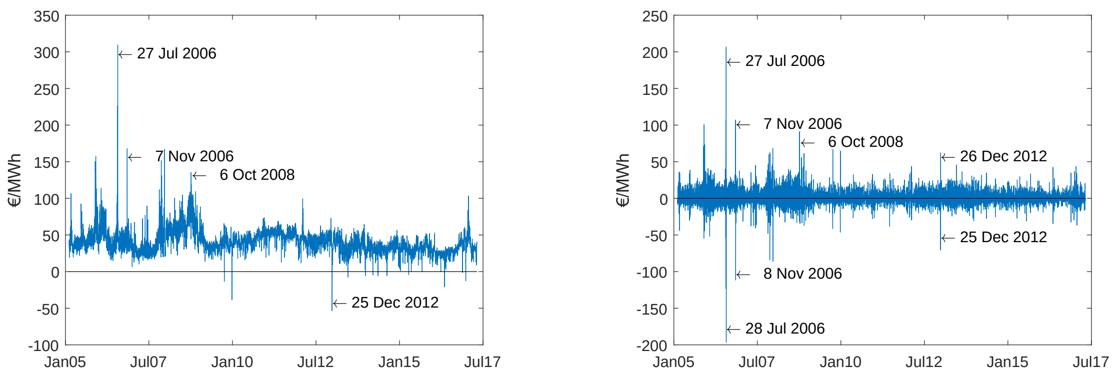

(a) Prices (b) First difference

Figure 1.1: Daily volume-weighted German electricity spot prices and first differences from 8 February 2005 to 28 April 2017.

Electricity as a commodity is almost unique in that it is widely used, but cannot be stored on any meaningful scale. While it is true that the source of electricity generation can be accumulated and inventoried[2], there are financial or physical restrictions which lead to bottlenecking in various stages from generation to delivery. Renewable energy constitutes a large proportion

[2]For instance, hydroelectric dams can collect water, and the resulting potential energy can then be converted into electrical energy (as well as by-products such as heat and noise). Similarly, coal, biomass fuels, and nuclear cores can also be accumulated.

of energy generated in certain markets. In Germany[3], wind is a prominent source of electricity; in South Australia, wind and solar together have contributed close to half of the energy in 2018, a proportion which has been increasingly yearly[4]. Energy generated by these methods, in particular, is entirely dependent on weather and cannot be stored as inventory.

The fact that electricity is non-storable, at least for now, manifests in prices in two ways. The most obvious consequence is that current supply and demand must match; prices cannot be smoothed by holding inventory. As such, wholesale electricity prices may be more volatile around certain periods, for example, when generators go into forced production (Simonsen, 2005). The plot of first differences in Figure 1.1(b) gives some indication of clustering. Large changes in prices tend to be followed by large changes, and vice versa. In addition, generation plants, especially those for renewables, incur a cost to be shut down during periods of low demand. Consequently, occasionally it is preferable for prices to fall into the negative region to allow for the generated power to be consumed, instead of shutting down generators altogether. These negative prices are seen clearly in Figure 1.1(a), with the largest of them occurring on Christmas day 2012. Perhaps the most striking feature of electricity prices is the presence of spikes. While many other asset prices have jumps in them, the spikes in electricity prices are typically short-lived, and can be much larger in magnitude. Although several explanations for the occurrence of spikes have been offered, there is still no obvious or accurate way to empirically identify spikes. However, upon inspection of Figure 1.1(a) it is tempting to declare that prices on 27th July 2006 and 7th November 2006 are examples of upward spikes, and that the negative price on Christmas day of 2012 was a downward spike. These three events are merely some of the most prominent, but other spikes are likely to have occurred also. Finally, the close connection between weather changes and electricity demand and prices is obvious; the population will demand cooling when it is too hot, and heating when it is too cold.

Electricity spot prices share certain characteristics with other commodities. They exhibit seasonality corresponding to, among others, yearly seasons, days of the week, and hour of the day (Koopman et al., 2007).

[3] To be precise, the data used here is for the joint bidding zone between Germany and Austria before the split in electricity price zones which occurred in 2018. However, the literature typically refers to it simply as "Germany". To be consistent with the literature and for simplicity, the data will also be referred to as German data in this thesis.

[4] See `https://opennem.org.au/energy/sa1/?range=all&interval=1y`

Volatility clustering

The presence of volatility clustering in electricity spot prices has been widely acknowledged. Unsurprisingly, spot price volatility is the key focus of a number of studies. Clustering is often placed under the umbrella of conditional heteroskedasticity, for which the autoregressive conditional heteroskedastic (ARCH) model of Engle (1982) and generalised ARCH (GARCH) model of Bollerslev (1986) offer typical estimation approaches. ARCH- or GARCH-type models were employed by Knittel and Roberts (2005), Goto and Karolyi (2004), and Escribano et al. (2011), among others.

It is worth noting that although Knittel and Roberts (2005), who evaluate an AR-EGARCH specification, conclude that volatility clustering (along with higher order autocorrelation) is a key feature in electricity pricing, their results show some variation when assessed over different periods in time. Similarly, Garcia et al. (2005) conclude that the improved performance derived from introducing GARCH residuals to an autoregressive integrated moving average (ARIMA) model is contingent on the choice of sample period.

Spikes

Spikes are characterised by one or more extreme upward (downward) jumps followed immediately by a sharp downward (upward) corrective move. These spikes are typically due to physical anomalies in either the supply side (*e.g.* generator shutdowns) or demand side (*e.g.* heatwaves). The presence of spikes has the capacity to bias the estimation of the conditional mean of prices due to their large magnitudes. There are three different treatments of spikes which are common in the literature. Lucia and Schwartz (2002) conduct their estimations by ignoring the issue altogether, treating all prices as if spikes were not present. Others transform the data so as to remove, or at least minimise, the effect of spikes. Examples include winsorising to remove spikes beforehand (see, *inter alia*, Cartea and Figueroa (2005), Gianfreda (2010), and Ketterer (2014)), or "stabilising" the variance of the dataset before applying the model of choice (Uniejewski et al., 2018). The third approach is to construct models which account for the presence of spikes. Of course, there are also those whose modelling focuses on the spikes themselves, largely ignoring base prices (Becker et al., 2007; Christensen et al., 2009, 2012).

Electricity price spikes are commonly modelled using jumps or jump-diffusion models (see, *inter* *alia*, Barone-Adesi and Gigli (2003), Escribano et al. (2011), Knittel and Roberts (2005), Weron and Misiorek (2005), and Weron et al. (2004)). Popular alternatives to jump-diffusion models are regime-switching models. This class of models is, perhaps, better-suited for spikes as they allow for the near-instantaneous mean-reversion which follows. Furthermore, there is persistence to the occurrence of spikes (Christensen et al., 2009) which tends to be better-captured by regime-switching models (Weron,2014). On this note, a number of different specifications have been proposed for regime-switching models. Two-state models typically contain one regime for the base trend, and another for spikes. For example, Huisman and De Jong (2003) approximate the spikes by a Gaussian distribution, whereas Weron et al. (2004) use a lognormal distribution, and Bierbrauer et al. (2004) specify a Pareto distribution. Huisman and Mahieu (2003) propose a model with three states, namely, the base, initial jump, and reversing jump regimes. In the three-regime model, the initial jump is followed immediately by a reversing jump. As a result, the states are not independent. Furthermore, the forced return to the base regime does not permit consecutive spikes.

Negative prices

It is common practice to transform the data from levels to logarithms when modelling prices of many other assets. This treatment has also been applied to electricity prices in many studies (see, *inter alia*, Bierbrauer et al. (2007), Janczura and Weron (2010), Knittel and Roberts (2005), Koopman et al. (2007), and Weron et al. (2004)). One of the benefits of doing so is that the effect of spikes will be reduced, allowing for more conventional modelling (Uniejewski et al., 2018). However, since late 2009 a number of negative prices have been observed, thereby making the log-transformation infeasible.

Erni (2012) argues that the motivations behind taking logarithms of prices of other assets do not apply to electricity prices. Working with logarithms guarantees models which produce positive price predictions. In this scenario, however, the imposition of positive prices would be inappropriate. In fact, the raw data itself contains negative prices, and it would be pertinent for a good model to be able to predict such instances. Furthermore, the fact that log-transforms smooth out high volatility potentially works against such treatment of the data. Doing so would significantly dampen the effects of spikes, which are a key feature in electricity prices.

Sensitivity to weather

The impact of weather on the demand for electricity has long been established. In essence, there is a non-linear relationship between temperature and electricity consumption (and hence prices) since the need for cooling would be greater on hot days and vice versa, and there is a comfortable range of temperatures in between which neither cooling nor heating would be particularly essential. Depending on the electricity-generation methods in a particular country, there may also be a causal link between climate and energy supply. For a recent review of some empirical literature related to this subject, see Auffhammer and Mansur (2014).

Some early studies used heating degree days (HDD) and cooling degree days (CDD) to capture the non-linearity (for examples, see Amato et al. (2005) and Sailor and Munõz (1997)). However, certain limitations of the HDD-CDD approach have been put forward. In determining HDDs and CDDs, threshold temperatures must first be chosen. There is no universally-accepted rule for determining these thresholds, and they vary by location[5][6][7]. Additionally, this method does not truly explore the extent of the non-linear relationship since temperature values are now transformed into a dummy or count variable.

An alternative to the HDD-CDD approach is to estimate the effects of temperature using semiparametric or nonparametric models (Engle et al., 1986; Henley and Peirson, 1997). Naturally, the nonparametric (or nonparametric portion of the semiparametric) specification is flexible to the non-linear relationship between temperature and electricity consumption. However Bessec and Fouquau (2008) suggest that such models are not adequate for the purpose of analysing the link between the two.

A third option is in the form of logistic smooth threshold regression (LSTR) models (Moral-Carcedo and Vicéns-Otero, 2005) and panel smooth transition regression (PSTR) models (Bessec and Fouquau, 2008).

Mean-reversion

Mean-reversion is a phenomenon which applies to many time series, and is not new to econometrics. In the context of electricity pricing, a common way of modelling mean-reverting prices

[5] European Union: `https://www.eea.europa.eu/data-and-maps/indicators/heating-degree-days-2`
[6] Australia: `http://www.bom.gov.au/climate/map/heating-cooling-degree-days/documentation.shtml`
[7] United States: `https://w1.weather.gov/glossary/index.php`

is by means Ornstein-Uhlenbeck processes and estimating the discrete data using autoregressive models (Bierbrauer et al., 2007; Benth et al., 2007; Knittel and Roberts, 2005).

Note, however, a cursory glance at Figure 1.1(a) suggests that the mean may be changing over time. Furthermore, some empirical studies, for example, Al-Faris (2002) and Dergiades and Tsoulfidis (2008, 2011), find their respective samples of annual average real residential electricity prices to be integrated of order 1. As such, care must be taken in ensuring that an appropriate time series model is used for the data which is being analysed.

Seasonality

The impact of weather on prices was discussed earlier in this section. There, the implication was that extreme temperatures would lead to changes in prices. However, the usual annual temperature changes also influence prices, particularly in subtropical and temperate countries, where the seasons may be vastly different from one another.

In addition, it is also common to find seasonal components in shorter frequencies. For instance, inspection of Figure 1.2(a) shows that prices are lower on weekends than weekdays, on average. One explanation for this pattern is simply that the demand for electricity is lower on weekends due to many workplaces being relatively unused on weekends, and also potentially families being out of their homes more on weekends. An interesting point to note is that when the sample is arbitrarily split into two sub-samples, as in Figure 1.2(b)–(c), weekend prices are still lower than on weekdays, but the mean prices are different between the two sub-samples. Once again, this finding supports the need for time-varying parameters in modelling electricity prices. For data at the hourly level, diurnal patterns have been observed (Karakatsani and Bunn, 2008b; Li and Flynn, 2004, 2006). The simple and obvious explanation for this pattern is that the population engages in more energy-consuming activities at certain times of the day.

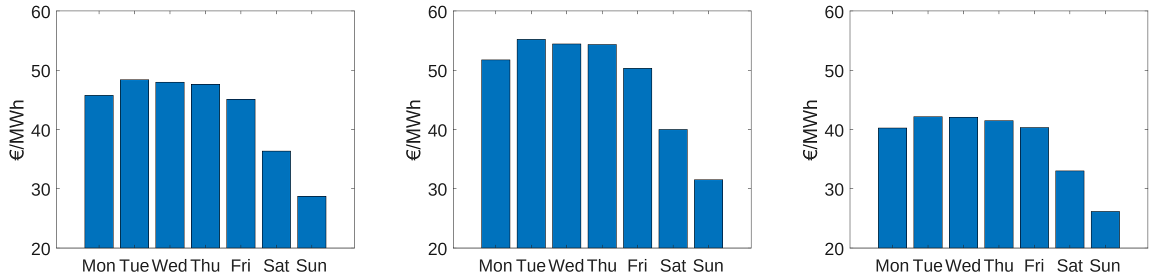

(a) 8th Feb 2005 – 28th Apr 2017 (b) 8th Feb 2005 – 31st Dec 2010 (c) 1st Jan 2011 – 28th Apr 2017

Figure 1.2: Mean spot electricity prices in Germany for each day of the week.

1.2 Sale and purchase curves

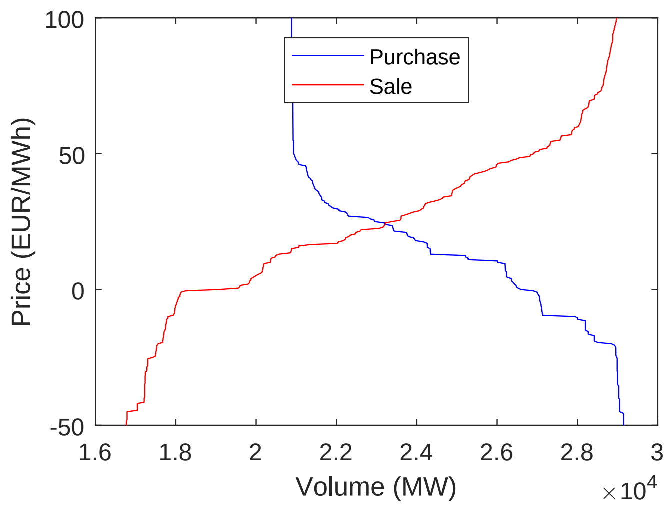

Figure 1.3: Sale and purchase curves within the price range -EUR 50 and EUR 100 on the first hour of 1 January 2016 in Germany.

Where data permits, individual bids from wholesale market participants can be combined to form sale (supply) and purchase (demand) curves, also referred to as "merit order curves". Take, for example, Figure 1.3, which shows a part of the curves for the first hour of 1 January 2016 in the German wholesale market. The price of wholesale electricity in that hour of EUR24.70/MWh corresponds to the intersection between the two curves. While the series in Figure 1.1(a)) provides some insight into the trajectory of prices over time, it can be argued that the sale and purchase curves contain more information within a single intra-day period. For example, if the intersection occurs at a point where both curves are relatively flat, then a small shift in either curve would not result in a big price change. Conversely, an intersection at a steep portion of either curve means that prices are very sensitive to small changes in volume. In other words, it is possible to use supply and demand curves to investigate or predict prices, volumes, and elasticities.

The literature on forecasting these curves is relatively new. Ziel and Steinert (2016) and Shah and Lisi (2020) use different approaches to forecast the sale and purchase curves, and use the intersection of the curves as forecasted prices. On the other hand, Mestre et al. (2020) propose a new model, and show that it produces good forecasts of the sale curves. While the curves themselves may exhibit some seasonality, using them to forecast prices leads to different considerations over the use of prices themselves. For example, a price spike in Figure 1.1 may be the result of a sudden and short-lived change in either the shape or location of one of the two curves, in which case an appropriate model would need to accommodate some sort of instantaneous shock to the system. On the other hand, a spike might simply be the result of a small, explainable shift in one of the curves when the intersection lies on a particularly steep part of either curve, which might already be accounted for in a relatively simple model.

1.3 Thesis structure

As with many other financial and economic time series, it is reasonable to expect the evolution of wholesale electricity prices to change over time. These changes may be due to regulatory adjustments, technological advancements, or other shifts in supply or demand. Without delving into the cause of any such variation over time, Chapter 2 introduces semiparametric models which can capture time-variation in regression model parameters.

Forecasts of supply and demand curves have been used to predict electricity prices. The first to do so was Ziel and Steinert (2016), who use forecasts of volume (y-axis in Figure 1.3) at a set of predetermined prices (x-axis in Figure 1.3), and construct sale and purchase curves by essentially joining the dots. Shah and Lisi (2020) use a nonparametric kernel-type approach to forecasts the curves. The former method is simple, since it is akin to estimating the same time series forecasting model multiple times, each with a different set of volume data. however, it implicitly assumes that the bid volumes at each price level are independent of those at other prices. On the other hand, the latter allows volumes of bids at each price to influence levels at other prices, but is requires somewhat-tedious numerical estimation. Chapter 3 uses a functional principal components analysis approach to forecast sale and purchase curves. This approach treats each curve as a datum point instead of as a collection of data points like Ziel and Steinert (2016). However, it is equally simple, while being quicker to implement than the method proposed by

Shah and Lisi (2020).

As indicated earlier in this introduction, the sale and purchase curves contain more information than just price at the intersection. Since these curves can be constructed for each intra-day interval, it becomes possible to analyse the behaviour of price elasticity of demand in wholesale electricity markets over time. This particular aspect of wholesale electricity markets has not been considered before, but provides further insight which may benefit market participants. Chapter 4 shows how intra-day price elasticity of demand can be computed, and also illustrates certain patterns in these quantities. The chapter also contains discussions on various ways in which existing regression-based estimates of price elasticity of demand may not be appropriate. The discussion does not seek to discredit or dismiss the usefulness of regression in this context, but highlights the importance of taking great care when interpreting estimates.

Chapter 5 briefly introduces some other research ideas which were hatched and explored at a very preliminary level. These ideas may form the basis of future research. Finally Chapter 6 summarises and concludes the thesis.

Chapter 2

Semiparametric point forecast of prices

The electricity market in the European Union underwent deregulation in 1996 as a direct result of the EU Directive 96/92/EC. One of the main objectives of this directive was to establish a competitive electricity market. As a consequence of the deregulation, the electricity spot market went from one which was extremely technical to one which is largely similar to other commodities. Today, spot prices are determined by matching supply and demand curves in daily auctions.

The measures adopted by Germany in response were "well beyond the requirements of the ... directives" (OECD, 2004). The *Erneuerbare-Energien-Gesetz* (EEG), or Renewable Energy Sources Act, which was introduced in 2000 and revised in 2014 and 2017, has proven to be effective in promoting production of electricity from renewable sources (Bode and Groscurth, 2006). Furthermore, the passing of legislation pertaining to *Energiewende*, or energy transition, in 2010 lay the groundwork for a shift towards more environmentally friendly, reliable, and affordable energy generation. In addition to the transition *towards* more sustainable generation, there was also a significant shift *away* from nuclear energy. In 2010 the Federal Government adopted energy concepts which outlined several targets for the German energy industry until 2050 (BMU and BMWi, 2010). These targets surpassed those of any other EU directives, and included reducing greenhouse gas emissions by 40% by 2020, and at least 80% by 2050, taking 1990 levels as the basis of comparison. Goals were also determined for the amount of renewable energy as a percentage of gross final energy consumption, share of heat consumption, *et cetera*

(BMWi, 2015). In 2011, in response to the nuclear disaster in Fukushima, the 2010 energy concepts were amended to accelerate the phasing out of nuclear energy (Grossi et al., 2017). A reasonable supposition to make is that each of these regulatory changes, in altering the structure of the energy market, will, in turn, affect the evolution of prices. Of course, there may also be other factors which could likewise lead to variations in the structure of prices.

2.1 Introduction

The dynamics of electricity prices in wholesale spot markets contain many characteristic features. As one may expect, prices exhibit strong seasonality. For example, demand for electricity in temperate countries tends to rise in winter when heating is required, driving prices up. In addition to this annual pattern, it is common to observe higher-frequency seasonal trends in the form of day-of-the-week effects; prices are typically lower on Saturdays, and even more so on Sundays, than they are on weekdays. Besides the annual seasonality in temperature and weather which is generally attributed to well-known patterns in climate, other weather-related shocks are also expected to influence prices. For example, on abnormally-hot summer days, the demand for cooling—and therefore electricity—increases significantly in countries like Australia. Conversely, extremely cold winter days in northern Europe and North America may drive prices up as a result of increased heating requirements. Although weather and temperature have been cited as being demand-side drivers, it is important to note that they can also influence supply in certain markets. For instance, countries which generate a large share of their energy from hydroelectric sources such as Norway, Paraguay, and Malaysia may experience increased (ability to) supply in months with more rainfall[8], and regions with significant amounts of solar energy, such as South Australia, may see higher supply on hotter days.

Prices are also generally regarded as mean-reverting, as alluded to in numerous studies including Bierbrauer et al. (2004), Bierbrauer et al. (2007), Weron et al. (2004), Knittel and Roberts (2005), and Eichler and Türk (2013), although Bosco et al. (2007), Bosco et al. (2010), and Lisi and Nan (2014) model electricity prices as nonstationary time series. Depending on the dataset and the sample period, electricity prices may also be censored from above or below. For instance, the Australian National Electricity Market (NEM) has an annually-reviewed market price cap

[8] Extended periods of greater rainfall does not automatically lead to larger quantities of generated electricity. Instead, more water can be collected in hydroelectric dams, thereby increasing a plant's capacity for electricity generation.

which was set at AUD14,700/MWh for the period 1 July 2019 to 30 June 2020[9]. In Germany, wholesale electricity prices in the day-ahead spot market were restricted to being non-negative until this condition was relaxed in 2008. Evidence of volatility clustering has been found in electricity prices, and treated using (generalised) autoregressive conditional heteroskedasticity models (Goto and Karolyi, 2004; Knittel and Roberts, 2005; Escribano et al., 2011). However, Knittel and Roberts (2005) and Garcia and Contreras (2005) report that the improvement in performance using these models is contingent on the choice of sample period. Perhaps the most pronounced feature of electricity prices is the presence of spikes. Although a definitive way of identifying price spikes has yet to be developed, they are typically described as being short-lived, unanticipated, extreme changes in the spot price. Here, the term "short-lived" refers not only to the fact that the large change occurs in very few (typically one, at the daily frequency) consecutive periods, but also that the extreme movement will be immediately followed by a sharp and near-equal reversal. In other words, the mean of prices before and after a spike are not significantly different, otherwise the phenomenon would be better referred to as a jump[10].

In essence, there is a mean-reverting base component containing the aforementioned annual seasonality, weekly trend, and reaction to weather changes, as well as a spike component. It is the existence of the latter which introduces most of the complexity to electricity price forecasting (EPF). The occurrence of spikes is uncontested and taken as matter of fact; the real question relates to when exactly they happen, and what the magnitude of these spikes may be. Furthermore, it is often useful to have some proxy of what the base price would have been at times when spikes occur. These values, too, are not known. The literature has explored a few ways in which to approach the concurrent existence of the two components. The simplest method would be to ignore spikes altogether (Lucia and Schwartz, 2002). However, a number of works argue that spikes are particularly hazardous and non-negligible (Cartea and Figueroa, 2005; Huisman, 2008; Clements et al., 2013). Most works on EPF acknowledge these spikes. After separating the two components, some research only models the base component (Cartea and Figueroa, 2005; Gianfreda, 2010; Ketterer, 2014) by replacing spikes with a proxy of base prices, or only models spikes (Becker et al., 2007; Christensen et al., 2009, 2012). Alternatively, both parts are modelled concurrently (Bierbrauer et al., 2007; Cartea and Figueroa, 2005; Huisman, 2008;

[9] This price cap was announced on 20 June 2019 on the Australian Energy Market Operator (AEMO) website as Market Notice number 68807 (Change Number: CHG0054651).

[10] Of course, in the absence of an undisputed method for identifying spikes, there is really no way to confirm or refute this claim or, indeed, any competing ones.

Karakatsani and Bunn, 2008a; Janczura and Weron, 2010, 2012; Janczura et al., 2013). Even when modelling both components, the norm when using regime-switching models, for example, is still to isolate the components before fitting models (Bierbrauer et al., 2007; Janczura et al., 2013). Of course, any method which separates the two components will likely be subject to some spike-identification algorithm. An examination and comparison of such techniques was conducted by Janczura et al. (2013).

With regards to the actual modelling of the components, a number of different, often complementary, approaches have been proposed. Surveys of relevant literature can be found in Aggarwal et al. (2009a,b), which have since been superseded by Weron (2014), Jiang and Hu (2018), and Nowotarski and Weron (2018). All of the three more recent survey papers group different EPF approaches into a few classes, namely, multi-agent models, fundamental models, reduced-form models, statistical models, and computational intelligence models. The methods used in this chapter fall under the categories of reduced-form and statistical models, as defined by the survey papers.

A number of techniques for modelling the trend in the first component, i.e. the base component, was studied in Lisi and Nan (2014). They further decompose the base component into long-term component and a "periodic and calendar" component. For the long-term trend, Lisi and Nan (2014) consider the use of polynomial and sinusoidal regression, local polynomial kernel regression, splines, spectrum analysis, various decomposition methods, and a number of filters such as the Hodrick-Prescott filter and Kolmogorov-Zurbenko filter. Sinusoidal functions and wavelet decomposition, in particular, are popular in the literature (Cartea and Figueroa, 2005; Geman and Roncoroni, 2006; Bierbrauer et al., 2004; Weron and Misiorek, 2008; Conejo et al., 2005; Janczura and Weron, 2010). Regardless of the specific method used, it is apparent from the results in Lisi and Nan (2014) that the trend in prices varies over time. They also investigate a few methods of modelling the shorter-term periodic and calendar component. Ultimately, the simple practice of using appropriate dummy variables seems to have gained some traction (Knittel and Roberts, 2005; Bierbrauer et al., 2007; Huisman, 2008; Lisi and Nan, 2014; Nowotarski and Weron, 2016).

The effect of exogenous variables such as weather has long been included in electricity modelling. Temperature, in particular, is often used in the literature. Although the manner in which this variable is introduced varies, the general consensus, dating back to Engle et al. (1986), is that its effect on electricity demand is nonlinear. Knittel and Roberts (2005) include the first, second, and third powers of this explanatory variable to capture the nonlinearity. In literature related to electricity demand or load modelling, Sailor and Munõz (1997) and Amato et al. (2005) include heating degree days and cooling degree days instead of temperature measurements. On the other hand, Henley and Peirson (1997) uses kernel smoothing techniques to estimate a nonparametric specification for the effect of temperature on demand. Finally, the persistence and mean-reversion in prices is often modelled by autoregressive-type specifications (Cuaresma et al., 2004; Bierbrauer et al., 2007; Misiorek et al., 2006).

Instead of attempting to filter out spikes before modelling either of the components separately, an alternative to isolation is to estimate the base and spike components concurrently. Bierbrauer et al. (2007), Huisman (2008), Karakatsani and Bunn (2008a,b), Janczura and Weron (2012), and Eichler and Türk (2013), among others, attempt this by means of regime-switching models with various specifications, which are now generally preferred over the jump-diffusion alternatives employed by Escribano et al. (2011), Cartea and Figueroa (2005), Knittel and Roberts (2005), Geman and Roncoroni (2006), *et cetera*. Another type of model which has been used, although seemingly less so, is the threshold autoregressive model (see, for example, Robinson (2000), Rambharat et al. (2005), Weron and Misiorek (2006), and Gaillard et al. (2016)).

At this point it is worth noting the differences in the implementation of modelling and forecasting. In both cases, the data used to estimate model parameters is available only up to the present day. However, when those estimates are used to forecast the variable of interest out of sample, only the deterministic variables are known in the future. Any random variables will need to be predicted with reasonable accuracy. In some cases, forecasts of such random variables are generally accepted to be quite good, depending on the intended forecast horizon. In other cases, the availability of such predictions is a considerable restriction, and those variables, informative as they may be, simply cannot be included. Additionally, reliable prediction of spikes in the future is questionable. While certain models have been shown to produce a reasonable interpretation of the probability and magnitude of spikes in-sample, works which forecast spikes are typically performed on high frequencies such as on intra-day (Christensen et al., 2009, 2012) or, at most daily (Mount et al., 2006; Huisman, 2008) data.

Finally, limited consideration has been given to the idea of time-varying parameters in electricity pricing models. This aspect of forecasting models is of particular importance due to the various regulatory changes discussed at the beginning of this chapter. For the purpose of this thesis, random coefficients-type models are excluded from this particular discussion. This is because their parameters vary in the sense that they are functions of other explanatory variables, those functions themselves having constant parameters. In this sense, "true" time-varying parameters are either those of state-space models (Karakatsani and Bunn, 2008a; Bordignon et al., 2013) in which the trend takes the form of a random walk, or nonparametric and semiparametric techniques in which one or more coefficients are smooth functions in time. The latter group is focused on in this report, noting that expanding the time-varying coefficients to include more than just the trend requires no more than a simple extension. In parallel literatures, Gao and Hawthorne (2006) fit a semiparametric model to estimate the trend of a global temperature series and Chen et al. (2018) use a nonparametric version of a heterogeneous autoregressive (HAR) model to forecast realised volatility. Both works find improvements in forecasting performance when coefficients are time-varying.

Note that the terms "nonparametric" and "semiparametric" have also been used in contexts other than that of time-varying coefficients, for instance, when a parametric form for the error distribution (Weron and Misiorek, 2008) or spike distribution (Eichler and Türk, 2013) is not assumed. A different semiparametric approach using kernel smoothing was employed by Clements et al. (2013), who model the probability of spikes with weights which are state-dependent, and obtained using multivariate kernel smoothing.

The time series models here are used to forecast daily electricity prices in the German wholesale market at different horizons, namely, 1, 7, and 30 day-ahead.

2.2 Models

The model developed here is, essentially, a reduced-form model. Whereas a structural model would simultaneously estimate a supply and a demand equation, the reduced form is more parsimonious and is simpler to estimate. Even so, it is imperative to include variables which represent supply-side or demand-side drivers, where available. To this end, seasonality and temperature are included as variables which influence demand. On the other hand, the variable which can be expected to have a significant influence on the demand side, particularly in Germany, is not suitable for our forecasting exercise. Recent years have seen a marked increase in the amount of wind energy generated in Germany, which serves as a substitute for the electricity source in our data set. It is known that on days when wind-generated energy is in excess, the price of electricity can drop drastically, sometimes even becoming negative. Unfortunately, this variable cannot be used for forecasting as predicted volumes of wind energy are not available early enough even for the 1 day-ahead forecast horizon, and are also generally very inaccurate.

As outlined in Section 2.1, a large number of different techniques have been considered for electricity price modelling and forecasting. In the time series modelling framework, the log transformation of prices is often used as a means of obtaining a more stable variance (Aggarwal et al., 2009a). Indeed, a large number of works analyse the logarithm of prices—deseasonalised or raw—including Lucia and Schwartz (2002), Bierbrauer et al. (2004, 2007), Weron et al. (2004), Cartea and Figueroa (2005), Bosco et al. (2010), Janczura and Weron (2012), Bordignon et al. (2013), and Ketterer (2014), to name a few. In more recent years, negative electricity prices are observed, which makes the practice of taking logarithms infeasible. Although models which were originally designed for log-prices can still be drawn upon for inspiration, an unfortunate consequence is that they cannot be used for direct comparison of the performance of any newly-proposed models.

2.2.1 Building blocks: Base regime

Methods used to account for the characteristics of base (non-spike) prices which were identified in Section 2.1, regardless of whether spikes are featured in the models, are presented here.

Seasonality

The long-term seasonal component of prices is sometimes represented using dummies for months or quarters (Knittel and Roberts, 2005). In the case of electricity prices, this seasonality is largely a proxy for changes in demand due to the change in seasons, therefore the change in average temperature. As such, it seems more realistic to model the annual seasonality as a smooth function instead of discrete steps. This can be achieved using a combination of sine and cosine functions (Bierbrauer et al., 2007). On the other hand, the shorter-term "seasonal" components may not have the same kind of wave form. Furthermore, the short-term seasonal effects typically manifest for the length of a day, which is exactly the frequency of our data. Therefore, there is no need to model them as functions, and using dummy variables for these effects such as day of the week or public holiday dummy variables (Bierbrauer et al., 2007; Escribano et al., 2011) is acceptable.

Temperature

Temperature clearly has a causal relationship with electricity price movements. Specifically, while seasonal average temperatures can be accounted for by other means, temperature anomalies should be included in a model for electricity prices. By this reasoning, it is clear that raw temperatures need to be preprocessed before entering into the model. A common treatment for temperature is to replace them with heating-degree days (HDD) and cooling-degree days (CDD). However, this practice can be expected to result in large overlaps and collinearity between long-term seasonality and the HDD and CDD variables. Using a treatment for temperature which has some similarities with Becker et al. (2007), this study calculates the deviation of daily temperature from its seasonal pattern. Furthermore, the impact of deviations from seasonal temperatures is permitted to be different in colder months and in hotter months.

First, the annual seasonality in temperature is removed, giving τ_t^* as the residuals from the ordinary least squares model

$$\tau_t^* = r_t - \hat{a}_r - \hat{b}_r t - \hat{\gamma}_{\cos,r} \cos\left(\frac{2\pi}{365}t\right) - \hat{\gamma}_{\sin,r} \sin\left(\frac{2\pi}{365}t\right), \tag{2.1}$$

where r_t is the recorded average temperature on day t, and $\hat{a}_r$, $\hat{b}_r$, $\hat{\gamma}_{\cos,r}$, and $\hat{\gamma}_{\sin,r}$ are estimated least-squares coefficients. Even after detrending, it can be expected that temperature anomalies of different directions will have different effects, depending on the season. In winter, a negative shock in temperature should increase electricity demand whereas a positive shock should decrease it. The converse should hold true for summer months. On the other hand, seasonal mean temperatures in spring and autumn are commonly known to fall within a comfortable range for humans, and a deviation from these temperatures, unless very extreme, would not necessitate heating or cooling. Assuming that abnormal temperatures play a less significant role in spring and autumn, τ_t^* is rescaled by a cosine function to give

$$\hat{\tau}_{t,w} = -\cos\left(\frac{t_{yr}(t)}{T_{yr}(t)} \times 2\pi\right) \times \tau_t^* \times \mathbb{1}\left(-\cos\left(\frac{t_{yr}(t)}{T_{yr}(t)} \times 2\pi\right) < 0\right), \tag{2.2}$$

$$\hat{\tau}_{t,s} = -\cos\left(\frac{t_{yr}(t)}{T_{yr}(t)} \times 2\pi\right) \times \tau_t^* \times \mathbb{1}\left(-\cos\left(\frac{t_{yr}(t)}{T_{yr}(t)} \times 2\pi\right) > 0\right), \tag{2.3}$$

where $\mathbb{1}(\cdot)$ is an indicator function which takes the value 1 if its argument is true and 0 otherwise, $t_{yr}(t)$ is how far into the year the date on t is, and $T_{yr}(t)$ is the total number of days in the year which t falls into. For example, on 4th January 2015, $t_{yr}(t)/T_{yr}(t) = 4/365$ and on 5th March 2016, $t_{yr}(t)/T_{yr}(t) = 64/366$. In other words, the cosine function takes a value close to -1 on 1st January and close to 1 on 1st July. This transformation inverts the sign of temperature anomalies in the colder months and gradually dampens the magnitude of anomalies the further away a date is from the peak of summer or winter. The two indicator functions further separate the transformed series into one for the colder months, $\hat{\tau}_{t,w}$, and one for warmer months, $\hat{\tau}_{t,s}$. The reason for splitting the series in two is that winter in Germany is generally cold enough to necessitate heating, while summer temperatures are often comfortably in the low 20°Cs. Increases in temperature during summer will not have as large an impact on demand for electricity as would a decrease in temperature during sub-zero temperatures in winter.

A benefit of including temperature in this manner over the use of, say, heating degree days and cooling degree days (Sailor and Munõz, 1997; Amato et al., 2005) is that it permits estimation the non-linear impact of the magnitude of temperature difference on prices instead of considering the extent of price changes as being conditional on the persistence in abnormal weather.

Meanwhile, the usual annual seasonality in temperature is not introduced as a variable in its own right, as it is captured in whichever method of accounting for long-term seasonality in prices is employed.

2.2.2 Building blocks: Spikes

Spikes are particularly tricky to model as little is known about them. Even so, it is important to properly acknowledge and account for them, either through an appropriate filtering method or a suitable model.

Filtering spikes

One of the aspects of the modelling which translates well between logs and levels is that of identifying spikes. Janczura et al. (2013) considered a variety of methods by which this exercise has been conducted. A number of these methods select spikes based on some fixed threshold. Some authors compare price levels against the predetermined threshold, others compare returns or first differences. While thresholding itself is a useful tool, there are several issues associated with the manner in which it is conducted, particularly when applied to electricity prices. The method proposed here addresses some of them, and has not been considered in the literature.

In devising the approach proposed here, the main objective is to strike a balance between aggressive and conservative identification. Overly-aggressive identification would result in filtering out prices which belong to the base regime, thereby excluding informative data which should be preserved. For example, even though Janczura et al. (2013) conclude that there is no single best method for outlier detection, the two methods which which result in identification of the most spikes in the German and New South Wales markets often led to parameter estimates of their simulated seasonal pattern which are less accurate compared to other approaches. On the other hand, incorrectly retaining spikes in the filtered series is also likely to lead to incorrect estimation of the model parameters, given the large magnitude of spikes.

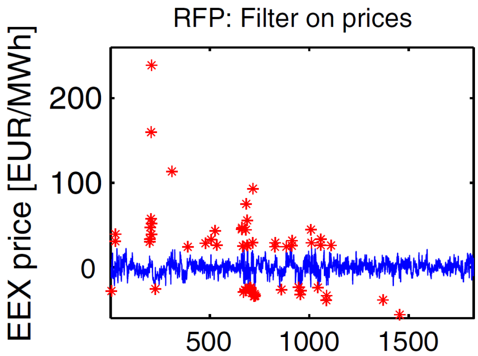

(a) Threshold for levels

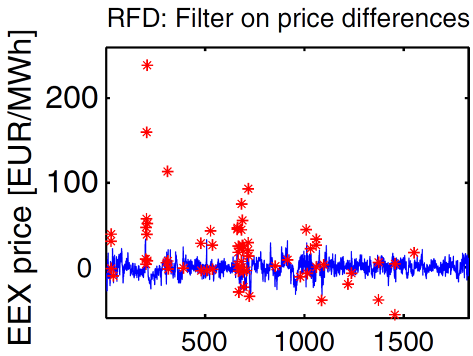

(b) Threshold for differences

Figure 2.1: A sample of filtered deseasonalised daily electricity prices in the German market using two different approaches. The red crosses indicate the deseasonalised levels which were identified as spikes using the respective methods. These figures are copied exactly from Fig. 2 in Janczura et al. (2013).

Similar to most existing methods, the proposed filtering process involves thresholds which are data-driven and are updated at each point in time. Another feature of spikes is that they occur for a brief period of time. Although the literature is still to come to a consensus on the length of the time interval over which spikes occur, it is noted that spikes may happen close to each other but typically do not occur on consecutive days[11]. This statement is easily justified by comparing the outcome from two of the "best" identification methods in Janczura et al. (2013)—one uses thresholds for levels and the other uses thresholds for first differences. For ease of comparison, two plots from Fig. 2 of Janczura et al. (2013) are copied exactly and included in Figure 2.1. These two plots illustrate two of the spike-identification methods considered by Janczura et al.

[11]This observation is based solely on daily-level frequencies. Intra-day electricity prices have different dynamics.

(2013) on wholesale electricity prices in Germany. Since the two methods are different, the outcomes will not be identical. However, upon closer inspection it can be seen the points that are identified as spikes in panel (a) which are also identified in panel (b) are accompanied by another "spike" in panel (b), which is around the mean of the series. This accompanying red cross is certainly not in concordance with the idea of spikes being abnormally large deviations from base prices. These accompanying "spikes" which are identified in Panel (b) are clearly the result of prices returning back to base prices *after* a spike has occurred.

On the other hand, a small number of red crosses which occur in panel (b) which *do not* turn up in panel (a) are strangely close to 0. These points are not accompanied by a second red cross since they are quite small anyway and do not require a large shift to return to base prices. Finally, focusing on panel (a), it can be seen that a small number of the red crosses, such as the one around approximately the 240th observation, may be outside some threshold level, but do not stray far from their respective neighbouring points. It is possible that instead of being spikes, these levels simply fall into a period where some variable exogenous to the seasonal trend has a prolonged effect on price levels. If so, it may be more appropriate to include the relevant variables in the model for base prices instead of treating those elevated prices as spikes.

The following is a proposed technique for filtering electricity prices at daily frequencies which addresses the issues highlighted above. Identification of spikes is based on three criteria. For each point in time, t, let $\mathcal{N}(t)$ be some neighbourhood around t. Then a spike is said to have occurred at t if

$$|\Delta p_t| > 2\sigma^{\Delta}_{\mathcal{N}(t)},$$

$$|\Delta p_{t+1}| > 2\sigma^{\Delta}_{\mathcal{N}(t+1)},$$

and

$$\operatorname{sgn}(\Delta p_{t+1}) = -\operatorname{sgn}(\Delta p_t),$$

where p_t is the observed spot price at time t, Δ is the first difference operator, $\operatorname{sgn}(\cdot)$ is the sign of its argument, and $\sigma^{\Delta}_{\mathcal{N}(t)}$ is the standard deviation of first differences of prices which fall within some neighbourhood of t, $\mathcal{N}(t)$. The first condition relies on the fact that a spike is an abnormally large movement in prices, the second exploits the nature of spikes to return to base prices the following day, and the third condition ensures that the large movement on the following day is, indeed, a reversal. Since this research lies in the context of forecasting, the neighbourhood $\mathcal{N}(t)$ will be backward-looking; it will include past observations up to, and including, t.

A drawback of this approach to filtering is its reliance on a future value to compute Δp_{t+1}. In the applications which follow, when the most current observation is p_t, then its state (spike or base) will depend only on the first condition $|\Delta p_t| > 2\sigma^{\Delta}_{\mathcal{N}(t)}$. On the following day, the data is re-examined to check whether or not the large shift is followed by a correction, and the assessment of the state at time t is amended if necessary. Prices which are identified as spikes are replaced with some proxy of their base values. In the study, the proxy is the base price from the most recent same weekday. In other words, the series of filtered prices, $\tilde{p}_t$, is

$$\tilde{p}_t = \begin{cases} \tilde{p}_{t-7} & \text{, if a spike is present at } t \\ p_t & \text{, otherwise} \end{cases} . \tag{2.4}$$

Note that the series $\tilde{p}_t$ must be populated recursively, since $\tilde{p}_{t-7}$ must be known at time t.

Although the norm is to conduct filtering on deseasonalised prices instead of raw prices because "the seasonal patterns substantially complicate the identification of price spikes and drops in raw data" (Janczura et al., 2013), this method, applied to the data set used in this report, identifies exactly the same points as being spikes regardless of whether deseasonalisation was conducted beforehand[12].

Including spikes

Regime-switching models are increasingly popular, and they have the added benefit of being able to provide an estimate of the probability of spikes as a function of some exogenous variable (Huisman, 2008). However, this thesis places emphasis on the definition of Janczura and Weron (2010), who label spikes as being "generally unanticipated". Accordingly, prediction of spikes in the future is not addressed here. Even so, the author recognises the possibility that there may be gains to be had from modelling base prices and spikes concurrently instead of filtering spikes out and modelling a proxy of the base regime.

[12] Janczura et al. (2013) deseasonalise the data using a wavelet approximation and subtracting the mean of deseasonalised prices corresponding to each day of the week. Instead, this study used local constant kernel smoothing.

2.2.3 Constructed models

With the building blocks established, the models are assembled additively. Once the basic model is constructed, it is further enhanced by modelling some of the coefficients as time-varying parameters instead of restricting them to be constant throughout each estimation window.

Constant-parameter model for base prices

The model for base (filtered) prices is given as

$$\textbf{Model 1:} \qquad \tilde{p}_t = \boldsymbol{w}_t' \boldsymbol{\phi} + \varepsilon_t, \tag{2.5}$$

where

$$\begin{aligned}
\boldsymbol{\phi} &= \left(\mu_0, \beta_1, \beta_2, \beta_3, \beta_4, \boldsymbol{\gamma}'\right)' \\
\boldsymbol{w}_t &= \left(1, \tilde{p}_{t-1}, \tilde{p}_{t-2}, \tilde{p}_{t-3}, \tilde{p}_{t-7}, \boldsymbol{x}_t'\right)', \\
\varepsilon_t &\sim N\left(0, \sigma^2\right).
\end{aligned}$$

The $\boldsymbol{x}_t = (x_{t,1}, x_{t,2}, \dots)'$ variables on the right-hand side of (2.5) explain some of the movement in the levels of $\tilde{p}_t$. First, trigonometric functions of time are introduced to account for the long-term seasonality in prices,

$$x_{t,1} = \cos\left(\frac{2\pi}{365}t\right) \qquad \text{and} \qquad x_{t,2} = \sin\left(\frac{2\pi}{365}t\right).$$

Lower prices on public holidays are accounted for by using dummy variables, $x_{t,3}, x_{t,4}$. Since these holidays only occur once a year, those days with somewhat similar effects on prices are arbitrarily grouped together into one variable, where possible, to reduce sparsity. Accordingly, $x_{t,3}$ takes the value 1 on New Year's day, Easter Sunday, May day, and Christmas day[13]. $x_{t,4}$ represents Unity day. The possibility that prices on different days of the week may vary is accommodated by means of day-of-the-week dummy variables, $x_{t,5}, \dots, x_{t,10}$, for Sunday to

[13]Other holidays including Corpus Christi, Good Friday, Easter Monday, Ascension Day, Whit Monday, and Boxing Day were also considered. They were excluded from the model due to statistical insignificance of their respective coefficients. Boxing Day on 2012 was the only instance in which the price of electricity was, upon visual inspection (see Figure 2.2), obviously different from the general trajectory of prices. However it appears to be a singular episode as the prices of electricity on Boxing Day of other years do not stand out.

Saturday, with Wednesday as the base day. Finally, variables for temperature anomalies, $(x_{t,11}, x_{t,12}, x_{t,13}, x_{t,14}) = \left(\hat{\tau}_{t,w}, \hat{\tau}^2_{t,w}, \hat{\tau}_{t,s}, \hat{\tau}^2_{t,s}\right)$, which are defined in Section 2.2.1, are also included.

Much of the literature which uses time series modelling assumes an additive form $p_t = S_t + X_t$, where p_t is the price, S_t represents a seasonal component, and X_t is a stochastic component. In the absence of spikes, the stochastic component is usually modelled as a time series, $X_t = \rho 0 + \sum_{i=1}^{k} \rho_i X_{t-i} + \eta_t$, and the additive model can be rearranged and represented by an autoregressive distributed lag (ARDL) model. However, the proposed model (2.5) leads to marginally lower mean squared forecast errors, even if the difference in accuracy is not statistically significant. More important, however, is the fact that (2.5) involves considerably fewer right-hand side variables than an ARDL model would, and is simply more compact and parsimonious without compromising forecasting ability.

Time-varying models for base prices

Given the changing conditions in the German electricity market, whether they be from the regulatory side or innovations in supply of renewable energy, there is almost surely a gradual change in the structure of electricity market and the impact of various inputs. Karakatsani and Bunn (2008a) and Bordignon et al. (2013) propose models which are time-varying in the sense that the trends are modelled as random walks. On the other hand, Troncoso et al. (2007) use weighted nearest neighbours for price prediction in electricity markets at the intra-day frequency. Their approach uses only historical prices as inputs for the smoothing, and does not incorporate any additional structure to the forecasting model. The proposed models in this chapter take a different route and use kernel smoothing techniques for selected parameters in (2.5) thereby imposing some structure while still permitting flexibility in coefficients across time[14].

First, consider the case where only the trend changes across time, and the effect of other parameters on prices is constant. Model TV-trend takes the form

$$\textbf{Model TV-trend:} \qquad \tilde{p}_t = g^a(t) + \boldsymbol{w}_t^{a'} \boldsymbol{\phi}^a + \varepsilon_t, \tag{2.6}$$

[14]An interesting point to note is that the proposed semiparametric models (TV-coef, and TV-double) meet the definitions for locally-stationary models (Dahlhaus, 2012).

where

$$\boldsymbol{\phi}^a = \left(\beta_1, \beta_2, \beta_3, \beta_4, \boldsymbol{\gamma}'\right)'$$

$$\boldsymbol{w}_t^{a'} = \left(\tilde{p}_{t-1}, \tilde{p}_{t-2}, \tilde{p}_{t-3}, \tilde{p}_{t-7}, \boldsymbol{x}_t'\right)',$$

and $g^a(t)$ is an unknown continuous smooth function in t.

In addition, it may also be useful to permit other coefficients to vary if the changing conditions around electricity pricing affect the manner in which prices are influenced by external factors. In the second version of a time-varying model, all the parameters except for holidays and day-of-week dummy variables are permitted to change with time. This model, Model TV-coef, is

$$\textbf{Model TV-coef:} \qquad \tilde{p}_t = \boldsymbol{v}_t^b \boldsymbol{\theta}^b(t) + \boldsymbol{z}_t' \boldsymbol{\zeta} + \varepsilon_t, \tag{2.7}$$

where

$$\boldsymbol{v}_t^b = (1, \tilde{p}_{t-1}, \tilde{p}_{t-2}, \tilde{p}_{t-3}, \tilde{p}_{t-7}, x_{t,1}, x_{t,2}, x_{t,11}, x_{t,12}, x_{t,13}, x_{t,14}) \tag{2.8}$$

$$\boldsymbol{z}_t' = (x_{t,3}, x_{t,4}, \ldots, x_{t,10}).$$

All the elements in $\boldsymbol{\theta}^b(t)$ are assumed to have the same degree of smoothness as one another, a concept which will become clearer when the estimation procedures are discussed. Occasionally there may be reason to expect different parameters to vary at different rates. For example, the coefficient for the vector of ones in (2.8) represents the overall price level, which may reasonably be expected to shift at a different rate to, say, the response of prices to the other variables such as temperature anomalies or lagged prices. In the next model, the trend, $g^c(t)$, is allowed to change with a different rate than the other time-varying parameters. The model, which has two different levels of smoothing, is

$$\textbf{Model TV-double:} \qquad \tilde{p}_t = g^c(t) + \boldsymbol{v}_t^{c'} \boldsymbol{\theta}^c(t) + \boldsymbol{z}_t' \boldsymbol{\zeta} + \varepsilon_t, \tag{2.9}$$

where

$$\boldsymbol{v}_t^{c'} = (\tilde{p}_{t-1}, \tilde{p}_{t-2}, \tilde{p}_{t-3}, \tilde{p}_{t-7}, x_{t,1}, x_{t,2}, x_{t,11}, x_{t,12}, x_{t,13}, x_{t,14}).$$

Although (2.7) and (2.9) may appear to be very similar, it will be convenient to have the two models distinguished in this manner when describing the estimation process.

2.3 Estimation

Each of the four models for base prices and the model which incorporates spikes described in Section 2.2.3 requires a different estimation procedure. The methods are described below, followed by a description of the method used to forecast prices q days ahead.

Model 1

Model 1 is the simplest as coefficient estimates can be obtained by simple linear regression under simple assumptions. The ordinary least squares estimator for $\boldsymbol{\phi}$ is given as

$$\hat{\boldsymbol{\phi}} = \left(\sum_{t=8}^{T} \boldsymbol{w}_t \boldsymbol{w}_t' \right)^{-1} \sum_{t=8}^{T} \boldsymbol{w}_t \tilde{p}_t.$$

Model TV-Trend

For the semiparametric models, the unknown time-varying parameters are estimated using kernel smoothing methods. First, define a probability weight function of the form

$$W_s(t) = W_{T,s}(t;h) = \frac{K\left(\frac{t-s}{Th}\right)}{\sum_{u=8}^{T} K\left(\frac{t-u}{Th}\right)},$$

where $K(\cdot)$ is a kernel probability function and $h = h_T$ is some bandwidth depending on T which satisfies $\lim_{T\to\infty} Th^2 = \infty$ and $\lim_{T\to\infty} h = 0$. It is commonly known that the choice of h has a significant impact on estimates, whereas the choice of the kernel function $K(\cdot)$ is less important as long as it is a bounded function that is symmetric around 0 and integrates to 1. The standard normal probability density function is used $K(x) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{x^2}{2}\right)$ throughout this chapter.

For Model TV-trend, the trend, $g^a(t)$, is estimated as a smooth function in time. The estimates $\hat{\boldsymbol{\phi}}$ and $\hat{g}^a(t)$ are obtained by the following three-step procedure, as outlined in Härdle et al. (2000) and Gao and Hawthorne (2006).

1. *Step one: Nonparametric estimation.* For every given $\boldsymbol{\phi}^a$, the estimator of the time-varying

trend is

$$\begin{aligned}
\tilde{g}^a(t) &= \tilde{g}^a_T(t; h, \boldsymbol{\phi}^a) \\
&= \sum_{s=8}^{T} W_s(t) \left(\tilde{p}_s - \boldsymbol{w}_t^{a'} \boldsymbol{\phi}^a \right) \qquad (2.10) \\
&= \sum_{s=8}^{T} W_s(t) \tilde{p}_s - \sum_{s=8}^{T} W_s(t) \boldsymbol{w}_t^{a'} \boldsymbol{\phi}^a.
\end{aligned}$$

2. *Step two: Parametric estimation.* Replace $g^a(t)$ by $\tilde{g}^a(t)$. Then approximate (2.6) by

$$\tilde{p}_t^a = \tilde{\boldsymbol{w}}_t^{a'} \boldsymbol{\phi}^a + \varepsilon_t,$$

where

$$\tilde{p}_t^a = \tilde{p}_t - \sum_{s=8}^{T} W_{T,s}(t) \tilde{p}_s \qquad \text{and} \qquad \tilde{\boldsymbol{w}}_t^{a'} = \boldsymbol{w}_t^{a'} - \sum_{s=8}^{T} W_{T,s}(t) \boldsymbol{w}_s^{a'}.$$

The least-squares estimator of $\tilde{\boldsymbol{\phi}}$ is then

$$\hat{\boldsymbol{\phi}}^a = \left(\sum_{t=8}^{T} \tilde{\boldsymbol{w}}_t^a \tilde{\boldsymbol{w}}_t^{a'} \right)^{-1} \sum_{t=8}^{T} \tilde{\boldsymbol{w}}_t^a \tilde{p}_t^a.$$

3. *Step three: Semiparametric estimation.* Finally, substitute $\hat{\boldsymbol{\phi}}^a$ into (2.10) to obtain an estimate of the nonparametric component,

$$\hat{g}^a(t) = \sum_{s=8}^{T} W_s(t) \left(\tilde{p}_s - \boldsymbol{w}_t^{a'} \hat{\boldsymbol{\phi}}^a \right).$$

Model TV-coef

In Model TV-coef, the variables corresponding to the nonparametric components are random variables instead of constants as in Model TV-trend. Consequently, Model TV-coef is estimated using a similar, albeit slightly more involved, three-step procedure.

1. *Step one: Nonparametric estimation.* For every given $\boldsymbol{\zeta}$, the estimator of the nonparamet-

ric parameters is

$$\tilde{\boldsymbol{\theta}}^b(t) = \left(\sum_{s=8}^{T} \boldsymbol{v}_s^b \boldsymbol{v}_s^{b'} W_s(t)\right)^{-1} \sum_{s=8}^{T} \boldsymbol{v}_s^b \left(\tilde{p}_s - \boldsymbol{z}_t' \boldsymbol{\zeta}\right) W_s(t) \tag{2.11}$$

$$= \left(\sum_{s=8}^{T} \boldsymbol{v}_s^b \boldsymbol{v}_s^{b'} W_s(t)\right)^{-1} \sum_{s=8}^{T} \boldsymbol{v}_s^b \tilde{p}_s W_s(t) - \left(\sum_{s=8}^{T} \boldsymbol{v}_s^b \boldsymbol{v}_s^{b'} W_s(t)\right)^{-1} \sum_{s=8}^{T} \boldsymbol{v}_s^b \boldsymbol{z}_t' \boldsymbol{\zeta} W_s(t)$$

2. *Step two: Parametric estimation.* Replace $\boldsymbol{\theta}^b(t)$ in (2.7) with $\tilde{\boldsymbol{\theta}}^b(t)$, giving the approximation

$$\tilde{p}_t^b = \boldsymbol{z}_t^{b'} \boldsymbol{\zeta} + \varepsilon_t,$$

where

$$\tilde{p}_t^b = \tilde{p}_t - \left(\sum_{s=8}^{T} \boldsymbol{v}_s^b \boldsymbol{v}_s^{b'} W_s(t)\right)^{-1} \sum_{s=8}^{T} \boldsymbol{v}_s^b \tilde{p}_s W_s(t)$$

$$\boldsymbol{z}_t^{b'} = \boldsymbol{z}_t' - \left(\sum_{s=8}^{T} \boldsymbol{v}_s^b \boldsymbol{v}_s^{b'} W_s(t)\right)^{-1} \sum_{s=8}^{T} \boldsymbol{z}_t' W_s(t).$$

Then the least squares estimator of $\boldsymbol{\zeta}$ is

$$\hat{\boldsymbol{\zeta}} = \left(\sum_{s=8}^{T} \boldsymbol{z}_t^b \boldsymbol{z}_t^{b'}\right)^{-1} \sum_{s=8}^{T} \boldsymbol{z}_t^b \tilde{p}_t^b.$$

3. *Step three: Semiparametric estimation.* Finally, replace $\boldsymbol{\zeta}$ with $\hat{\boldsymbol{\zeta}}$ in (2.11) to obtain the estimate

$$\hat{\boldsymbol{\theta}}^b(t) = \left(\sum_{s=8}^{T} \boldsymbol{v}_s^b \boldsymbol{v}_s^{b'} W_s(t)\right)^{-1} \sum_{s=8}^{T} \boldsymbol{v}_s^b \left(\tilde{p}_s - \boldsymbol{z}_t' \hat{\boldsymbol{\zeta}}\right) W_s(t).$$

Model TV-double

Now allow the same parameters as those in Model TV-coef to vary with time, but additionally permit the trend to have a different level of smoothness from the other coefficients. In practical terms, what this means is that different bandwidths, h_1 and h_2, will be used to estimate $g^c(t)$ and $\boldsymbol{\theta}^c(t)$, respectively. A larger bandwidth parameter places more equal weights on the observations in the estimation sample, whereas a smaller bandwidth attaches heavier weights to values which are closer in time to the observations of interest than those which are further away. The parameters in this model are estimated using a process adapted from Park et al. (2015) by adding an additional step to the TV-coef method.

1. *Steps one to three: Three-step estimation.* Let

$$\boldsymbol{\theta}(t) = \left(g^c(t), \boldsymbol{\theta}^c(t)'\right)',$$

   and $\boldsymbol{v}_t^b$ be defined as in (2.8). Then let $\tilde{\boldsymbol{\theta}}(t) = \left(\hat{g}^c(t;h_1), \tilde{\boldsymbol{\theta}}^c\left(t;h_1\right)'\right)'$ and $\hat{\boldsymbol{\zeta}}$ be the estimates of $\boldsymbol{\theta}(t)$ and $\boldsymbol{\zeta}$, respectively, using bandwidth $h = h_1$ in the three-step method for Model TV-coef.

2. *Step four: Semiparametric estimation (b).* In the final step, define

$$\tilde{e}_t = \tilde{p}_t - \hat{g}^c\left(t;h_1\right) - \boldsymbol{z}_t'\hat{\boldsymbol{\zeta}}.$$

   Then the estimator for $\boldsymbol{\theta}^c(t)$ is

$$\begin{aligned}\hat{\boldsymbol{\theta}}^c(t) &= \hat{\boldsymbol{\theta}}^c(t;h_2) \\ &= \left(\sum_{s=8}^{T} \boldsymbol{v}_s^c \boldsymbol{v}_s^{c'} W_{T,s}\left(t;h_2\right)\right)^{-1} \sum_{s=8}^{T} \boldsymbol{v}_s^c \tilde{e}_t W_{T,s}\left(t;h_2\right).\end{aligned}$$

Note that the estimates for all four models have closed-form solutions, and are simpler and quicker to obtain than other models and approaches which use numerical optimisation.

Forecasting

Let $\boldsymbol{\mathcal{C}}(t)$ and $\boldsymbol{\mathcal{X}}_t$ represent the coefficients and right-hand side variables, respectively, at time t for any of the above models. Then the q-day-ahead forecast on day t, $\hat{p}_{t+q|t}$ is calculated as follows.

1. Estimate $\hat{\boldsymbol{\mathcal{C}}}(t)$ using the past observations up to time t.

2. For the one-day-ahead forecast at time t, use $\hat{\boldsymbol{\mathcal{C}}}(t)$ and $\boldsymbol{\mathcal{X}}_t$ to predict $\hat{p}_{t+1|t}$.

3. For the two-day-ahead forecast at time t, note that since $\boldsymbol{\mathcal{X}}_{t+1}$ includes $\tilde{p}_{t+1}$, all the predictor variables are not observed. Instead, use $\hat{\boldsymbol{\mathcal{C}}}(t)$ and $\hat{\boldsymbol{\mathcal{X}}}_{t+1|t}$, in which $\tilde{p}_{t+1}$ is replaced with $\hat{p}_{t+1|t}$ and calculate $\hat{p}_{t+2|t}$.

4. For the q-day-ahead forecast, recursively replace $\tilde{p}_{t+1}, \ldots, \tilde{p}_{t+q-1}$ with the predicted values, $\hat{p}_{t+1|t}, \ldots, \hat{p}_{t+q-1|t}$, until $\hat{p}_{t+q|t}$ is obtained[15].

2.4 Data

Prices

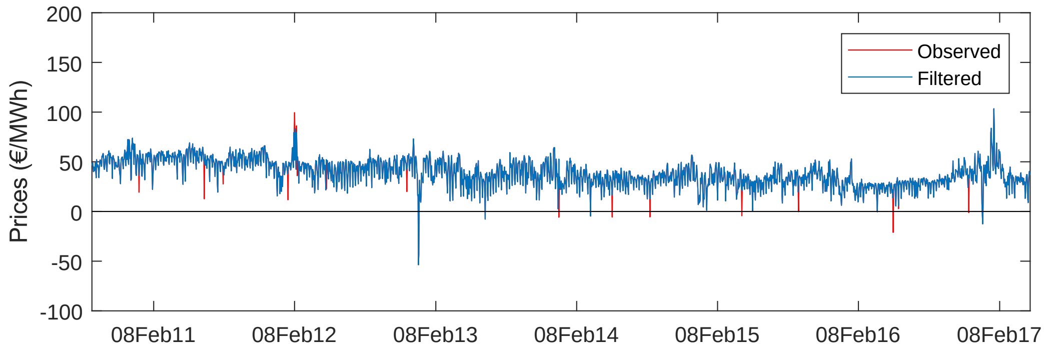

Figure 2.2: Plot of volume-weighted daily German electricity spot prices from 1 September 2010–28 April 2017. The red line represents observed prices, whereas the blue line corresponds to prices which have been preprocessed as described in Section 2.2.2.

Hourly electricity spot prices and volumes are obtained from the European Electricity Exchange (EEX). Daily prices are then calculated as the volume-weighted average price for each day. The time series of daily German electricity spot prices are presented in Figure 2.2, with summary statistics shown in Table 2.1.

Table 2.1: Summary statistics of daily German electricity spot prices and volumes from 1 September 2010–28 April 2017. There are 2432 observations in the sample.

	Min.	Max.	Mean	Std. dev.	Skew.	Kurt.
Bids (€/MWh) [raw]	-53.63	103.40	38.45	13.29	-0.20	4.96
Bids (€/MWh) [preprocessed]	2.82	74.88	38.61	12.15	0.03	2.51
Volume (MWh)	482657	1066710	669043	84151	0.73	4.13

German electricity spot prices have been investigated in numerous modelling and forecasting studies, using a variety of methods. Evidently, this particular electricity market has been of interest to researchers for many reasons including the regulatory background and natural features

[15]This method of forecasting produces what is referred to in Terasvirta et al. (2010) as a naïve forecasts. The errors will be cumulative as forecast horizon increases. However, this aspect of the naïve forecast does not affect this Chapter; the only instance in which forecast errors are used is in the tests for predictive accuracy or superiority. The Diebold-Mariano test Diebold and Mariano (1995) accounts for forecast horizon in its own way, and the model confidence set procedure uses empirical distributions obtained by a stationary bootstrap.

in Germany. More recently, however, the characteristics of spot price movements seem to be different from those in the 2000's. This sample differs from those considered in the vast majority of the literature in that it seems to exhibit a large number of negative spikes relative to the fairly small number of positive spikes. Of the papers surveyed for this chapter, only Grossi and Nan (2019) consider a data set (Italian market from 1 January 2013 to 1 January 2014) which has this feature.

Sudden and large price movements, or spikes, are a salient feature of electricity prices. Interestingly, the point in Figure 2.2 which seems to best fit this description falls on Christmas 2012, but is not retrospectively identified as a spike by the filtering method proposed in this chapter. This is because prices remained low on Boxing Day 2012, only returning to normal levels on 27 December 2012. Although this particular non-identification is almost certainly incorrect, it appears from the plots that, overall, the filtering method seems to work quite well. It is worth noting, however, that besides Christmas 2012, spikes that have occurred since 2010, although more frequent, may not be as hazardous as the ones which occurred before. The structure of prices in the German electricity market has clearly undergone some changes. As such, it would appear that a re-examination of German electricity spot prices and related modelling approaches is warranted. Note that all of the prices alluded to in the remainder of this chapter are filtered, since enough price data prior to this sample is available for the filtering process.

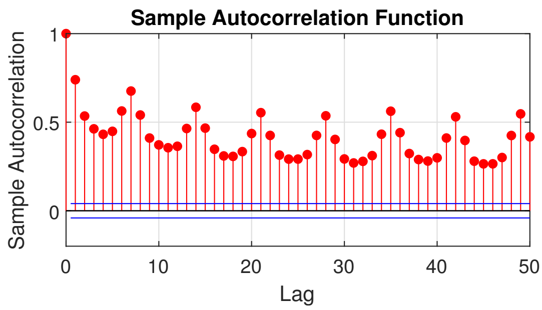

(a) Full sample

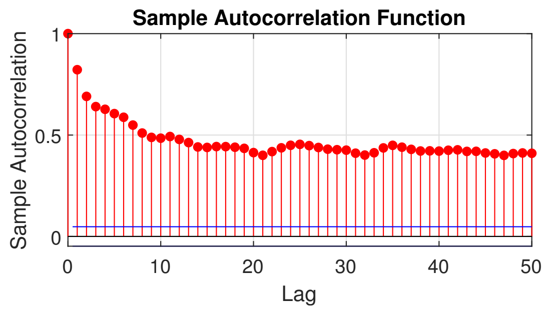

(b) Sample without weekends

Figure 2.3: Autocorrelation function of preprocessed prices

An interesting feature in the structure of electricity prices can be seen in Figure 2.3. From Panel (a), there are peaks in the sample autocorrelation function every increment of seven lags. However, when weekends are removed from the sample entirely, Panel (b) shows that the persistence is almost monotonically-decreasing in lags. The lots in Figure 2.3 suggest that prices on weekdays may be quite significantly different to those on weekends. This fact provides some justification for the specific inclusion of the seventh lag of prices, $\tilde{p}_{t-7}$, in specification (2.5).

Temperature

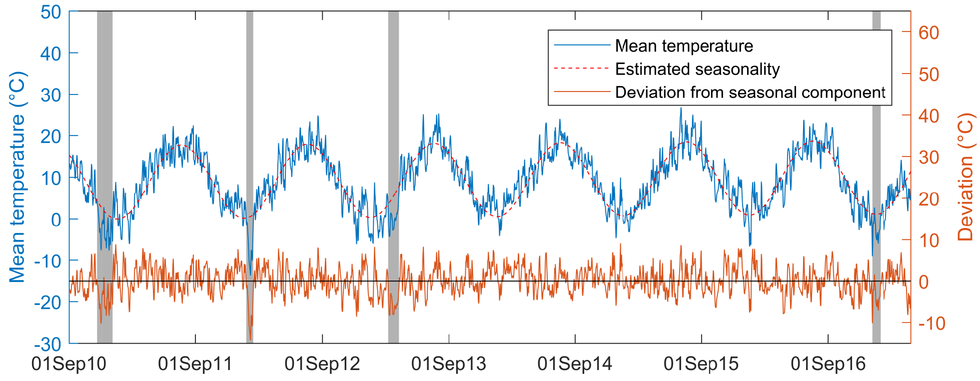

Figure 2.4: Average daily temperature (r_t), fitted seasonality, and deviation from seasonal component (τ_t^*) for the period 1st September 2010 to 28th April 2017. Several periods of persistent, larger-than-average negative temperature anomalies are highlighted.

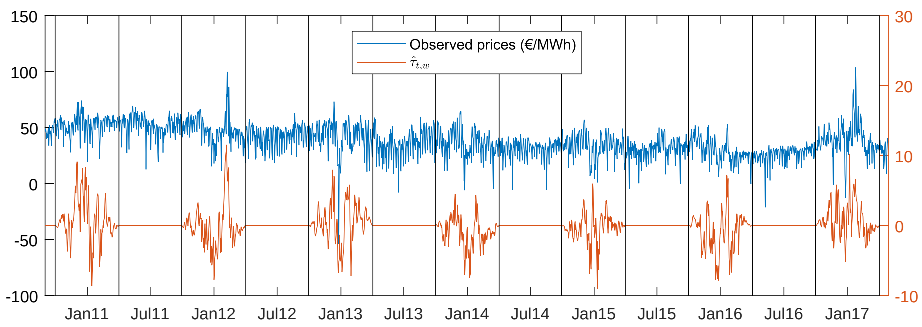

(a) $\hat{\tau}_{t,w}$

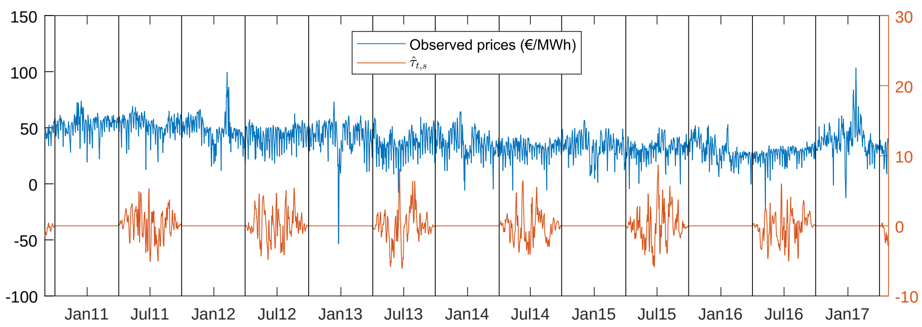

(b) $\hat{\tau}_{t,s}$

Figure 2.5: Plots of daily German electricity spot prices and transformed temperature anomalies. The black lines indicate where the cosine function in (2.2) meets zero.

Daily mean temperature data for 504 stations across Germany were obtained from DWD Climate Data Center (2018). The average daily German temperature used in this analysis is the mean of those temperatures recorded at all stations with available readings across the entire estimation period. It is plain to see from the plot in Figure 2.4 that there is a distinct annual cycle, as expected in a temperate country. The cyclical trend according to (2.1) fits the data quite well, revealing several large deviations. For instance, it is clear that dramatic reductions of up to 13.6°C from the seasonal average occurred during particularly cold winters in 2010–2011, early 2012, March 2013, and January 2017. Comparison between Figures 2.2 and 2.4 does not suggest a strong relationship between temperature anomalies and spiky behaviour. Note, however, that the transformed temperature series for colder months, $\hat{\tau}_{t,w}$, appears to have some relationship with p_t, particularly during some periods with higher prices, as seen in Figure 2.5(a). On the other hand, temperature anomalies in the warmer months, $\hat{\tau}_{t,s}$, shown in Figure 2.5(b) show a much weaker co-movement with prices. This is expected since a departure from average temperatures in summers, which are generally still comfortable, will have a much smaller impact than going from approximately 0°C to either -8°C or 8°C in winter. Data on actual temperatures are more easily obtained than a historical record of predicted temperatures. For example, 30-day-ahead temperature forecast are available each day, but historical records of these temperatures are not available. Observed temperatures have been used for retrospective forecasting under the assumption that temperature forecasts are reasonably accurate, and the discrepancy between using observed and predicted temperatures will be minimal.

2.5 Empirical results

A number of different ways in which to forecast electricity prices have been proposed here. First, it is prudent justify the proposed extensions to Model 1, that is, the decision to introduce Models TV-trend, TV-coef, and TV-double.

As a first point of comparison all the forecasts on the series of observed prices are superimposed to get some basic intuition of how the methods perform relative to one another. Then plots of the cumulative mean-square forecast errors (MSFE) of each method are shown. Finally, Diebold-Mariano statistics for relative prediction accuracy are presented.

All the coefficients are estimated using moving subsamples (rolling windows) of 1337 observations, or just over three and a half years.

2.5.1 Justification of time-varying extensions

Before delving into comparisons between the models, it is useful to first motivate the use of time-varying parameters. Figure 2.6 shows some of the estimated coefficients using Model 1 and

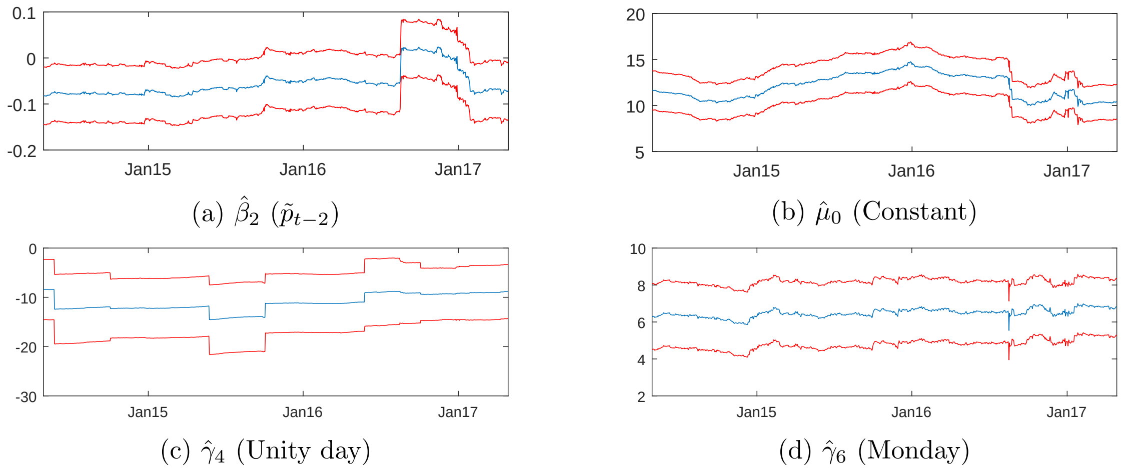

Figure 2.6: Selected plots of estimated coefficients (blue) and 95% confidence intervals (red) using Model 1.

rolling windows with 1337 observations over time. Even with the minor changes in the sample observations (excluding the oldest observation and including one new value) there is clearly some movement in the coefficient of $\tilde{p}_{t-3}$ and the intercept (panels (a) and (b)). The estimates of coefficients on some days lie outside the confidence band on other days. On the other hand, the coefficients on dummy variables for Unity day and Mondays in Panels (c) and (d) do not change much throughout the period. These plots clearly illustrate variation in coefficient estimates over time, as well as the relevance of restricting some other coefficients to be constant.

2.5.2 Price levels

Forecast accuracies are presented in a few different ways. First, plots of forecasts from each of the four models and the original observed prices are plotted in Figure 2.7. Days on which prices were filtered are also indicated with a red cross. On those days, errors are measured as the discrepancy between forecast and observed values, and are expected to be relatively large since those extreme values do not form part of the estimation sample.

There is little to distinguish between the four models when forecasting a single day into the future. When making 7-day-ahead forecasts, Model 1, which uses constant parameters, tends to perform rather poorly for most of the out-of-sample period. forecasts from Model 1 are greater than the observed prices as well as the forecasts from the other models until around October 2016. Similar comparisons can be made between the models with the 30-day-ahead forecast horizon in Figure 2.7(c).

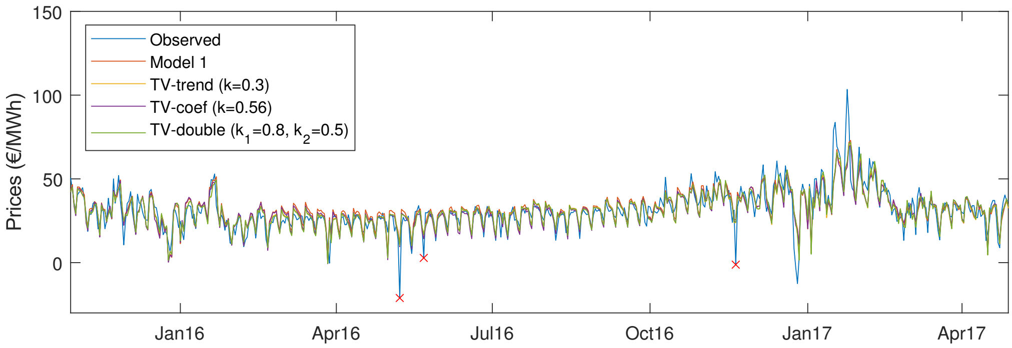

(a) 1-day-ahead

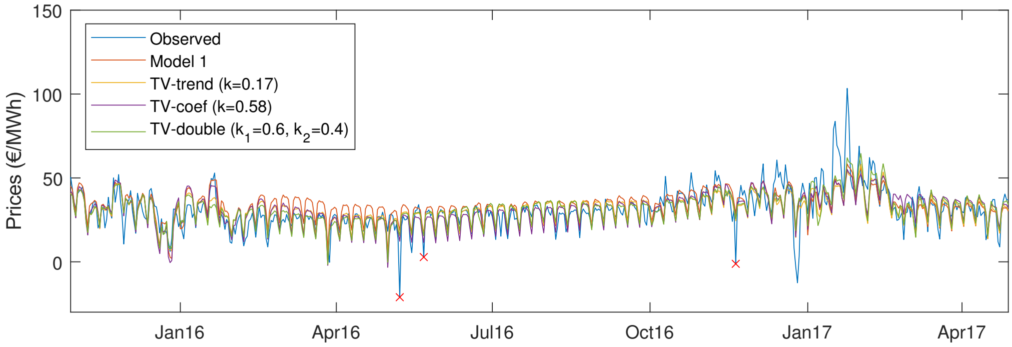

(b) 7-day-ahead

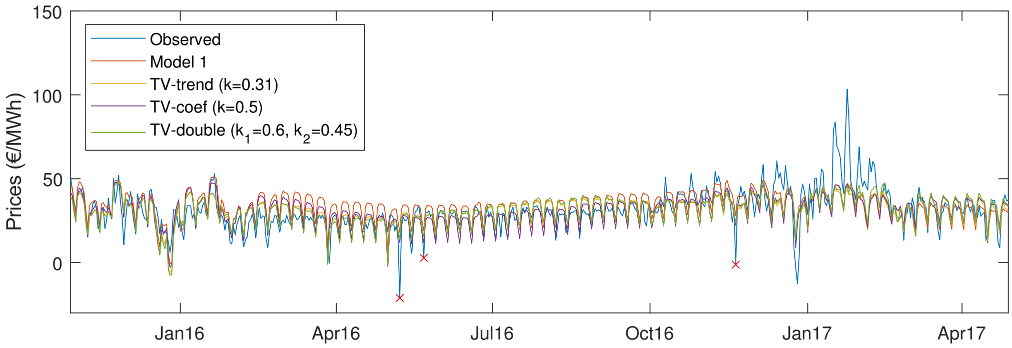

(c) 30-day-ahead

Figure 2.7: Comparison between observed prices and prices forecasted using the four models from 29th October 2015 to 28th April 2017. The red crosses indicate the prices which have been filtered out for the estimation process.

One thing which is quite clear from these plots, and is more obvious the longer the forecast horizon, is that the constant parameter model produces forecasts which track the average. Take, for example, the forecasts in Figure 2.7(b). Between April 2016 and October 2016, Model 1 forecasts are higher than all others. They are, in fact, closer to the average levels of prices in the period prior to February 2016, which forms a considerable part of the sample on which coefficients are estimated. Either by a turn of luck or by design, observed prices after March 2017 return to levels which are ever-so-slightly higher than those of the April–October 2016 period, and are closer to those of the pre-February 2016 portion of the sample, so Model 1 becomes reasonably accurate again. On the other hand, even though the forecast errors of Models TV-trend, TV-coef, and TV-double are generally in the same direction as those of Model 1, the magnitude of the error is smaller most of the time. To show this, two ratios are calculated for each of the models represented in Figure 2.7. The first is

$$\frac{\sum_{s=1}^{t} \hat{e}^2_{s|1:s-q}}{\sum_{s=1}^{t} \breve{e}^2_{s|1:s-q}}, \tag{2.12}$$

where $\hat{e}_{s|1:s-q}$ is the difference between the price forecasted at time s for horizon q using one of the three semiparametric models and the observed price. $\breve{e}_{s|1:s-q}$ is the corresponding forecast error from Model 1. Ratios smaller than 1 imply that, on average, Model 1 produces larger (squared) errors than the semiparametric competitors. The second ratio is

$$\frac{\sum_{s=1}^{t} \mathbb{1}\left(\hat{e}^2_{s|1:s-q} < \breve{e}^2_{s|1:s-q}\right)}{\sum_{s=1}^{t} \mathbb{1}\left(\hat{e}^2_{s|1:s-q} > \breve{e}^2_{s|1:s-q}\right)}, \tag{2.13}$$

where $\mathbb{1}(\cdot)$ is a function which takes the value 1 when its argument is true and 0 otherwise. In other words, (2.13) presents the ratio of the number of times a candidate semiparametric model has smaller squared error than Model 1 over the number of times the converse is true. A value larger than 1 implies that forecasts from the semiparametric approach in consideration are closer to true values more often than Model 1 forecasts. Calculated values from (2.12) and (2.13) are presented in Table 2.2. All ratios from (2.12) are less than 1, and all ratios from (2.13) are greater than 1.

Table 2.2: Ratio between the sum of squared forecast errors of the proposed semiparametric models and Model 1. Bandwidths for the semiparametric models are those represented in Figure 2.7.

Ratio from equation	(2.12)			(2.13)		
Forecast horizon, $h =$	1	7	30	1	7	30
TV-trend	0.9768	0.8851	0.8531	1.1490	1.0679	1.1158
TV-coef	0.9918	0.9221	0.8984	1.4356	1.4685	1.4356
TV-double	0.9709	0.8789	0.8627	1.4248	1.5607	1.6863

Overall, these results suggest that the semiparametric models do indeed adapt to changes better than the constant-parameter model.

Interestingly, even though actual temperatures were used instead of predicted values, the 7-day-ahead and 30-day-ahead forecasts do not seem to be able to accurately predict the increase in prices in early-2017. This suggests two things. First, there is possibly at least one omitted exogenous variable which may capture the unusual prices in January–February 2017, whose effect was transmitted through lagged prices in the 1-day-ahead forecast. Perhaps the duration of abnormal cold days in a similar vein to the cooling degree days (CDD) variable (Sailor and Munõz, 1997; Amato et al., 2005) might provide more explanatory power. Second, the persistence in prices, captured by the inclusion of lagged values, plays an important role in forecasting prices in that period. Note that the main difference between a 1-day-ahead forecast and a 7-day-ahead forecast in our analysis is simply that actual prices are used to forecast prices one day into the future, whereas we use six days' worth of forecasts based on our model estimates in order to forecast prices on the seventh day into the future.

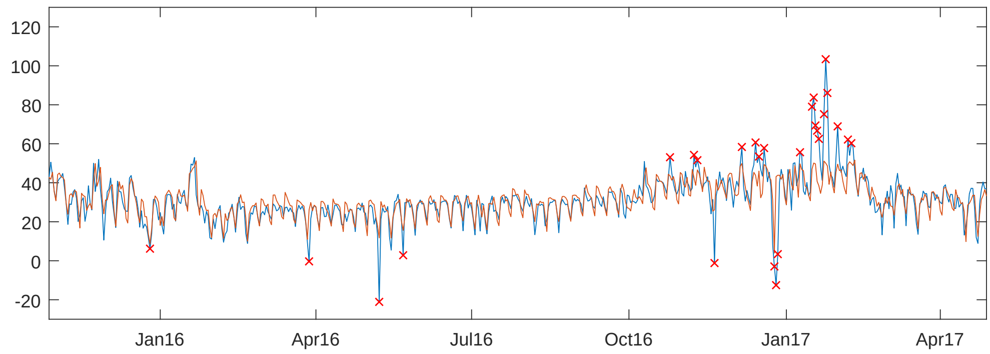

Figure 2.8: Comparison between observed prices and 1-day-ahead prices forecasts using Model 1 (constant parameters) from 29th October 2015 to 28th April 2017. The red crosses indicate the prices which have been filtered out for the estimation. Preprocessing in this series was conducted using a levels instead of first differences.

As a brief aside, an additional benefit of the proposed spike-filtering algorithm is highlighted. Figure 2.8 shows 1-day-ahead forecasts using Model 1 with the same inputs. The only difference is that instead of the algorithm described in Section 2.2.2, these prices are filtered based on their levels compared to a threshold, which is defined as three standard deviations from the mean of a neighbourhood around each point. In other words, the only criteria for spike identification is

$$\left| p_t - \frac{\sum_{s \in \mathcal{N}(t)} p_s}{\#\left(\mathcal{N}(t)\right)} \right| > 3\sigma_{\mathcal{N}(t)},$$

where $\#(\cdot)$ is the cardinality of its argument and $\sigma_{\mathcal{N}(t)}$ is the standard deviation of all prices in $\mathcal{N}(t)$. Note that a large number of prices in January 2017 were identified as spikes.

When data has been filtered, forecasts are expected to be muted or less volatile compared to the actual observation, as is the intention of the filtering process; we only want to model and forecast the predictable base prices, after all. However, since the 1-day-ahead forecasts in Figure 2.7(a) are able to track the prices in January 2017 with reasonable accuracy except for on the two days with the largest increases, those observed prices should not be identified as spikes, and therefore should not be excluded from the data set as in Figure 2.8. This example clearly illustrates the importance of a good filtering method; it is imperative that essential information in the data be preserved.

2.5.3 Cumulative mean-square forecast errors

Figure 2.9 plots the cumulative mean squared forecast error (MSFE) of each of the four models. Lower lines indicate smaller (cumulative) MSFE, and higher forecasting accuracy. For each $t \in [1, 2, \ldots, 548]$ and each forecast horizon q, the cumulative MSFE is defined as

$$cumMSFE(t) = cumMSFE(t; q) = \frac{\sum_{s=1}^{t} \hat{e}_{s|1:s-q}^2}{548}.$$

As noted when comparing the plots in Figure 2.7, the performance of the four base price models in Figure 2.9(a) is virtually indistinguishable when forecasting a single day into the future. However, the longer the forecast horizon, the larger the overall improvement in forecasting accuracy achieved by time-varying models relative to the constant-parameter model, particularly in the lead-up to October 2016. Note how the distance between cumMSFE for Model 1 and the other models increases between April 2016 to October 2016 in Figure 8(c), suggesting that in this period, the time-varying forecasts are consistently more accurate than those of Model 1.

One way to interpret the results of these comparisons is to consider the MSFE or cumulative MSFE to be a loss function. This loss function is symmetric in the sense that over-prediction and under-prediction are weighted equally. With this, the comparison is simply that the one with the lowest MSFE is best, and the bigger the difference between the MSFE between two methods, the greater the magnitude of improvement which stands to be gained. In this case, these results show that time-varying models can produce the largest gains when the forecasting horizon is longer, and are no less accurate than Model 1 even at the 1-day-ahead horizon.

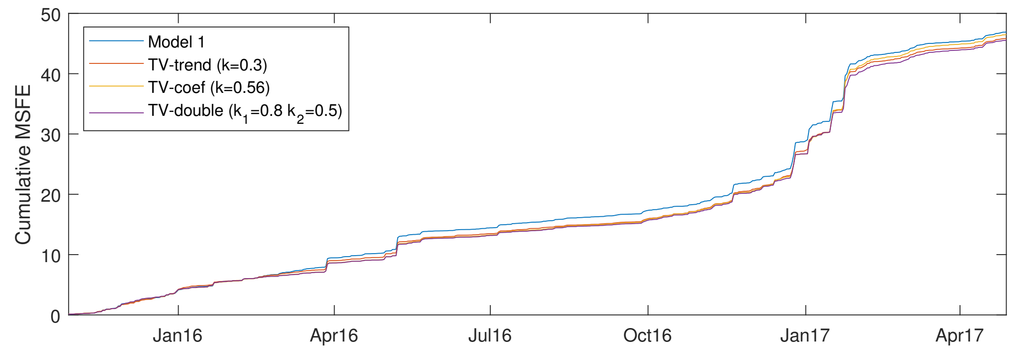

(a) 1-day-ahead

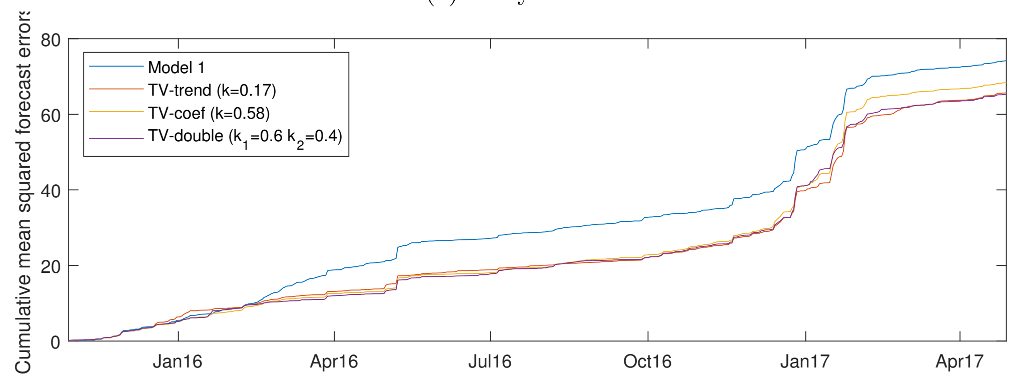

(b) 7-day-ahead

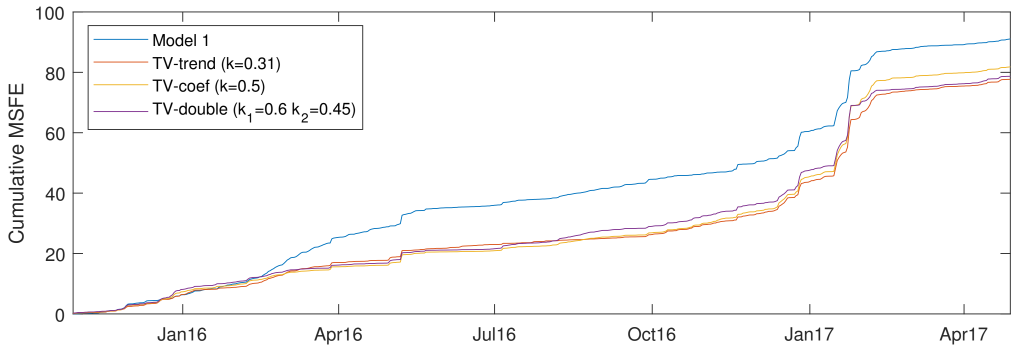

(c) 30-day-ahead

Figure 2.9: Cumulative MSFEs for all six models for the period 29th October 2015 to 28th April 2017.

2.5.4 Comparing predictive accuracy

The forecasting performance of each semiparametric model is compared against the parametric benchmark using Diebold-Mariano tests Diebold and Mariano (1995). Further, a model confidence set (MCS) as defined by Hansen et al. (2011) is constructed for each forecast horizon, identifying the set of "best" models. In this case, the candidate models are Model 1 (the parametric benchmark) and semiparametric models (TV-trend, TV-coef, and TV-double), each with a range of bandwidths.

While Figures 2.7 and 2.9 give some insight into the relative performance of the four models, it

Table 2.3: Diebold-Mariano test statistics (DM-stat) for comparing forecast accuracy between Models TV-trend and TV-coef to that of Model 1. For different forecast horizons and for each model, different ranges of bandwidths, $h = k\tilde{T}^{-1/5}$, are selected.

Horizon:	1-day-ahead				7-day-ahead				30-day-ahead			
Model:	TV-trend		TV-coef		TV-trend		TV-coef		TV-trend		TV-coef	
	k	DM	k	DM	k	DM	k	DM	k	DM	k	DM
	0.12	1.30	0.50	0.37	0.09	1.46	0.38	0.92	0.19	1.26	0.30	0.48
	0.14	1.35	0.52	0.40	0.11	1.64	0.40	1.04	0.21	1.31	0.32	0.59
	0.16	1.37	0.54	0.42	0.13	1.75	0.42	1.14	0.23	1.34	0.34	0.68
	0.18	1.37	0.56	0.42	0.15	1.80	0.44	1.22	0.25	1.36	0.36	0.74
	0.20	1.34	0.58	0.41	0.17	1.82	0.46	1.27	0.27	1.37	0.38	0.79
	0.22	1.31	0.60	0.40	0.19	1.82	0.48	1.29	0.29	1.37	0.40	0.72
	0.24	1.27	0.62	0.40	0.21	1.90	0.50	1.30	0.31	1.36	0.42	0.73
	0.26	1.23	0.64	0.39	0.23	1.77	0.52	1.29	0.33	1.34	0.44	0.72
	0.28	1.19	0.66	0.37	0.25	1.74	0.54	1.27	0.35	1.32	0.46	0.81
	0.30	1.15	0.68	0.36	0.27	1.70	0.56	1.24	0.37	1.30	0.48	0.79
	0.32	1.11	0.70	0.34	0.29	1.66	0.58	1.21	0.39	1.27	0.50	0.76

DM-stats greater than critical values indicate that the semiparametric forecast is significantly better. DM-stat$\sim N(0,1)$. The 10%, 5%, and 1% right-tailed critical values are 1.2816, 1.6449, and 2.3263, respectively. Statistical significance is indicated using cells shaded in magenta (10% level), yellow (5% level), and cyan (1% level).

Table 2.4: Diebold-Mariano test statistics (DM-stat) for comparing forecast accuracy between Models TV-double to that of Model 1.
For each forecast horizon, different ranges of bandwidths multipliers, k_1 and k_2, are selected for $h_i = k_i\tilde{T}^{-1/5}$, $i = 1, 2$.

Panel A: 1-day-ahead

k_1 \ k_2	0.40	0.50	0.60	0.70	0.80
0.50	0.59	1.07	0.98	0.78	0.65
0.60	0.70	1.32	1.39	1.31	1.26
0.70	0.73	1.39	1.53	1.51	1.48
0.80	0.73	1.39	1.53	1.52	1.49
0.90	0.71	1.35	1.48	1.45	1.40

Panel B: 7-day-ahead

k_1 \ k_2	0.20	0.40	0.60	0.80	1.00
0.30	-1.13	0.89	-0.37	-1.27	-1.54
0.40	-1.14	1.50	1.42	0.72	0.48
0.50	-1.15	1.74	2.08	1.71	1.60
0.60	-1.15	1.82	2.44	2.25	2.19
0.70	-1.15	1.83	2.59	2.46	2.39

Panel C: 30-day-ahead

k_1 \ k_2	0.30	0.45	0.60	0.75	0.90
0.40	0.60	0.85	0.76	0.68	0.71
0.50	0.71	1.00	0.91	0.86	0.91
0.60	0.73	1.11	1.07	1.05	1.11
0.70	0.71	1.13	1.12	1.10	1.15
0.80	0.66	1.11	1.08	1.03	1.05

DM-stats greater than critical values indicate that the semiparametric forecast is significantly better. DM-stat$\sim N(0,1)$. The 10%, 5%, and 1% right-tailed critical values are 1.2816, 1.6449, and 2.3263, respectively. Statistical significance is indicated using cells shaded in magenta (10% level), yellow (5% level), and cyan (1% level).

is important to also conduct statistical tests of the relative accuracies of Model 1 and each of the semiparametric alternatives. This comparison can be performed using the Diebold-Mariano test (Diebold and Mariano, 1995).

From Table 2.3, Model TV-trend outperforms Model 1 at the 10% level of significance for all three forecast horizons for a number of bandwidths, and at the 5% level of significance at the 7-day-ahead horizon. On the other hand, Model TV-coef does not appear to provide forecasts which are any different in accuracy from Model 1, in a statistically-significant sense. Another way to conceptualise Models TV-trend and TV-coef is to think of them both as Model TV-double, but in one case $h_2 = \infty$, and in the other $h_2 = h_1$. DM-stats in Table 2.4 are for cases where the restrictions on h_2 are relaxed. Remarkably, there appears to be some improvement in forecasting accuracy, in that the largest DM-stat is greater than those of Models TV-trend and TV-coef in Table 2.3 at the 1-day-ahead and 7-day-ahead forecast horizons. If one were to make a choice between the four model specifications in this paper based solely on the Diebold-Mariano test, then Models TV-trend and TV-double would be preferred. They are able to produce statistically-significant improvements in forecast accuracy at the 1-day-ahead and 7-day-ahead horizons, and Model TV-trend forecasts are also statistically more accurate at the 30-day-horizon.

In addition to the Diebold-Mariano tests for comparative predictive accuracy in Section 2.5.4, the model confidence set for each forecast horizon was also constructed[16]. Table 2.5 contains the models which are *excluded* from the MCS. The set of all candidate models (both included and excluded from the MCS) are those represented in Tables 2.3–2.4. These results suggest that Model 1 is always outperformed by some of the semiparametric models at the 1-day and 7-day forecast horizons. At the 30-day horizon, the MCS procedure is unable to exclude any model. However, it is clear that within the range of bandwidths selected for this study, some choice of k and (k_1, k_2) will lead to the TV-trend and TV-double approach being in the set of best models. This conclusion is in line with that of the Diebold-Mariano tests. On the other hand, when there is any significant difference in predictive ability, according to the MCS procedure, Model 1 is always excluded from the set of best models.

[16]The MCS procedure was executed using the MFE Toolbox in MATLAB by Kevin Sheppard, which is available at `https://www.kevinsheppard.com/code/matlab/mfe-toolbox/`.

Table 2.5: Models which are *excluded from* the confidence sets for each of the three forecast horizons. The confidence level is 0.9. The resulting MCS is robust to the stationary bootstrap with average window lengths indicated in the table. $B = 999$ bootstrap replications are performed.

Forecast horizon, $h =$	1	7	30
Model 1	Yes	Yes	No
TV-trend:	None	None	None
TV-coef	All	None	None
TV-double	$k_1 = 0.50$, $k_2 = 0.40$	$k_1 = 0.30$, $k_2 = 0.20$	None
	$k_1 = 0.60$, $k_2 = 0.40$	$k_1 = 0.40$,$k_2 = 0.20$	
	$k_1 = 0.50$, $k_2 = 0.50$	$k_1 = 0.50$,$k_2 = 0.20$	
	$k_1 = 0.50$, $k_2 = 0.60$	$k_1 = 0.60$,$k_2 = 0.20$	
	$k_1 = 0.50$, $k_2 = 0.70$	$k_1 = 0.70$,$k_2 = 0.20$	
	$k_1 = 0.50$, $k_2 = 0.80$	$k_1 = 0.30$,$k_2 = 0.40$	
		$k_1 = 0.30$,$k_2 = 0.60$	
		$k_1 = 0.40$,$k_2 = 0.60$	
		$k_1 = 0.30$,$k_2 = 0.80$	
		$k_1 = 0.40$,$k_2 = 0.80$	
		$k_1 = 0.30$,$k_2 = 1.00$	
		$k_1 = 0.40$,$k_2 = 1.00$	
Window lengths	52–58	30–40	10–59

An important point to discuss is the range of bandwidths reported in Tables 2.3–2.4. While there is some overlap in the bandwidths, it is clear that the suitable values of k for Models TV-trend and TV-coef are generally smaller than those of k_1 and k_2 Model TV-double. More interestingly, the ranges of values of k_1 and k_2 in Table 2.4 for which DM-stats are similar or better than the statistics in Table 2.3 are noticeably greater than the ranges of k in Table 2.3. This suggests that Model TV-double is more robust to deviations in bandwidths from the optimal value.

Note that despite the fact that the Diebold-Mariano test does not seem to recommend the use of Model TV-double at the 30-day horizon, these results do not necessarily contradict those of Section 2.5.3. The variance of the null distribution of the DM test statistic grows with the forecast horizon. Therefore even with larger magnitudes of improvement in prediction accuracy, the Diebold-Mariano test may have lower statistical significance, as seen here.

2.5.5 A note on bandwidth selection

Many other studies which use kernel estimation methods will employ some form of bandwidth selection process, be it the rule-of-thumb formula or cross-validation. These methods are less useful for dependent data. Since the purpose of this study is forecasting, a target function for optimisation can be based on some measure of in-sample forecasting inaccuracy. To that end, the sample is split in two. The final date in the earlier subsample is defined as the "current" date. Each of the semiparametric models is used to produce in-sample forecasts until the "current" day. The best bandwidth from this in-sample forecasting exercise is chosen as that which gives the smallest in-sample MSFE. Then, for robustness, some other bandwidths around the selected "best" value are also chosen. Out-of-sample forecasts are performed using this range of bandwidths. Naturally, the best bandwidth for the out-of-sample period may not necessarily coincide with that of the earlier period.

In this bandwidth-selection exercise, a general relationship between bandwidth and MSFE was noticed. As k (or h) approaches zero, the MSFE increases sharply. This result is to be expected, since an infinitesimal bandwidth would mean that very little information from the sample is incorporated in the estimation. On the other hand, as k approaches infinity, the MSFE approaches that of Model 1. This, too, is expected since an infinitely-large bandwidth means that equal weights are placed on all observations in the sample, which is no different from ordinary least squares estimation. Somewhere in between $k = 0$ and $k = \infty$ lies at least one local minimum for MSFE. For the wide range of bandwidths we considered, not all of which are reported in this paper, we find that the smallest local minimum is the first local minimum as k moves away from zero. Thus, should a search algorithm be applied in order to find the optimal bandwidth, a reasonable starting point would be to initialise the algorithm at a small bandwidth value. However, the definition of "small" is, of course, subjective. Ultimately, the choice of bandwidth and the frequency at which to update the chosen bandwidth are dependent on the econometrician and data set.

2.6 Discussion

The literature on electricity price forecasting has a long history. In light of changes to electricity markets, regulatory or otherwise, certain innovations to forecasting methods may be in order.

In this chapter, two contributions are made to the literature. First, a different treatment for the temperature variable is presented, which aims to capture the effect of deviations from seasonal temperature patterns, and also permits the effects of these deviations to depend on the season. The second and main contribution is in presenting a relatively simple model for forecasting daily electricity spot prices with some extensions, namely, allowing for time-varying parameters. The relative forecasting performance between the base model and its extensions was examined in Section 2.5. Throughout these comparisons, a few points of discussion spring to mind.

First, from visual inspection of Figures 2.7 and 2.9, it would seem that there may be economic gains to be had from using the semiparametric models over the constant-parameter base model, particularly at longer forecasting horizons. However the DM-test reveals that the forecast horizon most suited to the semiparametric models is, in fact, the shortest one Ultimately, the choice of modelling approach for some forecast horizon would depend on the econometrician's objective. What is empirically true from the work in this paper, however, is that for each semiparametric model at each of the three forecast horizons, there exists a range of bandwidth parameters such that the use of Models TV-trend, TV-coef, or TV-double will not lead to statistically poorer forecasts than those of Model 1.

It is also worth stating the fact that findings may differ for different sample periods. For instance, if the sample were to end on 1st October 2016 instead of 28th April 2017, large and clear gains would be found, both economically and statistically, from using any of the three semiparametric models over the constant-parameter version.

Chapter 3

Point forecast of prices using functional data analysis

In theory, the price of an item is determined as the intersection of supply and demand curves. This relationship is not always directly observable[17]. Wholesale electricity market data is almost unique in that prices are, in actuality, determined by the intersection of the two curves. German market data from the EEX contains disaggregated information about the sale and purchase bids made by all market participants. Using this data, the sale and purchase curves can be constructed. The plots in Figure 1.3 are of sale and purchase curves in the German wholesale market at the first hour of 1 January 2016. Every hourly auction in the wholesale market will have its own corresponding sale and purchase curves arising from the bids of market participants. The traded volume (MWh) and price (€/MWh) occurs at the intersection of the two curves. Instead of forecasting just the series of prices, functional data analysis techniques permit forecasting of supply curves and demand curves. Then, forecasts of prices can be obtained as the intersection of the two curves at any point in time. The extant literature using this approach is new and small, but promising.

[17] For example, the price of a tin of tuna in shops is subject to marketing decisions based on some market research into the availability of supply and the willingness of consumers to purchase at various prices. However, the exact supply and demand curves themselves cannot be drawn up.

3.1 Introduction

Ziel and Steinert (2016) were the first use sale and purchase curves to forecast wholesale electricity prices, although they did not employ functional data analysis methods. They create forecasts of volume bids at a selection of discrete points along each curve, then interpolate the values to construct forecasts of entire curves. Following that, Shah and Lisi (2020) treats each curve as a datum in itself, and forecasts entire curves using functional data analysis tools instead of as reconstructions from several points. In brief, functional data analysis (FDA) is the practice of analysing and modelling functional data, where each observation in the sample is viewed as a function over some set, as opposed to being a single point of datum. Comprehensive texts on functional data and its analysis include Ramsay and Silverman (2002, 2005), Horváth and Kokoszka (2012), and Kokoszka and Reimherr (2017). Section 3.2 outlines some of the FDA modelling approaches which have been widely used in practice.

FDA can be employed when there is sufficient reason to believe that the underlying data are continuous curves or can be approximated by such curves. Functional data can be found in various disciplines. A number of illustrative examples, ranging from handwriting samples to human growth, can be found in Ramsay and Silverman (2002). More recently, FDA has also been used to model or predict mortality rates (Hyndman and Ullah, 2007), biomechanics related to joint injuries (Hébert-Losier et al., 2015), disease characteristics (Kendrick et al., 2017), financial returns (Shang, 2017), electricity prices (Chen and Li, 2017; Gonzalez et al., 2018) and load (Goia et al., 2010), supply and demand curves (Lisi and Shah, 2020; Mestre et al., 2020), and many other types of data. Reviews and surveys related to functional data analysis in practice can be found in Rice (2004), Ullah and Finch (2013), Cuevas (2014), Shang (2014), and Wang et al. (2016).

The monograph by Bosq (2000) addresses the functional autoregressive (FAR) model, including relevant theory and estimation methods. At the risk of oversimplification, the FAR model may be viewed as a Hilbertian-space analogue to the scalar autoregressive (AR) model. While coefficients in the AR model are scalars, the FAR equivalents are continuous linear operators (or functionals). When certain conditions are met, the linear operators can be estimated by least squares (Bosq, 2000; Kokoszka and Reimherr, 2012). Functional time series models are also estimated by means of basis expansions (Hyndman and Ullah, 2007; Chiou, 2012; Wang et al., 2016; Shang, 2017).

The literature pertaining to functional time series estimation and forecasting of supply and demand curves is extremely small. In fact, only two papers to date use functional methods to forecast wholesale electricity market curves are Shah and Lisi (2020) and Mestre et al. (2020). Furthermore, of these two articles, only Shah and Lisi (2020) forecast both supply and demand curves in order to obtain the price at the intersection; Mestre et al. (2020) forecast only the sale curves. Despite the scarcity of immediately-relevant works, the theory, intuition, and techniques which are employed in functional forecasting of sale and purchase curves is the similar to that of any other application, for which there is a wealth of literature.

3.2 Functional time series in wholesale electricity markets

Functional models can be estimated using a variety of methods. Thus far, two approaches have been applied to modelling sale and purchase curves in wholesale electricity markets. Both methods are collectively referred to in this chapter using the term "kernel-type estimation". An alternative way of conducting FDA is through basis function decomposition, which has been widely used in other disciplines.

Kernel-type estimation

Shah and Lisi (2020) adopt the kernel-type estimator described in Ferraty et al. (2012), relaxing the linearity assumption on the operators. In essence, the procedure in Shah and Lisi (2020) applies larger weights to more recent data. Let $\xi_{t,h}$ represent the functional object at hour h on day t (for $t = 1, \ldots, T$ and $h = 1, \ldots, 24$) which, in this case, is either a supply or demand curve. Then the assumption underpinning this approach is that it follows the autoregressive process

$$\xi_{t,h} = m\left(\xi_h\right) + u_{t,h}, \tag{3.1}$$

where $m\left(\xi_h\right) \equiv m\left(\xi_{t-1,h}, \ldots, \xi_{t-d,h}\right) = \mathbb{E}\left[\xi_{t,h} | \mathcal{I}_{t-1}\right]$ is a mean function, with ξ_h and $\mathcal{I}_{t-1}$ referring to a vector of lagged functional variables and the available information up to and including time $t-1$, respectively. The maximum number of lags, d, is user-defined, and the term $u_{t,h}$ denotes a functional error. Out of the four econometric methods which were compared in Shah and Lisi (2020), the one which led to the best predictive performance most of the time was the nonparametric functional autoregressive (NPFAR) model. If information is available up to (and including) time $t-d-1$, the kernel regression estimator for NPFAR for the function $m(\cdot)$ evaluated at $\tilde{\xi}$ is given as

$$\hat{m}\left(\tilde{\xi}\right)=\frac{\sum_{t=d+1}^{T}\xi_{t,h}K_h\left[\lambda\left(\xi_{t-d,h},\tilde{\xi}\right)\right]}{\sum_{t=d+1}^{T}K_h\left[\lambda\left(\xi_{t-d,h},\tilde{\xi}\right)\right]}, \tag{3.2}$$

where $\xi_{t,h}(\cdot)$ is a functional random variable in some semimetric space (E,λ), $\tilde{\xi}$ is a given element of E, $K(\cdot)$ is a kernel function, h is the pre-determined bandwidth, and $K_h(a)=K(a/h)$. The lag, d, must be suitably chosen. The expression in (3.2) is, essentially, a functional version of the univariate local-constant Nadaraya-Watson estimator. In other words, $\hat{m}(\tilde{\xi})$ is the weighted sum of all observations in the sample, where data points closest (with distance measured in the time domain) to $\tilde{\xi}$ are given highest weighting.

Although Gonzalez et al. (2018) also use the "kernel" in referring to their operators, the term has a different meaning. Their kernel is a three-dimensional surface which permits every point in each functional explanatory variable (including lags of the dependent variable, lags of errors, and other exogenous variables) to influence the level at any point on the current curve. Instead of using a probability density kernel function, the bivariate "kernel" is represented by the sum of trigonometric functions, with parameters which are estimated. Mestre et al. (2020) further propose using a simpler version of the kernel surface, where the volume in a supply curve for some given price, say p_0, at time t is affected by the volumes of curves at other points in time, but only at the same price level, p_0. As a consequence, the kernel surface becomes two-dimensional or flat. A key difference between the methods proposed by Gonzalez et al. (2018) and Mestre et al. (2020) and the NPFAR approach of Shah and Lisi (2020) is that the former two also include other exogenous explanatory variables such as weather and temperature. Each of these exogenous variables, whether functional or scalar, will have a functional operator (kernel surface) associated with them, which can be viewed as functional versions of regression coefficients.

These methods were shown to lead to more accurate price forecasts that standard scalar (as opposed to functional) models. However, they are complex and somewhat complicated to implement[18].

[18] At time of writing, one of the authors of Gonzalez et al. (2018) had been contacted. An R program for estimating their model was under development but not published.

Basis function decomposition

The approach adopted in this chapter uses basis expansions to estimate a functional autoregressive model. The procedure will be discussed in order of the steps outlined in Hyndman and Ullah (2007). The discussion will include reasons that some steps are performed, how they are adopted in this chapter, and, where applicable, reasons why certain steps are not adopted for the analysis in this chapter.

Consider a data set with observations $\{p_i, v_t(p_i)\}$, $t = 1, \ldots, T$, $i = 1, \ldots, N$, where

$$v_t(p_i) = f_t(p_i) + \sigma_t(p_i)\varepsilon_{t,i}, \tag{3.3}$$

where $\varepsilon_{i,t}$ is an iid standard normal random variable and $\sigma_t(p_i)$ allows the amount of noise to vary with p. The general algorithm used by Hyndman and Ullah (2007) is as follows, with specific details provided in Section 3.3.

1. Smooth the data for each t using a nonparametric smoothing method to estimate $f_t(p)$ for $p \in [p_1, p_N]$ from $\{p_i, v_t(p_i)\}$, $i = 1, 2, \ldots, N$.

2. Decompose the fitted curves via a basis function expansion using the following model:

$$f_t(p) = \mu(p) + \sum_{k=1}^{K} \beta_{t,k}\phi_k(p) + e_t(p), \tag{3.4}$$

   where $\mu(p)$ is a measure of location of $f_t(p)$, $\{\phi_k(p)\}$ is a set of orthonormal basis functions and $e_t(p) \sim N(0, \psi(p))$.

3. Fit univariate time series models to each of the coefficients $\{\beta_{t,k}\}$, $k = 1, \ldots, K$.

4. Forecast the coefficients $\{\beta_{t,k}\}$, $k = 1, \ldots, K$, for $t = T+1, \ldots, T+h$ using the fitted time series models.

5. Use the forecast coefficients with (3.4) to obtain forecasts of $f_t(p)$, $t = T+1, \ldots, T+h$. From (3.3), forecasts of $f_t(p)$ are also forecasts of $v_t(p)$.

First, the data is smoothed. Rice and Silverman (1991) show that under the assumption that the mean function is smooth, estimates are more accurate when the raw data first undergoes a smoothing process. It is worth noting, however, that Cuevas (2014) and Gao et al. (2020) motivate this step as a method by which noise in data measurements is reduced. Furthermore, Bosq (2000) describes interpolation and smoothing as techniques which are "for converting discrete data into functional data". In the case of hourly electricity auction curves, sale and purchase curves are typically piece-wise constant, so are not naturally smooth. However, although the data is unevenly spaced on account of different bidding prices, these data points completely and exactly define the continuous curve over the entire domain of possible prices. As such, neither of these two reasons is applicable to the context of auction curves as functional data. Additionally, smoothing is not performed in recent works which use similar methods, such as Shang (2019).

The second step is to decompose the (fitted or smoothed) curves by basis function expansion. Here, the condition is that the set of basis functions are orthonormal. Functional principal components are often used as the basis functions as they capture the largest amount of variation in the data, by design. However, Ramsay and Silverman (2005) point out that there is no single prescribed basis which works well in all situations. Even so, functional principal components analysis (FPCA) has its charms. For example, the use of FPCA significantly reduces the number of arbitrary decisions which need to be made. The principal components are derived from observed data, and the order of the basis is a data-driven choice. It is worth noting, additionally, that in defining the basis functions for functional data, Ramsay and Silverman (2005, Chapter 8.2) and Horváth and Kokoszka (2012) do not require the mean or location, $\mu(p)$, to be isolated from the basis functions.

All subsequent steps in the algorithm are uncomplicated. In the third step, since the basis functions are orthogonal, the coefficients for each k are independent of each other and can therefore be modelled as univariate time series. However it may be useful to model the $\{\beta_{t,k}\}$, $k = 1, \ldots, K$, using AR models with exogenous regressors (ARX) if relevant regressors exist. Forecasting in the fourth and fifth steps are straightforward.

Beyond the algorithm

The algorithm above provides a method of producing forecasts for the functional data. However, it is important to then relate those forecasts with the data and setting at hand. Two additional considerations must be made, regardless of the chosen functional forecasting approach. The first is that sale and purchase curves are monotonic by construction. As such, the forecasted market curves are monotonically-constrained versions of the forecasts produced by the algorithm.

A predicted sale curve, $\{\hat{v}(p_i)\}$, is the result of minimising

$$\frac{1}{N}\sum_{i=1}^{N}\left(\tilde{v}(p_i)-\hat{v}(p_i)\right)^2$$

with respect to $\{\hat{v}(p_i)\}$, subject to constraints

$$0 \leq \hat{v}(p_j) \leq \hat{v}(p_{j+1}) \text{ for } j = 1, \ldots, N-1, \tag{3.5}$$

where $\tilde{v}_s(p)$ denotes the prediction obtained from the forecasting algorithm. For a forecast of a (monotonically increasing) purchase curve, the inequality constraints in (3.5) are inverted.

The second additional step which is required relates to the fact that the quantity of interest in this study is prices. Forecasts of prices are obtained from the intersection of forecasted curves. however, the predicted sale and purchase curves are defined on a discrete grid. The approximation of the intersection price, $\tilde{p}$ is simply determined to be the price which minimises the distance between the forecasted supply and demand curves,

$$\tilde{p} = \arg\min\left|\hat{v}^s(p) - \hat{v}^d(P)\right|,$$

where the superscript indicates a supply (s) or demand/purchase (d) curve.

Functional principal components analysis

The idea behind principal components (PC) in functional data is the same as in the multivariate case. The theory differs mainly in that in the functional case, the data, and therefore the PCs, are continuous. Following Ramsay and Silverman (2005, Chapter 8.4), a simple way to compute functional PCs is by discretising the continuous functions to a grid of equally-spaced values. Then, the PCs are computed using the discretised data exactly as they would have been in the multivariate setting. In brief, let $\boldsymbol{\mathcal{V}}$ be the $N \times T$ matrix consisting of the T functions discretised with N equally-spaced values, with sample variance-covariance matrix $\hat{\boldsymbol{\Sigma}} = N^{-1}\boldsymbol{\mathcal{V}}'\boldsymbol{\mathcal{V}}$. Then the (discretised) functional PCs are the eigenvectors of $\hat{\boldsymbol{\Sigma}}$ or, equivalently, the $\boldsymbol{U}$ matrix in the singular value decomposition $\boldsymbol{UDW'}$ of $\boldsymbol{\mathcal{V}}$. If continuous functional PCs were required, the discretised PCs could simply be interpolated. Of the four approaches to selecting the number of

PCs mentioned by Shang (2017), the easiest is by comparing the ratio of the eigenvalues (Chiou, 2012). The other three methods alluded to were through pseudo AIC or BIC (Yao et al., 2005), cross-validation (Rice and Silverman, 1991), and bootstrapping (Hall and Vial, 2006).

An alternative to using eigendecomposition or singular value decomposition is by basis function expansion of the functions. This can be viewed as a more general version of FPCA, since PCs satisfy the requirements to be basis functions. As mentioned in the previous section in reference to Ramsay and Silverman (2005), there is no rule stating that FPCA yields the best basis, but finding a specific set of basis functions which yield better performance is not always an easy task.

3.3 Empirical application

3.3.1 Data

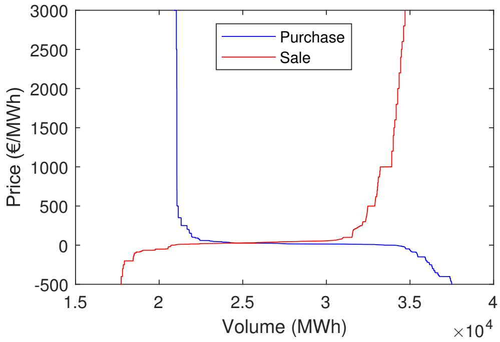

(a) Market curves at 00:00 1 February 2015 for the full range of prices

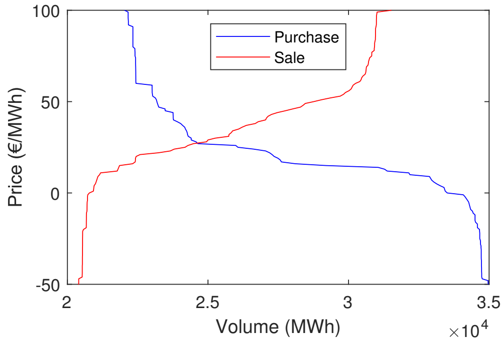

(b) Market curves at 00:00 1 February 2015 between -€50/MWh and €100/MWh

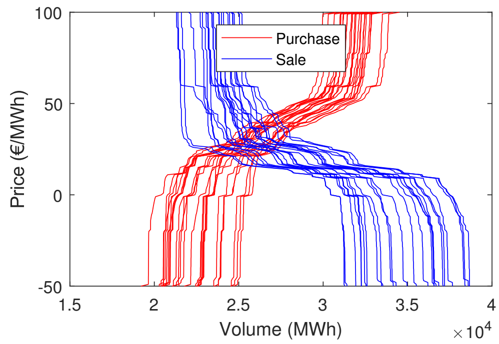

(c) Market curves for each of the 24 hours on 1 February 2015 between -€50/MWh and €100/MWh

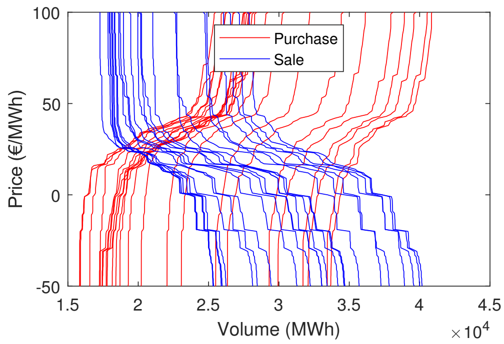

(d) Market curves for each of the 24 hours on 1 August 2015 between -€50/MWh and €100/MWh

Figure 3.1: Representation of sale and purchase curves in Germany.

German wholesale electricity market data for the period 1 February 2015–30 April 2016 (455 days) are obtained from the European Electricity Exchange (EEX). For this study, the various purchase and sale bids are aggregated to produce 10,920 (24×455) hourly sale and purchase curves. Observed sale and purchase curves from 1 February 2015 are shown in Figure 3.1. Panel (a) contains plots of the curves for the first hour, with Panel (b) focusing on the portion between the prices -€50/MWh and €100/MWh. To see what a small collection of these curves would look like, hourly purchase and sale curves for 1st February 2015 and 1st August 2015 are presented in Panels (c) and (d), respectively. As discussed, purchase (sale) curves are monotonically decreasing (increasing), The price at which the curves cross is the equilibrium price. For the first hour of 1 Feb 2015, this price is €27.04/MWh.

In the empirical application which follows, only the portion of the curves between the prices of -€100/MWh and €150/MWh are considered. The reason for this restriction is that all the prices in the sample lie between these bounds, and a smaller range of prices means that the characteristics of the curves within this region will be better captured in the modelling stage.

3.3.2 Forecasting

The basis function decomposition approach outlined in Section 3.2 was adopted with the described changes, namely, smoothing was not performed and the locality measure term, $\mu(p)$, in (3.4) is omitted from the model. Functional principal components, as described in Section 3.2, were used as the basis functions. For the purpose of forecasting, an in-sample period of 28 days (672 hours) was used for fitting model coefficients. This also means that the principal components were updated every day using the most current 28 days' worth of hourly data. Based on scree plots each day, the number of principal components included was consistently chosen to be $K = 2$ throughout the sample. Examples of principal components used in this analysis and the associate coefficient estimates are shown in Figure 3.2. The plots in Figures 3.2(c)–(d) clearly demonstrate a diurnal pattern in coefficients of the basis function decomposition.

Forecasts are performed in daily (24-hourly) blocks. In other words, fitting a forecasting model for any hour of a given day must be performed with data up to and including the 24-th hour of the preceding day only. Consequently, a two-step ahead (second hour of the day) forecast involves re-estimating the univariate model from Step 2 on page 49 using the forecasted coefficients in the basis function decomposition from the one-step ahead forecast. The process is then iterated until the forecast for the 24th hour. Naturally, any other random explanatory variables used

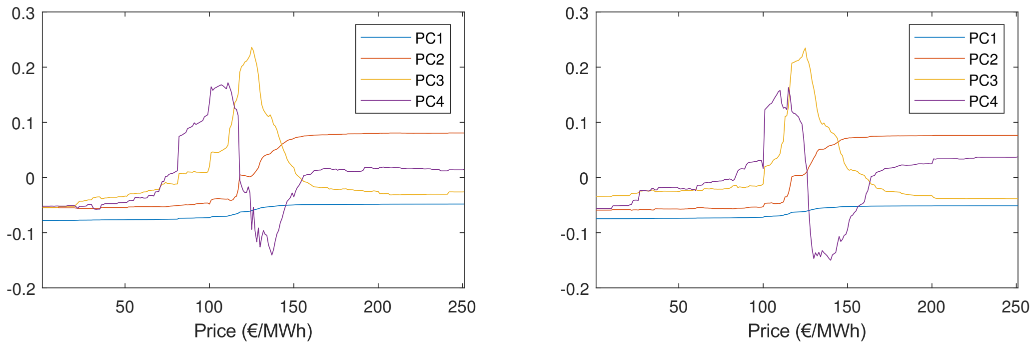

(a) Principal components for subsample ending 1 February 2015

(b) Principal components for subsample ending 4 March 2015

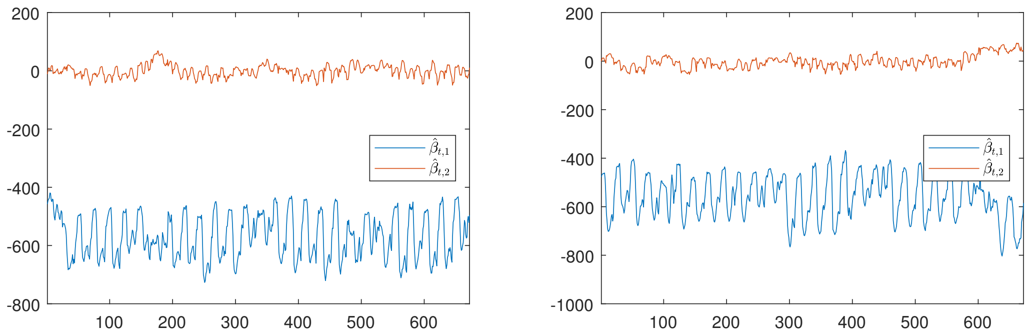

(c) Coefficient estimates for subsample ending 1 February 2015

(d) Coefficient estimates for subsample ending 4 March 2015

Figure 3.2: First four principle components used to decompose the demand curves on subsamples of 28 days or 672 hours (Panels (a)–(b)). Only the first two principle components were used. Time series of coefficient estimates are shown in Panels (c)–(d).

in Step 3 should also be forecasted. In this chapter, the univariate model in Step 3 is an ARX model

$$\hat{\beta}^a_{k,t} = \alpha_k + \delta_k \hat{\beta}^a_{k,t-1} + \boldsymbol{z}'_t \boldsymbol{\gamma} + \eta_{k,t}, \tag{3.6}$$

for $a \in \{s, p\}$ to indicate whether the model is for a sale (s) or purchase (p) curve. The control variables on the right-hand side, $\boldsymbol{z}'_t$, consist of hour-of-day dummies, day-of-week dummies, daily mean temperature, and volume of renewable energy generated at time t. Only the first lag, $\hat{\beta}_{k,t-1}$, was chosen as it led to the lowest forecast error for each $\hat{\beta}^a_{k,t}$.

For the sake of comparison, a univariate model is also used to forecast hourly prices. The benchmark model is

$$p_t = \alpha_k + \sum_{b \in \mathcal{B}} \delta_b p_{t-b} + \boldsymbol{z}'_t \boldsymbol{\gamma} + \varepsilon_t, \tag{3.7}$$

where p_t is the price at time t, the lags are $\mathcal{B} = \{1, 2, 3, 4, 24\}$, and exogenous variables z_t' are the same as in (3.6). The lags in $\mathcal{B}$ represent prices in the four preceding hours, as well as the price at the same hour on the previous day. The benchmark model (3.7) is based on existing established forecasting models, adapted to best suit the data at hand.

In both (3.6) and (3.7), lags are calibrated based on a training sample. Forecasts of prices were made using a range of different lag combinations for the sub-sample from 1 March 2015–31 July 2015[19]. The set of lags which minimised the sum of squared forecast errors $\sum(p_t - \tilde{p}_t)^2$ in this training sample is then chosen for the entire pseudo out-of-sample period of 1 August 2015–30 April 2016.

Once price forecasts from each method are obtained, their accuracy can be compared using the Diebold-Mariano test (Diebold and Mariano, 1995).

3.3.3 Results

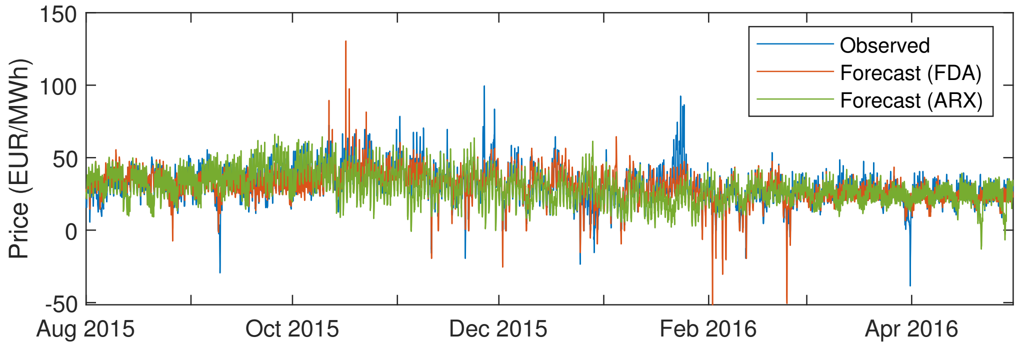

(a) Observed and forecasted prices

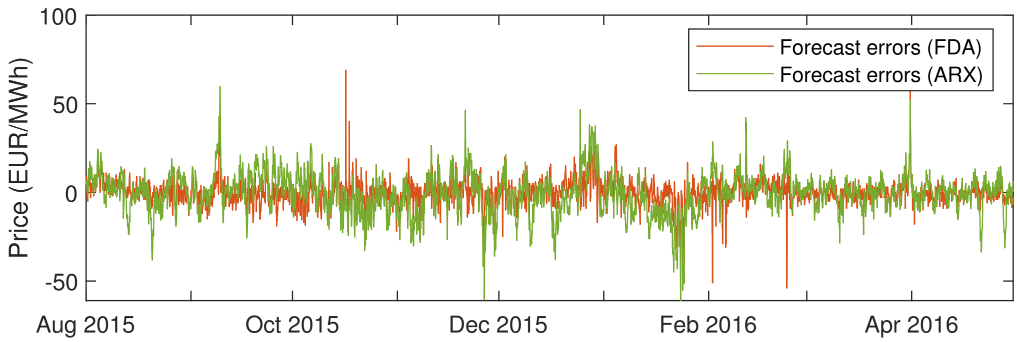

(b) Forecast errors prices

Figure 3.3: Hourly wholesale electricity price forecasts and forecast errors from 1st August 2015 to 30th April 2016. Forecasts are performed using either functional data analysis or an ARX model.

[19]Recall that model coefficients are fitted for subsamples of 28 days, so the first day with a forecast is 1 March 2015.

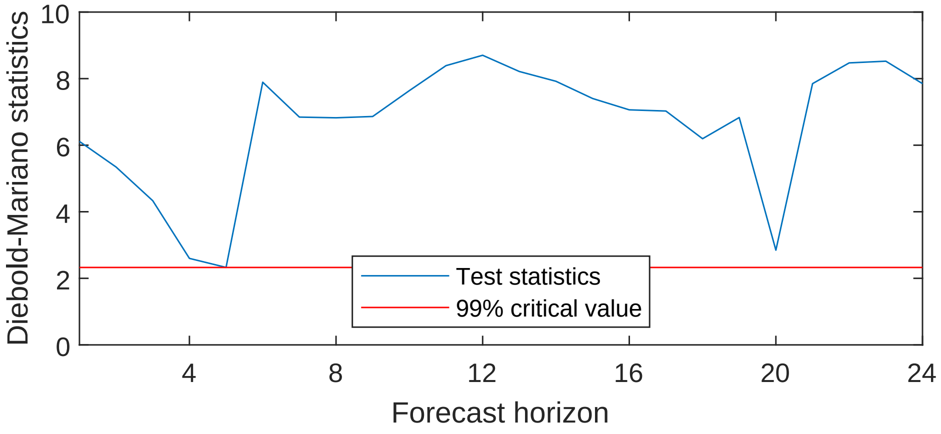

Figure 3.4: Diebold-Mariano test statistics for comparing predictive accuracy between the FDA and ARX approaches. The 99% critical value for a one-tailed test is 2.32.

Forecasts from the FDA approach and the ARX method are shown in Figure 3.3(a). In Figure 3.3(b), it is clear that FDA forecasts are closer to observed values, with errors that generally have smaller and more constant variance than the ARX forecast errors. The corresponding Diebold-Mariano test statistics for predictive accuracy are shown in Figure 3.4. Note that there are 24 statistics, each corresponding to a different hour of the day. The reason for the 24 Diebold-Mariano tests is that forecasts are conducted in blocks of 24 hours (one day), and there are therefore 24 forecasts horizons. Despite some erroneous spike predictions produced by the FDA approach (see Figure 3.3(a)), all of the Diebold-Mariano statistics are greater than the one-tailed 99% critical value, indicating that the FDA forecasts are statistically more accurate than the ARX forecasts at the 1% level of significance.

Note that although (3.7) is a very simple model, it is largely similar to the benchmark model used in Shah and Lisi (2020). While it is certainly possible that a more complex specification for an ARX—or some other univariate forecasting approach—may yield better results than (3.7), the fact is that even this simple specification for the basis function FDA method, (3.6), performs well.

3.4 Discussion and conclusion

Forecasting sale and purchase curves in order to obtain price forecasts is relatively new, the application of functional data analysis for this purpose even less explored. From the limited literature, this approach appears to have great potential for increasing the accuracy of wholesale price forecasts. The only other study to use FDA for this purpose is set in Italy, whereas the data studied in this chapter is from Germany. The findings suggest that improvements gained from adopting an FDA forecasting method is not particular to a single data set. Potential extensions and future research in this area could be from a methodological perspective, where different FDA techniques might be considered. Although not the focus of this chapter, it is worth noting that supply and demand curves contain other information which may be useful in making energy markets more efficient. Investigation into other uses of market curves may be an avenue for research as well. Other explanatory variables which may reasonably be assumed to contribute to changes in levels of shapes of supply and demand curves could also be explored.

Certain limitations and benefits of the model employed in this study warrant a discussion. The basis function decomposition method advocated for in this chapter is described as being easier to estimate than the kernel estimation alternatives. Although the kernel approach was not conducted here, Shah and Lisi (2020) emphasise that it takes only 0.5 seconds to conduct a one-step ahead prediction with their NPFAR model. In contrast, the basis function approach here takes, on average, under 0.1 seconds to forecast all 24 hours' prices for one day[20].

In addition to savings in computing complexity and time requirements, a benefit of the basis decomposition approach over the NPFAR method of Shah and Lisi (2020) is that it easily admits other explanatory variables which explain movements in wholesale market curves and prices. However, it must be noted that the variables used for forecasting in this study are observed, not forecasted. For some variables such as temperature, this practice is not likely to significantly affect the outcome since hourly day-ahead temperature forecasts are very accurate. However, this study also uses the volume of renewable energy to predict sale and purchase curves. Using observed values of this variable instead of forecasts may artificially inflate the predictive accuracy since forecasts of renewable energy in Germany are typically inaccurate[21]. At least, the comparison between FDA and the benchmark is made on equal grounds since the benchmark also uses observed as opposed to forecasted values.

In conclusion, using functional data analysis tools to forecasts supply and demand curves has been shown to be effective at producing more accurate forecasts of their intersections than merely forecasting the intersections themselves. However, this avenue is still very new to the literature.

[20]Shah and Lisi (2020) use an Intel Core i7-4510U CPU at 2.60GHz, whereas the computer used in this chapter contains an AMD Ryzen 7 4800H CPU at 2.90GHz.

[21]The forecasts from this chapter were repeated without renewables as explanatory variables. Diebold-Mariano statistics in Appendix A show that the FDA approach still outperforms the benchmark.

As such, there are various other research opportunities in related areas.

Chapter 4

Real-time price elasticity of wholesale electricity demand

Price elasticity of electricity demand is a topic which has received increasing attention over the years. A large and growing number of studies use econometric techniques to estimate the value of price elasticity, typically by regressing quantity on price and other variables, and using the coefficient of price. This chapter explores a number of ways in which such an approach might lead to inaccurate inference, and proposes a novel way of computing intra-day price elasticity of wholesale demand. The discussion and empirical results serve as a reminder that regression-based estimates must be used with care.

4.1 Introduction

Since the deregulation of electricity markets across the world, the study of various electricity-related quantities has become increasingly relevant and important. The fact that prices in many countries are no longer determined exogenously by a regulatory body but are a product of market forces means that a better understanding of the various quantities, be it price, load, usage, or elasticities, has the potential to lead to more efficient electricity markets. Any savings in the wholesale market could be passed on to end-users, resulting in higher welfare for the general public. Efficient bidding by wholesale market participants is one of the ways in which prices can be optimised. To this end, a growing body of literature focuses on empirically estimating the price elasticity of electricity demand (PEDE) in various energy markets around the world, although price elasticity of supply in electricity (PESE) seems to be less-studied.

Two of the key factors which drive advancements in empirical research are methodological innovations and data availability. Thus far, articles on price elasticities in electricity markets contribute to methodology. Broadly speaking, the accuracy of predictions or estimates produced by new estimation methods can be evaluated by comparing fitted values against their observed counterparts. Occasionally, such comparison is not possible because the quantity of interest is inferred but not observed. Until now, the empirical literature on price elasticity of demand in electricity has fallen under the latter category. Consequently, studies have introduced various ways in which price elasticity can be estimated from the available data without being able to truly verify the accuracy of the results.

Through access to highly-detailed data, this chapter is the first to measure true, real-time price elasticities in the wholesale market. The two key words here are "real-time", reflecting the fact that elasticities are measured for every traded time interval[22], and "true", indicating that these elasticities are computed from actual market supply and demand curves instead of being reverse-engineered from transaction price, quantity, and other data over a period of time. In other words, it is the first to provide a verifiably-good measure of observed price elasticities based on the most fundamental definition.

The focus on price elasticity of electricity demand, while relatively recent compared to studies on prices and loads, is certainly gaining momentum in recent years. Selected works in the growing body of related literature can be found tabulated in Lijesen (2007), Fan and Hyndman (2011), and Inglesi-Lotz (2011), often together with the geographic region, the reported values of price elasticities, as well as the methodology employed. Since then, of course, more studies have contributed to the literature on estimating price elasticity, including, but not limited to, Feehan (2018), Tiwari and Menegaki (2019), Dong et al. (2020), and Yu and Xin (2020). Each of these empirical studies computes price elasticities using some estimated quantity which represents some average over the entire sample. Typically, a single value for the coefficient of price in a regression model is estimated over the sample, and is then used to compute price elasticities. Two notable exceptions to this practice are Inglesi-Lotz (2011) and Tiwari and Menegaki (2019), who employ a state-space approach in their models, thereby allowing the relevant coefficient to be time-varying.

[22] The term "real-time price elasticity" has recently been used by Lijesen (2007). However, a term which might better-represent where their approach lies in relation to the rest of the literature is "very short term elasticity".

As indicated earlier in this section, the key factor differentiating this chapter from others is the data which is used. The existing literature examines various types of data, mainly focusing on residential, industrial, or other non-wholesale data. Additionally, most of the existing work uses lower frequencies such as daily- or annually-aggregated prices and volumes. On the other hand, this chapter focuses solely on the wholesale electricity market at the intra-day frequency. Most important, however, is that all bids in the intra-day market are examined here in their most raw, disaggregated form, which permits construction of exact supply and demand curves in each intra-day period. This is in stark contrast to other studies, which consider only the *traded* price and volumes, and are therefore unable to take into account the collective nature of all the bids made in intra-day electricity auctions. Access to raw data permits exact measurement of price elasticity of demand. The study in this chapter is the first to compute price elasticity of demand directly from raw hourly supply and demand curves.

There are a few studies which are close in spirit to this present work, albeit in different ways. Lijesen (2007) was the first to use intra-day data from the wholesale market to compute price elasticity of demand. However, the data used in that paper is the time series of prices and loads rather than sale and purchase curves. Consequently, the regression model in that study produces only a single value for elasticity over a period of one year, representing an average measure of intra-day real-time elasticities. Knaut and Paulus (2016) adopt a similar framework on hourly data, but use dummy variables to obtain twenty four coefficients on price, one for each hour of the day. They are therefore able to show an intra-day pattern for price-elasticity of electricity demand. However, these hourly price elasticities still represent average values—one for each intra-day interval—over the sample period.

To date, the only study to investigate price elasticity at the hourly frequency, as opposed to using hourly data to estimate price elasticity, is Kulakov and Ziel (2019). The method employed by Kulakov and Ziel (2019) is similar to the approach here in that demand elasticity is computed using slopes and values along the demand curve. However, they use supply and demand curves constructed from a decomposition of observed wholesale curves. Rather than studying price elasticity of demand of the wholesale market, their model, which they refer to as the "Fundamental Model", is designed with the intention of estimating the response in *overall consumption* to prices of an aggregation of all generators. In contrast, the study in this chapter has a simple, singular goal: To compute the *wholesale demand* response to a change in the price of wholesale electricity, that is, to compute the price elasticity of demand in the wholesale electricity market.

A justification for the focus of this chapter is that consumers are only very indirectly affected by wholesale prices. In fact, the prices which consumers are charged by their retailers are very stable, unlike the prices observed in wholesale markets. As such, it is unintuitive to investigate their response to changes in wholesale prices, especially at high frequencies. Therefore, attempting to evaluate consumers' demand response to a volatile price to which they are not directly exposed, and of which they often have no current knowledge, is unrealistic. However, in the same way that wholesale electricity prices indirectly, but surely, affect end-users, any increased clarity about the behaviour of wholesale market participants is almost certain to indirectly benefit consumers.

Most existing empirical works, with the exception of Kulakov and Ziel (2019), compute price elasticity using an estimated coefficient in a regression model. Under certain model specifications, the same economic reasoning which underlies this estimation also permits the construction of a time series of price elasticities of demand. Such time series of PEDEs are constructed as functions of an estimated regression coefficient and the time series of prices and quantities[23]. The fact that the coefficient is usually constant over the sample, save for a few studies which permit time-varying parameters (Inglesi-Lotz, 2011; Tiwari and Menegaki, 2019), imposes certain unrealistic implicit assumptions on the shape of the demand curve. The degree to which such assumptions will affect the quality of the PEDE estimates can be assessed by comparing those estimates against true values, which are approximated with a high degree of accuracy here. In this regard, the present study finds that a regression-based method is unable to provide estimates which would accurately represent average PEDE over a sample.

Section 4.2 provides a brief overview of the existing literature on econometric estimation of price elasticities. The implicit assumptions about demand curves which are imposed by regression estimates will also be discussed in the same section. Following that, Section 4.3 describes the novel approach for measuring real-time, true price elasticity of supply and demand for electricity, as well as an existing regression based approach which is to be evaluated. Section 4.4 introduces the data used in this chapter. In this section real-time price elasticities are shown, and some of their characteristics are discussed. Furthermore, the accuracy of elasticities estimated using a

[23] Admittedly, since the coefficient estimate represents an average effect over the whole sample, using it to construct a time series of price elasticities may not be ideal. However, the average of these elasticities, which is equivalent to computing the average elasticity using mean price and volume, should still coincide with the average of elasticities computed using the proposed approach.

regression approach is evaluated. Finally, Section 4.5 concludes with a discussion of modelling considerations and empirical performance of the regression-based method, as well as some closing remarks.

4.2 Existing works

4.2.1 Literature overview

Price elasticity of demand (supply) is the percentage change in the quantity demanded (supplied) of a commodity in response to a percentage change in its price, and can be represented by the formula

$$\mathcal{PE} = \frac{dQ/Q}{dP/P} = \frac{dQ}{dP} \times \frac{P}{Q}, \tag{4.1}$$

where $\mathcal{PE}$ is price elasticity, and Q and P are the quantity and price of the good.

The formula in (4.1) is applied in a few ways in empirical studies which use econometric methods. For example, Lijesen (2007) fits a linear regression model of the form

$$Q_t = c_P P_t + \boldsymbol{z}_t' \boldsymbol{\beta} + \varepsilon_t, \tag{4.2}$$

in which quantity (Q_t) is the dependent variable and price (P_t) is one of the explanatory variables. The vector of variables $\boldsymbol{z}_t$ controls for any other external factors which influence the quantity of electricity demanded so that the estimated coefficient c_P captures only the effect of price. The estimated coefficient on price, c_P, represents dQ/dP in (4.1), and (average) elasticity is computed by multiplying the c_P by (the mean of) prices divided by (the mean of) quantities. Other models use log-transformations of both quantity and price, often referred to as log-log models, and can be expressed as

$$q_t = c_p p_t + \boldsymbol{z}_t' \boldsymbol{\beta} + \varepsilon_t, \tag{4.3}$$

where $q_t = \log Q_t$ and $p_t = \log P_t$. This model is used by Lijesen (2007), Fan and Hyndman (2011), and many others. Elasticity in log-log models is simply the coefficient on the log of price,

c_p. On the other hand, for log-linear models such as

$$q_t = c_P P_t + \boldsymbol{z}_t' \boldsymbol{\beta} + \varepsilon_t, \tag{4.4}$$

in which quantity is log-transformed but price is not, elasticity is given by

$$\mathcal{PE} = [\exp(c_P) - 1] \, P,$$

where c_P is the coefficient of price, and P is the price level (Fan and Hyndman, 2011). These three approaches are among the most popular methods which have been employed.

This general idea is used, albeit with different model specifications, in the existing empirical literature. Selected recent works are summarised in Table 4.1. References to less-recent, but no-less-important, related studies can be found within the works contained in Table 4.1 and in the meta-analysis by Labandeira et al. (2017).

Note that since electricity prices are allowed be negative, the sign for the estimated elasticities may occasionally be different from those which would traditionally be expected. Furthermore, a log-log model would no longer be feasible in general since the logarithm of a negative number is not defined.

4.2.2 Implicit assumptions in regression-based methods

For decades, various established econometric models have been used to provide best estimates of PEDE based on the available data. The results may then be used to inform policy or adjust consumer and supplier behaviour. However, the accuracy of the estimates and the validity of the methods have never been clearly evaluated. Furthermore, this approach assumes or imposes certain restrictions on the shape of demand curves. Using the three examples presented in (4.2)–(4.4), demand curves are implicitly assumed to be of the forms expressed in Table 4.2. If the true demand curves are vastly different from these forms, then price elasticity, which is a function of the slope of the curve at a given point, will be inaccurately estimated.

In truth, there is no reason to believe that demand or supply curves adhere to any of these three forms. In fact, upon inspecting the superimposition of the implied linear, log-linear, and log-log

[24] Although Ziramba (2008) claims to use a cointegration approach, the method is more accurately described as one which exploits a level relationship. The study in question invokes the test for level relationships by Pesaran et al. (2001), but avoids any unit root tests. Cointegration is merely an element in the set of level relationships.

Table 4.1: Summary of recent literature on price elasticity for electricity demand.

Researcher	Year	Type of model	Account for endogeneity	Elasticity (LR: Long run; SR:Short run)	
Alberini and Filippini	2011	Dynamic panel	Yes	LR:	-0.43 to -0.73
				SR:	-0.08 to -0.15
Al-Faris	2002	ECM | Cointegration	-	LR:	-0.82 to -3.39
				SR:	-0.04 to -0.18
Boogen et al.	2017	Dynamic panel	Yes	LR:	-0.6
				SR:	-0.3
Burke and Abayasekara	2018	Panel data	Yes	LR:	-0.95 to -1.01
				SR:	-0.06 to -0.24
Dergiades and Tsoulfidis	2008	ARDL | Cointegration	Yes	LR:	-1.07
				SR:	-0.39
Dergiades and Tsoulfidis	2011	ARDL | Cointegration	Yes	LR:	-0.61 to -0.67
				SR:	-0.09
Dilaver and Hunt	2011a	STSM	-	LR:	-0.38
				SR:	-0.09
Dilaver and Hunt	2011b	STSM	-		-0.11
Dong et al.	2020	TSLS	Yes	LR:	-0.51
				SR:	-0.68
Filippini and Pachauri	2004	OLS	-		-0.29 to -0.51
Inglesi-Lotz	2011	State-space	-		-1.08 to -0.05
Lijesen	2007	TSLS	Yes		-0.001 to -0.004
Nakajima	2010	Panel cointegration	-		-1.127
Nakajima and Hamori	2010	Panel cointegration	-		-0.12 to -0.33
Narayan and Smyth	2005	ARDL | Cointegration	Yes		-0.47 to -0.54
Silva et al.	2018	(Pseudo-)Panel	-		-0.59 to -0.84
Tiwari and Menegaki	2019	State-space	-		-0.21
Yu and Xin	2020	Panel data	-		-0.47 to -0.80
Zhou and Teng	2013	OLS	-		-0.35 to -0.50
Ziramba[24]	2008	ARDL	Yes	LR:	-0.04
				SR:	-0.02

*ARDL: Autoregressive distributive lag, ECM: Error-correction model, OLS: Ordinary least squares, STSM: Structural time series model, TSLS: Two-stage least squares.
†All the papers in this list use log-log models. Lijesen (2007) also uses a linear model.
‡Note that different papers use different terminology for variable transformations. In particular, the term "loglinear" is sometimes used to refer to log-log transformations, and should not be confused with the term "log-linear", which describes (4.4).

Table 4.2: Regression model specifications, functional form of the demand curve, and equation for price elasticity. In each of the following, c represents a coefficient estimated from the relevant regression equation in each case, and α is an intercept which is some function of all other explanatory variables included in the regression.

Model	Demand curve	Elasticity at time t
Linear	$Q = \alpha + cP$	$\mathcal{PE} = c\,(P_t/Q_t)$
Log-linear	$Q = \exp(\alpha + cP)$	$\mathcal{PE} = [\exp(c) - 1]\,P_t$
Log-log	$Q = \alpha P^c$	$\mathcal{PE} = c$

curves on the actual demand curve (see Figure 4.1), it is quite clear that none of the functional forms match reality. In Figure 4.1(a), it is plain to see that the shapes of the implied curves are only similar to the actual demand curve in some small neighbourhood around the point at which they are designed to meet. It is important to emphasise that c and α were chosen so that the tangents of the implied curves are all equal to the calculated slope of the smoothed demand curve at the intersection. In Panel (b), α and c are calibrated to match the intersection quantity and price at 12:00 on 18 January 2015 (the next day).

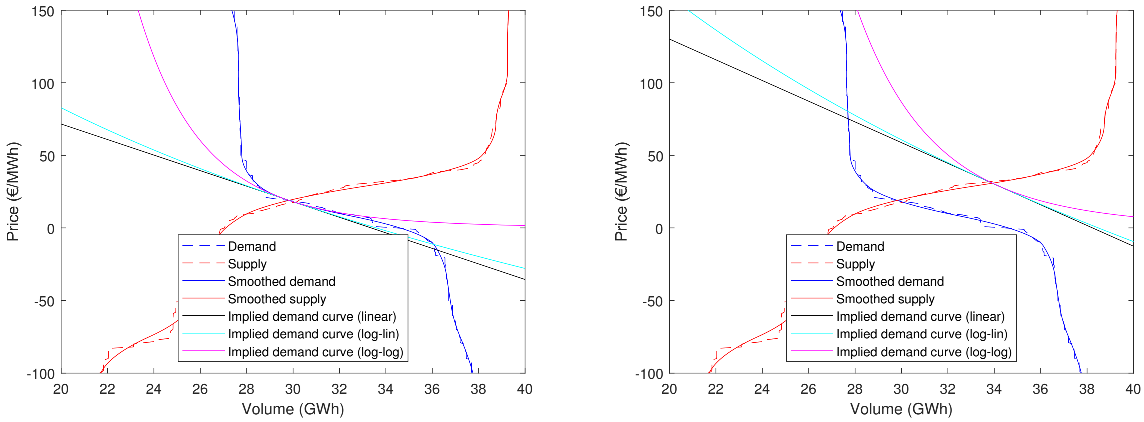

(a) Implied curves

(b) Implied curves for next day's equilibrium

Figure 4.1: Actual and smoothed supply and demand curves at 12:00 17 January 2015. Also shown are the demand curves implied by a linear, log-linear, or log-log (only for positive prices) model. In Panel (a), the values of c and α are computed purely from smoothed curves, and are chosen to exactly match their intersection. Panel (b) assumes the correct value of c is implied by a regression model, but chooses α, where applicable, to match the equilibrium price and quantity for the same hour of the next day.

Each of these plots illustrates a possible situation with regards to a regression-based approach. Recall that regression estimates are based on an entire sample instead of data from just a single intra-day market period. As a result, the coefficient may coincidentally be exactly the value which corresponds to the slope of the demand curve at the equilibrium point for a specific time of a specific day. More often than not, the coefficient estimate will differ from the the required value. Panel (a) shows that even in a rare and unlikely event when regression estimates yield the correct estimate of price elasticity for a particular point in time, they would still be unable to provide any accurate information pertaining to the rest of the demand curve. In Panel (b), the estimated coefficients are correct for a different day, but do not provide any information about the market at 12:00 on 17 January 2015. Clearly, attempting to accurately measure real-time price elasticities requires that certain parameters be allowed to vary.

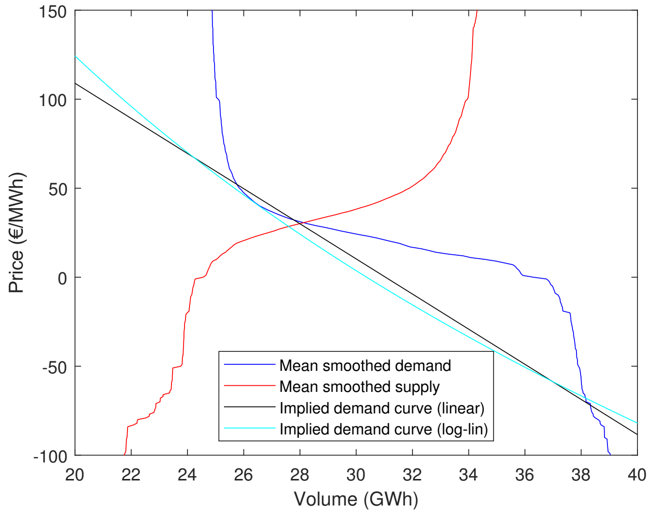

Figure 4.2: Mean smoothed supply and demand curves for the entire sample and the curves implied by coefficients from a simple linear regression of volume on price.

One way in which the discussion around Figure 4.1 might be rebutted is the claim that regression estimates give an idea of the average price elasticity of demand over the entire sample, therefore considerations relating to a single observation are invalid. To address this view, Figure 4.2 contains the mean of the smoothed supply and demand curves over the entire sample. These curves represent average bidding behaviour of all wholesale market participants throughout the sample. Superimposed upon these plots are the implied demand curves from fitting a simple linear regression of equilibrium quantity on price (in either the linear or the log-linear specifications) and, where necessary, using the sample means of price and volume. Even in this mean case, the implied demand curves do not match the true curve. Furthermore, the slope of the implied curves in the vicinity of the equilibrium point are different from the true slope, which would lead to an inaccurate estimate of price elasticity of demand.

In addition to the incongruity between implied and observed demand curves, many studies also ignore the possibility of endogeneity in the model. In particular, prices can be expected to be correlated with the error term due to simultaneity between price and quantity, leading to biased coefficient estimates and, therefore, biased estimates of price elasticity. Just under half of the references in Table 4.1 account for endogeneity in the model.

Attention must also be paid to the meaning of the expression dQ/dP. In regression equation (4.2), the quantity $c_P = dQ/dP$ captures the expected difference in quantity, Q_t, to various price levels, P_t, in a given set of time series data. On the other hand, dQ/dP in the classical expression for price elasticity in (4.1) represents the change in quantity demanded or supplied for a change in price in the same market for the same period. Indeed, many introductory econometrics texts refer the the quantity $(dy/dx) \times (x/y)$ as elasticity. However, electricity price and volume data are recorded at different points in time. In order to use the same logic to estimate price elasticity as some interpretation of regression coefficients, an assumption must be made that all market participants behave similarly across the time domain. This assumption simply does not hold true for the wholesale electricity market since we can see from observed supply and demand curves that market participants behave differently in each hour. In summary, even though both quantities have the same expression, dQ/dP, they represent different things.

An additional comment about regression-based estimation of price elasticity must be made. This observation depends on the data which analysed. A number of studies use some variation of the regression methods in Section 4.2.1 to investigate the response of consumers (residential, industrial, etc.) to wholesale electricity prices. This practice posits a direct relationship between wholesale markets and consumers which may not exist. Wholesale market prices are highly volatile, whereas consumers typically pay rates which are considerably more stable. In other words, consumers are not directly financially exposed to wholesale market prices. As such, any statistically-estimated relationship between the two may be spurious. Of course, empirical methods are not just limited to the three approaches highlighted in this section. However, other techniques, such as those used by the studies in Table 4.1, can mostly be viewed as extensions of these. Any shortcomings of these models will likely also affect their extensions.

Since Kulakov and Ziel (2019) propose an approach which can be viewed as a more complex extension of the method proposed here, some remarks on their paper are warranted. Although the technique is similar, it would appear that the central focus of Kulakov and Ziel (2019) is different. The "fundamental model" they propose seems to perform some transformations so as to estimate the demand response of market *participants*, recognising that some participants place bids to both buy and sell electricity in the market. On the other hand, this present study is not concerned with the activities of each market participant. The only quantity of interest here is what percentage change in quantity can be expected for a percentage change in bidding price. More crucially, there appears to be a flaw in the motivation for their fundamental model. Specifically, their explanation of how a "utility" in the market would arbitrage from placing sale and purchase bids at the same time would, in fact, lead to utilities being in a situation where they may be contractually obligated to generate a quantity which exceeds their capacity instead of making an easy profit.[25]

4.3 Real-time price elasticity of demand and supply

4.3.1 Description

Price elasticity in wholesale electricity markets can be computed at every intra-day trading interval (hourly, in Germany) using information contained within that interval alone. All the bids, consisting of price (€/MWh) and quantity (GWh) pairs, in a certain hour can be used to construct a demand or a supply curve. Consider, for example, Figure 4.1. The dashed blue line represents the demand curve at noon on 1st January 2016 in the German wholesale electricity market. The full range of permissible prices is between -€500€/MWh and €3,000/MWh. However, the plot in Figure 4.1 focuses on a subset of prices so as to better display the curvature of the market curves around the intersection. By construction, the curve is piecewise-constant. Price elasticity can be computed at any point, (P, Q) on the smoothed demand curve, represented by the solid blue line in Figure 4.1. The inverse of the gradient at the chosen point is dQ/dP which is then substituted into (4.1). The same procedure can be carried out for the curves observed at 1pm, 2pm, 3pm, and so forth. If elasticity is computed at equilibrium prices and quantities for each hour, a time series of real-time realised price elasticities of demand for electricity, say $(\mathcal{PE}_1, \mathcal{PE}_2, \ldots, \mathcal{PE}_T)$, is produced. The same process, when applied to smoothed supply curves, produces a time series of price elasticities of electricity supply.

4.3.2 Estimating price elasticities

In the empirical component of this chapter, realised price elasticity of supply and demand are estimated by first smoothing the supply and demand curves using cubic B-splines, constructed using the Cox-de Boor recursion outlined in Hastie et al. (2009). Choose a grid of N prices, P_i, $i = 1, \ldots, N$. Let $Q_t^s(P_i)$ and $Q_t^d(P_i)$, $i = 1, \ldots, N$, represent the smoothed supply and demand curves at time t, respectively. Realised price elasticity of demand (supply) is then estimated using (4.1), with equilibrium price and quantity (P_t^* and Q_t^*) as well as an approximation of

[25] An email sent to the corresponding author of Kulakov and Ziel (2019) requesting clarification did not receive a response.

dQ/dP for the demand (supply) curve at that point, $\hat{g}_t^{d*}$ ($\hat{g}_t^{s*}$). These values are computed as

$$P_t^* = \arg\min_{P_i} \left| Q_t^s(P_i) - Q_t^d(P_i) \right| \tag{4.5}$$

$$Q_t^* = \frac{Q_t^s(P_t^*) + Q_t^d(P_t^*)}{2} \tag{4.6}$$

$$\hat{g}_t^{d*} = \frac{Q_t^d(P_t^* + \delta) - Q_t^d(P_t^* - \delta)}{2\delta} \tag{4.7}$$

$$\hat{g}_t^{s*} = \frac{Q_t^s(P_t^* + \delta) - Q_t^s(P_t^* - \delta)}{2\delta}, \tag{4.8}$$

where δ is some small interval. In this chapter, smooth supply and demand curves are defined at €1/MWh intervals, so $\delta = 1$.

4.3.3 Comparing results

This study is able to compute price elasticity of demand in wholesale electricity from fundamental concepts on an hour-by-hour basis instead of using some aggregate or average value estimated over a period of time. Furthermore, the only error that exists in this approach arises from the choice of smoothing algorithm of demand curves, and not from any other assumptions or misspecification that could be present in econometric models. This means that the intra-day wholesale PEDE values produced by this approach are arguably the most accurate. It may be enlightening to compare these values to those empirically estimated using other methods.

In selecting an existing approach for comparison, the main consideration was that the method be used here with minimal modification, and is therefore as close to the original proposed method as possible. The models in Lijesen (2007) were chosen for a number of reasons. First, one of the two specifications used in Lijesen (2007) does not take the logarithm of prices as a variable. This is important because wholesale electricity prices around the world are now permitted to drop below zero. Second, the study in Lijesen (2007) investigates intra-day data, which is the frequency of the data used here. Third, the models in Lijesen (2007) were proposed with wholesale electricity prices in mind, whereas others are tailored to disaggregated data such as residential or industrial usage only. Finally, the two-stage least squares approach addresses the endogeneity in prices, correcting for biased coefficient estimates. It is simple enough to replace hourly electricity load (the dependent variable of choice in Lijesen (2007)) with wholesale quantity traded as the variable representing quantity demanded, with some justification for doing so provided in Section 4.1. In short, the linear model from Lijesen (2007) which is evaluated here takes the form

$$\begin{aligned} Q_t = cP_t + \alpha Trend_t + \sum_{r=1}^{10} \beta_{H,r} x_{H,r} + \sum_{s=1}^{2} \beta_{D,s} x_{D,s} + \sum_{u=1}^{6} \beta_{M,u} x_{M,u} \\ + \sum_{v=1}^{5} \beta_{W,v} x_{W,v} + \sum_{\ell=1}^{2} \beta_{B,\ell} x_{B,r} + \varepsilon_t, \end{aligned} \tag{4.9}$$

with twenty seven explanatory variables. Since the data is at the hourly frequency, the subscript t represents the hour in the sample. For example, observation at $t = 23$ occurs at 11pm on the first day, and $t = 25$ refers to 1am on the second day. The two key variables of interest are Q_t and P_t, which are demand and price, respectively. The ten dummy variables $(x_{H,1}, \ldots, x_{H,10})$ indicate readings which occur at each of the ten hours between 9am and 6pm, inclusive. Dummy variables $(x_{D,1}, x_{D,2})$ indicate observations on Thursdays and Fridays, respectively. Month-of-year dummies, $(x_{M,1}, \ldots, x_{M,6})$, indicate observations occurring in January, February, April, July, September, and October, respectively. Variables $(x_{W,1}, \ldots, x_{W,5})$ are weather variables: $x_{W,1}$ is the maximum temperature on a given day, $x_{W,2}$ is the interaction term between maximum temperature and the time falling between 12pm and 4pm, inclusive; $(x_{W,3}, x_{W,4}, x_{W,5})$ are the daylight variable times a 9am dummy, a 5pm dummy, and a 6pm dummy, respectively. The holiday (break) dummy variable $x_{B,1}$ indicates when most of the country is on summer holidays, and $x_{B,2}$ indicates observations falling in the week from Christmas Day to New Year's day.

The daylight variable used in $(x_{W,3}, \ldots, x_{W,5})$ is defined in Lijesen (2007) as the quadratic difference, measured in days, from the longest day of the year, or the summer solstice. The interaction terms are designed to accommodate differences in electricity usage around changes in sunrise and sunset times. $x_{B,2}$, which is called "Week 53" in Lijesen (2007), accounts for the fact that the week between Christmas and New Year's is typically a holiday[26]. Additionally, where Lijesen (2007) has two variables for summer holidays, the adapted model which is used in this chapter has only one. This is because the data used here is from Germany instead of the Netherlands. Whereas summer breaks in the Netherlands are determined separately in the north and south regions (hence two variables), Germany has sixteen regions which may all have slightly different summer breaks. Instead of including a dummy variable for each region, the

[26]Note that this is in contrast with the study in Chapter 2, where only Christmas and New Year's days are holidays instead of the whole week. The choice in Chapter 2 was motivated by observations in the data (see Figure 1.1), whereas the use of "Week 53" in this study is intended to match the variable defined in Lijesen (2007)

variable $x_{B,1}$ consists of the two months in which most regions are on summer holidays. In order to account for endogeneity in P_t, lagged prices, P_{t-1} are used as instruments for P_t in two-stage least squares regressions.

Additionally, it is worth noting that Lijesen (2007) limit their study to peak usage periods, which is defined as working days from 9am to 6pm. They further restrict their analysis to one calendar year. Therefore, part of the comparison in this chapter will also be restricted to peak hours in one calendar year. However, a second round of comparison will extend the data to every single hour in a year, as well as permitting the sample to be longer than one calendar year.

In the event that the sample does not contain negative prices, then the log-log specification may still be used. It turns out that prices at peak hours in this data do not fall below zero. In this case, the model is adapted from (4.9) as

$$\begin{aligned} q_t = cp_t + \alpha Trend_t + \sum_{r=1}^{10} \beta_{H,r} x_{H,r} + \sum_{s=1}^{2} \beta_{D,s} x_{D,s} + \sum_{u=1}^{6} \beta_{M,u} x_{M,u} \\ + \sum_{v=1}^{5} \beta_{W,v} x_{W,v} + \sum_{\ell=1}^{2} \beta_{B,\ell} x_{B,r} + \varepsilon_t, \end{aligned} \tag{4.10}$$

where q_t and p_t are the natural logs of Q_t and P_t, respectively. For two-stage least squares, instruments for p_t are p_{t-1}.

Where applicable, additional dummy variables are added to capture more hours of the day as well as Saturdays and Sundays.

Minor deviation from the original model

The construction of the "daylight" variable specified in Lijesen (2007) is not completely representative of daylight hours since it does not consider that certain days in the year may be closer to the summer solstice in an adjacent year than in the current year, and may under-represent the daylight hours on those days. This issue becomes obvious when trying to construct a daylight variable using the given definition for samples which span more than one year. Instead of taking quadratic differences, an alternative is to fit an interpolating cubic spline to points which take the value 1 on winter solstice and 0 on the summer solstice of each year. The reason that splines are chosen here instead of some parametric curve, e.g. a sinusoidal function, is simply because summer and winter solstices may occur on different days each year, leading to slightly different spacing between the peaks and troughs.[27]

4.4 Results and discussion

4.4.1 Data

Hourly sale and purchase curves are constructed using intra-day bid and offer data in the German wholesale electricity market, obtained from the European Power Exchange (EPEX). The data spans from 00:00 on 1st April 2015 to 23:00 on 31st March 2017, encompassing two years' (731 days') worth of hourly bids. Daily maximum temperatures are from DWD Climate Data Center 2018. The summer solstice in 2015, 2016, and 2017, used in computing the variable "daylight", occurred on 21st, 20th, and 21st of June, respectively. Winter solstices for the same years occurred on 22nd, 21st, and 21st of December, respectively. Summer holidays for the variable $x_{B,1}$ in (4.9) and (4.10) were the entire months of July and August.

4.4.2 Real-time price elasticity in the German spot market

The series of hourly realised price elasticity of supply and demand from 00:00 1st April 2015 to 23:00 31st March 2017 are presented in Figure 4.3[28]. The real-time (very short-term) PEDEs are larger than the averages reported by Knaut and Paulus (2016), who estimate the response of consumer demand to wholesale prices in Germany, albeit for a different sample period. This outcome is not surprising as changes in wholesale prices do not affect consumers as directly as they do participants in the wholesale market. Therefore it is only to be expected that consumer usage, represented by load in Knaut and Paulus (2016), is less responsive to such changes. At a relatively small number of hours, the signs of price elasticity of supply and demand are the opposite of what might traditionally be expected. These are due to the fact that the equilibrium electricity price at those times was negative. There do not appear to be any obvious changes in the unconditional mean level of price elasticity of supply or demand across these two years. Summary statistics for these prices, demand, and price elasticities are reported in Table 4.3.

[27] One set of results which uses the original quadratic difference definition for "daylight" is included in Appendix B, and are similar to those estimated using the spline definition in Section 4.4.

[28] Smoothing of the supply and demand curves was conducted using cubic B-splines with 24 equally-spaced interior knots, giving 28 orthonormal spline functions.

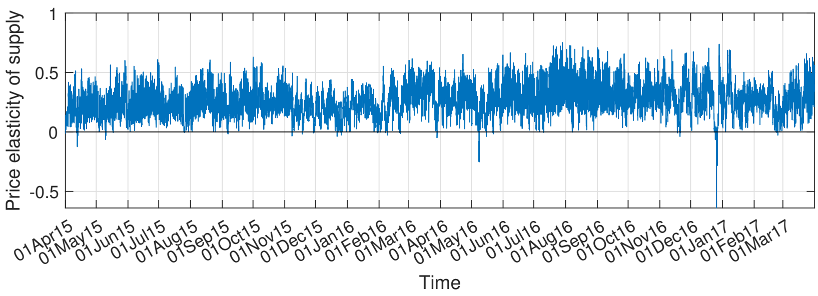

(a) Price elasticity of supply

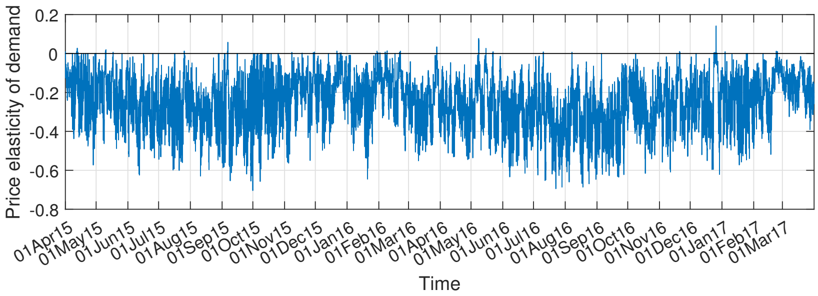

(b) Price elasticity of demand

Figure 4.3: Hourly price elasticity of supply and demand from 00:00 on 1st April 2015 to 23:00 on 31st March 2017.

Table 4.3: Summary statistics of hourly prices, demand, and real-time price elasticities.

	Panel A: 01:00 1st January 2016–23:00 31st December 2016							
	Whole sample				Peak hours only			
	Mean	Std. dev.	Min	Max	Mean	Std. dev.	Min	Max
Price (€/MWh)	29.31	12.44	-135.00	107.00	35.32	11.23	3.00	107.00
Demand (GWh)	26.75	5.06	17.21	47.97	30.22	4.66	18.19	46.05
Real-time PEDE	**-0.27**	0.13	-0.69	0.14	**-0.30**	0.11	-0.62	-0.00
Real-time PESE	0.28	0.14	-0.64	0.75	0.22	0.07	0.01	0.66
	Panel B: 00:00 1st April 2015–23:00 31st March 2017							
	Whole sample				Peak hours only			
	Mean	Std. dev.	Min	Max	Mean	Std. dev.	Min	Max
Price (€/MWh)	31.59	14.05	-135.00	164.00	38.41	14.06	3.00	158.00
Demand (GWh)	27.86	5.59	17.21	51.47	31.79	5.44	17.88	51.47
Real-time PEDE	**-0.25**	0.13	-0.70	0.14	**-0.28**	0.11	-0.62	-0.00
Real-time PESE	0.25	0.13	-0.64	0.75	0.20	0.08	0.01	0.66

Patterns in price elasticities

A different visualisation of realised PEDE and PESE for the sample is shown in Figure 4.4. These two plots are constructed using exactly the same information as Figure 4.3, but presented such that any daily or annual patterns are more shown. In both plots of Figure 4.4, the bottom left point represents price elasticity at 01:00 on 1st April 2015. Going horizontally across corresponds to different times on the same day. On the other hand, moving vertically from any given point means viewing the price elasticity on a different day in the sample, but at the same time of the

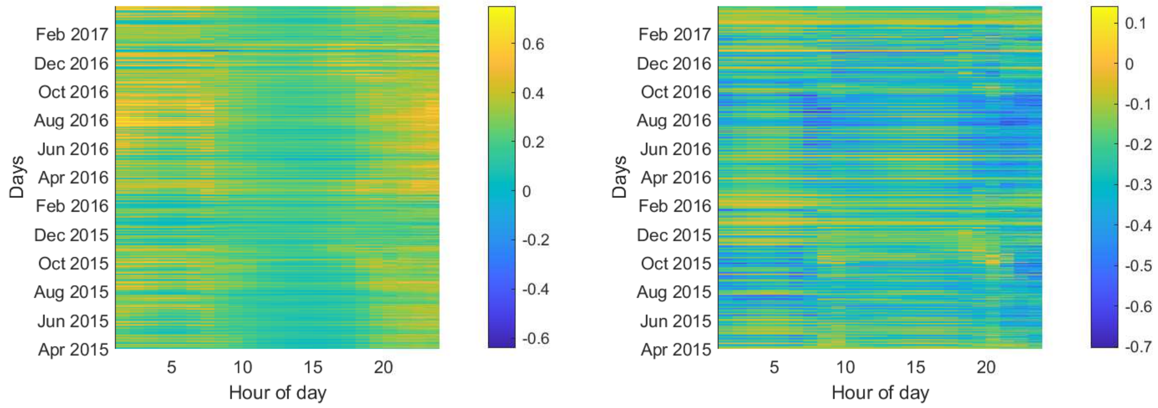

(a) Price elasticity of supply

(b) Price elasticity of demand

Figure 4.4: Hourly price elasticities from 1st April 2015 to 31st March 2017. The y-axis indicates the date, and each point on the x axis corresponds to an hour of the day, from 1–24.

day. It seems that both price elasticity of demand and supply are typically close to zero around midday, and furthest from zero when it is dark. This diurnal pattern is in accordance with the patterns found by Knaut and Paulus (2016) and Kulakov and Ziel (2019) in the German market, even though the market participants represented in their studies are different from those in this study. A simple explanation for this pattern is that during waking and working hours, a certain minimum amount of electricity is required to power factories, machinery, infrastructure in office buildings, and lights and air-conditioning all over. This required energy must be delivered regardless of the price, so buyers do not have much flexibility to reduce the quantities bid when prices rise. Consequently, elasticity is lower during these hours. In the evenings and early mornings, particularly after midnight at which point most of the population is resting, electricity usage is more flexible and variable, granting more flexibility with regards to wholesale bids.

In addition to the diurnal pattern, the plots in Figure 4.4 also reveal that there is an annual seasonality associated with those hours with higher price elasticity, i.e. between approximately 7pm and 9am. For instance, it is quite clear from Panel (b) that price elasticity of demand of electricity at 8pm is lowe (closest from zero) in the coldest months (October–February) than in warmer months (June–September). A likely cause of this pattern is similar in spirit to that of the diurnal pattern. In general, locations with cold winters which have long nights will require more heating and lighting during the winter months. The increased need for electricity will lead to lower bargaining power of buyers in the wholesale market, so to speak, making demand less elastic to price.

4.4.3 Regression-based estimates of price elasticity

Different subsamples of the data are investigated. In order to compare results from the method proposed by Lijesen (2007), albeit on a different dataset to that on which it was originally used, the first set of estimates is restricted to a single year from 00:00 1st January 2016–23:00 31st December 2016, and only for peak hours as determined by Lijesen (2007), that is, weekdays between 9am and 6pm, inclusive. In this subsample, all prices are positive. Hence, the two-stage least squares linear and log-log models from (4.9) and (4.10), respectively, are both feasible, and will be estimated[29].

Market behaviour in peak hours may be expected to be more homogeneous than if every hour in the year was included. To examine the possibility of this potential phenomenon artificially inflating the accuracy of the regression-based estimates, the linear model was estimated on the sample consisting of every single hour in 2016. Since this sample contains negative prices, only the linear model was estimated.

While the span of one year seems to be an obvious choice of sample length, it is also interesting to investigate the effect of a longer period. In particular, if the coefficient on the price variable is unstable (is different for different sample periods), then estimates of price elasticity will also change with sample choice. To this end, estimation of all the models and specifications was repeated for the full sample of two years from 00:00 1st April 2015–23:00 31st March 2017.

Coefficient estimates and t-statistics for the ordinary least squares and two-stage least squares linear regression model may be viewed in Appendix C. The 99% confidence intervals of the key coefficients, that is, the coefficients on price (or the log of price) estimated by two-stage least squares, are shown in Table 4.4. Also included are the values of the coefficients which would be necessary to produce the corresponding mean price elasticity of demand values in Table 4.3. For the linear model, this value is calculated as $\overline{\mathcal{PE}} \times \bar{Q}/\bar{P}$, where a bar, $\bar{\cdot}$, denotes the mean of its argument over the sample and $\mathcal{PE}$ is price elasticity approximated using the proposed method in Section 4.3.2. In the log-log model, the required coefficient is simply the mean of price elasticities, $\overline{\mathcal{PE}}$. Of the six confidence intervals, only one contains the required value for correct estimation of the average price elasticity for the sample. Considering that these models were originally designed with only peak hours in mind, it is also interesting to find that none of the

[29] Augmented Dickey-Fuller (ADF) tests (Dickey and Fuller, 1979; Said and Dickey, 1984) find that the series of prices and volumes are both stationary, so any estimated relationship between the two is not spurious.

peak-hour-only samples give accurate estimates. Furthermore, the regression-based estimates do not appear to over- or underestimate the coefficient in a consistent manner, since the "true" coefficient lies on either side of the confidence interval, depending on which sample is chosen.

Table 4.4: 99% confidence intervals of coefficients on the price variable from estimated models and the value of the coefficient which would yield the true sample average price elasticity of electricity demand.

	00:00 1st Jan 2016–23:00 31st Dec 2016	
	Confidence interval	Implied true coefficient
Linear (Whole)	(-0.2597,-0.2411)	**-0.2479**
Linear (Peak)	(-0.2897,-0.2695)	-0.2631
Log-log(Peak)	(-0.3662,-0.3443)	-0.3045
	00:00 1st Apr 2015–23:00 31st Mar 2017	
	Confidence interval	Implied true coefficient
Linear (Whole)	(-0.1828,-0.1704)	-0.2216
Linear (Peak)	(-0.1736,-0.1623)	-0.2284
Log-log(Peak)	(-0.2675,-0.2532)	-0.2752

Instead of an average price elasticity of demand over the entire year, it is possible to construct real-time estimates as a function of price, quantities, and the estimated coefficient on price at each hour using the formula for elasticity with the linear model in Table 4.2.[30] The differences between regression-based real-time estimates using this approach and the true values in Figure 4.3 are represented by the blue lines in Figure 4.5. Panels (a) and (b) correspond to the subsample of one year, whereas Panels (c) and (d) are for the full sample consisting of two years. Unsurprisingly, the shape of the plots in Figures 4.5(a) and (c) are similar to their counterparts in Figures 4.5(a) and (d), respectively, since they are essentially just scaled differently.

Overall, it appears that regression-based estimates of real-time PEDE from the one-year sample in Panels (a) and (b) are reasonably accurate in spring and summer, but deviate from true values by a significant amount in the colder months. When the sample is increased to two years in Panels (c) and (d), more incidents of estimated PEDE being further than two standard deviations from true values occur, even in the warmer months of 2016 which were reasonably accurate in Panels (a) and (b). This suggests that although sample selection may not impact the estimated average PEDE, the real-time PEDEs estimated from this regression-based approach are sensitive to the choice of sample period.

[30] It is acknowledged, at this point, that the original model in Lijesen (2007) was never intended for this purpose. Nevertheless, it is simple to execute and interesting for the sake of comparison.

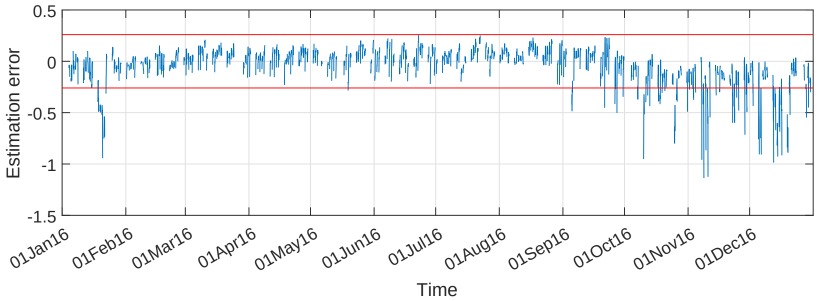

(a) Peak hours in 2016.

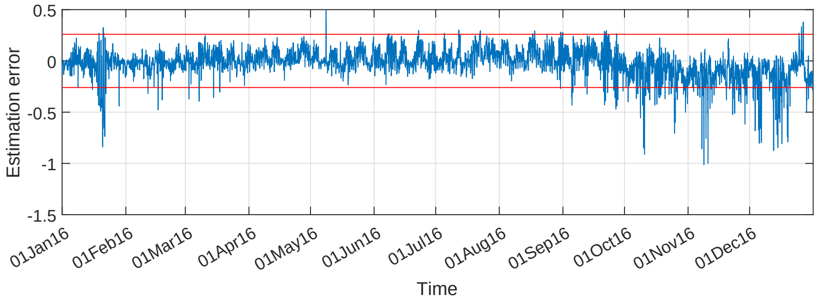

(b) All hours in 2016.

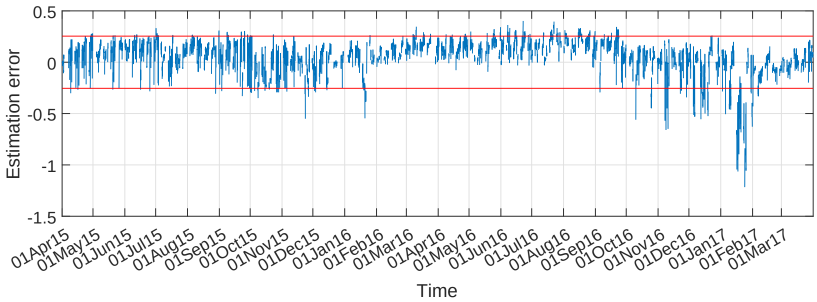

(c) Peak hours for the two years from 1st April 2015 to 31st March 2017.

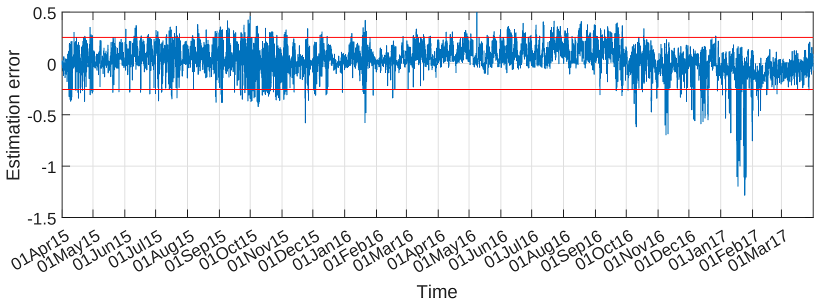

(d) All hours in the two years from 1st April 2015 to 31st March 2017.

Figure 4.5: Estimation error of regression-based real-time price elasticity of demand for peak hours from the linear model estimated by two-stage least squares. The red lines demarcate 2 standard deviations of true PEDE from Table 4.3 from zero.

Of course, the regression-based approaches in Lijesen (2007) were not intended to be used for estimating real-time price elasticities (in the sense of this chapter as opposed to the definition in Lijesen (2007)). However this comparison serves to illustrate the benefits of being able to investigate price elasticity of electricity demand at a higher granularity instead of just as a sample average. Even the hourly-averages presented by Knaut and Paulus (2016) are unable to capture and display other forms of seasonality beyond the intra-day pattern.

4.5 Discussion

This study presents a novel method of measuring price elasticity of supply and demand for wholesale electricity based on a highly-detailed data set. These real-time true elasticities are computed by making full use of intra-day auction market data, and are not subject to exogenous variables. Price elasticities computed in this manner are more-or-less exact, as they use a fundamental approach, namely, computing price elasticities from supply and demand curves. The only error in measurement arises from the choice and calibration of a smoothing algorithm for these curves. The series of real-time price elasticities of supply and demand were shown, and some properties were discussed. It was observed that price elasticities in electricity markets exhibit some diurnal pattern in which they are most inelastic during daylight hours, as well as annual seasonality.

Before conducting any empirical analysis, a few potential drawbacks of regression-based estimation were suggested. These drawbacks include the imposition of an unrealistic functional form on the demand curve and the presence of endogeneity. Furthermore, an argument is also made against the use of regression-based estimation of elasticities with time series data. The empirical section of this chapter investigates the accuracy of a relevant regression-based model. The regression model in question was adapted from Lijesen (2007), and was largely unmodified in this study. Estimates from the regression methods over different sample periods were compared to the values which would have been necessary in order to produce accurate measures of average price elasticities, and were mostly found to be significantly different from the required values.

Ultimately, this chapter does not seek to discourage the use of regression models to estimate price elasticity of electricity demand. In fact, depending on data limitations for specific research questions, regressions may still be the best available technique. However, results from such approaches should always be used with caution, since their accuracy is generally unverified. Accordingly, any policy recommendations related to price elasticity of demand in electricity markets must allow for the likely event that estimates from regression models are inaccurate.

In addition to the analysis presented in this study, it may be of interest for future research to consider price elasticity of supply and demand at prices and quantities other than the equilibria. Such analysis may help prepare market participants for situations which may shift one or both curves such that a different price is traded.

In summary, this study introduces a new way of accurately measuring intra-day price elasticity of demand, and opens up new avenues for examining various electricity market quantities in real-time for very-short-term analysis. The same method can also be used to compute price elasticity of electricity supply, although that quantity seems to be of less interest in the literature at the moment.

Chapter 5

Other ideas and future research

In addition to the completed projects in Chapters 2–4, some other research ideas were conceived, with varying levels of results and completion. The first was a planned exploration of counterfactuals about the operation of grid-scale battery storage systems in South Australia. The second is an investigation into the whether the tools used in some studies on probabilistic modelling and forecasting of electricity prices are appropriate.

5.1 Battery storage and wholesale price

Historically, energy has not been economically storable. As a consequence, changes in the demand of electricity across every single second must be met perfectly by adjusting supply accordingly. As with any other commodity, periods of high demand are expected to be associated with higher prices, all else being constant. However, even if all generators are functioning properly, there may be instances when demand is so high that currently-operating generators are unable to deliver enough energy, and more electricity sources need to be commissioned. This type of occurrence leads to an even sharper price increase due to the extra cost of activating additional—typically less-efficient—generators, sometimes resulting in positive spikes. Since the widespread adoption of renewable energy sources, a new situation has arisen in which energy needs to be balanced out. Instead of having to increase supply to meet demand, an excess of generated energy may lead to the need for increased demand. For example, wind and solar energy are intermittent, highly-dependent on environmental factors outside of human control, and, more importantly, not biddable in response to price (Giulietti et al., 2018). Wind turbines and solar farms remain active in periods of prolonged strong winds[31] or intense sunlight so large quantities of renewable energy are generated. Any excess energy must have an outlet, leading to prices dropping quite drastically in order to boost quantity demanded.[32] The apparent consensus in the literature on electricity pricing is that non-storability of energy has led to prices being more volatile and more prone to large spikes than most other commodity markets (see, inter alia, Higgs and Worthington (2008) and Weron (2014)).

If this explanation for large and frequent electricity price movements holds, the solution for smoothing prices is to introduce batteries capable of storing sufficient energy. Such facilities could help by taking on excess energy, and releasing it in periods of high demand. The necessary technology has recently become available, with a number of grid batteries having been installed around the world. In particular, the Hornsdale Power Reserve (HPR) in South Australia ranks among the biggest in the world in terms of energy storage capacity, and is second in terms of power (rate of delivery of energy). Furthermore, it has been operational since December 2018, giving at least two full years' worth of data. The South Australian electricity market is therefore an ideal setting in which to study the effects of having a large battery. A growing number of studies (Giulietti et al., 2018; Davies et al., 2019) study the economic benefits, or revenue, which could be enjoyed by the battery operator. This situation is exemplified by the HPR, which is owned by a foreign private company, and therefore operates in a manner which maximises the profits of its parent company and its shareholders. On the other hand, it would be of academic and public interest to investigate the possibilities associated with a hypothetical similar battery which was state-owned, and would therefore aim to improve the utility of the general public by smoothing or lowering wholesale prices.

Installing grid batteries can be expected to reduce volatility and the occurrence spikes. On the other hand, in 2018, Matt Canavan, the Australian Minister for Resources at the time, remarked of the HPR that "this big battery is the Kim Kardashian of the energy world—it's famous for being famous. It really doesn't deliver much". Ultimately, it seems that the impact of such technology on energy prices boils down to two things, namely, whether the battery is capable of contributing to lower prices, and whether it is utilised with such a goal in mind. To illustrate this point, the battery in South Australia is able to supply up to 129MWh at 100MW (currently

[31]Wind turbines are shut down only under extreme conditions when wind speeds exceed the furling speed, or the maximum speed at which a turbine can safely operate.

[32]The technical aspects of electricity generation, modulation, and safety requirements are beyond the scope of this thesis, but basic information is available on the website of the Australian Energy Market Commission (AEMC) and Australian Energy Market Operator (AEMO).

undergoing expansion to 194MWh at 150MW) to a market in which half-hourly demand is around 1,290MWh, on average. However, the contribution from batteries is 0.4% on average, instead of the maximum of approximately 2% average contribution[33].

The fact that the maximum capacity of the HPR is able to account for only a small proportion of demand may be discouraging, but it is worth considering that the battery may not need to be pushed to its limits in order to smooth out prices, depending on the circumstances. The wholesale electricity demand curve in South Australia is perfectly inelastic. If the intersection occurs within a region in which the supply curve is relatively flat, then a small change in the shape of the curve induced by different bidding behaviour by the HPR would not significantly affect prices. On the other hand, when the intersection of supply and demand curves occur at steep points in the two curves, then a small shift in part of the supply curve could potentially lead to a large change in equilibrium price.

The HPR is operated by NEOEN (the operator of the neighbouring Hornsdale Wind Farm), under certain agreements with the South Australian government. Of the battery's capacity, 119MWh of energy storage capacity and 30MW of the discharge capacity are permitted to be used for energy arbitrage (AEMO, 2018). This means that most of the storage capacity is used to buy and store energy at low prices and sell it to the market when prices are high. Clearly, this procedure does not help smooth or lower electricity prices. The remaining capacity is reserved for power system reliability purposes, and participates in the Frequency Control Ancillary Services (FCAS) markets[34].

So far, it seems that the HPR could conceivably be capable of contributing to smoother electricity prices, although this point is not certain, given the small energy storage capacity relative to demand in South Australia. It is clear, though, that the HPR is not currently being utilised with price smoothing as a goal.

The proposed study would attempt to answer two questions. First, a counterfactual analysis can be conducted to find out exactly how much of an impact the HPR could have had since beginning operations in 2018 if the 119MWh storage and 30MW discharge capacity was used solely for the purpose of smoothing prices instead of energy arbitrage. The second question

[33] This value is calculated under some very naive and simplistic assumptions: (a) the full discharging capacity of 100MW and charging capacity of 80MW are utilised, (b) the battery is in operation 24 hours a day, with no breaks, (c) the battery always fully charges and discharges, and (d) there is no energy loss.

[34] HPR is reported to have reduced FCAS costs by AUD116 million in 2019, but there is no evidence of these savings being passed on in any form to consumers. FCAS is not a vital part of this proposed study, but the reader is referred to AEMO (2015) for a brief explanation of the process.

considers a similar situation as the first, but from the alternative perspective: For a given target reduction in either price volatility or spike frequency, how large would the capacity of a battery in South Australia need to be? Answers to questions such as these may serve as a guide to the optimal usage of grid batteries.

The AEMO makes public a vast amount of data. Unfortunately, the author of this thesis was unable to use this data to reconstruct supply curves which corresponded to realised price and quantities traded, despite reaching out to various experts on the Australian wholesale electricity markets. Nevertheless, these questions are interesting, and alternative approaches may be devised to examine the potential for large scale batteries to influence electricity prices.

5.2 Methods in probabilistic price forecasting

Following the 2014 edition of the Global Energy Forecasting Competition (GEFCom), the recent trend in electricity price forecasting is producing probabilistic, instead of point, forecasts. A review of probabilistic forecasting is presented by Nowotarski and Weron (2018). One of popular techniques employed in this area is quantile regression (QR). Put simply, QR models allow the econometrician to model or forecast different expected quantiles of the response variable (electricity price) as functions of explanatory variables. Two different applications of QR have been used in probabilistic electricity price forecasting. The first estimates usual models, with price as the dependent variable and various quantities which can be expected to affect price levels as explanatory variables on the right-hand side. Hagfors et al. (2016b) and Hagfors et al. (2016c) fit quantile regression models to wholesale electricity market prices in order to obtain a better understanding of the effects of various price drivers. The quantiles τ selected in these studies range from 1% to 99%, covering some of the most extreme cases. The other application of QR is called Quantile Regression Averaging (QRA), where right-hand side variables are point forecasts from various other models (Maciejowska et al., 2016; Nowotarski and Weron, 2015). For both applications of QR, the quantile regression model is estimated in the same way, the only difference being the choice of right-hand side variables.

Consider a classical linear model

$$y_t = \boldsymbol{z}_t' \boldsymbol{\alpha} + e_t, \tag{5.1}$$

where $\boldsymbol{z}_t$ is a vector of regressors and e_t are i.i.d. mean zero errors independent of $\boldsymbol{z}_t'$. The conditional τ-th quantile of the distribution of y_t is given by

$$Q_\tau\left(y_t\right) = \boldsymbol{z}_t'\boldsymbol{\alpha}^{(\tau)}, \tag{5.2}$$

with parameter estimates obtained as

$$\hat{\boldsymbol{\alpha}}^{(\tau)} = \underset{\boldsymbol{\alpha}^{(\tau)}}{\arg\min} \sum_t \left(\tau - \mathbb{1}_{y_t < \boldsymbol{z}_t'\boldsymbol{\alpha}^{(\tau)}}\right)\left(y_t - \boldsymbol{z}_t'\boldsymbol{\alpha}^{(\tau)}\right),$$

where $\mathbb{1}$ is an indicator function taking the value 1 if its subscript argument is true and 0 otherwise.

The first potential problem with a direct application of QR methods to electricity price data is in the specification of the model in the presence of price spikes. For comparison, consider Markov regime-switching (MRS) models, in which price spikes are treated as originating from a different data generating process than that which produces the base prices. This approach is reasonable since the effect of factors which drive normal prices (such as seasonality and lagged prices) will have a negligible effect on extreme prices. The classical linear model in (5.1), on the other hand, makes no such accommodation.

Second, it has been shown that asymptotic approximation for distributions of QR estimates in extreme quantiles (close to either 0 or 1) is different from those in more central quantiles (Chernozhukov et al., 2017). There is no clear threshold quantile beyond which this problem is present, but given that Hagfors et al. (2016a) include 1% and 99% quantiles, it is possible that inference based on asymptotic approximation are not appropriate.

A potential solution to the second problem is to use an appropriate bootstrap method. Indeed, Hagfors et al. (2016c) obtain standard errors of their coefficient estimates by bootstrapping. However, they make no mention of which bootstrapping technique is used. In mean (least squares) regressions, the choice of bootstrap is relevant as it depends on the type of data. This is also the case in the quantile regression setting (Koenker, 2005; He, 2017). With electricity prices in particular, the choice of an appropriate bootstrap procedure may be even more crucial owing to the presence of extreme values.

The quantile crossing problem, where predicted values resulting from the estimates of a lower value of τ end up being larger than predictions from a higher τ value is a known issue with QR in practice (He, 1997; Chernozhukov et al., 2010; Liu and Wu, 2011; Schmidt and Zhu, 2016). Of all the papers which use quantile regression for electricity prices, only Haben and Giasemidis (2016) explicitly acknowledge this problem, addressing it using the rearrangement or bootstrap method of Chernozhukov et al. (2010). Even so, they do not explore the extent to which this problem occurs, or the conditions under which it happens.

Preliminary Monte-Carlo simulations have been conducted using two data generating processes (DGPs). The first one is an autoregressive model with an exogenous variable, without spikes. The second introduces randomly-occurring spikes represented by a normally-distributed random variable with a large mean. Results from these simple simulations suggest that parameter estimates for extreme quantiles under the spike DGP are distorted. Furthermore, quantile crossing occurs with these simulations, although the problem seems to be reduced when extreme values are introduced under the second DGP.

There is clearly a need to re-evaluate the manner in which probabilistic electricity price forecasting is carried out in order to ensure that the techniques employed are both relevant and appropriate.

Chapter 6

Conclusion

Research into energy is so important that it spans multiple disciplines. The technical aspect will be of interest to engineers and physicists, whereas the impact of generation methods and emissions are of concern to climate scientists and the public at large. No matter the angle taken, any advancements in research can better inform policy-makers. The studies in this thesis contribute to general understanding and clarity of wholesale markets in Germany from a quantitative perspective.

6.1 Findings and contributions

In the first study in Chapter 2, a semiparametric model for time-varying coefficients was proposed for forecasting daily volume-weighted spot prices in the German wholesale market. The motivation behind this type of model is that various external factors may lead to one-off changes in the long-term structure or trajectory of prices. Such effects are posited to affect the value of model coefficients, but are also expected manifest gradually, as opposed to instantaneously. For example, a new policy to discontinue the use of nuclear power plants would have to allow enough time for facilities and equipment of alternative energy sources to be constructed and installed. However, such a change can be expected to lead to a change in electricity prices due to the cost associated with alternative generation methods and, potentially, generation capacities at different times of the day. Results from this study show that the time-varying coefficient models produce more accurate forecasts than their constant-parameter counterpart. This outcome is verified using the Diebold-Mariano test.

The study in Chapter 3 focuses on predicting day-ahead prices at the intra-day frequency by forecasting sale and purchase curves using functional data analysis (FDA) techniques. This approach to price forecasting is relatively new. The proposed method differs from the only other FDA approach in the literature in its estimation method. The functional principal components technique used here is far simpler, quicker, and less computationally-intensive than the kernel-type method proposed by Shah and Lisi (2020). Despite the fact that estimation and forecasting with the FDA method is only slightly more involved—and no more complicated—than a regular time series benchmark, Diebold-Mariano tests reveal that it clearly outperforms the benchmark for every hour of the day.

Departing from electricity price forecasting, Chapter 4 proposes a new method of estimating price elasticity in the wholesale market which is possible because German wholesale electricity market data is detailed enough to permit construction of supply and demand curves. This approach is simple, but appeals to the most fundamental concept of price elasticity, which is percentage change in quantity demanded or supplied for a percentage change in price. The elasticities approximated here are at the intra-day level, providing an insight into diurnal and annual seasonality in the behaviour of wholesale market participants. This chapter also includes an extensive discussion on other ways in which the literature estimates price elasticity of demand using regression estimates, and highlights a number of ways in which those methods may be inappropriate. Finally, a relevant approach from the literature is adapted for the wholesale market, with its estimates compared to those from the proposed technique. Estimated price elasticity of demand from the adapted model are shown to be different from the values produced by the proposed method, which is claimed here to be very accurate. Ultimately the message from this study is not that regression-based estimates of price elasticity should be avoided in favour of the proposed approach. Instead, regression for this branch of research should be performed carefully and with certain additional considerations.

6.2 Future direction

The scope for research in energy markets is wide. A few research ideas could not be included in this thesis due to data or time constraints. In addition to possible extensions to Chapters 2–4 which were included in the conclusion of the relevant chapters, other suggestions were introduced and discussed in Chapter 5 as possible future research directions.

The first of these ideas is particularly germane at a time when energy demands have changed due to lockdowns associated with the Covid-19 pandemic, and while global temperatures and carbon dioxide levels continue to rise in spite of international agreements. Increased implementation of grid-scale batteries would facilitate wider adoption of renewable electricity generation, which is likely to significantly reduce emissions. The proposed study would investigate the potential impact of large batteries on wholesale prices.

The second direction for future research is to inspect and dissect some econometric approaches which have been adopted in the probabilistic electricity price forecasting literature. Specifically, quantile regression is increasingly popular in this area. However, none of the studies to date appear to account for the fact that electricity prices are not well-behaved, in that certain assumptions underlying quantile regression techniques may not be satisfied. A valuable contribution to the literature would be to investigate this question by simulations as well as empirical applications, and to suggest ways in which any shortcomings may be addressed.

Bibliography

AEMO (2015). Guide to ancillary services in the National Electricity Market. Media release, Australian Energy Market Operator. `https://aemo.com.au/-/media/files/pdf/guide-to-ancillary-services-in-the-national-electricity-market.pdf`.

AEMO (2018). Initial operation of the Hornsdale Power Reserve Battery Energy Storage System. Media release, Australian Energy Market Operator. `https://www.aemo.com.au/-/media/Files/Media_Centre/2018/Initial-operation-of-the-Hornsdale-Power-Reserve.pdf`.

Aggarwal, S. K., Saini, L. M., and Kumar, A. (2009a). Electricity price forecasting in deregulated markets: A review and evaluation. *International Journal of Electrical Power and Energy Systems*, 31:13–22.

Aggarwal, S. K., Saini, L. M., and Kumar, A. (2009b). Short term price forecasting in deregulated electricity markets. A review of statistical models and key issues. *International Journal of Energy Sector Management*, 3:333–358.

Al-Faris, A. R. F. (2002). The demand for electricity in the GCC countries. *Energy Policy*, 30:117–124.

Alberini, A. and Filippini, M. (2011). Response of residential electricity demand to price: The effect of measurement error. *Energy Economics*, 33:889–895.

Amato, A. D., Ruth, M., Kirshen, P., and Horwitz, J. (2005). Regional energy demand responses to climate change: Methodology and application to the commonwealth of Massachusetts. *Climatic Change*, 17:175–201.

Auffhammer, M. and Mansur, E. T. (2014). Measuring climatic impacts on energy consumption: A review of the empirical literature. *Energy Economics*, 46:522–530.

Barone-Adesi, G. and Gigli, A. (2003). Managing electricity risk. *Economic Notes*, 32:283–294.

Becker, R., Hurn, S., and Pavlov, V. (2007). Modelling spikes in electricity prices. *The Economic Record*, 83:371–382.

Benth, F. E., Kallsen, J., and Meyer-Brandis, T. (2007). A non-Gaussian Ornstein-Uhlenbeck process for electricity spot price modeling and derivatives pricing. *Applied Mathematical Finance*, 14:153–169.

Bessec, M. and Fouquau, J. (2008). The non-linear link between electricity consumption and temperature in Europe: A threshold panel approach. *Energy Economics*, 30:2705–2721.

Bierbrauer, M., Menn, C., Rachev, S. T., and Trück, S. (2007). Spot and derivative pricing in the EEX power market. *Journal of Banking and Finance*, 31:3462–3485.

Bierbrauer, M., Trück, S., and Weron, R. (2004). Modeling electricity prices with regime switching models. *Lecture Notes on Computer Science*, 3039:859–867.

BMU and BMWi (2010). *Energy concept for an environmentally sound, reliable, and affordable energy supply*. Federal Ministry for the Environment, Nature Conservation and Nuclear Safety (BMU) and Federal Ministry of Economics and Technology (BMWi), Berlin.

BMWi (2015). *Fourth "Energy Transition" monitoring report*. Federal Ministry for Economic Affairs and Energy (BMWi), Berlin.

Bode, S. and Groscurth, H. (2006). The effect of the German Renewable Energy Act (EEG) on "the electricity price". HWWA Discussion Paper No. 358, Hamburg Institute of International Economics, Hamburg.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31:307–327.

Boogen, N., Datta, S., and Filippini, M. (2017). Dynamic models of residential electricity demand: Evidence from Switzerland. *Energy Strategy Reviews*, 18:85–92.

Bordignon, S., Bunn, D. W., Lisi, F., and Nan, F. (2013). Combining day-ahead forecasts for British electricity prices. *Energy Economics*, 35:88–103.

Bosco, B., Parisio, L., Pelagatti, M., and Baldi, F. (2010). Long-run relations in European electricity prices. *Journal of Applied Econometrics*, 25:805–832.

Bosco, B. P., Parisio, L. P., and Pelagatti, M. M. (2007). Deregulated wholesale electricity prices in Italy: An empirical analysis. *International Advances in Economic Research*, 13:415–432.

Bosq, D. (2000). *Linear Processes in Function Spaces.* Springer New York.

Burke, P. J. and Abayasekara, A. (2018). The price elasticity of electricity demand in the United States: A three-dimensional analysis. *The Energy Journal*, 39:123–145.

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

Chen, X. B., Gao, J., Li, D., and Silvapulle, P. (2018). Nonparametric estimation and forecasting for time-varying coefficient realized volatility models. *Journal of Business & Economic Statistics*, 36:88–100.

Chen, Y. and Li, B. (2017). An adaptive functional autoregressive forecast model to predict electricity price curves. *Journal of Business & Economic Statistics*, 35:371–388.

Chernozhukov, V., Fernández-Val, I., and Kaji, T. (2017). Extremal quantile regression. In *Handbook of Quantile Regression*, pages 333–362. Chapman and Hall/CRC.

Chernozhukov, V., Fernández-Van, I., and Galichon, A. (2010). Quantile and probability curves without crossing. *Econometrica*, 78:1093–1125.

Chiou, J.-M. (2012). Dynamical functional prediction and classification, with application to traffic flow prediction. *The Annals of Applied Statistics*, 6:1588–1614.

Christensen, T., Hurn, S., and Lindsay, K. (2009). It never rains but it pours: Modelling the persistence of spikes in electricity prices. *The Energy Journal*, 30:25–48.

Christensen, T. M., Hurn, A. S., and Lindsay, K. A. (2012). Forecasting spikes in electricity prices. *International Journal of Forecasting*, 28:400–411.

Clements, A., Fuller, J., and Hurn, S. (2013). Semi-parametric forecasting of spikes in electricity prices. *Economic Record*, 89:508–521.

Conejo, A. J., Contreras, J., and Espínola, R. (2005). Forecasting electricity prices for a day-ahead pool-based electric energy market. *International Journal of Forecasting*, 21:435–462.

Cuaresma, J. C., Hlouskova, J., Kossmeier, S., and Obersteiner, M. (2004). Forecasting electricity spot-prices using linear univariate time-series models. *Applied Energy*, 77:87–106.

Cuevas, A. (2014). A partial overview of the theory of statistics with functional curves. *Journal of Statistical Planning and Inference*, 147:1–23.

Dahlhaus, R. (2012). Locally stationary processes. In Subba Rao, T., Subba Rao, S., and Rao, C. R., editors, *Handbook of statistics*, volume 30, pages 351–413. Elsevier.

Davies, D. M., Verde, M. G., Mnyshenko, O., Chen, Y. R., Rajeev, R., Meng, Y. S., and Elliott, G. (2019). Combined economic and technological evaluation of battery energy storage for grid applications. *Nature Energy*, 4:42–50.

Dergiades, T. and Tsoulfidis, L. (2008). Estimating residential demand for electricity in the United States, 1965–2006. *Energy Economics*, 30:2722–2730.

Dergiades, T. and Tsoulfidis, L. (2011). Revisiting residential demand for electricity in Greece: new evidence from the ARDL approach to cointegration analysis. *Empirical Economics*, 41:511–531.

Dickey, D. A. and Fuller, W. A. (1979). Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American statistical association*, 74:427–431.

Diebold, F. X. and Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business and Economic Statistics*, 13:134–144.

Dilaver, Z. and Hunt, L. C. (2011a). Modelling and forecasting Turkish residential electricity demand. *Energy Policy*, 39:3117–3127.

Dilaver, Z. and Hunt, L. C. (2011b). Turkish aggregate electricity demand: An outlook to 2020. *Energy*, 36:6686–6696.

Dong, Z., Liu, Z., Liu, J., Li, L., and Zhao, J. (2020). Estimating the price elasticity of electricity of urban residential consumers in eastern China. In *2020 Asia Energy and Electrical Engineering Symposium (AEEES)*, pages 881–885.

DWD Climate Data Center (CDC) (2018). Historical daily station observations (temperature, pressure, precipitation, sunshine duration, etc.) for germany. version v006.

Eichler, M. and Türk, D. (2013). Fitting semiparametric Markov regime-switching models to electricity spot prices. *Energy Economics*, 36:614–624.

Engle, R. F. (1982). Autoregressive conditional heteroskedasticity with estimates of the variance of U.K. inflation. *Econometrica*, 50:987–1008.

Engle, R. F., Granger, C. W. J., Rice, J., and Weiss, A. (1986). Semiparametric estimates of the relation between weather and electricity sales. *Journal of the American Statistical Association*, 81:310–320.

Erni, D. (2012). *Day-ahead electricity spot prices - fundamental modelling and the rold of expected wind electricity infeed at the European Energy Exchange*. Phd dissertation, University of St. Gallen.

Escribano, A., Peña, J. I., and Villaplana, P. (2011). Modelling electricity prices: International evidence. *Oxford Bulletin of Economics and Statistics*, 73:622–650.

Fan, S. and Hyndman, R. J. (2011). The price elasticity of electricity demand in South Australia. *Energy Policy*, 39:3709–3719.

Feehan, J. P. (2018). The long-run price elasticity of residential demand for electricity: Results from a natural experiment. *Utilities Policy*, 51:12–17.

Ferraty, F., Van Keilegom, I., and Vieu, P. (2012). Regression when both response and predictor are functions. *Journal of Multivariate Analysis*, 109:10–28.

Filippini, M. and Pachauri, S. (2004). Elasticities of electricity demand in urban Indian households. *Energy Policy*, 32:429–436.

Gaillard, P., Goude, Y., and Nedellec, R. (2016). Additive models and robust aggregation for GEFCom2014 probabilistic electric load and electricity price forecasting. *International Journal of Forecasting*, 32:1038–1050.

Gao, J. and Hawthorne, K. (2006). Semiparametric estimation and testing of the trend of temperature series. *Econometrics Journal*, 9:332–355.

Gao, Y., Shang, H. L., and Yang, Y. (2020). Modelling functional data with high-dimensional error structure. In Aneiros, G., Horová, I., Hŭsková, M., and Vieu, P., editors, *Functional* *and High-Dimensional Statistics and Related Fields*, chapter 14, pages 99–106. Springer International Publishing.

Garcia, R. C. and Contreras, J. (2005). A GARCH forecasting model to predict day-ahead electricity prices. *IEEE Transactions on Power Systems*, 20:867–874.

Geman, H. and Roncoroni, A. (2006). Understanding the fine structure of electricity prices. *The Journal of Business*, 79:1225–1261.

Gianfreda, A. (2010). Volatility and volume effects in European electricity spot markets. *Economic Notes*, 39:47–63.

Giulietti, M., Grossi, L., Trujillo-Baute, E., and Waterson, M. (2018). Analyzing the potential economic value of energy storage. *The Energy Journal*, 39:101–122.

Goia, A., May, C., and Fusai, G. (2010). Functional clustering and linear regression for peak load forecasting. *International Journal of Forecasting*, 26:700–711.

Gonzalez, J. P., San Roque, A. M., and Perez, E. A. (2018). Forecasting functional time series with a new Hilbertian ARMAX model: Application to electricity price forecasting. *IEEE Transactions on Power Systems*, 33:545–556.

Goto, M. and Karolyi, G. A. (2004). Understanding electricity price volatility within and across markets. Working paper, Ohio State University and Central Research Institute of Electric Power Industry of Japan.

Grossi, L., Heim, S., and Waterson, M. (2017). The impact of the German response fo the Fukushima earthquake. *Energy Economics*, 66:450–465.

Grossi, L. and Nan, F. (2019). Robust forecasting of electricity prices: Simulations, models and the impact of renewable sources. *Technological Forecasting and Social Change*. Forthcoming.

Haben, S. and Giasemidis, G. (2016). A hybrid model of kernel density estimation and quantile regression for GEFCom2014 probabilistic load forecasting. *International Journal of Forecasting*, 32:1017–1022.

Hagfors, L. I., Bunn, D., Kristoffersen, E., Staver, T. T., and Westgaard, S. (2016a). Modeling the UK electricity price distributions using quantile regression. *Energy*, 102:231–243.

Hagfors, L. I., Kamperud, H. H., Paraschiv, F., Prokopczuk, M., Sator, A., and Westgaard, S. (2016b). Prediction of extreme price occurrences in the German day-ahead electricity market. *Quantitative Finance*, 16:1929–1948.

Hagfors, L. I., Paraschiv, F., Molnar, P., and Westgaard, S. (2016c). Using quantile regression to analyze the effect of renewables on EEX price formation. *Renewable Energy and Environmental Sustainability*, 1. Article 32.

Hall, P. and Vial, C. (2006). Assessing the finite dimensionality of functional data. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 68:689–705.

Hansen, P. R., Lunde, A., and Nason, J. M. (2011). The model confidence set. *Econometrica*, 79:453–497.

Härdle, W., Liang, H., and Gao, J. (2000). *Partially Linear Models.* Springer Series in Contributions to Statistics. Springer-Verlag, New York.

Hastie, T., Tibshirani, R., and Friedman, J. (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction.* Springer Series in Statistics. 2nd ed. edition.

He, X. (1997). Quantile curves without crossing. *The American Statistician*, 51:186–192.

He, X. (2017). Resampling methods. In *Handbook of Quantile Regression*, pages 7–20. Chapman and Hall/CRC.

Hébert-Losier, K., Pini, A., Vantini, S., Strandberg, J., Abramowicz, K., Schelin, L., and Häger, C. K. (2015). One-leg hop kinematics 20 years following anterior cruciate ligament rupture: Data revisited using functional data analysis. *Clinical Biomechanics*, 30:1153–1161.

Henley, A. and Peirson, J. (1997). Non-linearities in electricity demand and temperature: Parametric versus non-parametric methods. *Oxford Bulletin of Economics and Statistics*, 59:149–162.

Higgs, H. and Worthington, A. (2008). Stochastic price modeling of high volatility, mean-reverting, spike-prone commodities: The Australian wholesale spot electricity market. *Energy Economics*, 30:3172–3185.

Horváth, L. and Kokoszka, P. (2012). *Inference for Functional Data with Applications.* Springer New York.

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

Huisman, R. and De Jong, C. (2003). Option formulas for mean-reverting power prices with spikes. *Energy Power Risk Management*, 7:12–16.

Huisman, R. and Mahieu, R. (2003). Regime jumps in electricity prices. *Energy Economics*, 25:425–434.

Hyndman, R. J. and Ullah, M. S. (2007). Robust forecasting of mortality and fertility rates: A functional data approach. *Computational Statistics & Data Analysis*, 51:4942–4956.

Inglesi-Lotz, R. (2011). The evolution of price elasticity of electricity demand in South Africa: A Kalman filter application. *Energy Policy*, 39:3690–3696.

Janczura, J., Trück, S., Weron, R., and Wolff, R. C. (2013). Identifying spikes and seasonal components in electricity spot price data: A guide to robust modeling. *Energy Economics*, 38:96–110.

Janczura, J. and Weron, R. (2010). An empirical comparison of alternate regime-switching models for electricity spot prices. *Energy Economics*, 32:1059–1073.

Janczura, J. and Weron, R. (2012). Efficient estimation of Markov-switching models: An application to electricity spot prices. *Advances in Statistical Analysis*, 96:385–407.

Jiang, L. and Hu, G. (2018). A review on short-term electricity price forecasting techniques for energy markets. In *2018 15th International Conference on Control, Automation, Robotics and Vision (ICARCV)*, pages 937–944. IEEE.

Karakatsani, N. V. and Bunn, D. W. (2008a). Forecasting electricity prices: The impact of fundamentals and time-varying coefficients. *International Journal of Forecasting*, 24:764–785.

Karakatsani, N. V. and Bunn, D. W. (2008b). Intra-day and regime-switching dynamics in electricity price information. *Energy Economics*, 30:1776–1797.

Kendrick, S. K., Zheng, Q., Garbett, N. C., and Brock, G. N. (2017). Application and interpretation of functional data analysis techniques to differential scanning calorimetry data from lupus patients. *PLOS ONE*, 12:e0186232.

Ketterer, J. C. (2014). The impact of wind power generation on the electricity price in Germany. *Energy Economics*, 44:270–280.

Knaut, A. and Paulus, S. (2016). Hourly price elasticity pattern of electricity demand in the German day-ahead market. Working Paper No 16/07, EWI Working Paper.

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

Koenker, R. (2005). *Quantile Regression*. Econometric Society monographs ; no. 38. Cambridge University Press, New York.

Kokoszka, P. and Reimherr, M. (2012). Determining the order of the functional autoregressive model. *Journal of Time Series Analysis*, 34:116–129.

Kokoszka, P. and Reimherr, M. (2017). *Introduction to Functional Data Analysis*. Taylor & Francis Ltd.

Koopman, S. J., Ooms, M., and Carnero, M. A. (2007). Periodic seasonal Reg-ARFIMA-GARCH models for daily electricity spot prices. *Journal of the American Statistical Association*, 102:16–27.

Kulakov, S. and Ziel, F. (2019). Determining the demand elasticity in a wholesale electricity market. arxiv preprint arxiv:1903.11383.

Labandeira, X., Labeaga, J. M., and López-Otero, X. (2017). A meta-analysis on the price elasticity of energy demand. *Energy Policy*, 102:549–568.

Li, Y. and Flynn, P. C. (2004). Deregulated power prices: comparison of diurnal patterns. *Energy Policy*, 32:657–672.

Li, Y. and Flynn, P. C. (2006). Electricity deregulation, spot price patterns and demand-side management. *Energy*, 31:908–922.

Lijesen, M. G. (2007). The real-time price elasticity of electricity. *Energy Economics*, 29:249–258.

Lisi, F. and Nan, F. (2014). Component estimation for electricity prices: Procedures and comparisons. *Energy Economics*, 44:143–159.

Lisi, F. and Shah, I. (2020). Forecasting next-day electricity demand and prices based on functional models. *Energy Systems*, 11:947–979.

Liu, Y. and Wu, Y. (2011). Simultaneous multiple non-crossing quantile regression estimation using kernel constraints. *Journal of Nonparametric Statistics*, 23:415–437.

Lucia, J. J. and Schwartz, E. S. (2002). Electricity prices and power derivatives: Evidence from the Nordic Power Exchange. *Review of Derivatives Research*, 5:5–50.

Maciejowska, K., Nowotarski, J., and Weron, R. (2016). Probabilistic forecasting of electricity spot prices using Factor Quantile Regression Averaging. *International Journal of Forecasting*, 32:957–965.

Mestre, G., Portela, J., San Roque, A. M., and Alonso, E. (2020). Forecasting hourly supply curves in the Italian Day-Ahead electricity market with a double-seasonal SARMAHX model. *International Journal of Electrical Power & Energy Systems*, 121:106083.

Misiorek, A., Trueck, S., and Weron, R. (2006). Point and interval forecasting of spot electricity prices: Linear vs. non-linear time series models. *Studies in Nonlinear Dynamics & Econometrics*, 10. Article 2.

Moral-Carcedo, J. and Vicéns-Otero, J. (2005). Modelling the non-linear response of Spanish electricity demand to temperature variations. *Energy Economics*, 27:477–494.

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

Nakajima (2010). The residential demand for electricity in Japan: An examination using empirical panel analysis techniques. *Journal of Asian Economics*, 21:412–420.

Nakajima, T. and Hamori, S. (2010). Change in consumer sensitivity to electricity prices in response to retail deregulation: A panel empirical analysis of the residential demand for electricity in the United States. *Energy Policy*, 38:2470–2476.

Narayan, P. K. and Smyth, R. (2005). The residential demand for electricity in Australia: an application of the bounds testing approach to cointegration. *Energy Policy*, 33:467–474.

Nowotarski, J. and Weron, R. (2015). Computing electricity spot price prediction intervals using quantile regression and forecast averaging. *Computational Statistics*, 30:791–803.

Nowotarski, J. and Weron, R. (2016). On the importance of the long-term seasonal component in day-ahead electricity price forecasting. *Energy Economics*, 57:228–235.

Nowotarski, J. and Weron, R. (2018). Recent advances in electricity price forecasting: A review of probabilistic forecasting. *Renewable and Sustainable Energy Reviews*, 81:1548–1568.

OECD (2004). *OECD Reviews of Regulatory Reform: Regulatory Reform in Germany*. Paris: OECD.

Park, B. U., Mammen, E., Lee, Y. K., and Lee, E. R. (2015). Varying coefficient regression models: A review and new developments. *International Statistical Review*, 83:36–64.

Pesaran, M. H., Shin, Y., and Smith, R. J. (2001). Bounds testing approaches to the analysis of level relationships. *Journal of Applied Econometrics*, 16:289–326.

Rambharat, B. R., Brockwell, A. E., and Seppi, D. J. (2005). A threshold autoregressive model for wholesale electricity prices. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 54:287–299.

Ramsay, J. O. and Silverman, B. W. (2002). *Applied Functional Data Analysis: Methods and Case Studies*. Springer, New York.

Ramsay, J. O. and Silverman, B. W. (2005). *Functional Data Analysis*. Springer Series in Statistics. Springer.

Rice, J. A. (2004). Functional and longitudinal data analysis: Perspectives on smoothing. *Statistica Sinica*, 14:631–647.

Rice, J. A. and Silverman, B. W. (1991). Estimating the mean and covariance structure non-parametrically when the data are curves. *Journal of the Royal Statistical Society. Series B (Methodological)*, 53:233–243.

Robinson, T. (2000). Electricity pool prices: A case study in nonlinear time-series modelling. *Applied Economics*, 32:527–532.

Said, S. E. and Dickey, D. A. (1984). Testing for unit roots in autoregressive-moving average models of unknown order. *Biometrika*, 71:599–607.

Sailor, D. J. and Munõz, J. R. (1997). Sensitivity of electricity and natural gas consumption to climate in the USA – Methodology and results for eight states. *Energy*, 22:987–998.

Schmidt, L. D. W. and Zhu, Y. (2016). Quantile spacings: A simple method for the joint estimationof multiple quantiles without crossing.

Shah, I. and Lisi, F. (2020). Forecasting of electricity price through a functional prediction of sale and purchase curves. *Journal of Forecasting*, 39:242–259.

Shang, H. L. (2014). A survey of functional principal component analysis. *AStA Advances in Statistical Analysis*, 98:121–142.

Shang, H. L. (2017). Forecasting intraday S&P 500 index returns: A functional time series approach. *Journal of Forecasting*, 36:741–755.

Shang, H. L. (2019). A robust functional time series forecasting method. *Journal of Statistical Computation and Simulation*, 89:795–814.

Silva, S., Soares, I., and Pinho, C. (2018). Electricity residential demand elasticities: Urban versus rural areas in Portugal. *Energy*, 144:627–632.

Simonsen, I. (2005). Volatility of power markets. *Physica A*, 355:10–20.

Terasvirta, T., Tjostheim, D., and Granger, C. W. J. (2010). *Modelling Nonlinear Economic Time Series*. OUP Oxford.

Tiwari, A. K. and Menegaki, A. N. (2019). A time varying approach on the price elasticity of electricity in India during 1975–2013. *Energy*, 183:385–397.

Troncoso, A., Riquelme, J. C., Riquelme, J. M., Martínez, J. L., and Gómez, A. (2007). Electricity market price forecasting based on weighted nearest neighbors techniques. *IEEE Transactions on Power Systems*, 22:1294–1301.

Ullah, S. and Finch, C. F. (2013). Applications of functional data analysis: A systematic review. *BMC Medical Research Methodology*, 13:43.

Uniejewski, B., Weron, R., and Ziel, F. (2018). Variance stabilizing transformations for electricity spot price forecasting. *IEEE Transactions on Power Systems*, 33:2219–2229.

Wang, J.-L., Chiou, J.-M., and Müller, H.-G. (2016). Functional data analysis. *Annual Review of Statistics and Its Application*, 3:257–295.

Weron, R. (2014). Electricity forecasting: A review of the state-of-the-art with a look into the future. *International Journal of Forecasting*, 30:1030–1081.

Weron, R., Bierbrauer, M., and Trück, S. (2004). Modeling electricity prices: jump diffusion and regime switching. *Physica A*, 336:39–48.

Weron, R. and Misiorek, A. (2005). Forecasting spot electricity prices with time series models. *Proceedings of the European Electricity Market EEM-05 Conference*, pages 133–141.

Weron, R. and Misiorek, A. (2006). Short-term electricity price forecasting with time series models: A review and evaluation. HSC Research Report HSC/06/01, Hugo Steinhaus Center, Wrocław University of Technology, Poland.

Weron, R. and Misiorek, A. (2008). Forecasting spot electricity prices: A comparison of parametric and semiparametric time series models. *International Journal of Forecasting*, 24:744–763.

Yao, F., Müller, H.-G., and Wang, J.-L. (2005). Functional data analysis for sparse longitudinal data. *Journal of the American Statistical Association*, 100:577–590.

Yu, H. and Xin, X. (2020). Demand elasticity, ramsey index and cross-subsidy scale estimation for electricity price in China. *Sustainable Production and Consumption*, 24:39–47.

Zhou, S. and Teng, F. (2013). Estimation of urban residential electricity demand in China using household survey data. *Energy Policy*, 61:394–402.

Ziel, F. and Steinert, R. (2016). Electricity price forecasting using sale and purchase curves: The X-Model. *Energy Economics*, 59:435–454.

Ziramba, E. (2008). The demand for residential electricity in South Africa. *Energy Policy*, 36:3460–3466.

Appendix

A FDA forecasts without renewables

The use of observed—as opposed to forecasted—values of renewable energy volumes in Chapter 3 is not strictly appropriate. Diebold-Mariano statistics shown in Figure A.1 are the analogue to those in Figure 3.4 when forecasting is performed without the inclusion of wind or solar energy. The outcome is the same; forecasts using FDA techniques outperform the benchmark.

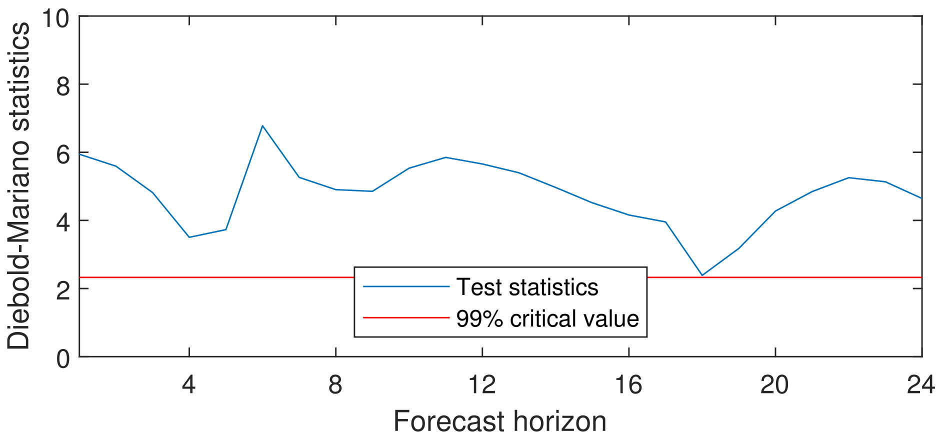

Figure A.1: Diebold-Mariano test statistics for comparing predictive accuracy between the FDA and ARX approaches. The 99% critical value for a one-tailed test is 2.32. All models estimated for this figure omit renewables as explanatoty variables

B Regression with quadratic daylight term

This section presents regression estimates (Table B.1) and estimation errors of regression-based real-time PEDE (Figure B.2) for a one-year sample when the "daylight" variable is constructed using the quadratic distance definition from Lijesen (2007). The results here are almost identical to those in the main body of the text. These results are included to show that the new definition used in this thesis does not materially change the outcome of the models.

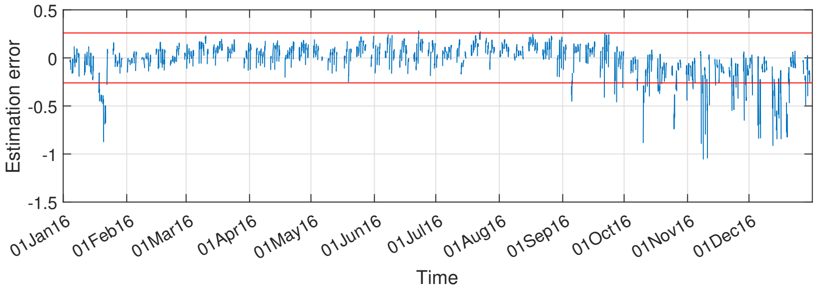

(a) Peak hours in 2016.

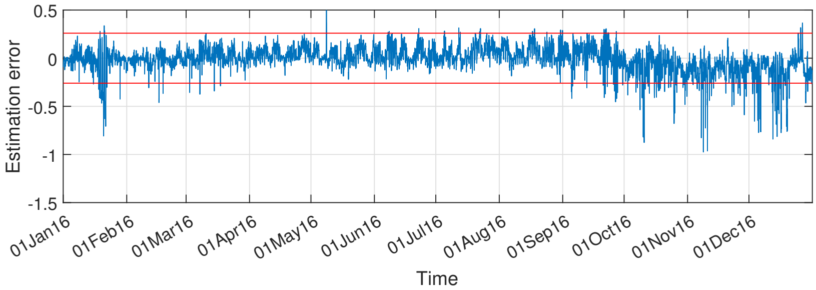

(b) All hours in 2016.

Figure B.2: Estimation error of regression-based real-time price elasticity of demand for peak hours from the linear model. The red lines demarcate 2 standard deviations of true PEDE from Table 4.3 from zero.

Table B.1: Results from estimating model (4.9) with daylight computed as the quadratic difference from the longest day of the year.
00:00 1st January 2016–00:00 31st December 2016.

	Two-stage least squares				Ordinary least squares			
	Peak		Whole sample		Peak		Whole sample	
Variable	Coef.	*t*-stat	Coef.	*t*-stat	Coef.	*t*-stat	Coef.	*t*-stat
Price (€/MWh)	-0.28	-72.04	-0.25	-69.72	-0.26	-67.33	-0.24	-67.29
Trend	0.01	12.82	0.01	15.98	0.01	10.14	0.01	15.02
Time of day dummies, hour starting at:								
1am			27.46	94.19			27.23	93.44
2am			28.33	123.29			28.15	122.57
3am			28.20	123.36			28.04	122.68
4am			30.02	122.44			29.90	121.97
5am			29.87	121.64			29.74	121.16
6am			30.14	122.20			30.00	121.70
7am			31.64	126.01			31.46	125.32
8am			34.47	133.48			34.22	132.61
9am	35.02	140.90	34.00	137.29	34.39	138.55	33.71	136.18
10am	37.35	178.26	35.67	144.86	36.80	175.87	35.39	143.77
11am	38.54	188.29	37.08	152.46	38.04	186.06	36.82	151.38
12pm	38.11	186.12	38.33	158.11	37.61	183.92	38.07	157.10
1pm	38.21	190.55	38.61	161.53	37.75	188.51	38.38	160.59
2pm	37.69	189.47	38.01	160.48	37.25	187.50	37.78	159.59
3pm	36.55	184.56	36.83	156.22	36.12	182.62	36.61	155.35
4pm	35.32	155.30	35.53	149.78	34.89	153.64	35.30	148.88
5pm	35.91	179.30	34.52	144.19	35.45	177.21	34.28	143.23
6pm	33.88	145.82	34.05	137.92	33.42	144.00	33.77	136.83
7pm			33.43	133.22			33.12	132.05
8pm			33.07	131.44			32.76	130.27
9pm			31.70	128.83			31.42	127.74
10pm			29.33	97.54			29.04	96.61
11pm			29.98	125.52			29.74	124.57
12am			28.53	96.89			28.28	96.07
Day of week dummies								
Thursday	-0.05	-0.73	0.07	0.73	-0.04	-0.61	0.06	0.68
Friday	-0.16	-2.15	-0.30	-3.18	-0.13	-1.82	-0.30	-3.15
Saturday			-2.54	-26.54			-2.50	-26.08
Sunday			-3.20	-31.31			-3.10	-30.40
Month-of-year dummies								
January	4.24	24.62	3.31	19.00	3.98	23.13	3.25	18.64
February	3.03	20.94	3.57	22.87	3.02	20.93	3.58	22.95
April	0.18	1.52	0.21	1.62	0.23	1.95	0.21	1.65
July	1.07	7.78	1.24	8.19	1.01	7.38	1.22	8.06
September	-0.29	-2.27	-0.95	-6.60	-0.22	-1.70	-0.95	-6.66
October	-0.88	-7.50	-0.43	-3.42	-0.85	-7.31	-0.44	-3.55
Weather variables								
Max. temperature	0.05	7.33	0.03	4.93	0.06	8.43	0.04	5.58
Max. temperature × 12–4pm dummy	0.12	15.42	-0.16	-17.09	0.12	15.53	-0.16	-17.35
Daylight × 9am dummy	0.00	9.39	0.00	6.78	0.00	9.77	0.00	7.00
Daylight × 5pm dummy	-0.00	-0.75	0.00	7.02	-0.00	-0.99	0.00	7.26
Daylight × 6pm dummy	0.00	15.72	0.00	4.45	0.00	15.01	0.00	4.72
Holiday dummies								
Summer holidays	-1.41	-11.36	-1.70	-12.25	-1.33	-10.72	-1.69	-12.22
Week 53(dummy)	-2.05	-8.29	-1.05	-4.19	-1.79	-7.26	-0.88	-3.50

C Regression model estimates for different specifications

Tables C.1–C.2 section presents regression estimates for the models estimated in Section 4.4. In summary, TSLS and OLS estimates are largely similar. Estimates of the coefficient on price are largely unaffected by the distinction between samples containing onle peak hours, or those which include all hours in the period. However, the coefficients vary considerably when estimated over a one-year period compared to a two-year period which overlaps the shorter sub-sample.

Table C.1: Results from estimating model (4.9) over one year.
00:00 1st January 2016–23:00 31st December 2016.

	Two-stage least squares				Ordinary least squares			
	Peak		Whole sample		Peak		Whole sample	
Variable	Coef.	*t*-stat	Coef.	*t*-stat	Coef.	*t*-stat	Coef.	*t*-stat
Price (€/MWh)	-0.28	-71.49	-0.25	-69.47	-0.26	-66.77	-0.24	-67.06
Trend	0.01	13.61	0.01	16.32	0.01	10.88	0.01	15.38
Time of day dummies, hour starting at:								
1am			26.87	82.52			26.63	81.82
2am			28.25	122.89			28.07	122.17
3am			28.12	122.96			27.96	122.28
4am			29.98	122.49			29.85	122.01
5am			29.82	121.69			29.69	121.20
6am			30.10	122.23			29.95	121.70
7am			31.59	126.00			31.40	125.31
8am			34.41	133.42			34.18	132.56
9am	34.40	127.87	33.91	136.76	33.76	125.68	33.62	135.66
10am	37.15	177.05	35.57	144.33	36.60	174.68	35.30	143.25
11am	38.35	187.02	36.99	151.89	37.84	184.80	36.83	150.86
12pm	37.94	185.60	38.24	157.59	37.44	183.41	37.99	156.58
1pm	38.05	190.04	38.53	161.03	37.60	188.01	38.29	160.08
2pm	37.53	188.96	37.92	159.99	37.09	186.99	37.70	159.10
3pm	36.39	184.05	36.74	155.74	35.96	182.11	35.52	154.87
4pm	35.36	143.13	35.44	149.30	34.94	141.51	35.22	148.40
5pm	35.72	178.00	34.43	143.69	35.26	175.94	34.19	142.74
6pm	33.57	133.01	33.96	137.40	33.12	131.41	33.68	135.31
7pm			33.33	132.68			33.03	131.53
8pm			32.97	130.91			32.67	129.74
9pm			31.61	128.32			31.33	127.23
10pm			28.89	86.32			28.58	85.44
11pm			29.89	125.05			29.66	124.10
12am			28.06	85.34			27.80	84.57
Day of week dummies								
Thursday	-0.05	-0.72	0.07	0.72	-0.04	-0.59	0.06	0.68
Friday	-0.15	-2.10	-0.30	-3.18	-0.13	-1.77	-0.30	-3.15
Saturday			-2.54	-26.52			-2.49	-26.07
Sunday			-3.19	-31.27			-3.10	-30.37
Month-of-year dummies								
January	4.35	25.36	3.34	19.24	4.08	23.84	3.27	18.89
February	3.05	21.05	3.55	22.79	3.04	21.03	3.56	22.86
April	0.17	1.47	0.20	1.58	0.22	1.90	0.20	1.61
July	1.12	8.09	1.27	8.40	1.06	7.68	1.25	8.28
September	-0.39	-3.04	-1.00	-7.05	-0.32	-2.45	-1.01	-7.13
October	-0.99	-8.43	-0.48	-3.89	-0.96	-8.22	-0.20	-4.03
Weather variables								
Max. temperature	0.06	7.66	0.04	5.42	0.06	8.74	0.04	6.07
Max. temperature × 12–4pm dummy	0.11	14.89	-0.16	-17.33	0.11	15.01	-0.16	-17.60
Daylight × 9am dummy	2.99	10.32	3.42	7.58	3.09	10.66	3.52	7.81
Daylight × 5pm dummy	-0.59	-2.07	3.20	7.07	-0.64	-2.25	3.31	7.32
Daylight × 6pm dummy	3.83	13.13	2.38	5.25	3.64	12.47	2.50	5.52
Holiday dummies								
Summer holidays	-1.45	-11.73	-1.72	-12.48	-1.37	-11.09	-1.72	-12.44
Week 53(dummy)	-1.84	-7.48	-0.96	-3.82	-1.59	-6.46	-0.78	-3.13

Table C.2: Results from estimating model (4.9) over two years.
00:00 1st April 2015–23:00 31st March 2017.

	Two-stage least squares				Ordinary least squares			
	Peak		Whole sample		Peak		Whole sample	
Variable	Coef.	*t*-stat	Coef.	*t*-stat	Coef.	*t*-stat	Coef.	*t*-stat
Price (€/MWh)	-0.17	-76.80	-0.18	-73.81	-0.17	-75.69	-0.17	-72.99
Trend	-0.01	-64.76	-0.01	-41.74	-0.01	-64.87	-0.01	-41.84
Time of day dummies, hour starting at:								
1am			29.59	110.85			29.52	110.58
2am			32.14	176.26			32.08	175.94
3am			32.02	176.72			31.96	176.41
4am			34.38	175.56			34.33	175.32
5am			34.25	174.59			34.20	174.35
6am			34.44	175.42			34.39	174.16
7am			36.00	178.14			35.94	177.84
8am			38.70	185.31			38.60	184.85
9am	36.40	159.61	38.18	191.82	36.27	129.07	38.09	191.48
10am	40.00	230.90	39.76	201.49	39.87	230.23	39.68	201.06
11am	41.36	243.24	41.26	211.63	41.25	242.60	41.17	211.22
12pm	41.00	244.09	4249	218.90	40.89	243.45	42.41	218.49
1pm	41.29	250.69	42.89	224.32	41.19	250.08	42.81	223.93
2pm	40.83	249.83	42.38	223.76	40.73	249.23	42.30	223.38
3pm	39.68	243.48	41.19	218.53	39.58	243.18	41.12	218.15
4pm	37.85	179.23	39.85	210.15	37.85	178.79	39.78	209.76
5pm	38.61	232.07	38.59	201.39	38.51	231.46	38.51	201.00
6pm	35.76	164.36	37.97	191.77	35.66	163.91	37.88	191.33
7pm			37.36	185.30			37.27	184.84
8pm			37.03	183.13			36.93	182.67
9pm			35.36	178.90			35.28	178.47
10pm			31.86	115.80			31.77	115.47
11pm			33.67	176.07			33.60	175.67
12am			30.51	113.20			30.43	112.91
Day of week dummies								
Thursday	-0.08	-1.30	0.04	0.45	-0.08	-1.25	0.04	0.46
Friday	-0.44	-6.66	-0.42	-5.34	-0.43	-6.59	-0.42	-5.32
Saturday			-2.44	-30.02			-2.43	-29.86
Sunday			-2.75	-32.29			-2.73	-32.01
Month-of-year dummies								
January	0.63	5.60	-0.35	-3.06	0.61	5.38	-0.36	-3.15
February	1.60	15.35	1.40	12.92	1.61	15.47	1.41	12.98
April	0.72	7.12	-0.46	-4.45	0.73	7.25	-0.45	-4.41
July	2.56	21.09	2.00	15.81	2.56	21.04	2.00	15.79
September	-2.24	-22.19	-1.83	-17.11	-2.25	-22.27	-1.84	-17.20
October	-1.08	-11.05	-0.50	-4.97	-1.10	-11.23	-0.52	-5.12
Weather variables								
Max. temperature	0.10	14.93	0.01	1.75	0.10	15.17	0.01	1.95
Max. temperature × 12–4pm dummy	0.11	15.68	-0.18	-23.21	0.11	15.70	-0.18	-23.28
Daylight × 9am dummy	4.42	18.50	5.40	15.03	4.43	18.52	5.42	15.09
Daylight × 5pm dummy	0.82	3.47	4.74	13.19	0.81	3.41	4.77	13.25
Daylight × 6pm dummy	4.44	18.45	4.68	13.02	4.41	18.33	4.71	13.08
Holiday dummies								
Summer holidays	-3.10	-28.36	-2.17	-18.84	-3.11	-28.47	-2.18	-18.93
Week 53(dummy)	-1.38	-6.23	0.07	0.38	-1.37	-6.17	0.10	0.49